The background of the cover features a grid of interlocking puzzle pieces. Several pieces are highlighted in color: red, orange, purple, green, and blue. In the center, a grayscale image of a young child's face is visible, with the puzzle pieces overlaid on it. The child has long hair and is looking down with a slight smile.

Robots for Autistic Children

**Understanding and Facilitating Predictability
for Engagement in Learning**

Bob R. Schadenberg

ROBOTS FOR AUTISTIC CHILDREN

UNDERSTANDING AND FACILITATING PREDICTABILITY
FOR ENGAGEMENT IN LEARNING

by

BOB R. SCHADENBERG

Graduation committee

Chairman and Secretary:

Prof. dr. J. N. Kok University of Twente

Supervisors:

prof. dr. Vanessa Evers University of Twente & Nanyang Technological University

prof. dr. Dirk K. J. Heylen University of Twente

Co-supervisor:

dr. ir. Dennis Reidsma University of Twente

Committee members:

prof. dr. Elizabeth Pellicano Macquarie University

prof. dr. Adriana Tapus ENSTA Paris

prof. dr. Mark A. Neerincx Delft Technical University & TNO

prof. dr. Gerben Westerhof University of Twente

prof. dr. Carolien Rieffe University of Twente & Leiden University

Paranymphs

Tim P. Schadenberg & dr. Jeroen M. Linssen

DSI Ph.D. Thesis Series No. 21-006 (ISSN: 2589-7721)

Digital Society Institute

P.O. Box 217, 7500 AE Enschede, the Netherlands.

The research reported in this dissertation was carried out at the Human Media Interaction group of the University of Twente.

SIKS Dissertation Series No. 2021-13

The research reported in this dissertation was carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

This work was made possible through funding from the European Union's Horizon 2020 research and innovation program under grant agreement no: 688835 (DE-ENIGMA).

ISBN: 978-90-365-5164-9

DOI: 10.3990/1.9789036551649

Typeset with $\text{\LaTeX}$. Printed by Gildeprint.

Cover design and layout design by Bob R. Schadenberg. Photo by Max Vakhtbovych, puzzle pieces by freepik.

©2021 Bob R. Schadenberg, the Netherlands. All rights reserved. No parts of this thesis may be reproduced, stored in a retrieval system or transmitted in any form or by any means without permission of the author. Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd, in enige vorm of op enige wijze, zonder voorafgaande schriftelijke toestemming van de auteur.

ROBOTS FOR AUTISTIC CHILDREN

UNDERSTANDING AND FACILITATING PREDICTABILITY
FOR ENGAGEMENT IN LEARNING

DISSERTATION

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus
prof. dr. ir. A. Veldkamp
on account of the decision of the Doctorate Board
to be publicly defended
on Friday 30 April, 2021 at 12:45 hours.

by

Bob Rinse Schadenberg
born on the 10th of May, 1987
in Assen, the Netherlands

This dissertation has been approved by:

Supervisors:

prof dr. V. Evers	University of Twente & Nanyang Technological University
prof dr. D. K. J. Heylen	University of Twente

Co-supervisor:

dr. ir. D. Reidsma	University of Twente
--------------------	----------------------

For my parents, Harm and Jacky.

Abstract

“A novel tool for delivering learning content to autistic children in a manner that keeps the children engaged, is tailored towards their strengths, and improves learning gains”; this is the promise of using robots to enhance interventions for autistic children. While robots have been found to pique the children’s interest and improve engagement in interventions, designing robots to sustain long-term engagement that leads to learning is difficult. The children are very different from each other in how autism affects the development of their cognitive, language, and intellectual ability, which needs to be taken into account for the child-robot interaction. How this can and should be done is still an open question — one that will be addressed in this dissertation.

In the first part of this dissertation, I describe our research in developing a novel robot-assisted intervention for teaching the basics of emotion recognition to autistic children. Our research was part of the EU-project DE-ENIGMA. A key research question that we address is how we can design a robot-assisted intervention to engage autistic children in learning. To address this broad research question, we report a descriptive study where autistic children interacted with our initial prototype of the “DE-ENIGMA robot-assisted intervention”. In this study, we report on how the children’s individual differences are correlated with how the children interacted within the intervention. Furthermore, we conducted a literature study to assess the user needs and user requirements that are relevant for developing robot-assisted interventions for autistic children, and describe how we developed the DE-ENIGMA intervention.

The second part of this dissertation is related to a concept that is central to autism: *predictability*. Autistic children are believed to generally favour predictable environments, and contemporary Bayesian accounts of autism place the inability to effectively deal with unpredictability at the core of the condition. Because robots are programmable, we could, in theory, program them to be highly predictable. By doing so, we could address this need for predictability and possibly improve the engagement of autistic children to the intervention. In fact, the highly predictable nature of robots is a commonly used argument for why robots may be promising tools for those working with autistic children. Predictability, however, is poorly defined in current literature and lacks an operationalisation that we can use for manipulating a robot’s predictability. We therefore provide a novel formalisation and operationalisation of predictability, and how it relates to human-robot interaction, based on the predictive processing framework. Furthermore, we report on two experimental studies where we investigated the effect of a robot’s predictability on the studies’ participants. In

one study, we specifically look at people's social perception of a robot in relation to its predictability. In the other study, we investigated the effect of a robot's predictability on the engagement of autistic children to the robot-assisted intervention.

The work described in this dissertation is a step towards better understanding the concept of predictability and its effects on human-robot interaction, as well as how we can design robot-assisted intervention for autistic children that may sustain engagement.

Samenvatting

Een nieuw stuk gereedschap dat leermateriaal kan presenteren aan autistische kinderen op een manier waardoor de kinderen betrokken blijven, aangepast aan hun sterkte punten, en wat uiteindelijk tot beter leren leidt. Dit is de belofte van het gebruik van een robot in een interventie voor autistische kinderen. Hoewel robots de interesse kunnen wekken van autistische kinderen en hun betrokkenheid bij het leren vergroten, is het nog altijd lastig om robots zo te ontwikkelen dat ze een interactie voor de langere termijn kunnen onderhouden die tevens tot leren kan leiden. De kinderen verschillen erg van elkaar in hoe hun autisme hun ontwikkeling (op gebied van cognitie, taal, en intelligentie) beïnvloedt. Dit moet meegenomen worden in het ontwerp van de robot. Hoe we dit moeten bewerkstelligen is nog een grote vraag. In dit proefschrift ga ik hier dieper op in.

In het eerste deel van dit proefschrift beschrijf ik ons onderzoek naar het ontwikkelen van een nieuwe robot-geassisteerde interventie voor autistische kinderen. Hiermee kunnen zij de basis leren van het herkennen van emoties. Ons onderzoek is deel van het Europese project genaamd DE-ENIGMA. Een van de hoofdvragen die we onderzoeken is hoe we een robot-geassisteerde interventie zo kunnen ontwikkelen waardoor de kinderen betrokken blijven bij het leren. Om antwoorden te vinden op deze brede onderzoeksvraag hebben we een beschrijvende studie uitgevoerd waarin autistische kinderen interacteerden met een vroeg prototype van de “DE-ENIGMA robot”. We rapporteren over hoe de individuele verschillen van de kinderen gecorreleerd zijn met de typen van spontane interacties van de kinderen binnen de interventie. Daarnaast hebben we een literatuurstudie uitgevoerd waarin we naar *user needs* en *user requirements* hebben gekeken die relevant zijn voor het ontwikkelen van een robot-geassisteerde interventie. Als laatste in dit eerste deel van het proefschrift beschrijf ik hoe we de DE-ENIGMA robot-geassisteerde interventie hebben ontwikkeld.

Het tweede deel van dit proefschrift gaat over een concept wat centraal staat binnen autisme, namelijk “voorspelbaarheid”. Autistische kinderen zouden over het algemeen voorkeur hebben voor voorspelbare omgevingen. Hedendaagse Bayesiaanse theorieën over autisme verklaren deze conditie aan de hand van een onvermogen om efficiënt met onvoorspelbaarheid om te gaan. Omdat robots programmeerbaar zijn, zouden we deze ook zo kunnen programmeren dat ze erg voorspelbaar zijn. Hierdoor zouden we aan de behoefte van autistische kinderen kunnen voldoen naar voorspelbare omgevingen. En mogelijk zijn de kinderen daardoor dan ook meer betrokken bij het leren in een robot-geassisteerde interventie. De voorspelbaarheid van robots is dan ook een argument dat vaak wordt gebruikt om aan te geven dat robots zo veelbe-

lovend zijn als het om toepassingen gaat voor autistische kinderen. Echter, het concept voorspelbaarheid is slecht gedefinieerd in de huidige literatuur. Tevens ontbreekt er een operationalisering van dit concept die toepasbaar is zodat de voorspelbaarheid van een robot te manipuleren en programmeren is. Om dit probleem op te lossen hebben wij een nieuwe formalisering en operationalisatie van voorspelbaarheid bedacht, die betrekking heeft op het vakgebied van mens-robot interactie, en gebaseerd is op het *predictive processing* raamwerk. Met deze formalisatie en operationalisatie hebben wij twee experimentele studies gedraaid waarin we naar het effect van de voorspelbaarheid van een robot hebben gekeken op de onderzoeksdeelnemers. In één studie hebben we specifiek gekeken naar hoe de voorspelbaarheid van de robot de sociale perceptie van mensen beïnvloedt. In de andere studie hebben we onderzocht wat het effect is van de voorspelbaarheid van de robot op de betrokkenheid van autistische kinderen in het leren binnen een robot-geassisteerde interventie.

Het werk dat ik in dit proefschrift beschrijf brengt ons een stap dichterbij het begrijpen van het concept van voorspelbaarheid en de effecten hiervan op de mens-robot interactie. Tevens begrijpen we door ons werk nu beter hoe we robot-geassisteerde interventies kunnen ontwerpen die de betrokkenheid van autistische kinderen vast kan houden.

Acknowledgements

The journey of PhD started after finishing my master's. Not because I managed to get a coveted PhD position, but because I said to myself that I would never do a PhD. While conducting research was a passion of mine, writing up the results was not. The writing of my master's thesis turned out to be a demotivating experience. Day in day out, the only task you had was writing up your results. There was no variation in what I needed to work on, no meetings, there was just writing. I would be damned if I considered tormenting myself for (at least) another four years when doing a PhD. Hence, I took a job as a researcher in industry. There I talked with various colleagues who had finished a PhD and actually loved it. Turns out I knew nothing of what it was to be a PhD student. Maybe I was too quick to judge? Later that year, after finishing an assessment for a job interview, the psychologists, Nico, asked me out of nowhere whether I had ever thought of pursuing a PhD. What made him ask this question¹? After all, it was totally unrelated to the job I was applying for. He was right though. At that time, I was wondering whether academia was the right place for me to be. Now that my PhD journey is coming to an end, I can say that it is all I could ever wish for. And it were these small nudges from my former colleagues (Marion, Pieter, and Wilco) and Nico which created a butterfly effect that eventually led me to the point I am now at. For this, I am grateful.

While this dissertation has my name on it, the research I describe is the product of various collaborations and was made possible by all the people who supported me over the years. First of all, I would like to thank my supervisors. Dennis, as my daily supervisor, we worked closely together and you were instrumental to this dissertation. You made time to discuss our work in depth, even when you had no time to spare. You taught me how to do proper science, how to write articles that people could understand, how to “kill your darlings”, and to live up to the values and responsibilities of being a researcher. Your enthusiasm and endless positivity kept me going, even when I was inclined to let my negative valence get the better of me. You have been an incredible mentor to me, for which I am eternally grateful. I strive to pass down all you have taught me to those I may supervise or coach myself one day. Vanessa, as my promotor, you have steered me away from the cliffs more than once. In the first year, you urged me to find a worthy research question. One that must be answered. I thought I had (several times), but you gently disagreed. Looking back at those research questions, you were right, and prevented me from pursuing less important

¹I did ask him after accepting a different job offer. Turns out, Nico asked me this questions because my personality test showed that I had a strong focus on analysing and conceptualising.

questions. Overall, you were a great help in guiding me on this PhD journey. You also allowed me the freedom to investigate my own research interests throughout my PhD, which I appreciate very much. I was lucky to have not one, but two promoters. Dirk, thank you for believing in me, giving me the opportunity to start the PhD at the Human Media Interaction group, and for supervising me. Your philosophical remarks helped me think beyond the status quo. You also helped me deal with the struggles of my project. Thank you for all of this.

Conducting research with autistic children can be challenging. I was fortunate that I was part of a large research consortium, funded by the European Union, who enabled my research. Our partners at the University College London and the Serbian Society for Autism had good connections with various special education schools, and knew many autistic children whom we could approach to participate in our studies. Their efforts enabled the work described in this dissertation. I am grateful to all the children, their families, the schools that hosted our studies, and the school staff members, who all generously gave their time for our studies. Liz, you have greatly enhanced my understanding of autism and autistic individuals. Your theory on autism are central to this dissertation. You also taught me how we can, and should, listen to autistic individuals and to be mindful of their views on autism and society. It still saddens me that I did not manage to visit your lab in London or Australia as a visiting scholar. Lynn en Bridgette, thank you for managing the DE-ENIGMA project, as well as listening to me complaining and offering your support. Furthermore, I would like to thank both of you, as well as Betsy, for the proofreading of my articles. Pauline, Daniel, Jamy, and Vicky, thanks for all your great work on the DE-ENIGMA project. It was a great pleasure to work with you.

I would like to thank all my committee members for taking the time to review this dissertation. I am honoured to have you as my reviewers.

When you start your PhD at a university, you become part of a research group. For me, this was the Human-Media Interaction (HMI) group; an incredibly welcoming group of people. Jeroen, my former colleague, dear friend, and paranymp, you had a large part in this. Because of your warm and cheerful personality, and our shared interests, we instantly got along with each other. It was not long before we went to concerts (metal matters), bouldering, gaming, and sharing dinners. Let's continue doing so in the future! Jan, thank you for being my roommate, putting up with me, all your puns, and our discussions on statistics. You are not average, nor are you mean! Charlotte and Alice, you are the foundation of HMI. Thank you for all your work to keep the department running. Khiết, thank you for being the voice of our little robot. And of course, to all my former and current colleagues at HMI whom I have not yet mentioned, thank you for providing such a cosy and warm environment to work in. I have learned lots from our many discussions. Whether these were about work, or other topics.

Doing a PhD is not without its hardships. Facing harsh reviews, realising they are right, going back to a paper for the seventh time, dealing with the politics of academia, ... (the list goes on). One way of dealing with all this was going hiking with friends. These long-distance hikes have the tendency to put life back in perspective. Suddenly, you do not have to worry about what you need to do on a certain day; you only need

to walk from A to B in one piece. The latter was sometimes challenging, as my feet have the tendency to disintegrate after day one. Nevertheless, your company keeps me going. Whether it be hiking, sipping a glass of beer, gaming, or other activities, it is always “gezellig” when we spend time together.

And last, but certainly not least: my family. Paps en mams, thanks you for everything. I am truly blessed to have you as parents. At an early age, you always pushed me to ask questions and to try out new things, which sparked by curiosity. Our long-distance hikes taught me to persevere and endure hardships. That doing so can get you to places that are otherwise not accessible (both literally and figuratively), and that one should focus on the positive aspects of such experiences. These are the qualities that enabled me to pursue this PhD. The way you raised me shaped me in the person I am today, for which words fail to express my gratitude. I dedicate this dissertation to you. I am immensely proud of you, and hope you are proud of me as well. Tim, my dear brother and paranymp, you are as defining for my development as our parents. As my big brother, you are showing me the way in life, and you are always there for me. Your talk about innovation during our early twenties also got me interested in it. This eventually resulted in my interest in robots. Thank you for paving the way for me. Susanne, thank you for your advice on the design of this dissertation’s cover. And finally, Dorien, my loving girlfriend. You saved me from the confines of Hengelo, which surely made the end of my PhD more bearable. Thank you for sticking with me during my PhD, and for supporting me. Thank you for being at my side.

Bob Schadenberg
Looking out over Arnhem, on the 16th of April, 2021

Contents

1	Introduction	1
1.1	Introduction	1
1.2	Autism Spectrum Condition	2
1.3	Research context and scope	4
1.4	DE-ENIGMA database	5
1.5	Research questions and structure of the dissertation	6
1.6	Contributions	8
I	BACKGROUND	11
2	Autistic children interacting with robots	13
2.1	Robots and interventions for autistic children	14
2.2	Engagement definitions and measures	16
2.3	The interactions of autistic children with robots and objects	19
2.4	The idiosyncrasies of autistic children	22
2.5	Conclusion	23
II	TOWARDS AN INTERVENTION	27
3	User needs and requirements for robot-assisted interventions	29
3.1	Introduction	29
3.2	Materials and methods	32
3.3	User needs and discussion thereof	34
3.4	User requirements and discussion thereof	38
3.5	General discussion and conclusion	42
4	Spontaneous interaction of autistic children and the role of autistic traits	47
4.1	Introduction	47
4.2	Materials and methods	49
4.3	Results	57
4.4	Discussion	64
4.5	Conclusion	70

5	The evolution of DE-ENIGMA intervention	73
5.1	Introduction	73
5.2	Exploratory study, v1	74
5.3	Exploratory study, v2	78
5.4	From v3 to the final version of the DE-ENIGMA intervention	85
5.5	The DE-ENIGMA intervention	86
III	PREDICTABILITY AND ITS EFFECTS ON HUMAN-ROBOT INTERACTION	95
6	Autism, robots, and the role of predictability	97
6.1	Introduction	97
6.2	Autism, predictability, and individual differences	99
6.3	The role of predicting in human perception and implications for robot predictability	101
6.4	Conceptualising predictability within the context of HRI	104
6.5	Operationalisation of behavioural predictability as variance in a robot .	107
6.6	Reducible and irreducible unpredictability in robot behaviour	108
6.7	Conclusion	110
7	Measuring robot predictability	113
7.1	Measuring the predictability of a robot	113
7.2	Materials and methods	114
7.3	Results	115
7.4	Final multi-item scale for attributed unpredictability and its limitations	116
8	Understanding people's perception of a robot and its predictability	119
8.1	Introduction	119
8.2	Research questions and hypotheses	121
8.3	Materials and methods	124
8.4	Results	128
8.5	Discussion	132
8.6	Conclusion	137
9	Robot predictability and the engagement of autistic children	141
9.1	Introduction	141
9.2	Research questions and hypotheses	143
9.3	Materials and methods	144
9.4	Results	157
9.5	Discussion	165
9.6	Conclusion	168
IV	DELIBERATION	171
10	Discussion and conclusion	173
10.1	Main findings and contributions	173

10.2 Reflection and future work	177
10.3 Closing remarks	181
About the author	183
Bibliography	187
SIKS dissertation list	213

1

Introduction

1.1 Introduction

HUMAN beings have evolved as creatures who can make and use simple tools to tools that can act on their own accord and behave in sometimes humanlike fashion. This dissertation relates to such tools, namely to *robots*. Specifically at robots which are primarily designed to work in close collaboration with people. These are the type of robots we often see in science-fiction media. Robots such as R2D2 and BB8 (Star Wars), Baymax (Big Hero 6), Number Six (Battlestar Galactica), HAL (2001: A Space Odyssey), or Marvin (The Hitchhiker's Guide to the Galaxy). They have captured our imagination for ages, and their potential to transform the world (for good or bad, depending on the media) is immense. Video demonstrations of current robots do a good job at showing that they already live up to their potential, showing incredible feats in their interaction with people, making it easy to believe that science fiction future is already here today. The reality, however, is often less auspicious. Developing robots that can meaningfully interact socially with people for longer periods of time is hard. In the demonstrations, such robots are therefore often operated by a human controller, who creates the illusion that the robot can hold its own while interacting with people. One of the reasons why it is hard to develop a robot that can hold its own in social interactions with people is because robots often fail to meet our expectations. Initial expectations set by robots we know from science fiction ([Kriz et al., 2010](#)), but also expectations that stem from seeing a robot and interacting with it. If a robot can speak, it does not mean it can also hear and understand you when you talk to it, even though we would expect that it does. For most use-cases that revolve around robots that add value through social interaction, the technological requirements for such robots are very demanding, and are often too demanding ([Fresh Consulting, 2020](#)).

This dissertation relates to a promising use-case for robots that interact with other people, but for which the technological requirements are possibly much lower for it to add value. That is, robots which are designed for autistic children. For these children, the limited capabilities of robots may actually be beneficial, rather than

to the detriment of the robot. This statement is based on the idea that simple and highly predictable interactions may make it easier for autistic children to focus on the content, rather than being distracted or overloaded by too much sensory information. Thus, using robots for autistic children is a use-case where very simple interactions may be of great value, and thus hold great potential for becoming reality in the near-future.

Work on robots for autistic children started with the turtle robot (designed by BBN engineer Paul Wexelblat), equipped with a light, horn, and a pen, that facilitated a Logo learning environment in which Piagetian learning can occur and is supported (Papert, 1973). This robot-assisted learning programme was introduced to an autistic child by Emanuel and Weir (1976). The child could control the robot through a button box, where each button corresponded to a command (e.g. move forward, backward, left, right, hoot) (Perlman, 1974). Emanuel and Weir (1976) concluded:

"In the process of acting-out, David [the child] seems to be both telling himself and telling us what he understood — his monologue trails into dialogue spontaneously. Overt non-verbal and verbal social gestures and an increasing willingness to commit himself followed from the reality of his being a free agent of his own actions and learning, and of the self validating effect of understanding and being understood. The Logo environment served as a catalyst in developing our relationship with David precisely because he was able to actively control and understand an object of common interest, the turtle." [p. 128]

Much has changed since 1976, but the idea that a robot can elicit interaction between an autistic child and an adult remains relevant today.

This dissertation is dedicated to further our understanding of developing effective robot-assisted interventions for autistic children. Specifically, I will focus on understanding the concept of predictability in the context of human-robot interaction (HRI); what is this concept about, how to measure it, and what are its effects on (autistic) HRI.

1.2 Autism Spectrum Condition

In this dissertation, we focus specifically on children with an Autism Spectrum Condition (hereafter referred to as “autism”) and how we may promote their engagement in robot-assisted interventions to sustain long-term engagement. Autism is a lifelong neurodevelopmental condition that affects the way an individual interacts with others and experiences the world around them. Throughout this dissertation, we will refer to individuals with autism as autistic individuals², compared to non-autistic, or typically developing individuals. According to the DSM-V (American Psychiatric Association, 2013), autism is a spectrum condition, which means that while there is wide variation

²Work by Gernsbacher (2017) suggests that person-first language may be stigmatising, and autistic adults prefer the use of disability-first terms, rather than person-first terms because they feel that being autistic is central to their identity (Kenny et al., 2016). We will therefore refer to these individuals/children throughout this dissertations as *autistic individuals/children* (or just as individuals/children).

in the type and severity of symptoms autistic individuals experience, these symptoms are believed to result from the same underlying mechanism. Diagnostic criteria for autism, defined in the DSM-V, include two core features, namely (a) difficulties in social interaction and communication (so-called “social features”), and (b) the presence of rigid and repetitive patterns of behaviours and limited personal interests (so-called “non-social features”). Both sets of features must be present from early in development and cause significant difficulties to the individual and/or to the people around the individual (American Psychiatric Association, 2013). Current prevalence estimates of autism found that around 1 in every 100 individuals is on the autism spectrum (Brugha et al., 2011; Elsabbagh et al., 2012), many of whom struggle to find and retain employment, to live independently, and to sustain friendships and intimate relationships (Howlin et al., 2004). For Europe, this means approximately 7 million autistic individuals. If you include their families, autism is a part of daily life for more than 24 million individuals.

Autism was first termed described by Leo Kanner in 1943, based on his observations of eleven children (Kanner, 1943). Since then, there have been numerous attempts to explain the behavioural features of autism. The most prominent attempts have included the theory of mind hypothesis (Baron-Cohen et al., 1985), the executive dysfunction hypothesis (Ozonoff et al., 1991; Hill, 2004) and weak central coherence theory (Frith and Happé, 1994; Happé and Frith, 2006). More recent accounts of autism, include the empathising-systemising account (Baron-Cohen, 2002, 2009) and Bayesian accounts of autism (Pellicano and Burr, 2012; Van de Cruys et al., 2014; Lawson et al., 2014; Sinha et al., 2014), which we describe briefly in turn.

According to Baron-Cohen (2002), the core features of autism can be explained as individual variation on the dimensions of *empathising* and *systemising*. In this, empathising refers to the drive to identify other people’s emotions and thoughts (cognitive empathy) and the appropriate emotional response (affective empathy). Systemising is the tendency to analyse and explore a system and extract underlying rules that govern its behaviour. In this, a system can be any kind of system that follows specific rules, whether mechanical (e.g. a robot), abstract, or any other type. According to the empathising-systemising theory, the social features of autism can be explained by having difficulties with empathising, while the non-social features can be explained as a high tendency to systemise. This could also explain why many autistic children are drawn to technology such as a robot.

The Bayesian accounts of autism attempt to identify the mechanisms underlying the difficulties autistic individuals face within a Bayesian computational model of perceptual inference (Pellicano and Burr, 2012; Van de Cruys et al., 2014; Lawson et al., 2014; Sinha et al., 2014). These Bayesian models of perception include predictive coding/predictive processing and other generative models (Rao and Ballard, 1999; Bar, 2007; Friston, 2010; Clark, 2013; Hohwy, 2013), which all assume that perception is an optimised combination of external sensory data (the likelihood) and an internal model on what sensory information is expected (the prior). For autistic individuals, the process of matching sensory data with the internal model is assumed to be atypical by these Bayesian accounts of autism. As a result, autistic individuals have difficulty predicting sensory information. Social environments are highly un-

predictable, and therefore rely more strongly on the internal model for generating predictions, which can explain the social features of autism. In turn, the non-features are the result of trying to maintain a highly predictable environment. We will go more in depth into these Bayesian accounts of autism in Chapter 6, as they play a central role in explaining how we predict a robot's behaviour.

1.3 Research context and scope

The work presented in this dissertation is part of the DE-ENIGMA project³, which was funded by the European Union's Horizon 2020 programme. The DE-ENIGMA consortium was made up of five scientific institutions (University of Twente, University College London, Imperial College London, Institute of Mathematics of the Romanian Academy, and the University of Passau, for whom the members later moved to the University of Augsburg), two associations (Serbian Society of Autism, and Autism-Europe), and one industrial company (IDMind). Each partner took responsibility for one aspect in the development of a novel robot-assisted intervention. The University of Twente was primarily tasked with creating an understanding of the context of use for the DE-ENIGMA system and the design and development of the DE-ENIGMA intervention. The research reported in this dissertation is therefore related to this task.

The aim of the DE-ENIGMA project was to develop a novel intervention for teaching emotion recognition to autistic children with the help of a humanoid robot. Autistic children often have difficulties recognising, interpreting, and producing facial expressions and emotions, across modalities (face, voice, body) (for a review, see [Uljarević and Hamilton, 2013](#)). Recognising emotions is central to success in social interaction ([Halberstadt et al., 2001](#)), and due to impairment in this skill, autistic individuals may fail to accurately interpret the dynamics of social interaction. Learning to recognise the emotions of others may provide a toehold for the development of more advanced emotion skills ([Strand et al., 2016](#)), and ultimately improve social competence ([Denham et al., 2003](#)).

The target audience for the DE-ENIGMA robot-assisted intervention are autistic children who are “less cognitively able”, as research for this group of autistic children is relatively limited — most research focuses on more cognitively able autistic children ([Tager-Flusberg and Kasari, 2013](#)). Many of the “less cognitively able” children have limited receptive language and lower intellectual ability, and often require much higher levels of support from specialist teaching than regular, mainstream schools can typically provide. For the design of the autistic child-robot interaction, the needs of these children bring additional challenges to providing interactions where those individuals involved and the robot understand each other. Some of these challenges are addressed in this dissertation, such as providing non-verbal ways to interact with the robot.

The robot that was used in the DE-ENIGMA project is RoboKind's R25 humanoid robot called “Zeno” or “Milo” (see Figure 1.1). This robot was used in all studies reported in this dissertation. The main feature of this robot is its expressive face,

³See <http://de-enigma.eu/>.

Figure 1.1: Robokind's R25 humanoid robot “Zeno” used in the DE-ENIGMA project.

which can be used to display facial expressions of emotion. It has five degrees of freedom in its face and two in its neck, which allowed us to design facial expressions of emotion for expressing joy, sadness, anger, fear, surprise, and disgust (Schadenberg et al., 2018). For the DE-ENIGMA intervention to be meaningful, learning to recognise the facial expressions of Zeno will need to generalise to humans. The expressive face of Zeno resembles that of a human (although the proportions of the face are off to make it slightly cartoonish). This made the robot particularly suitable for the DE-ENIGMA intervention, as gap between Zeno’s facial expressions and human facial expressions is much lower than when using robots that are less humanlike.

1.4 DE-ENIGMA database

One of the major outcomes of the DE-ENIGMA project is the development of the DE-ENIGMA database⁴ — a publicly available multi-modal database of autistic children’s interactions in the setting of a robot-assisted intervention. The database allows researchers to train autism-specific algorithms, based on the audio, video, and 3D video recordings that we collected, as well as to conduct behavioural analyses on the children’s behaviour. To develop the DE-ENIGMA database, we (the DE-ENIGMA consortium) conducted a large data collection study in the first year of the project. The database includes 121 autistic children aged 5 to 12, from Serbia ($n = 59$) and the United Kingdom ($n = 62$). The recordings in the United Kingdom were held at three separate special education schools, where some of the children of those schools participated. For the Serbian recordings, the Serbian Society of Autism — our partner in the DE-ENIGMA project — invited parents of autistic children from their network to participate in the study, which was held in a rented apartment.

For the data collection study, we developed an initial robot-assisted activity, which was based on the teaching programme developed by Howlin et al. (1999). This teach-

⁴See <https://deenigmadb.wordpress.com/>.

ing programme focuses on teaching perception, expression, understanding, and social imagination related to the affective states happiness, sadness, anger, and fear. We adapted the teaching programme to include a robot, which would perform the facial expressions of emotions. Each child was randomly assigned to either the robot-assisted or adult-assisted activities. For the latter, the activity was simply to engage in the [Howlin et al. \(1999\)](#) teaching programme.

1.5 Research questions and structure of the dissertation

The initial goal of the DE-ENIGMA project was to develop a robot-assisted intervention for teaching emotion recognition to autistic children. For any robot-assisted intervention that is to lead to learning, it is essential that the children engage and stay engaged with the tasks in the intervention. This is because engagement is a necessary prerequisite for learning ([McCormick et al., 1998](#)), where higher engagement results in more opportunities for cognitive and social skill learning ([Greenwood, 1991](#); [Fredricks et al., 2004](#)). Developing a robot that can sustain long-term engagement is difficult even for typically developing children (e.g. [Kanda et al., 2004](#); [Leite et al., 2013](#)), and autistic children have additional (support) needs and desires that need to be addressed before they can engage within the intervention. For instance, autistic children can have cognitive difficulties that may limit their language understanding or intellectual ability. They might also be preoccupied with certain senses that then dictate their behaviour, or are more easily overloaded by sensory information compared to typically developing children. These examples can all prevent an autistic child from engaging in the intervention, or disrupt an ongoing interaction. To make matters more difficult, autism is a spectrum condition. Their atypicalities in cognitive, emotional, behavioural and social functioning caused by autism therefore manifest very differently in both quality and quantity ([Happé et al., 2006](#)). The first research question that I will address is therefore:

Research question 1: How can a robot-assisted intervention be designed to engage autistic children in learning?

To answer this broad research question, I will describe the concept of engagement in the context of within Human-Robot Interaction (HRI) in Chapter 2, as this concept can be (mis)interpreted in various ways ([Azevedo, 2015](#)). In Chapter 3, I report on a literature analysis that we conducted in order to assess what was already known for developing effective robot-assisted interventions. This resulted in a list of general user requirements for robot-assisted intervention. In this chapter, I also report on two studies that we conducted to find current needs of educators and autism professionals. In Chapter 4, I report on our study where we explored how autistic children interacted within the initial prototype of the DE-ENIGMA robot-assisted intervention. In this study, we also look into how differences in the interaction are correlated with autism-specific traits (e.g. autism severity scores), and investigate how autistic children spontaneously interact with objects and robots. Lastly, we describe the development of the DE-ENIGMA intervention prototypes in Chapter 5, as well as how each prototype was tested.

In the design of our final version of the DE-ENIGMA intervention, one concept stood central for facilitating engagement, namely the concept of *predictability*. In our description of autism in Section 1.2, we briefly touched on the Bayesian accounts of autism that explain the condition through atypicalities in the mechanism to generate and/or use predictions, which results in a desire for predictability. In the context of social skill learning, experiencing discomfort due to dealing with unpredictability is problematic, because it could prevent children from being in a mental state where learning can occur. For instance, dealing with unpredictability can cause anxiety (Paulus and Stein, 2006). In current educational practices, predictability is therefore accentuated at schools (e.g. through the TEACCH approach (Mesibov and Shea, 2010)), so that autistic children know what to expect during the day, increasing their engagement in learning (MacDuff et al., 1993; Bryan and Gast, 2000; O'Reilly et al., 2005). Incorporating a robot in social skill learning can be helpful in that it can provide a highly predictable manner of learning social skills, as we can systematically control the predictability of the robot's behaviour (as we will see later in this dissertation). Indeed, the predictability of a robot is a commonly used argument for why robots may be promising tools for autism professionals working with autistic children (e.g. Dautenhahn and Werry, 2004; Duquette et al., 2008; Thill et al., 2012; Huskens et al., 2013; Sartorato et al., 2017; David et al., 2020). While the concept of predictability is often mentioned in literature on autism and robots for autism (or HRI for that matter), a conceptualisation of predictability is missing. This prevents us from effectively taking predictability into account when developing robots for autistic children. This leads to our second research question:

Research question 2: What is predictability in relation to people interacting with robots?

To this end, we look at robot predictability from the perspective of predictive processing in Chapter 6. Predictive processing is a framework from cognitive neuroscience, which provides an account of how the brain makes sense of sensory information. In this chapter, we make the distinction people's ability to predict robot behaviour and people's attribution of predictability to a robot. For the latter, we developed a new measurement scale, which we report on in Chapter 7.

With a better understanding of what robot predictability is, we then explored whether this concept is actually relevant for HRI, and in specific for autistic children interacting with robots. For autistic children, a highly predictable robot may be beneficial in facilitating engagement (for reasons stated above), but there is currently no experimental evidence that this is actually the case. The third and final research question that we address in this dissertation is therefore:

Research question 3: How does robot predictability influence human-robot interaction?

To address this research question, we conducted two experimental studies and investigated how a robot's predictability influences people's perception of that robot, and how it influences the engagement of autistic children. The first of those two studies was a video-HRI study with non-autistic individuals to assess how a certain

operationalisation of robot predictability influences people's social perception of the robot. I report on this study in Chapter 8. The reason for testing with non-autistic individuals is that our measures for robot predictability were too complex to be used with autistic children, many of whom have high support needs and limited language understanding. In Chapter 9, I report on a joint-study conducted by the DE-ENIGMA consortium. At this time, the aim of the project had shifted from evaluating the DE-ENIGMA intervention in a randomised control trial study to assessing the effect of the robot's predictability on autistic children. In our analysis, we investigate how robot predictability influences the (two types of) engagement of autistic children to the DE-ENIGMA intervention.

To conclude the dissertation, we discuss in Chapter 10 how our work addressed the three overarching research questions outlined above, and the implications of our work. In this chapter, we draw our final conclusions on engaging autistic children in a robot-assisted intervention and the role of robot predictability therein.

1.6 Contributions

This dissertation provides the following contributions:

An initial list of user requirements for developing robot-assisted interventions (Chapter 3). To develop a novel robot-assisted intervention, we have identified a number of user requirements based on a literature search and two studies where we interview our users on what they find important requirements. The list that we present provides an initial set of user requirements that can be used for a user-centered development of robot-assisted interventions.

New insights into how autistic children interact in robot-assisted settings (Chapters 4 and 5). The studies presented in this dissertation resulted in new insights in how autistic children interact with robots as well as interact with an adult in the presence of a robot.

- *Spontaneous interactions of autistic children within a robot-assisted intervention.* While there are several reports that describe how autistic children interact within robot-assisted settings, these are limited to qualitative reports with small sample sizes. Translating insights from these reports to design is difficult due to the large individual differences among autistic children in their needs, interests, and abilities. To address these issues, we conducted a descriptive study and report on quantitative and qualitative analyses of how 31 autistic children spontaneously interacted with a humanoid robot and an adult within the context of a robot-assisted intervention, as well as which autistic traits were associated with the observed interactions.
- *Facilitating communication with a robot through tangibles.* Autistic children will need to be able to communicate with a robot in a way that both parties can understand each other. This can be difficult, as autistic children may have limited language use and understanding. Through a number of small exploratory

studies, we investigated how we could facilitate engagement with learning in a robot-assisted intervention through the use of tangibles.

Formalisation and operationalisation of the concept of predictability as it relates to HRI (Chapter 6). Our current conceptual understanding of the concept predictability in the context of HRI, is too limited for understanding how robots can facilitate predictability and how this shapes people's perception of robots. This limits us in effectively taking predictability into account in the design of robot behaviour. We therefore provide a novel operationalisation and formalisation of predictability, as it relates to HRI, based on the predictive processing account of human cognition. Our operationalisation and formalisation now allows us to study the robot predictability, and investigate its effects on people's interaction with robots.

A new scale for measuring the unpredictability that is attributed to a robot by a person (Chapter 7). Previous studies measured to what extent people attribute predictability to a robot as an attribute either through a single item, or multiple items specific to predictability of robot motion. Single-item measures are more vulnerable to random measurement errors and unknown biases in the meaning and interpretation of that item. With multiple-item scales, the random measurement error is more likely to be cancelled out. Moreover, they cover a broader range of meanings of a construct, which can reduce the effect of differently interpreting an item. We therefore developed a new multi-item scale that measures to what extent unpredictability is attributed to a robot, and which is not restricted to the predictability of robot motions.

New insights into how a robot's predictability influences people (Chapters 8 and 9). Based on our operationalisation and formalisation of robot predictability, we conducted two experimental studies to assess how this concept influences HRI.

- *People's social perception of a robot in relation to its predictability.* A degree of unpredictability in robot behaviour may be desirable for social interactions in facilitating engagement and increasing the attribution of mental states to the robot, but can also negatively affect the social perception of the robot. We carried out a video HRI study where we manipulated the robot's predictability, and measured people's social perception of the robot, their ability to predict the robot's behaviour, and to what extent they attributed predictability as an attribute of the robot.
- *Autistic children's engagement in a robot-assisted intervention in relation to a robot's predictability.* The effectiveness of robot-assisted interventions designed for social skill learning presumably depends on the interplay between robot predictability, engagement in learning, and the individual differences between different autistic children. To better understand this interplay, we report on an experimental study where 27 autistic children participated in the DE-ENIGMA robot-assisted intervention. We manipulated the robot's predictability and measured the children's engagement, visual attention, as well as several individual factors.

Through these contributions, we can formulate an answer to the research questions posed in the previous section.

Part I

BACKGROUND

“The curse of climbing is discovering how great the distance yet to climb.”

Steven Erikson, Gardens of the Moon

2

Autistic children interacting with robots

This chapter is based on the literature sections of following articles:

Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., & Evers, V. (2020). Differences in Spontaneous Interactions of Autistic Children in an Interaction With an Adult and Humanoid Robot. *Frontiers in Robotics and AI*, 7, 19 pages. ([Schadenberg et al., 2020b](#))

Schadenberg, B. R., Reidsma, D., Evers, V., Li, J. J., Davison, D. P., Heylen, D. K. J., Neves, C., Alvito, P., Shen, J., Pantić, M., Cummings, N., Schuller, B. W., Olaru, V., Sminchisescu, C., Babović Dimitrijević, S., Petrović, S., Williams, A., Alcorn, A. M., & Pellicano, E. (under revision). Predictable Robots for Autistic Children — Variance in Robot Behaviour, Idiosyncrasies in Autistic Children's Characteristics, and Child-Robot Engagement. *ACM Transactions on Computer-Human Interaction*. 39 pages. ([Schadenberg et al., under revision](#))

In this chapter, we will give a theoretical background on engaging autistic children in child-robot interactions. To do so, we will start by explaining the use-case of using robotic technology for autistic children (Section 2.1). Why do we bother these children with a robot? What central need of them could a robot hope to solve? Well, a recurring pattern in earlier literature on robots for autistic children is that the children often showed increases in engagement when a robot was incorporated in the interaction, compared to the regular adult-child interactions without a robot involved. As the concept of engagement plays a central role in this dissertation, we then provide a brief explication of the concept of engagement as it relates to learning (Section 2.2). Next, we discuss what the engagement of autistic children with robots looks like, and what (dis)similarities there are with how they interact with inanimate objects (Section 2.3). After all, there is no evidence that autistic children will always perceive robots as social agents, and they may simply treat them as an inanimate object. As autistic children are known to be a highly

heterogeneous group of children, we then discuss several key idiosyncrasies that may influence how autistic children engage within a robot-assisted activity (Section 2.4). We end this chapter with a conclusion on the previously discussed sections (Section 2.5).

2.1 Robots and interventions for autistic children

To help autistic individuals lead lives of their own choosing, various interventions have been developed that try to teach certain social, behavioural, or cognitive skills. These include Applied Behavior Analysis (Baer et al., 1968) interventions, social skills training (McConnell, 2002), occupational therapy (Case-Smith and Arbesman, 2008), physical therapy (Srinivasan et al., 2014), and sensory integration therapy (Lang et al., 2012). To improve the efficiency and effectiveness of such interventions, they have been enhanced with Information Communication Technology (ICT) (Boucenna et al., 2014; Grynszpan et al., 2014), such as robots, interactive environments, computers, or touch screens (Boucenna et al., 2014). Autistic individuals are reported to often have affinity with ICT (Bernard-Opitz et al., 2001; Silver and Oakes, 2001), and have been found to prefer interacting with media over other play activities (Shane and Albert, 2008).

ICT can also provide interactions that are specifically designed to address the needs of autistic children, potentially creating more understandable and engaging interventions. The autistic children's need for sameness can be addressed by designing ICT to be highly predictable. For robots, this argument is commonly used to explain why robotic technology in specific may be promising tools for autism professionals working with autistic children (e.g. Dautenhahn and Werry, 2004; Duquette et al., 2008; Thill et al., 2012; Huskens et al., 2013; Sartorato et al., 2017; David et al., 2020). Furthermore, ICT is free of social demands that are often challenging for autistic individuals in human-human interactions.

Another factor that can be taken into account when designing ICT are the strengths of autistic children. For instance, they often have relatively strong visual processing skills (Shah and Frith, 1983), and show a strength in understanding the physical world, compared to understanding the human social world (Klin et al., 2009). This could be taken into account in the interaction design, for example, by reducing distracting stimuli to increase attention, and use visually cued instructions. All-in-all, ICT can be more easy to understand and more motivating to autistic children, improving their engagement in the ICT-enhanced intervention. As engagement is considered to be a necessary prerequisite for learning (McCormick et al., 1998), where higher engagement results in more opportunities to gain knowledge for cognitive and social development (Greenwood, 1991; Fredricks et al., 2004), improving the child's engagement within the intervention can improve its efficiency.

In a similar vein, robotic technology has been used to promote engagement of autistic children in interventions. Generally speaking, incorporating a robot in an intervention for autistic children appears to have a positive effect on the child's engagement and attention to the learning task (Scassellati et al., 2012; Simut et al., 2016), as shown by increases in positive affect (Costescu et al., 2015; Kim et al., 2015), communication (Kim et al., 2013; Wainer et al., 2014), and in attention (Dautenhahn and

Werry, 2004; Tapus et al., 2012; Simut et al., 2016). Importantly, the engagement that is observed is often social in nature and is directed not only at the robot, but also at other people near the robot (Robins et al., 2005; Duquette et al., 2008; Feil-Seifer and Matarić, 2009; Kozima et al., 2009; Kim et al., 2013). The latter is significant, because from a pedagogical point of view, it does not necessarily matter whether the child interacts with the robot or whether the robot elicits interaction between the child and adult, as learning can occur in either case.

Next to having a positive effect on engagement, robots are also thought to be less complex in terms of perceptual processing, where a robot's behaviour does not have the richness of social cues of human behaviour (Duquette et al., 2008; Sartorato et al., 2017). Robots could also deliver 'on demand' social skill learning, and provide quantified metrics of the child that can be used by an adult to further tailor the learning content to the child (Scassellati, 2007). Furthermore, some studies suggest that robots may be more engaging when compared to virtual agents. The robot's presence has been found to elicit more speech (Kim et al., 2013), and results in more social initiations of the child (Pop et al., 2013). As such, robots have been used to enhance interventions aimed at teaching various key behavioural skills, such as imitation (Tapus et al., 2012; Warren et al., 2015a), joint attention (Bekele et al., 2014; Warren et al., 2015b; David et al., 2018; Zheng et al., 2020), collaborative play (Wainer et al., 2014), turn-taking (David et al., 2020; Kostrubiec and Kruck, 2020), perspective taking (Scassellati et al., 2018; Wood et al., 2019), and emotion recognition and expression (Chevalier et al., 2017a; Scassellati et al., 2018). Learning such skills takes time, and as such, robots need to be capable of long-term interactions.

Scassellati et al. (2018) deployed a home-based robot-assisted intervention for one month, and showed that the robot was used for around 23 sessions of nearly 30 minutes long. In this intervention, the learning content was displayed on a touch-screen monitor in the form of games the child could play. The robot operated fully autonomously, where its role was to give feedback and keep the child motivated. The results showed positive effects on learning, in particular that the learned skills seemed to generalise to interactions with other people. However, the experimental design did not allow for the authors to draw firm conclusions regarding learning, as there was no control group, or randomisation of the design (e.g. an ABAB reversal design). Furthermore, and critically, the impact of the robot on the learning could not be assessed in this study. A similar research design and setup was used in Clabaugh et al. (2019), who also found that the robot-assisted intervention could sustain engagement over a month-long period of use (~14 sessions), but did report large differences in engagement between autistic children.

In a similar vein, Syrdal et al. (2020) reported on their experiences with deploying the KASPAR robot in a nursery school. The robot could perform some scenario's autonomously, whereas others required input from a school staff member. Contrary to Scassellati et al. (2018) and Clabaugh et al. (2019), the autistic children learned through interacting with the robot and the school staff member, rather than through computer-displayed learning content. On average, the children interacted with the robot 27 times, but again, there were large differences in engagement between children (ranging from a handful to nearly 70 interactions). Over time, the children

improved in the sensory and communication domains, although the research design did not allow the authors to conclude that these domains improved because of their interaction with the robot.

To summarise, the current state of the art robot-assisted interventions can maintain engagement for a month and operate autonomously (Scassellati et al., 2018; Clabaugh et al., 2019). To what extent the robot contributes to the sustained engagement over longer periods of time, however, is unclear, as the learning content (i.e. digital games) and the accompanying adult could also be motivating factors. Studies also reported large differences between autistic children in their engagement within robot-assisted interventions (Clabaugh et al., 2019; Syrdal et al., 2020). Nevertheless, initial results from studies that assess whether their robot-assisted intervention taught the children the targeted skill are promising (Scassellati et al., 2018; David et al., 2020; Syrdal et al., 2020), yet inconclusive. The reported increases in learning could also be the result of the passing of time or learning outside the intervention (Scassellati et al., 2018; Syrdal et al., 2020), or might be no different from learning in the intervention without a robot (David et al., 2020). Taken together, it is still unclear how we should design robot-assisted interventions to sustain long-term engagement that leads to learning, and where the robot is providing a benefit over other similar (technology-assisted) interventions. Given the central role of engagement in learning for the effectiveness of robot-assisted intervention (and for this dissertation), we will start with a brief explication of the concept of engagement.

2.2 Engagement definitions and measures

Azevedo (2015) discusses how the concept of engagement in learning has been (mis)interpreted in various ways in the literature leading to different definitions and meanings of this concept. Within the field of Human-Robot Interaction (HRI), engagement is often used as an outcome measure to say something about the quality and length of a participant's interaction with the robot. Engagement is also a concept used for the development of robots that can detect various stages of the user's engagement, such as the intention to engage, being engaged, or being disengaged. An often used definition of engagement in HRI literature is that of Sidner et al. (2005), who define engagement as “a process by which individuals in an interaction start, maintain and end their perceived connection to one another”. Other definitions emphasise that engagement is an affective process formalised as the degree to which an individual wants, or chooses, to engage with a system (e.g. Bickmore et al., 2010; O'Brien and Toms, 2008). While there is no consensus on a single definition for engagement, it is generally viewed as a multi-dimensional concept, including a behavioural, cognitive, and affective component (Connell and Wellborn, 1991; Fredricks et al., 2004). Note that these components of engagement are overlapping and can sometimes be difficult to disentangle (Sinatra et al., 2015).

2.2.1 Behavioural engagement

We are interested in engagement with a robot as it relates to the children's state in which learning can occur, as this is what we are trying to achieve with robot-assisted

interventions. In such a context, *behavioural engagement* refers to the child's participation in learning activities and involves on-task behaviour. Overall, studies on the engagement in learning of autistic children mostly relate to behavioural engagement (Keen, 2009), often referred to as “social engagement” when the task is to engage with another person. As with the definition of engagement, there are also various approaches to measuring the behavioural engagement of autistic children in their interaction with a robot. It can be directly assessed through observing the behaviour of the child, annotating the level of behavioural engagement on a macro-behavioural level. For instance, Kim et al. (2012) developed a compliance-based coding scheme for measuring (behavioural) engagement, where the speed of the autistic child's reaction to instructions or requests is indicative of the child's level of behavioural engagement. In this coding scheme, spontaneous engagement is the highest level of behavioural engagement, in contrast to a child refusing to comply to the robot or adult's request and walking away, which is the lowest level of behavioural engagement. Other studies report on micro-behavioural interactions, which are used to code the type of engagement, to get a deeper insight into how autistic children engage (e.g. Kostrubiec and Kruck, 2020; Schadenberg et al., 2020b). The choice for a certain measurement of behavioural engagement appears to be influenced by the purpose of the study and the type of data being gathered.

2.2.2 Affective engagement

Affective engagement is about the child's (inferred) interest in the learning activity and how much the child enjoys it. It is generally assessed through observing the autistic child's emotions from which the underlying affect is inferred in terms of valence and/or arousal (e.g. Kim et al., 2012; Rudovic et al., 2017, 2018). Correctly inferring the affect from the emotional expressions of autistic children can be difficult, however, as they can produce unique and unusual facial expressions, including blends of incompatible emotions that are not seen in typically developing children or children with Down syndrome (Yirmiya et al., 1989). Also, the vocal intonation of autistic children can be atypical when expressing emotions (Macdonald et al., 1989; McCann and Peppé, 2003; Paul et al., 2005). Notwithstanding, valence and arousal can be successfully annotated for autistic children with sufficient agreement between coders (Kim et al., 2012; Rudovic et al., 2017). Furthermore, using a machine learning approach trained on data of autistic children, valence and arousal can be detected using facial, body-pose, and audio features, and heart rate (Rudovic et al., 2018).

2.2.3 Cognitive engagement

Cognitive engagement refers to the quantity and quality of the child's psychological investment in learning (i.e. use of cognitive effort in order to understand). The cognitive engagement of autistic children is difficult to measure, as the current measures for cognitive engagement (Azevedo, 2015) overlap with the other components of engagement (Miller, 2015), or can be too complex to be used by autistic children, such as self-report questionnaires. Task-evoked pupillary responses have long been associated with attentional engagement (Beatty, 1982) and cognitive activity (Hess

and Polt, 1964; Kahneman and Beatty, 1966), as well as emotional arousal (Bradley et al., 2008), and have been used by researchers to measure the cognitive/affective engagement of autistic children (e.g. Frost-Karlsson et al., 2019). However, measuring pupil dilation requires carefully controlled experiments and experiment environment to control for other factors that influence pupil dilation, such as the pupillary reflex to changes in illumination (Beatty and Lucero-Wagoner, 2000). This makes pupil dilation difficult to use in real-world settings where the illumination cannot be fully controlled. A concept that is more easily measured — and is related to cognitive engagement — is that of attention, which is often viewed as a necessary component for basic forms of engagement to occur (Corrigan et al., 2016). Attention has both a covert and overt component, where overt visual attention is relying on the gaze fixation on a certain location, and covert attention involves cognitive processes for paying attention to something without the movement of the eyes (Wu and Remington, 2003). Indeed, gazing at a particular object is not always indicative of the person paying attention to that object (Posner, 1980). Nonetheless, measuring visual attention through gaze is a commonly used proxy for cognitive engagement in the field of HRI (Rich et al., 2010; Anzalone et al., 2015), and is also used with autistic children engaging with robots (e.g. Tapus et al., 2012; Anzalone et al., 2015; Javed et al., 2019, 2020; Kostrubiec and Kruck, 2020).

2.2.4 Engagement as a holistic concept

Clearly, the various measures and components of engagement also show overlap. For example, as Miller (2015) noted, gazing at a certain location may also indicate affective engagement, as people look more at what they like (Maughan et al., 2007). Altogether, engagement (the concept as a whole) is a fusion of behavioural, affective, and cognitive components of a person's involvement with a robot. Each of the components of engagement are considered essential for learning, as children are able to learn more effectively this way than when passively observing or receiving information (King et al., 2014). As such, considering all three components can provide a richer characterisations of a child's engagement than any single component. Sometimes all three components of engagement are combined into one bespoke measure for engagement (e.g. Simpson et al., 2013; Javed et al., 2019; Jain et al., 2020). For example, through measuring social signals such as eye gaze, vocalisations, smiles, spontaneous interactions, and imitation, which can then be converted into an engagement score by adding one point of engagement for each social signal that is present in a certain segment.

2.2.5 Conclusions

To conclude, engagement is a very broad concept and is formalised and operationalised in HRI studies in various ways. In spite the differences, studies commonly refer to the concept simply as “engagement” and do not make a distinction between the various components of engagement that is study relates to. This makes it difficult to generalise findings from one study that investigated engagement to other settings. As we are interesting in engagement as it relates to learning, we ascribe to the view that

engagement consists of a behavioural, affective, and cognitive component (Connell and Wellborn, 1991; Fredricks et al., 2004), since this distinction is widely used in educational sciences. These components can also be applied to HRI, and have also been used in this context, although to a lesser extent. Our own research will therefore also distinguish between these three components of engagement where necessary, or otherwise consider the concept holistically when we refer to it as “engagement”. Next, we will look at how autistic children engage within robot-assisted activities, and with robots and objects in specific.

2.3 The interactions of autistic children with robots and objects

2.3.1 Autistic children interacting with robots

The type of robots that are used in robot-assisted interventions are referred to as socially assistive robots. The main feature of these robots is that they interact socially with the user as a means of helping them in some way (Feil-Seifer and Matarić, 2011b). What this interaction looks like, and thus what the design of the robot’s behaviour should try to achieve, depends on how the robot is positioned within an intervention.

In a review, Diehl et al. (2012) identified three types of socially assistive robot applications in interventions for autistic children. Firstly, the robot can be used to elicit a target behaviour. This can then create a situation that can be utilised by an adult – or the robot – to promote prosocial behaviour. An example of this application is the intervention described by David et al. (2018), where the robot tries to elicit joint attention and provides feedback. Secondly, the robot can be used as a tool for learning and practising a target behaviour. For instance, in Chevalier et al. (2017a), the robot mimics the child’s facial expressions and serves as a mirror for the child to enable playful practice with facial expressions. In the intervention, the robot is a tool used by the adult who asks the child to make specific facial expressions of emotion. The resulting situation can then be utilised by the adult to teach more about the recognition and expression of that emotion. Lastly, the robot can provide encouragement and promote interaction with another person. An example of this approach is the intervention reported by Huskens et al. (2015), where the robot encouraged an autistic child and that child’s sibling to cooperate with each other in a Lego construction task.

From a pedagogical point of view, it does not necessarily matter whether the child interacts with the robot directly or whether the robot promotes interaction between the child and another person, as learning can occur in either case. For example, teaching joint attention can be done through using the robot as an object of shared attention between adult and child (e.g. Robins et al., 2004), or the robot itself could direct the child’s attention elsewhere by saying “Look!” and pointing (e.g. David et al., 2018). While many robot-assisted interventions seek to actively design robot behaviours for promoting interaction between the child and another person, there is a plethora of studies reporting that these interactions also occur spontaneously (Robins et al., 2004, 2005; Duquette et al., 2008; Feil-Seifer and Matarić, 2009; Kozima et al., 2009; Kim et al., 2013; Costa et al., 2015). For instance, Kozima et al. (2009) reported autistic children turning to the adult and sharing their enjoyment after the robot

responded unexpectedly to the child's touch. When the child initiates such a behaviour, it is commonly referred to as a *social overture* (e.g. Lord et al., 2012), which is a behaviour whose purpose is to communicate social intent. Teaching autistic children to spontaneously initiate interactions with others can also be a goal on its own (e.g. in Pivotal Reponse Treatment, Koegel and Koegel, 2006). Autistic children often have difficulty initiating social interactions (Stone et al., 1997; Koegel et al., 1999b), which may limit their ability of self-learning (Koegel et al., 1999a) and eliciting teaching interactions from their environment (Koegel et al., 2003). For this reason, researchers are also looking into enhancing interventions aimed at promoting social initiation skills with a robot (Huskens et al., 2013).

The above outlines the various kinds of social interactions between autistic child, robot, and other persons, robot-assisted interventions aim to achieve. However, from the child's perspective, the interaction "intended" by the designers may not be of interest or even be a logical response to the robot's morphological and behavioural cues. How these morphological and behavioural cues of a robot are processed depends on the child's cognitive ability (Johnson, 2003). Given that autism affects the cognitive development, these cues can be processed very differently from one autistic child to another. One may interpret these cues and consider the robot to be a *social actor*, whereas another child may come to the conclusion that the robot is an *inanimate non-social object*. As a result, there are large individual differences in how autistic children interact with a robot. While the cognitive processes underlying the perception of robots by autistic children remain unclear, there is behavioural evidence of autistic children interacting with robots in object-like manners, as well as interactions where they may consider the robot a social actor (Short et al., 2017). Moreover, compared to typically developing children who readily attribute human-like characteristics to a robot (Beran et al., 2011), there is preliminary evidence for autistic users, where this tendency was found to be reduced for autistic children (Chaminade et al., 2015) as well as autistic adults (Bird et al., 2007). Understanding how individual characteristics impact the interaction between an autistic child and robot is essential to effectively designing robot behaviour to engage these children and in choosing the robot morphology that is best suited to facilitate these robot behaviours.

Thus, while robot-assisted interventions are typically designed to elicit social interaction, this does not mean that autistic children also consider the robot to be a *social actor*. Because we can also expect object-like interactions with a robot, we will briefly discuss related work on how autistic children interact with inanimate objects.

2.3.2 Autistic children interacting with inanimate objects

While robots are relatively novel technology that autistic children interact with, their interactions with regular, inanimate objects – such as toys – have been studied extensively. Much of the research on how autistic children interact with objects is conducted in a play setting, where researchers study how the complexity of play develops as the development of the child progresses. In particular, the children's ability to focus their attention, motivation, and representational capacities, play an important role in their interaction with an object (Vig, 2007; Lifter et al., 2011). Different developmental stages of play can be distinguished, which generally include sensorimotor

or *exploratory play*, *relational play*, conventional or *functional play*, and representational or *symbolic play* (Libby et al., 1998; Casby, 2003; Naber et al., 2008). Children start out with exploratory play with objects and gradually develop the ability to create cognitive representations of objects and events required for more sophisticated types of play, such as symbolic play (Stagnitti, 2004). These developmental stages are not mutually exclusive, and children can exhibit a variety of play types. Autistic children seem to follow the same developmental trajectory of the play types as typically developing children (Vig, 2007), but may show the less sophisticated play types due to developmental delays.

Exploratory play is the earliest type of interaction with objects, starting to emerge at around three to four months of age during typical development, and are marked by oral or manual manipulation of objects, such as spinning, smelling, or mouthing (Williams, 2003). Through this type of play, children learn about the properties of different objects and how they relate to the world around them. As children begin to understand how objects relate to each other, they start showing *relational play*. This is play where a child uses different objects and relates them to each other in a way that does not indicate functional use of the object (i.e. using an object for its intended purpose). For example, nesting one object in another, or stacking objects. When children become aware of and show attention to the different properties of objects and their uses, objects start to be used in a conventional manner, which is called *functional play*. For example, children may push a toy car, or put a telephone to their ear. This type of play requires a first-order representation of the object. As children start to develop the cognitive capacity for second-order – or meta – representations of objects, they gain the ability to decouple mental representations of objects from reality (Leslie, 1987). Objects can then be used by pretending it is something else, attributing false properties to the object (e.g. the robot is ill), or referencing to an object as if it were present. This is called *symbolic play*.

Compared to typically developing children, or children with other developmental difficulties, the manifestation and frequency of the play types for autistic infants (Naber et al., 2008; Pierucci et al., 2015) and autistic children (Libby et al., 1998; Wilson et al., 2017) is markedly different. During exploratory play, autistic children seem to prefer using proximal senses of touch, taste, and smell to explore objects, rather than using vision (Williams, 2003; Naber et al., 2008). When they do visually inspect an object, they may place the objects close to their eyes, or they may focus on one aspect for an extended period of time (Williams, 2003). Libby et al. (1998) reported observing fewer instances of relational play in autistic children than typically developing or children with Down syndrome. Similarly, autistic children engage less frequently in functional play compared to other children (Williams et al., 2001; Christensen et al., 2010). Furthermore, their functional play tends to be less varied, integrated, and complex than that of other children (Williams et al., 2001; Christensen et al., 2010). Symbolic play is the play type with which most autistic children have difficulty (Jarrold, 2003). According to Jarrold (2003), autistic children may have the underlying capacity to understand symbolic play, but are less inclined to spontaneously engage in symbolic play. One explanation for this is that autistic children are more tied to the properties of an object and may have difficulty overriding

these properties by pretending it is something it is not.

In summary, the four play types describe how children may engage with an object. When applied to the interaction with a robot, child-robot interactions where the child explores the robot's materials through any of the senses would classify as an exploratory interaction. Relational interaction with the robot are interactions where the child uses additional objects with the robot in a non-functional manner. Functional interactions are interactions for which the robot was designed. In most cases of robot-assisted interventions, these would be social interactions. The distinction between functional and symbolic interaction is more delicate when it comes to SAR, as they are designed to elicit social interaction. This is only possible when the robot is viewed as a *social actor* to a certain extent, rather than a bunch material wired together to form a *non-social object*. Therefore, the appearance of robot and their behaviour purposefully create the illusion of animacy (Castro-González et al., 2016). Attributing animacy to a robot can then be considered a false belief. Whether the autistic child-robot interaction stems from the child's belief that the robot is a living entity, or merely because they learnt that this is how you should interact with a robot (e.g. by observing an adult interacting with the robot) is difficult to determine. Before an interaction can be classified as symbolic, there needs to be clear evidence that the child is aware of attributing a non-existing property to an object (Lillard, 2001).

2.4 The idiosyncrasies of autistic children

As of the fifth edition of the DSM (American Psychiatric Association, 2013), autism is considered a spectrum condition, combining autism-subtypes such as classical autism, Pervasive Developmental Disorder-Not Otherwise Specified, and Asperger's syndrome, into one classification. This is to account for the *wide* variety in the type and severity of the symptoms experienced by autistic individuals. There is a saying that "if you've met one person with autism, you have met one person with autism" (attributed to Stephen Shore), which embodies the immense variety between autistic individuals that supersedes natural variations in interests. As a neurodevelopmental condition, autism affects the developmental speed and trajectory of children. This results in certain abilities being developed late, or even not at all (e.g. some autistic individual never develop expressive speech). This also means that the child's age is not as good an indicator of the child's abilities as with typically developing children. Given that autistic children are a heterogeneous group of individuals, we need to take this into account when designing robot-assisted interventions. After all, the child-robot interactions are also different when the robot is to interact with a typically developing 2-year old, compared to a typically developing 10-year old child. In this section, we will therefore briefly discuss several autism-related aspects in which autistic children differ that can influence the autistic child-robot interaction.

Autism is a condition that causes atypicalities in cognitive, emotional, behavioural and social functioning. However, how these are manifested is different in both quality and quantity (Happé et al., 2006). In social communication, autistic children can have difficulty with understanding and using language (both verbal and non-verbal). This ranges from individuals who never develop spoken language to those who speak

fluently but for whom their use of language to communicate with others can seem awkward and often one-sided to non-autistic individuals. Some children may also echo what another person has said, which is called echolalia. In social interaction, autistic children often have difficulty with recognising or understanding others' feelings and intentions (Baron-Cohen, 2009), and with expressing their own emotions (Uljarević and Hamilton, 2013). In turn, this can result in various difficulties when navigating the social world (Halberstadt et al., 2001).

For the non-social features of autism, autistic children differ in their repetitive and restrictive behaviours. To deal with the unpredictability of the world, the children can develop certain routines so that they know what will happen — an *insistence of sameness*. Routines such as wearing the same clothes, or walking the same road to school. Autistic children may also repeat various behaviours that generate predictable sensory input, such as flapping with their hands, or rocking in a chair. Any changes in such routines or behaviours can distress the children and make them anxious (White et al., 2009). The majority of autistic children have unusual sensory responses (O'Neill and Jones, 1997), where some have certain sensory sensitivities, such as being sensitive to sounds, touch, tastes, smells, light, colours, temperatures or pain. For instance, an unexpected sound can feel like a needle piercing your eardrums, which is why some autistic children wear headphones. In contrast, rather than being sensitive to these modalities, other autistic children can be under-sensitive and may instead seek stimulation. For example, some children may wear a pendant for chewing, which can calm them. And finally, many autistic people have intense and highly focused interests (Klin et al., 2007). Being highly focused on certain interest to subjective well-being and satisfaction across specific life domains including social contact and leisure (Grove et al., 2018). However, it can also lead to being engrossed by a certain interest and result in neglecting other aspects of life.

Next to variability in how the core features of autism manifest themselves, there are also large individual differences in symptoms that are associated, but not core, to autism. In terms of cognitive functioning, autistic individuals range from having severe intellectual disability to superior intelligence (Grzadzinski et al., 2013). Autism is also associated with increased comorbidity for other conditions/disorders, such as ADHD, anxiety disorders, tic disorders, learning disabilities, and epilepsy (Matson and Nebel-Schwalm, 2007). Next to their difficulties, autistic individuals can also show areas of strength, such as in visuo-spatial skills (Shah and Frith, 1983), memory (Plaisted et al., 1998), or musical ability (Heaton, 2003).

2.5 Conclusion

In summary, studies on robots for autism show promising results in eliciting social interaction between the autistic child and the robot or another person (Scassellati et al., 2018; Clabaugh et al., 2019; David et al., 2020; Syrdal et al., 2020), and facilitating learning of a targeted skill (Scassellati et al., 2018; David et al., 2020; Syrdal et al., 2020). Through the use of robots, several key needs of autistic children could also be addressed, such as the need for predictability (Dautenhahn, 1999) and a reduced complexity of social interaction (Duquette et al., 2008; Sartorato et al.,

2017). Generally speaking, autistic children appear to be engaged in robot-assisted activities, both on a behavioural and affective dimension.

When we look closer at how these interactions between an autistic child and robot manifest, it is unclear to what extent some of the children view the robot as a social actor (Short et al., 2017). This may have to do with the child's (atypical) cognitive development (Johnson, 2003). From a developmental point of view, interactions with objects gradually increase in complexity as their cognition develops. Children first start with exploratory interactions, which gradually starts to involve relational, functional, and eventually symbolic interactions. Anecdotally, these types of interactions have been reported with robots as well (Scassellati, 2005; Kozima et al., 2009; Short et al., 2017). The differences how autistic children tend to interact with objects ties in with the idiosyncrasies of autistic children in general. As these children are so different from each other, it is unlikely that there is a one-size fits all solution to keeping autistic children engaged in robot-assisted interventions whilst providing meaningful learning.

Together, these works form the basis on how we can engage autistic children in the DE-ENIGMA intervention. In the next part of this dissertation, I will report on the major studies that shaped our intervention. In Chapter 3, we assessed what role robot-assisted interventions can play in education, and what is required to fulfil those roles, we conducted user needs and requirements analyses. The users targeted by the DE-ENIGMA project are autistic children, as well as those who are to use the robot: educators, therapists, or parents. As autistic children are so different from each other, we specifically focused on those with lower cognitive ability and speech, since there are not many tools available for helping these children. These children have high support needs not found in regular elementary schools, and instead go to special education schools. In Chapter 4, I report on a descriptive study that aimed to get a better understanding of what type of interactions might be interesting to autistic children and what role the idiosyncrasies of these children play in this. This will allow us to better develop a robot-assisted intervention that can engage and sustain engagement of autistic children in learning. Finally, in Chapter 5, I will describe how the DE-ENIGMA intervention evolved as the project progressed, and how this was influenced by the results of our studies.

Part II

TOWARDS AN INTERVENTION

*“My switch clicks on,
My eyes light up,
My motor starts the rumbles.
My cogs and gears all creak and crank.
My radar turns and tumbles.
My program data bank boots up,
It tells me what to do ...
I’m WL-3-61.
Now, how can I help you?”*

Mark Carthew, Vroom, Vroom! Poems About Things with Wheels

3

User needs and requirements for robot-assisted interventions

This chapter is partly based on the following paper:

Schadenberg, B. R., Li, J. J., Petrović, S., Reidsma, D., Heylen, D. K. J., & Evers, V. (2020). Helping Educators Monitor Autistic Children’s Progress Across Sessions. In *the Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 436–438). New York, NY, USA: ACM. ([Schadenberg et al., 2020a](#))

and references [Alcorn et al. \(2019\)](#), where I am one of the authors.

To develop the DE-ENIGMA intervention, we adopted a user-centered design approach. This involves assessing the needs of our users and defining user requirements for our system so that it could be successfully used by our users. For the DE-ENIGMA project, we focus the development on three user groups, namely the autistic children, their parents, and the autism professionals who are to use the robot (e.g. educators, or occupational therapists). In this chapter, we will describe how we assessed the needs of our users, and how we defined the user requirements. Note that the development of the DE-ENIGMA intervention was a collaboration between all consortium partners, and that the two need-finding studies described in this chapter should also be considered as such.

3.1 Introduction

FOR robots to be a useful addition to an autism intervention, they need to provide some benefit to the users. In this chapter, we will be looking at what benefits robots are anticipated to provide to our end-users, which are the autistic children, as well as their parents. We also consider the educators and occupational ther-

apists who work with the children as end-users⁵. Specifically, we will explore whether the anticipated benefits are addressing real problems our end-users face, as well as what our users would expect from a robot that is assisting in the intervention.

Researchers have argued that robots could provide various benefits to autistic children. As robots are programmable, they can be designed to provide highly predictable interactions (e.g. Dautenhahn and Werry, 2004; Duquette et al., 2008; Thill et al., 2012; Huskens et al., 2013; Sartorato et al., 2017; David et al., 2020), which autistic people are often said to favour (Baron-Cohen, 2009; Pellicano and Burr, 2012; Lawson et al., 2014). Robots may therefore “mediate between the (from an autistic child’s perspective) widely unpredictable world of ‘people’ and the much more predictable world of machines.” (Dautenhahn, 2007, p. 693). Another benefit of robots that is often mentioned is that robots can show less complexity in their behaviour and physical appearance than humans, which could make robots easier for autistic children to comprehend (e.g. Duquette et al., 2008; Cabibihan et al., 2013). This may also allow robot behaviour to be used as a stepping stone to teaching more complex behaviour (Robins et al., 2005). Taken together, the increased predictability and reduced complexity may explain the positive effects in terms of the children’s engagement that have been reported by researcher (see Chapter 2.1).

The addition of a robot to an intervention may also have benefits for those working with the autistic children. Parents of autistic children might be faced with challenges that can prevent them from accessing care for their autistic child, such as high costs of interventions, limited availability of providers, or geographic isolation (Zheng et al., 2020). These issues could be addressed by providing on-demand learning for autistic children through a robot-assisted intervention that is designed for at-home use by an autism professional. A robot could also alleviate the workload of autism professionals by providing an extra hand in an intervention (Huijnen et al., 2019).

Next to providing certain benefits through aiding in the delivery of the intervention, we believe robots may also provide unique opportunities to gather and present information on the child’s progression. This information is important for understanding whether an intervention is working, but can be difficult for autism professionals to assess, as the changes can be minor and are easily missed (Kientz et al., 2007). A robot could then serve as an additional set of sensors that provide information on the child’s progression. We consider this as a (possible) *secondary* purpose of the robot. Current methods for monitoring autistic children’s behaviour predominantly rely on embedded cameras (e.g. Kientz et al., 2005; Hayes et al., 2008) or wearables which record movement patterns (e.g. Nazneen et al., 2010). However, the children may not be aware that their behaviour is being monitored, as they do not directly engage with these systems, nor are these systems transparent to the children, which raises concerns on whether the children would assent to the monitoring of their behaviour (Spiel et al., 2019). Robots could alleviate this concern, as they already observe and analyse the child’s behaviour to facilitate a degree of autonomy in their interaction with the child. Moreover, robots can be designed to be highly transparent about their analysis of the child’s behaviour by communicating this explicitly to the child.

⁵In the remainder of this chapter, we will refer to the educators, therapists, and other professional users of the DE-ENIGMA intervention as “autism professionals”, when addressing the whole group.

As shown above, various benefits are anticipated for embedding robots in interventions for autistic children. Realising those benefits and addressing real problems experienced by our end-users, however, is no straightforward task (Diehl et al., 2012; Begum et al., 2016). While various research groups are studying and developing robot-assisted interventions, to date, there is no clinical evidence that they lead to learning. Thus, what is required of a robot to be successfully embedded into an autism intervention is not clear. These requirements likely differ per setting in which the robot is to be used, and how the robot is positioned within the intervention. For instance, the needs of autistic children educated in special education and the strategies to support them can be very distinct from autistic children educated within mainstream settings (Eaves and Ho, 1997). A robot that is to be used within either setting will likely require different designs. For the DE-ENIGMA intervention to be successful, we then first need to get a better understanding of our end-users — what are the problems they are facing in current practices, and what do they require from a robot that is embedded in an autism intervention? With a better understanding, we can deliberately design our intervention to align to current needs and meet the expectations of our end-users.

In this chapter, we will investigate the current unmet needs of our end-users, in relation to conducting, or participating in, autism interventions, and the monitoring and reporting of the autistic children's progression. Given that the focus of the DE-ENIGMA project is on autistic children with limited spoken communication and high support needs, we focus on them and the autism professionals who work and care for these children. To assess the needs of our end-users, we utilise three sources of information from which we extract user needs. The first source is a systematic literature search that we conducted to find literature on the needs of our end-users within a robot setting. The second and third source both relate to needfinding studies that were carried out by the DE-ENIGMA consortium. We conducted interviews and focus groups with educators in England, working in specialised settings for autistic children, to assess how they view the use of a humanoid robot in their schools — the second source⁶ The third source that we use relates to a small set of interviews that we conducted with autism professionals in the Netherlands and Serbia to assess their needs for monitoring and reporting the progression of autistic children (Schadenberg et al., 2020a). To distinguish the second and third source from the literature found through our systematic literature search, we refer to Alcorn et al. (2019) as the ALCORN ET AL. STUDY and Schadenberg et al. (2020a) as the SCHADENBERG ET AL. STUDY.

The rest of this chapter is structured as follows. First, we discuss the methodology used for each of the sources in Section 3.2. From these three sources of information, we then distil and discuss the user needs (Section 3.3), as well as assess the requirements for robot-assisted interventions according to autism professionals and discuss these (Section 3.4). In the last section (Section 3.5), we discuss and conclude on the overarching topics that need to be addressed in robot-assisted interventions.

⁶While I made a significant contribution to this needfinding study, I did not lead it. I will therefore only describe the article that came out of this study — Alcorn et al. (2019) — rather than provide a full report of the study.

3.2 Materials and methods

3.2.1 Method for the user needs and requirement analyses

User needs define what Through understanding the needs of our users, we can start to design our robot-assisted intervention in such a way that it can address those needs. User needs relate to *what* we want to achieve, but do not specify how they can be achieved using a robot. They are written in a manner that describes the problems that the user needs addressed rather than describe the solution for these needs. The user requirements specify *what* the users require from the solution for it to be successful. Note that user requirements do not specify *how* the robot should address those requirements (this are specified by system requirements); only that they should be addressed. The user requirements can relate to the user needs, but usually also relate to much broader and generic topics. In our discussion of the user requirements, we also provide some potential solutions for meeting these requirements, based on other literature.

We employed three sources to triangulate to the user needs and requirements: existing literature, an interview and focus groups conducted in England, and need-finding interviews carried out in the Netherlands and Serbia. When combined, they give us a better understanding of the user needs and requirements that are currently reported and how much evidence there is for those needs and requirements. In our analysis, we will go over each source and list the reported needs and requirements, assess which needs or requirements overlap with those mentioned in other sources, and finally summarise and discuss them in relation to the literature. As some studies report results that are specific to the robot(s) they showed to the participants, we also assessed whether the results could be generalised to other robot platforms. If so, they are included in our results. Below we will briefly explain the methods used for each of the sources.

3.2.2 Source 1: Existing literature

To support the identification of user needs and requirements, we first conducted a systematic literature review. The goal of this review was to assess what user needs and requirements were already reported. The databases that were accessed for the literature review included Scopus, Web of Science, and Google Scholar. For our electronic search, we used the following keywords: autism, autistic, ASD, ASC, robots, requirements, and needs. This resulted in 44 papers on requirements and 221 papers on needs. The selection of papers and articles was based on five additional criteria:

1. The paper should present a study involving end-users (autistic children, their parents, or autism professionals), or discuss user requirements or user needs in relation to robots.
2. The goal of the study should be to find user needs or requirements.
3. The user needs and requirements should relate to robot-assisted interventions for autistic children.

4. The end-users should either be the autism professionals or parents who are to carry out a robot-assisted intervention, or the autistic children themselves.
5. Only full papers are included in the analysis — extended abstracts were omitted because these often contained preliminary findings.

From the search, we excluded our own articles — the ALCORN ET AL. STUDY and SCHADENBERG ET AL. STUDY — as we discuss these as separate sources. We also excluded Robins et al. (2007), as the results reported in this paper are also reported, and expanded upon, in Robins et al. (2010b), which is included. The following papers, or articles, matched the selection criteria and were included in our analyses: Giullian et al. (2010); Robins et al. (2010a,b); Coeckelbergh et al. (2016); Huijnen et al. (2016); Zubrycki and Granosik (2016); Huijnen et al. (2017); Martin-Ortiz et al. (2017); Putnam et al. (2019).

3.2.3 Source 2: The views of educators on using robots

The DE-ENIGMA consortium carried out interviews and focus groups with educators who worked in specialised settings in England to assess their views and perspectives on the potential use of humanoid robots for autistic children. This is the ALCORN ET AL. STUDY. The goal was to get a better understanding of the factors perceived to be important for deploying robots in these settings. Below, I will briefly summarise how this study was conducted.

Fourteen educators were interviewed, and seventeen educators participated in one of three focus groups. All participants worked in a specialist setting (special education or as an autism resource base at a mainstream school) in England. The procedure was the same for the interviews and focus groups. First, the participants viewed six example images of existing humanoid robots (Zeno, KASPAR, Nao, Flobi, Parlo, and Pepper), but were not given any other information regarding these robots. They were then asked to consider the potential uses, as well as any potential benefits or concerns, of humanoid robots for autistic children through various probe questions. The interviews and focus groups were transcribed verbatim and analysed using thematic analysis (Braun and Clarke, 2006). We used the thematic analysis as a realist method, meaning we report on the experiences, meaning, and reality of participants. The themes were identified using a bottom-up (data-driven) approach. Two researchers familiarised themselves with the transcripts, reviewed the identified themes, and concluded on the final definitions of themes.

3.2.4 Source 3: Educators' views on the reporting and monitoring of progression

We aim to utilise the potential of social robots to provide information on an autistic child's progression, but we do not yet know well enough what information is actually useful to educators. Therefore, we conducted interviews to assess the needs of educators for monitoring and reporting the progression of autistic children. This study (SCHADENBERG ET AL. STUDY) is highly preliminary, given that we only interviewed nine educators. While we consider this source in conjunction with the other sources,

this study would be excluded based on the criteria of the systematic literature search, because the study was published as a late-breaking report. However, as we conducted the study ourselves and it influenced the DE-ENIGMA project, I do report this study in this chapter.

Participants. Nine educators (female: $n = 8$), between 32 and 54 years old ($M = 40$, $SD = 9$), participated in semi-structured needfinding interviews. Two were from the Netherlands (NL), and seven from Serbia (SRB). Participants were recruited from a local search of autism-related institutes (NL) and from a network of autism educators (SRB). Their titles are art therapist ($n = 1$, NL), head psychologist ($n = 1$, NL), special educator ($n = 6$, SRB) and psychomotor educator ($n = 1$, SRB). The educators had between 5- and 29-year's of experience ($M = 17$, $SD = 9$) working with autistic children. All participants were actively employed as educators working with autistic children or adolescents at the time of the interviews.

Procedure. The interviews were conducted over the phone (NL) or in person (SRB) and lasted approximately 30 to 45 minutes. They were structured as follows. Firstly, the interviewer introduced herself and obtained consent to record audio. The participants were then asked demographic and professional questions (e.g., type of children they work with, types of therapies they use). Next, the interviewer would ask the participants to describe actual positive and negative events that occurred in a session with an autistic child, how they recognised and dealt with the event. This was followed by the interviewer asking the participants to describe their current practices with measuring and reporting on the progression of an autistic child. Afterwards, participants were asked several questions regarding how they would like to measure and report on the child's progression if it were up to them.

Analysis. Similar to the interviews and focus group in the ALCORN ET AL. STUDY, we transcribed the audio-recordings verbatim and analysed using thematic analysis (Braun and Clarke, 2006). In this case, three researchers familiarised themselves with the transcripts, reviewed the identified themes, and decided on the final definitions of themes.

3.3 User needs and discussion thereof

In this section, we will discuss the user needs that we found through our analysis. Papers from which we extracted the user needs are: [Giullian et al. \(2010\)](#); [Zubrycki and Granosik \(2016\)](#); [Huijnen et al. \(2017\)](#); [Kim et al. \(2019\)](#); [Putnam et al. \(2019\)](#); ALCORN ET AL. STUDY; SCHADENBERG ET AL. STUDY. We will also discuss if and how a robot could potentially fill a role in addressing a user need by drawing upon the broader literature in the field of Human-Robot Interaction. We thus cite more papers than the seven listed above.

The user needs are divided into two subsections, where one subsection relates to the user needs in relations to carrying out autism interventions, and the other subsection relates to the user needs in relation to the monitoring and reporting of

autistic children's progression across sessions. Note that in both cases, the users are the autism professionals. None of the papers that we found addressed the user needs of the autistic children. Their needs were only indirectly assessed by interviewing autism professionals.

3.3.1 Needs of autism professionals in carrying out interventions

Need for a more comfortable and less threatening environment. For learning, it is important for autistic children to feel safe and secure in their environment. Educators mentioned that this may not always be the case [ALCORN ET AL. STUDY]. Because these children have difficulties with social interactions, the social nature of educational settings can be problematic. In particular, the *unpredictable* nature of people can make it difficult for autistic children to make sense of them (Huijnen et al., 2016, ALCORN ET AL. STUDY). After all, people behave in all sorts of different manners and ways (e.g. change of tone of voice or change of volume, use of unexpected gestures, or unclear (non)verbal messages). People can also look and smell differently each day, which some find difficult to deal with. For instance, an educator mentioned that the inconsistency in people's appearance may stop certain autistic children from communicating with a person whose look or smell they do not like that day. Overall, when the environment and its people are too unpredictable to a child, it can lead to anxiety (Paulus and Stein, 2006) and prevent the child from learning. Another factor that plays a role is that the social presence of an educator can make the children feel threatened [ALCORN ET AL. STUDY]. The example educators gave was one where an autistic child might feel threatened by the authority of the educator. This might lead to the child not trying out new things, because they are afraid of making a mistake in front of the educator.

Robots may be uniquely positioned to address this need as a technology that can provide both a highly predictable interaction, because we could program the robot's behaviour to be so, as well as elicit social interactions with the child (Robins et al., 2005; Duquette et al., 2008; Feil-Seifer and Matarić, 2009; Kozima et al., 2009; Kim et al., 2013). In other words, robots can provide *highly predictable* social interactions unlike human-human social interaction, which is typically unpredictable. The educators in the ALCORN ET AL. STUDY believed that autistic children might feel more at ease, and less threatened, when interacting with a highly predictable robot, relative to how they feel in other school activities or human interactions — they know what to expect from the robot, unlike from other people. In turn, this could help them focus their attention on the learning material.

Need for more scaffolding. Providing autistic children with the appropriate assistance and tools that they need to accomplish a new task or skill is not always easy. These “temporary adaptive supports” (i.e. certain activities, instructions, tools, or resources) are known as scaffolding (Wood and Middleton, 1975)⁷. In the ALCORN ET AL. STUDY and Kim et al. (2019), autism professionals reported that they constantly need to adjust and personalise the scaffolding for the children, as they differ in their

⁷Although the origins of the concept of scaffolding can be traced back to Lev Vygotsky, Nikolai Bernstein, and Alexander Luria (Shvarts and Bakker, 2019).

skills, as well as their preferences and needs. In some instances, providing scaffolding is problematic, as the support that can be given is either too easy or too difficult for the child — there is no support available in between.

Addressing the need for more scaffolding through the use of robots might be possible, as for typically developing children, a robot can be equipped with the required sensors and algorithms to detect how much support a child needs and personalise its behaviour accordingly (Jones et al., 2018; Jones and Castellano, 2018). In fact, autism professionals and educators suggested to use robots as tools to provide scaffolding (e.g. in Huijnen et al., 2016, ALCORN ET AL. STUDY). For instance, a child may understand pictorial version of human emotions, or photographs of human facial expressions, but the dynamic facial expressions of a person standing in front of them may still be too difficult. A robot's simplistic representation of human emotions (e.g. as in the facial expressions of Robokind's Zeno, see Schadenberg et al., 2018), may serve as the missing scaffold here. The children would first practice the targeted skill with a robot, which is simpler (e.g. simplistic representations of human emotions), and may also be less threatening, than practising the targeted skills with another person. In the next phase, the educator tries to generalise the targeted skill to human complexity. This way, the robot serves as a stepping stone for applying skills to the more complex and difficult human-setting. A robot may also be able to (re)create situations that are otherwise not possible (Giullian et al., 2010), or not responsible (Huijnen et al., 2017, ALCORN ET AL. STUDY), and bridge the gap between applying a skill in an abstract educational setting to applying it in a real-world setting. For instance, a robot could serve as an interaction partner to practice and recreate difficult situations or interactions with. These interactions could potentially result in uncomfortable situations, such as when addressing aggressive behaviour, where recreating the situation with another person may not be desirable.

3.3.2 Needs of autism professionals in reporting and monitoring

Need for moment-by-moment monitoring. In the SCHADENBERG ET AL. STUDY, the majority of the autism professionals mentioned that detailed, moment-by-moment monitoring is currently challenging. Some professionals work with multiple autistic children at the same time, which makes it impossible to always notice antecedent behaviours that lead up to significantly positive or negative events. Behaviours such as when a child gets upset or leaves. Participants mentioned that this could be because of something that happened during the session, or because of something that happened outside the session — the context is important. In one-on-one sessions it may also not be possible to pay close attention to the child at every moment. In Zubrycki and Granosik (2016), autism professionals said that they sometimes have difficulty interpreting the child's emotions and may act too late to address these, resulting in negative situations where they potentially need to call in help from colleagues. Other autism professionals mentioned the high workload during a session in which they need to keep track of various things; this can prevent them from paying close attention to the child at times [SCHADENBERG ET AL. STUDY].

Through moment-by-moment monitoring, it could also be possible to identify and understand antecedent behaviours that lead up to positive or negative events, which

may benefit the broader team that assists the child (i.e. parents and various care-takers) (Putnam et al., 2019, SCHADENBERG ET AL. STUDY). When these behaviours are communicated with this team, observing the positive behaviours could allow the team to better understand what the child may find particularly motivating, whereas the negative behaviours can signal that a team member needs to intervene. Thus, this could allow the team to better communicate with each other and assist the child.

A robot could relatively easily address the need for moment-by-moment monitoring, as any (semi) autonomous robot will already have the required sensors and use them to perceive and act upon the environment. These sensors can be utilised to continuously monitor and analyse the child's behaviour. In turn, this analysis could be used to mark salient moments (i.e. significantly positive or negative events) that occurred during the intervention in a recording, which autism professionals could then use to quickly review the interaction (in a similar fashion as proposed in Nazneen et al. (2012)). This could then aid in identifying and understanding the antecedent behaviours that lead up to positive or negative events.

Need for standardised measures of social skills. Six autism professionals in Serbia mentioned a need for standardised, validated measures of child progress in social skills learning [SCHADENBERG ET AL. STUDY]. Currently, they use tests that were developed for assessing typically developing children, because the tests for autistic children are not available, or not validated, in Serbian. Other Serbian educators mentioned that they report on the child's progression in social skills purely based on their observations, as they do not have access to any standardised tests. The dependency on a certain language seems to be a key barrier for the adoption of current standardised tests, but is likely to be restricted to certain countries.

Robots could possibly be beneficial in addressing this need however, as they can create standardised situations and provide stimuli that are always the same (Scassellati, 2007). The robot's analysis of the child's affective and behavioural response to these stimuli could be further developed into standardised measures that are language-independent.

Need for understanding child behaviour across settings. Most autism professionals in the SCHADENBERG ET AL. STUDY mentioned that they would like a better understanding of children's progress *outside* of their sessions to assess whether progression within sessions also generalises to other settings. For instance, the child could perform well at school, but abreact at home or vice versa. A professional cannot grasp the whole picture, because they do not see all aspects of the child's life (at school, at home, or at sports), and instead need to rely on what parents tell them. In Serbia, two educators reported a similar need that they partially address with school trips, where they assess how well the contents that they learned at school generalise to a setting outside the school. Nevertheless, the autism professionals felt that their current understanding of the child's behaviour across settings is insufficiently met with the tools and methods they have available today.

The need to assess progression across settings may be more difficult to address with a robot, as it would require that the robot observes and analyses the child's be-

haviour in different settings. This poses serious issues regarding the child's privacy, as well as the privacy of others who are accidentally observed and analysed. Furthermore, the child and any others may not be aware that the robot is observing and analysing his or her behaviour, which leads to questions whether they would give their assent or consent for the robot to do so.

3.4 User requirements and discussion thereof

In this section, we will elaborate on the user requirements resulting from our analysis. We excluded generic user requirements, such as that the robots they are working with should be safe to use and cause neither physical nor mental harm (clearly the most critical requirement), or that the user interface should be easy to use, in order to keep this section specific to the context of robots being used in autism interventions.

3.4.1 Autism professionals in control over the intervention and robot

Autism professionals expressed firmly that autism interventions should remain a human activity. A robot can possibly assist the professional, but not take over the role of the professional (Coeckelbergh et al., 2016; Martin-Ortiz et al., 2017, ALCORN ET AL. STUDY). Human-robot interaction should not replace human-human relationships for the autistic child, but autism professionals warn that this can be the case when robots are not used correctly — in a manner where the autism professional does not have control over the intervention. Interacting with a robot may be highly engaging to autistic children, and educators warned that the robot may have the properties to turn into an obsession for certain children [ALCORN ET AL. STUDY]. Furthermore, the children could also trust the robot and connect with it emotionally, which then could lead to becoming overly dependent on the robot and reduce the child's interaction with people. Similarly, Putnam et al. (2019) reports that one of the reasons why parents who avoided technology for their child did so because they were worried that it might contribute to isolation of their child. On the other hand, in the same study, participants also mentioned that a robot could become like a friend to the child, which seems to be at odds with the belief that technology can lead to isolation from other people.

Autism professionals expressed not only that they want to remain in control of the intervention (Coeckelbergh et al., 2016, ALCORN ET AL. STUDY), but also have (a degree of) control over the robot (Huijnen et al., 2016; Zubrycki and Granosik, 2016; Martin-Ortiz et al., 2017, ALCORN ET AL. STUDY). A robot should fill in for the weaknesses of autism professionals, and not replace their strengths. Educators are trained to assess the varying needs of the children, support those needs using (creative use of) distinct strategies [ALCORN ET AL. STUDY], and, in general, are particularly proficient into “reading the mind of the autistic children” (Huijnen et al., 2016; Martin-Ortiz et al., 2017) (i.e. noticing subtle changes in their emotional well-being). They fear that overly relying on a robot's senses and analyses thereof may deteriorate the quality of an intervention, because robotic technology was judged not to be up to this task currently, nor did they believe that it would be in the future [ALCORN ET AL. STUDY]. The educators also foresee that they need to adapt the robot's behaviour on the fly

when the situation demands it, as the children can behave unpredictably and have dynamic needs (Giullian et al., 2010; Huijnen et al., 2017). Control over the robot's behaviour is therefore necessary, so that the autism professionals have the last say in how the child is likely to feel, and what strategies are likely to be most appropriate.

To address the requirement for being in control over the robot's behaviour, autism professionals will need to be able to interface with the robot. When asked, the professionals said they likely would prefer to interface with the robot through speech (Martin-Ortiz et al., 2017). Alternatively, they could interface with the robot through a remote control, gestures, or touch, although these were judged less favourably. It is also important that the addition of a robot to the intervention, and control thereof, does not increase the workload of the professional, or make it more complicated, as this will likely decrease the adoption of such interventions (Giullian et al., 2010; Huijnen et al., 2017).

3.4.2 Providing a comfortable and safe learning environment

As we have seen in the previous chapter, being able to provide a comfortable and safe learning environment for the autistic children can sometimes be difficult, and a robot was perceived as a possible solution to some instances where the learning environment was not comfortable for the child. While the other user needs that we reported related to specific usages of the robot (i.e. for monitoring and reporting of progression, or for scaffolding), this user need applies to any robot-assisted intervention, regardless of its goal. It is therefore also a user requirement.

The autism professionals mentioned two aspects for addressing the user requirement of providing a comfortable and safe learning environment. Firstly, the unpredictability and complexity of people's behaviour and appearance can make it difficult for autistic children to understand them and can induce anxiety. Secondly, the (high) social demands experienced by the child of having to perform in the intervention can prevent the child from learning and also cause anxiety. To address the former, a robot can be highly predictable when it is programmed to do the same behaviour over and over again, in exactly the same manner, and look exactly the same every day. However, this may not be a very useful contribution to the intervention — the robot will likely have to do more. Current literature, nor the experiences of autism professionals, tell us how to design robots to be both highly predictable as well as provide meaningful interactions⁸.

To improve the simplicity of a robot, professionals noted that presenting multi-modal robot behaviour could cause an information overload (Huijnen et al., 2017), and that a simplistic appearance of it might make it easier for children to interact with them (Giullian et al., 2010; Huijnen et al., 2016). To address the social demands of the intervention placed on the child, the environment should allow for making mistakes and still be supportive (Robins et al., 2010b, ALCORN ET AL. STUDY). Rather than stating that the child's answer is incorrect, the robot could encourage the child to try again, or praise the child on the effort he or she is putting in. In Giullian et al. (2010), it is mentioned that a robot should also not be too large, as this may be too intimidating for the child.

⁸Spoiler alert: read Chapters 6, 8, and 9, to see how this can be achieved.

3.4.3 Familiarising the child with the robot

Needing to interact with an unfamiliar object, or change in general, can be unsettling to autistic children (Robins et al., 2010b; Huijnen et al., 2017). Similarly, interacting with a robot can be an unsettling experience for autistic children, when they do not know what to expect from the robot. What will it look like, what will it do, or how will it sound? To prevent this from happening, the autistic children will need time to get accustomed to the robot.

Three possible solutions on how to address the requirement of familiarising the child with the robot were put forward. First, an autistic adult mentioned that it would be beneficial if the children themselves could freely explore the robot, and become familiar with it, as they know best what they do and do not want (Huijnen et al., 2017). Second, in the ALCORN ET AL. STUDY, educators mentioned that creating a social story around the robot that shows what it looks like, explains what it is going to do, and when it is coming, may help the children anticipate and prepare for the robot's arrival. And lastly, providing some familiarity in appearance or behaviour might also put the children more at ease (Robins et al., 2010b).

3.4.4 Sensory hyper- and hypo-sensitivity

Unusual responses to sensory information are included as one of the non-social symptoms of autism (American Psychiatric Association, 2013). These responses vary widely between autistic individuals (Hazen et al., 2014). Some are hyper-sensitive to certain sensory experiences (e.g. strong reactions to loud or unpredicted sounds, or lights) which then cause great discomfort (e.g. feeling like a sharp needle pierces your eardrums), others are hypo-sensitive and react very slowly to certain sensory information, or are unaware of it, and there are also some that actively seek out certain sensory experiences. Autism professionals reported that taking these unusual responses to sensory information into account in the interventions could be important for keeping the children engaged (Giullian et al., 2010; Robins et al., 2010b; Huijnen et al., 2017). On the one hand, we do not want to create a sensory experience that triggers hypo-sensitivity, which is very unpleasant and may lead to disengagement. On the other hand, for those who actively seek out such sensory experiences, they may be particularly motivating.

To address sensory hyper- and hypo-sensitivities, personalising the sensory experiences is required. If the robot has lights, then it should be possible to both use them or deactivate them. If children want to feel the material of the robot, then this can be utilised to increase motivation, but it should not be required in order for the child to engage in the intervention, as this would prevent those who are sensitive to the robot's materials from engaging. Because it is often not possible to tell exactly *what* sensory experiences the child has unusual responses to, the professional using the robot will also need to pay close attention to this and intervene when necessary. In case a sensory sensitivity is triggered by something the robot is doing, then it should immediately stop doing it.

3.4.5 Personalising content and the robot

Personalising the autistic child-interaction in relation to sensory sensitivities of the children is one form of personalising that we believe is particularly important. However, personalising the autistic child-robot interaction *in general* will be important according to autism professionals. In autism education, personalising content is an essential task for autism professionals to adjust learning material to a specific child. Parents and educators noted that autistic children can have strong interests, and that utilising these interests could draw the attention of the child and keep the child engaged (Putnam et al., 2019). Not only in personalising the robot's behaviour [ALCORN ET AL. STUDY], but also its appearance (Huijnen et al., 2017). Some children may enjoy certain robot behaviours, which are particularly motivating for that child, while other children may enjoy different robot behaviours (Robins et al., 2010b). It should also be possible for autism professionals to personalise the learning content, as each child has an individual learning plan (Huijnen et al., 2017, ALCORN ET AL. STUDY) — what may be relevant for one child, is not so for another.

In robotics, much of the personalising that is done relates to personalising the difficulty of games through intelligent tutoring systems (e.g. Schadenberg et al., 2017; Ramachandran et al., 2017), which may partly address the need for personalising content and the robot. Similar systems have been applied to personalising learning content in robot-assisted interventions in terms of difficulty (Scassellati et al., 2018; Clabaugh et al., 2019) and feedback (Clabaugh et al., 2019). However, as we explained in the previous paragraph, much more personalisation, and different kinds of personalisation, are needed to effectively support the autistic children in a robot-assisted intervention. While some of the personalisation can be done by the robot (autonomously), other forms of personalisation will require the input from the autism professional (e.g. adjusting the learning content to the child's individual learning plan). The latter could be facilitated by providing the professionals with the ability to program the robot's behaviour (e.g. Barakova et al., 2013).

A different solution to the above is to give more control over the robot to the children themselves. This way, the children can choose what they enjoy (Robins et al., 2010b). Moreover, it lets the children be active participants, where *they* shape the interaction. In Robins et al. (2010b), educators suggested that simple controls on a toy (i.e. the robot) could provide the children with the means to explore the robot and control its behaviour. For instance, the KASPAR robot features several sensors on its body, which when pressed, cause the robot to react (Robins et al., 2010a).

3.4.6 Generalisation of learned skills to humans

As robots are not humans, it is essential that the targeted skill learned through interacting with the robot also generalises to human interaction. While a robot may provide certain benefits, these benefits may not contribute to this goal. For instance, educators in the ALCORN ET AL. STUDY said that robots that are *too* predictable, or *too* engaging, could potentially hinder the child's progress in learning to navigate in social environments.

To address this requirement, using a *humanoid* appearance for a robot may be

particularly successful in facilitating the generalisation of the learned skills to humans (Ricks and Colton, 2010; Scassellati et al., 2012, ALCORN ET AL. STUDY). However, for some children, “a robot is still a robot, even when it looks like a person” [ALCORN ET AL. STUDY, P.8], and this solution may not work. Instead, a more successful approach may be to actively embed the generalisation of learned skills from robots to humans into the intervention. Two approaches for doing so have been proposed in literature. One approach is where the robot is only used for eliciting certain social behaviours from the child that are directed at the professional, and learn a skill through this process (Colton et al., 2009). The skill is then already applied in the interaction with another person, circumventing the need for generalisation from robots to humans. A different approach to generalisation within intervention is by gradually fading the role of the robot in the intervention (Goodrich et al., 2012; Begum et al., 2015). The child may then first learn the skill through interacting with the robot, but later on in the interventions learns to apply this skill in the interaction with the professional.

3.4.7 Safety and robustness of the robot

Next to the safety of the people involved in the intervention, the robot itself should also be safe (Giullian et al., 2010). Autistic children may enjoy taking objects apart (Baron-Cohen, 2009), or may handle objects roughly, and a robot is unlikely to be an exception. As robots are often expensive and difficult to repair for a layperson, the robot should not be damaged during the intervention. As such, this type of behaviour needs to be accounted for through design, or through protocol, to address this requirement of having a sufficiently robust platform. Possible solutions include using a highly robust robot that cannot be taken apart without using tools, or simply dissuading the child to handle the robot roughly and intervening when this happens. Alternatively, scheduling a fixed period of time in the intervention where the children can engage in the tactile exploration of the robot could give them the satisfaction of doing so without further disrupting the rest of the intervention.

3.5 General discussion and conclusion

In this chapter, we have analysed what the current needs are of autism professionals in carrying out interventions for autistic children. These needs provide opportunities for robots that may be positioned to address those needs, and thus solve a real problem. Our analysis on user requirements shows what the important aspects are for successful interactions with autistic children, based on the expertise of autism professionals. Taking these requirements into account will be important for realising effective robot-assisted interventions. We summarise and rephrase all these requirements in Table 3.1 and provide the sources from which the requirements were extracted.

Through our analyses, three major themes emerged. Firstly, the need for a more predictable and simplistic interaction with a social actor. Robotic technology may be uniquely positioned by being able to provide more predictable, and less complex interactions, as well as being a social actor that elicits social interaction. This opens up various promising avenues for embedding a robot in interventions that target learning social skills. Not only could a predictable and simplistic robot be more easy to

Table 3.1: Summary of the user *requirements*. URP refers to User Requirement for the autism Professional, whereas URC refers to the User Requirement for the autistic Child. The sources are the articles from which the requirement was extracted. Note that the SCHADENBERG ET AL. STUDY is not included in this table, as it only related to user *needs*.

ID	User requirement	Sources
<i>As an autism professional, I require ...</i>		
[URP 1]	to be in control over the robot-assisted intervention.	d, h, j
[URP 2]	(partial) control over the robot.	a, e, f, h, j
[URP 3]	to personalise the robot-assisted intervention to accommodate the child's individual learning goals.	g, j
[URP 4]	to personalise the robot-assisted intervention to the child's interests.	c, g, i, j
[URP 5]	that the learned skills generalise from the robot to humans.	j
[URP 6]	that the learned skills generalise to settings outside the robot-assisted intervention.	j
[URP 7]	that the robot is sufficiently robust to withstand harsh handling.	a
<i>As an autistic child, I require ...</i>		
[URC 1]	to be in a predictable environment.	e, j
[URC 2]	not to feel pressured to perform.	j
[URC 3]	to get accustomed to the robot before I start with the learning content.	c, g
[URC 4]	my sensory hyper-sensitivities not to be triggered by the robot.	a, c, g
[URC 5]	my sensory hypo-sensitivities not to distract me during the robot-assisted intervention.	a, c, g

^aGiullian et al. (2010); ^bRobins et al. (2010a); ^cRobins et al. (2010b);

^dCoeckelbergh et al. (2016); ^eHuijnen et al. (2016); ^fZubrycki and Granosik (2016); ^gHuijnen et al. (2017); ^hMartin-Ortiz et al. (2017); ⁱPutnam et al. (2019);

^jALCORN ET AL. STUDY.

understand and comfortable for autistic children, it could also be less threatening, as it may be perceived by the children as being less socially demanding than a person. Future research will have to investigate whether this is the case, as well as how we can design robots to be highly predictable and still add meaning to an intervention.

A second theme that emerged was that autistic children are very different from each other. Accounting for these differences will be essential, but also challenging, because it is not always clear how the robot-assisted intervention should be adapted to the child. For instance, the use of lights could be anywhere from being highly motivating to causing great discomfort. It is unlikely that there is a one-size-fits-all robot for autistic children. Some children may not enjoy interacting with the robot, as they may be fearful of it (Putnam et al., 2019), or may not think the robot is “cool” (Huijnen et al., 2017). Others may be too aggressive to interact with a robot. In the end, the autism professional will need to decide how to use the robot and for which

children.

This brings us to our last theme, which is that the robot is to be a tool for autism professionals, who can use the robot in certain scenarios. For instance, by using the robot as a scaffold, to bridge the gap between learning with current materials and learning with people. This also means that the professionals should be empowered by having a robot at their disposal, which requires that they should remain in control over the intervention and be able to use the robot as they see fit. To enable the professionals to adjust and customise the robot's behaviour, they will need easy-to-use tools to program the robot. What these tools should look like, and how the professionals can control the robot in a session — taking into account that it should not cause additional workload during sessions — are questions that will need to be addressed in future research.

3.5.1 Limitations

In the papers that came out of the literature search, none of the papers involved autistic children in their search for user needs or requirements. This is unfortunate, as autism interventions are designed for *them*, but we do not know what needs and requirements they themselves report. Including autistic children in the design process would allow us to create more suitable and acceptable technologies, as well as allow the children to aid in shaping the robot-assisted intervention according to their needs and desires (Spiel et al., 2019). Frauenberger et al. (2011) argues that autistic children are rarely involved in the design process, because either the researchers have limited access to the target group, or hold views that autistic children may be impaired in their creative and communicative skills, limiting their potential to provide feedback. While it may seem difficult to involve autistic children with low language and cognitive ability, there are ways for doing so. For instance, researchers could use a combination of ethnography and structured observations to understand their experiences (e.g. Pellicano et al., 2014). Alternatively, autistic children who are farther along in their development could be involved in a study — for whom there are many methods to engage them in research (see Frauenberger et al., 2017; Spiel et al., 2017) — and could represent autistic children with more difficulty in their communication and cognition.

To limit the scope of our literature review, we only included studies that were related to *robot-assisted* interventions for autistic children. However, user needs do not specify through what means they can be addressed; only that users would like them to be addressed. A broader search for the user needs of autistic children and autism professionals would therefore likely result in more articles describing their needs. These may also be relevant for developing robot-assisted interventions.

While each paper reported on the user needs of end-users of robot-assisted interventions, some did so only briefly and provided little background regarding the user needs (Giullian et al., 2010; Zubrycki and Granosik, 2016; Kim et al., 2019). This resulted in various needs being reported, but not clearly explained what they were based on, nor why they cannot be met, or are met insufficiently so, through the current means available to autism professionals. We therefore left out several user needs that were only briefly mentioned in one of the papers, and instead focused on user

needs which were supported by statements from multiple papers. Given that we only found a few studies that investigated user needs, of which some were poorly reported, the validity of the user needs can be questionable. We believe it is likely that they are highly susceptible to regional and cultural differences. For instance, highly developed English-speaking countries, will likely have less difficulty with finding and applying standardised tests. Further research is required to verify whether the user needs are more widely shared between autism professionals and between countries. Nevertheless, for the DE-ENIGMA consortium, we believe the user needs are sufficiently valid, as they specifically relate to, and involved, our end-users (educators from special education schools, occupational therapists, and parents of autistic children), who were based in countries that are directly relevant for the testing of the DE-ENIGMA intervention (United Kingdom, Serbia, and the Netherlands).

3.5.2 Focus of the DE-ENIGMA intervention

These user needs and requirements form the basis of how we developed the DE-ENIGMA intervention. In particular, we aimed to position the DE-ENIGMA intervention to address the need for scaffolding and the need for a safe and comfortable learning environment. For the latter need, we specifically focused on developing a robot that was highly predictable in its behaviour. The DE-ENIGMA intervention also addressed the need for reporting the progression of autistic children, but I only briefly touch on this aspect of the intervention in this dissertation (Chapter 5.5), as I was not involved in the design of the solution. While all the user requirements that we established are likely to be important for an effective and engaging robot-assisted intervention, we focused on achieving a couple of them. These include URP 1 to 3, and URC 1 to 5 (see Table 3.1). In the next chapter, I will report on a descriptive study where we investigated how autistic children spontaneously interact within a robot-assisted intervention setting. Identifying such interactions can help us choose between the various ways of addressing the user requirements we identified.

4

Spontaneous interaction of autistic children and the role of autistic traits

This chapter is based on the following article:

Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., & Evers, V. (2020). Differences in Spontaneous Interactions of Autistic Children in an Interaction With an Adult and Humanoid Robot. *Frontiers in Robotics and AI*, 7, 19 pages. ([Schadenberg et al., 2020b](#))

It is important to note that the user requirements from the previous chapter do not state how they can be translated into design, or what the autistic child-robot interaction in general should look like. While we elaborated on possible solutions, there are more ways in which each user requirement can be addressed. Given the wide variety in abilities, interests, and needs between autistic children, the most effective way for addressing the user requirements likely also varies between children. While there are a few studies that provide qualitative descriptions of how autistic children interact with robots — or in the presence of a robot — and how these differ between children, these studies are with small sample sizes. This makes the results difficult to generalise to broader subgroups of autistic children. Our limited understanding of how individual differences affect the interaction within a robot-assisted intervention and what autistic children find interesting types of interaction formed the basis of the descriptive study that is presented in this chapter.

4.1 Introduction

INCORPORATING a robot in an intervention for autistic children appears to have a positive effect on the child's engagement and attention to the learning task ([Scassellati et al., 2012](#); [Simut et al., 2016](#)). Yet despite the promising findings related to autistic children engaging in robot-assisted activities, achieving sustained

engagement in a robot-assisted intervention that can lead to learning remains challenging. While studies on robots for autistic children generally report a positive effect of the robot on engagement, they also report on children who show very low levels of engagement or are not engaged (Rudovic et al., 2017; Desideri et al., 2018), or quickly lose attention within a session (Tapus et al., 2012). Moreover, sustaining high levels of engagement over multiple sessions is difficult (Desideri et al., 2018), where initially engaging interactions can become boring and too repetitive for autistic children over time (Srinivasan and Bhat, 2013). Additionally, some interactions may be very rewarding to the child and keep them engaged, but do not facilitate learning a targeted skill. Indeed, even when the children are engaged, current robot-assisted interventions do not necessarily lead to learning (Tapus et al., 2012; Kim et al., 2013; Pop et al., 2014; Simut et al., 2016; Desideri et al., 2018; Zheng et al., 2020). In part, these issues can be explained by individual differences among autistic children in what they do and do not consider interesting, as well as the idiosyncrasies in autism features that we described in the previous chapter (Section 2.4). Altogether, a personalised approach to design of robot-assisted intervention is essential for autistic children to engage in it and learn; the robot's behaviour needs to be in line with the interests, needs, and abilities of the child.

To design deliberate robot behaviour, aimed at engaging autistic children and facilitating learning, we then need to understand what interactions are interesting to them and how their individual characteristics play a role in this. There is a lot of work that provides qualitative descriptions of how autistic children interact with a specific robot (e.g. Kozima et al., 2009; Feil-Seifer and Matarić, 2009; Robins et al., 2009; Tapus et al., 2012; Costa et al., 2015; Boccanfuso et al., 2016) which can provide some guidance in the design of a robot-assisted intervention. For instance, in their study with the machine-like robot Sphero, which looks like a billiard ball, Boccanfuso et al. (2016) observed the responses of autistic children in relation to the robot's various expressions of emotions. The robot's behaviour was limited to rolling as a way to move, using its LEDs to change colour, and playing music. Responses varied from pushing, kicking, dropping, holding, and picking up the robot. Tapus et al. (2012) used the humanoid robot NAO and reported a detailed description of the interactions of four autistic children. Spontaneous interactions between the child and robot were observed, such as the child touching the robot. NAO also elicited interactions between the child and the experimenter, where children requested certain robot behaviours, or shared their enjoyment with the experimenter. We consider spontaneous behaviours, such as those reported above, to be indicative of potential for engagement. The behaviours are spontaneous, which means that there is no observable prompt that led to the behaviour (contrary to responsive behaviours which require a prompt), which can indicate that the children are intrinsically motivated to engage in such behaviour (Deci and Ryan, 1985). That is, they are motivated to perform the behaviour for its inherent satisfaction rather than for some separable consequence (Ryan and Deci, 2000). In turn, designing deliberate robot behaviour to support intrinsically motivating interactions may be particularly promising in keeping them engaged. However, our current understanding is insufficient to translate reported insights into design. Studies reporting on specific autistic child-robot interactions generally have less than

10 participants (Begum et al., 2016), and are qualitatively rich, but largely not quantitative. Moreover, participant characteristics are often not, or insufficiently, reported in studies on robots for autistic children (Diehl et al., 2012; Begum et al., 2016). This makes it difficult to ascertain the abilities and difficulties of the participants and generalise findings to more specific subgroups of autistic children to whom we can tailor the robot's behaviour.

In the descriptive study described in this chapter, we report analyses on the behaviour of autistic children in the context of a robot-assisted intervention. The interactions were collected as part of the DE-ENIGMA database⁹. This database is one of the major outcomes of the DE-ENIGMA project and is a publicly available multi-modal database of autistic children's interactions. This database hosts the largest set of recordings to date of autistic children ($N = 121$) engaged in a robot-assisted intervention featuring a humanoid robot. The analysis reported in this chapter specifically reports on interactions that were *spontaneously* initiated by the children. We aim to address the following two research questions:

Research question 1: What interactions did the autistic children spontaneously engage in within the DE-ENIGMA robot-assisted intervention?

Research question 2: Which autistic traits predicted these various interactions?

Specifically, we analysed the children's spontaneous interactions with the robot, where we classified the interactions as either exploratory, relational, or functional/symbolic interactions. We also analysed the children's spontaneous interactions towards the adult — social overtures — that made a reference to the robot in some way. To investigate whether certain interactions are of interest to particular groups of autistic children, we used the scores of various autism diagnostic assessments and demographics to differentiate among children, described in Section 4.2.6. By answering the above two research questions, we hope to deliver concrete insights into what robot behaviours to design for engaging specific groups of autistic children in a robot-assisted intervention.

4.2 Materials and methods

4.2.1 Dataset

DE-ENIGMA database. As we participated in the collection of audio and video recordings to develop the DE-ENIGMA database, we used the data collected for the database for our analyses (for a general description of the DE-ENIGMA database, see Chapter 1.4). The data collection for the DE-ENIGMA database took place in Serbia and the United Kingdom. In our analyses, we only consider the recordings from the United Kingdom to keep the cultural context similar between the recordings. The children were recruited from three special education settings in the United Kingdom, which were also the locations where the sessions took place. All of the children had

⁹Paper in preparation. See <https://deenigmadb.wordpress.com/>

Figure 4.1: Screenshot from one the sessions in the United Kingdom, where the adult is using the robot for the DE-ENIGMA learning task. The adult at the back is a school staff member who accompanies the child.

received an independent clinical diagnosis of autism according to criteria of the ICD-10 ([World Health Organization, 1992](#)), DSM-IV-TR ([American Psychiatric Association, 2000](#)), or DSM-V ([American Psychiatric Association, 2013](#)). The majority of the autistic children had additional intellectual disabilities and language challenges.

Ethical approval for the data collection in the United Kingdom for the DE-ENIGMA database was reviewed and approved by the ethics committee of the University College London, Institute of Education, and is registered under reference number “REC 796”. For some of the autistic children who participated, the parents only granted consent for using their child’s data within the project, and not for inclusion in the publicly available database.

Data selection. The video recordings used for the analyses presented in this chapter are a *subset* of the video recordings that were collected for the DE-ENIGMA database. The subset of the video recordings used in our analyses only includes the recordings of the autistic children from the *United Kingdom* who participated in the *robot-assisted condition* (see Figure 4.1 for a screenshot of one of the video recordings). Our subset includes the data of three autistic children for whom the parents only granted consent for their data to be used within the project and are not included in the database. Furthermore, our analyses use standardised tests for assessing autistic traits that were part of the data collection, but could not be included in the database due to the ethical constraints — they can only be used within the DE-ENIGMA project.

The children interacted with the robot for several sessions every or every other

Figure 4.2: The facial expressions of emotion of Zeno designed by Robokind.

day. The number of sessions for the children depended on their availability and progression through the learning content. For each child, one of their sessions was randomly selected for the analysis. A description of the children included in our sample is reported in Section 4.3.1.

4.2.2 Robot behaviour

To express the emotions, the robot used the default facial expressions designed by Robokind (see Figure 4.2). To make it easier to recognise the facial expressions, we augmented them with affect bursts. These are “short, emotional non-speech expressions, comprising both clear non-speech sounds (e.g. laughter) and interjections with a phonemic structure (e.g. “Wow!”), but excluding “verbal” interjections that can occur as a different part of speech (like “Heaven!”, “No!”, etc.)” (Schröder, 2003, p. 103). To assess whether this actually led to improved recognition rates, we conducted a small study with 28 typically developing children. For them, combining the affect bursts with the facial expressions led to an increase in recognition rate of 15% on average (Schadenberg et al., 2018).

In addition to the facial expressions, the robot had several gestures which could be used to (attempt to) elicit joint attention, various behaviours to draw the attention or to reward the child, and behaviour for saying “hi” and “goodbye”. When idle, the robot showed life-like behaviour by moving its wrists, blinking its eyes, and turning its head every now and then.

The Wizard-of-Oz paradigm was used for operating Zeno, where the adult who was giving the intervention also controlled the robot. The Wizard-of-Oz interface was a small keypad that was hidden underneath the table on which Zeno was standing. A cloth over the table further obscured the keypad from the child’s sight. The keypad contained controls for all of the robot behaviours, except the life-like behaviour which ran autonomously.

4.2.3 The activities used for the DE-ENIGMA data collection study

Prior to attending the first session, the autistic children received a social story, which are generally used as a narrative that is made to illustrate certain situations and prob-

lems and how people deal with them. In this case, the social story illustrated the session, which included specific information about what to expect (e.g., who they would see, what they would do). This helped the children prepare for what would otherwise have been an unfamiliar situation. The DE-ENIGMA sessions took place at the child's special education school in a separate room that was available for all sessions. The children were accompanied by their school staff member, whose main role was to provide additional support if needed but not to participate in the teaching. Each child started the first session with a free-play activity with toys to help them become comfortable with the unfamiliar adult and setting. The interaction with the robot started with a brief introduction to Zeno, where the adult would display the various behaviours of Zeno to familiarise the child with the robot. Next, the adult would demonstrate each of Zeno's dynamic facial expressions of emotion and label them. The children were then guided to work through an adapted version of the six steps defined by [Howlin et al. \(1999\)](#):

1. *Recognising static emotional expressions*: recognising Zeno's emotional facial expressions as depicted on laminated cards.
2. *Recognising abstract static emotional expressions*: recognising emotions expressed by emoticons on laminated cards.
3. *Recognising dynamic emotional expressions*: matching Zeno's dynamic emotional facial expressions with static emoticons on laminated cards.
4. *Recognising dynamic emotional expressions*: similar to step 3, but the child was also asked to express the shown emotion.
5. *Recognising dynamic emotional expressions*: similar to step 3, but the child was also asked to express the same emotion as the robot (no label of the emotion was provided).
6. *Recognising situation-based emotions*: recognising how the character (the robot, the child, or another child) in a social story feels. The social stories start out as simple situation-based stories, and gradually move towards incorporating desires and beliefs.

These steps were adapted to incorporate Zeno in the following way: instead of photographs of real people for step 1, photographs of Zeno were used; and instead of the adult showing dynamic facial expressions for step 3-6, Zeno was dynamically animated to display these expressions with its face. In each of the steps, Zeno expressed the emotions and provided positive feedback (either for the correct answer, or for the child's effort). At the end of the interaction with the robot, Zeno would say "goodbye" to the child.

The children engaged in the robot-assisted intervention for several sessions, scheduled over multiple days. The exact number and length of the sessions depended on the child's progress through the intervention, and additional factors such as their attention span. The intervention described above varied somewhat among the children, as the learning content was adapted to the child's behaviour and language by

the adult. The children were finished with the intervention when they had completed all six steps, or when they were not able to successfully complete a step after three separate attempts.

4.2.4 Coding scheme

For the analysis, we observed the interactions of the autistic children during the sessions with the robot. These were then annotated with the ELAN transcription software¹⁰, developed by the Max Planck Institute for Psycholinguistics in Nijmegen, the Netherlands.

We annotated interactions that met the following two criteria: firstly, the interaction should be *spontaneous*, which means that the child initiated the interaction without a prompt from the adult. For example, the child could ask the robot a question, or reach for the robot to explore its physical properties. Any responses or answers given by the children in response to a specific prompt by the adult were excluded. For example, the children were frequently required to choose an emotion they recognised, and they learned this through prior prompts by the adult. Secondly, the interaction should be *directed* towards either the robot, or the adult. For the latter, the interaction should also be related to the robot in some form (i.e. the behaviour would not have occurred if no robot was present). For example, asking the adult a question about the robot. Any spontaneous interactions directed towards the school staff member were not annotated, as they were asked not to be actively involved in the intervention. Potentially relevant interactions can occur at any point during the session, not only during the child-robot interaction (Dickerson et al., 2013). We therefore annotated the whole recording, which included parts where the robot was still covered by a blanket, as well as when the adult and child were engaged in free-play.

We used a grounded theory approach to design a coding scheme with which we analysed the observed child behaviours (Saldaña, 2015). The coding scheme that we used for the analysis to annotate the spontaneous interactions of the children can be seen in Table 4.1. All observations were placed in a behavioural unit, which describes how the interaction manifested itself. In turn, each behavioural unit is part of a categorical unit (in bold in Table 4.1), which describes what type of interaction the manifestation belongs to. For the robot directed interactions, the categorical units were based on the developmental types of object play (Libby et al., 1998; Casby, 2003; Naber et al., 2008). They include exploratory, relational, and functional/symbolic interactions. We combined functional and symbolic interactions into one categorical unit because the distinction between the two is precarious when it comes to robot (see Section 2.3.2). The exploratory interactions with the robot were restricted, due to the fragility of the robot. Beyond gently touching and visually inspecting the robot, other tactile interactions were actively prevented by the adult and the school staff member. These include possible exploratory interactions such as banging the robot on the table, or mouthing the robot. In addition to the child interactions with the robot, we also annotated any the children's interactions towards the adult which related to the robot in some way. Such interactions were categorised as social overtures – a spontaneous social initiation of the child towards the adult.

¹⁰<https://tla.mpi.nl/tools/tla-tools/elan/>

Table 4.1: Description of the coding scheme for annotating the autistic child's interactions with the robot and robot-mediated interactions with the adult.

<i>Interaction type</i>	<i>Description and examples</i>
Exploratory interactions with the robot	
Touching the robot	Touching part of the robot. This also includes attempts at touching Zeno, as children were prevented from actually touching it. An attempt is classified as moving towards Zeno and reaching for it, or when the intent is clear from previous attempts at touching Zeno.
Explicit visual inspection of the robot	Moving closer to the robot, or leaning over the table, and visually inspecting the robot.
Relational interactions with the robot	
Using additional objects with the robot	Relating one or more objects to the robot in a way that does not indicate functional or symbolic play. For instance, putting a laminated card on top of the robot or in its mouth.
Functional interactions with the robot	
Imitation of the robot	Child imitates Zeno's behaviour or the sounds it makes without being prompted. This excludes echolalic responses. For instance, imitating the robot's waving gesture.
Joint attention initiated by the robot	Gaze shift to area where Zeno is looking/pointing at and is saying "look".
Talking to the robot	Any questions, comments or vocalisations directed at Zeno. For instance, asking the robot about its favourite food.
Physical behaviours directed at the robot	Physical interactions with Zeno, which are not exploring behaviours. For instance pushing Zeno, blowing air in Zeno's face, or dancing with Zeno.
Controlling the robot	Making the robot do a certain behaviour through the Wizard-of-Oz keypad that controls the robot. This includes reaching for, or pressing, controls on the keypad.
Social overtures directed at the adult and related to the robot	
Social reference	A social reference refers to utilising the adult's interpretation of a novel situation to formulate one's own interpretation (Feinman, 1982). For instance, looking at the adult after Zeno did something unexpected.
Shared enjoyment	Indicating and communicating pleasure to the adult regarding something the robot did. For instance, looking and smiling at adult after Zeno finished its dance.
Directing attention	Getting the adult to direct her attention to Zeno. For example, the child may be pointing towards Zeno and saying "look".
Helping	Helping the adult with something that has to do with Zeno. For instance, covering or uncovering Zeno with a blanket.

Table 4.1: Table continued.

<i>Interaction type</i>	<i>Description and examples</i>
Requesting	(non) Verbally requesting the adult to have Zeno do certain behaviours. For example, the child could ask the adult to see the robot's happy facial expression by saying "do happy", or point to a laminated card that shows a certain facial expression of Zeno.
Conversing with the adult	Talking with the adult about the robot. For instance, asking questions about Zeno, commenting on certain features of Zeno.

4.2.5 Annotation procedure

For all observations, we annotated the start and end time of the observations. Observations that occurred within 2 seconds of each other were considered as the same observation. A single main coder annotated all the recordings. To calculate the reliability of these annotations, a second coder annotated a random selection of 20% of the recordings, which contained 19% of the annotations of the first coder. There were 21 segments in the recordings that were annotated by both coders. To determine the agreement between the two coders, Cohen's κ statistic was used. The agreement between the two coders on the behavioural units was good (Cohen's $\kappa = .83$, 95% CI [.66, 1.00], $p < .001$). However, there was a difference in sensitivity between the two coders, as an additional 24 segments were annotated by only one of the two coders. Of these additional segments, 21 were coded only by the main coder and 3 were coded only by the second coder.

As an additional check, the second coder annotated the 21 segments that were only annotated by the first coder. For these additional observations, the second coder also had to judge whether the child's behaviour was "spontaneous", "might have been spontaneous", or "was not spontaneous". The reason for the additional annotation of these 21 segments was given afterwards to the second coder, to prevent socially desirable responses. Of the 21 additional segments, 18 were judged to be "spontaneous", 3 as "maybe spontaneous", and none as "was not spontaneous". The second coder further mentioned that the child's behaviour in these segments were more subtle than in the segments that both coders initially annotated. For the additional segments, agreement was good (Cohen's $\kappa = .72$, 95% CI [.51, .93], $p < .001$); the aggregated Cohen's κ over all 42 segments is .78 (95% CI [.65, .92], $p < .001$), which we consider to be sufficiently high to proceed with.

We conclude that there was only a difference in sensitivity between the coders, where the main coder was more sensitive than the second coder, and not in differences in labelling segments. Both coders agreed that most of the additional segments were indeed spontaneous interactions and should be annotated. Difficulties with sensitivity in the behavioural analysis of autistic children is a known and common issue. For instance, in the development of diagnostic tools for autism, where subtle but potentially meaningful social communication behaviours are difficult to identify, which makes it difficult to develop a measure that is sensitive enough to account for these

behaviours (Grzadzinski et al., 2016; Anagnostou et al., 2015).

We also carried out a more detailed analysis of the coder disagreements through inspection of the confusion matrices. This showed two behavioural units that deserve special mention, namely “conversing with the adult” and “requesting”. These two units had a relatively high confusion (4 out of 15). This means that we should be careful about the distinction between these two particular units. Further results presented in this article are based on the annotations of the main coder.

4.2.6 Individual characteristic measures

To investigate which individual characteristics were associated with the observed interactions, we used three diagnostic measures and child demographics. These measures were collected for the children who participated in the recording session for the DE-ENIGMA database, but are not publicly available as part of the database as it is released.

Autism Diagnostic Observation Schedule. The Autism Diagnostic Observation Schedule - second edition (ADOS-2, Lord et al., 2012) is a structured play session conducted by a professional, and was administered to assess the autistic features. Each child is given one of five modules, each with their own activities for the play session. The module is primarily selected on the basis of the child’s expressive language capabilities and secondarily on the child’s chronological age. Module 1 is used for children older than 31 months who do not consistently use phrase speech. To account for the differences in cognitive and adaptive functioning (Bal et al., 2016), Module 1 distinguishes between two expressive language levels, namely “few to no words” (hereafter, “Module 1, FNW”) for children who used no words or fewer than five words during the ADOS administration, and “some words” (hereafter, “Module 1, SW”) for those who used more than five words up to those who used simple phrases (Gotham et al., 2007). Module 2 is used when children can use phrase speech, but are not yet verbally fluent, and Module 3 is for verbally fluent children and young adolescents. The other two modules were not applicable to our sample given the chronological age requirements.

The ADOS-2 Calibrated Severity Score (ADOS-2 CSS, Gotham et al., 2009) is the raw ADOS-2 score controlled for the chronological age and language skills. The CSS is therefore a more meaningful score for comparing scores across modules. A score of 1-2 is interpreted as minimal-to-no evidence, 3-4 as low, 5-7 as moderate, and 8-10 as high in autistic features.

Childhood Autism Rating Scale. The Childhood Autism Rating Scale - second edition (CARS-2, Schopler et al., 2010) is a 15-item autism screening and diagnostic tool and was administered to obtain a general measure of characteristics of autism. It is completed based on direct behaviour observation by a professional as well as reports from parents, teachers, or caretakers. The total score on the CARS-2 reflects the severity of autistic features with scores of 15.0-29.5 indicating minimal-to-no evidence, 30.0-36.5 is mild-to-moderate severity, and 37.0 and higher is severe autistic features.

Vineland Adaptive Behavior Scales. The Vineland Adaptive Behavior Scales - second edition (VABS-2, Sparrow et al., 2005) is a standardised measure of an individual's adaptive behaviour – the ability to undertake daily activities. Adaptive behaviour is a composite of five domains, namely the communication, daily living skills, socialisation, motor skills, and maladaptive behaviour domains. In this article, we report on the communication and socialisation domain scores. The former addresses receptive as well as expressive language usage, and the latter reflects functioning in social situations. The scores on the domains are standard scores ($M = 100$, $SD = 15$). For descriptive purposes, we also report on the Adaptive Behaviour Composite, which reflects an individual's overall adaptive behaviour, and is calculated using the domain scores.

Child demographics. For the child demographic characteristics, the children's chronological age and sex were included in the analyses.

4.3 Results

4.3.1 Description of the sample

The sample for the analysis consisted of 31 (84% male) autistic children from the United Kingdom, between the chronological age of 5 to 12 years. They were randomly assigned to the robot-assisted condition for the DE-ENIGMA database recordings. The average characteristics of the children can be seen in Table 4.2.

The ADOS-2 assessment was completed for all but one of the children included in the sample. For the ADOS-2 assessment, Module 1 was used for 24 children, of which 10 used few-to-no words and 14 used some words. There were six children who had phrase speech, but were not yet verbally fluent, for whom Module 2 was used. There were no children for whom Module 3 was deemed appropriate. The child for whom

Table 4.2: Average characteristics of the thirty-one autistic children who were included in the analysis.

Measure	<i>n</i>	Mean (<i>SD</i>)	Range
Age (years:months)	31	8:4.4 (2:2.7)	5:1 - 12:2
Sex			
Male	26	-	-
Female	5	-	-
ADOS-2 CSS			
Module 1, FNW	10	6.60 (1.84)	3 - 9
Module 1, SW	14	6.29 (0.83)	5 - 8
Module 2	6	6.33 (1.86)	3 - 8
CARS-2	31	33.68 (4.52)	24.5 - 45.0
VABS-2			
Adaptive Behaviour Composite	20	55.45 (8.48)	46 - 74
Communication domain	20	58.00 (12.55)	36 - 79
Socialisation domain	20	55.60 (7.34)	43 - 68

there is no ADOS-2 score was unable to participate in the ADOS-2 play session as he would not engage with the examiner that conducted the assessment. The CARS-2 assessment was completed for all children. All children scored above the autism cutoff on either the CARS-2 (30 or higher) or the ADOS-2 (4 or higher). On average, the children had moderate autistic features. The VABS-2 was assessed through survey interviews with the child's parents, and was completed for 20 children.

4.3.2 Observed interaction types

The analysis led to a total of 225 annotations in 450 minutes of video recordings. The sessions that were randomly selected for the analysis lasted from 5 min 55 s up to 37 min 12 s. On average, the sessions lasted 14 min 31 s ($SD = 8$ min 02 s). In our sample, the session ranged from the 1st to 7th sessions ($M = 3.03$, $SD = 1.43$). Twelve sessions contained the free-play activity, which lasted 6 min 27 s on average ($SD = 3$ min 05 s). Of the annotations, only seven occurred during the free play prior to the intervention with the robot, when the robot is still covered by a blanket.

The frequency and distribution of the observed interaction types can be seen in Figure 4.3. Most of the spontaneous interactions with the robot classify as functional interactions ($n = 71$), and were observed for 58% of the children. Exploratory interactions with the robot ($n = 57$) were observed for 42% of the children. For two children (6%), relational interactions with the robot were observed for a total of six annotations. A total of 91 social overtures were observed, spread out over 16 children (52%).

Notably, for eight children we did not observe *any* spontaneous interactions. They were all male and had moderate to severe autistic features. Of these eight children, two had an aversive reaction to the robot and left the room shortly thereafter. Both were assessed with ADOS-2 Module 1, SW. For one it was the first session, and for the other it was the second session. Four children did not seem to understand nor engage in the learning task. The adult spent most time trying to explain the task and focus the child's attention to the task. All of the four children were non-verbal. Two children engaged with the learning task and interacted with the robot and the adult when prompted, but did not initiate any interactions themselves. Both children were assessed with ADOS-2 Module 1, where one used some words and the other used few-to-none.

4.3.3 Individual differences in interaction types

Correlations between the autistic traits and interaction types. To calculate the association between the measured autistic traits and interaction types we use Kendall's Tau-b partial correlation, given the large amount of tied ranks, skewed distribution, and outliers (see Figure 4.3). The partial correlations were controlled for differences in session length and can be seen in Table 4.3. The ADOS-2 Module used during the ADOS-2 assessment showed positive associations with spontaneous functional interactions ($\tau_p = .39$, 95% CI [.16, .58], $p = .003$), social overtures ($\tau_p = .63$, 95% CI [.45, .76], $p < .001$), and the total number of spontaneous interactions ($\tau_p = .41$, 95% CI [.18, .60], $p = .002$). The correlations for the communication domain score

Figure 4.3: Raincloud plots (scatter, box, and density plot) that show the frequency of the interaction types observed per autistic child.

of the VABS-2 were not significant, but did show a similar trend to the ADOS-2 Module correlations. We found no evidence of an association between either the ADOS-2 Module or VABS-2 CD and spontaneous exploratory or relational interactions.

The ADOS-2 CSS was negatively associated with spontaneous functional interactions ($\tau_p = -.35$, 95% CI $[-.55, -.11]$, $p = .008$), while the CARS-2 score showed a negative association with social overtures ($\tau_p = -.28$, 95% CI $[-.49, -.04]$, $p = .027$). No evidence was found for either score regarding an association with exploratory or relational interactions, nor with the total number of spontaneous interactions.

The child's functioning in social situations, as measured through the VABS-2, showed a positive association with the spontaneous functional interactions ($\tau_p = .33$, 95% CI $[.02, .58]$, $p = .047$), social overtures ($\tau_p = .35$, 95% CI $[.04, .60]$, $p = .036$), and total number of spontaneous interactions ($\tau_p = .34$, 95% CI $[.03, .59]$, $p = .040$).

To calculate whether the child's sex was associated with initiating different types of interactions, we used the Mann-Whitney U test. Being male was coded as 0 and female as 1. We found no evidence of an association between sex and exploratory interactions ($U = 47.00$, $p = .313$, $r = .17$), relational interactions ($U = 60.00$, $p = 1.000$, $r = -.19$), functional interactions ($U = 56.00$, $p = .636$, $r = .09$), social overtures ($U = 63.00$, $p = .934$, $r = -.02$), or the total number of spontaneous interactions ($U = 48.00$, $p = .376$, $r = .17$). To account for differences in session length, we took the number of interactions per minute.

Table 4.3: Kendall's Tau-b partial correlations that show the strength of the association between the frequency of interaction types and the total number of spontaneous interactions with the children's chronological age, ADOS-2 Module, ADOS-2 Calibrated Severity Score (CSS), CARS-2 score, and the VABS-2 communication (CD) and socialisation domain (SD). The partial correlations were controlled for differences in session length.

Measure	Interaction type				Interactions total
	Exploratory	Relational	Functional	Social overture	
Age	-.12	-.05	-.13	.08	-.11
Language ability					
ADOS-2 Module	-.05	.08	.39**	.63***	.41**
VABS-2 CD	-.16	.11	.32	.31	.21
Autism spectrum severity					
ADOS-2 CSS	.14	-.01	-.35**	-.22	-.19
CARS-2	.06	-.03	-.18	-.28*	-.21
Social functioning					
VABS-2 SD	-.06	.02	.33*	.35*	.34*

* Correlation is significant at the .05 level (2-tailed).

** Correlation is significant at the .01 level (2-tailed).

*** Correlation is significant at the .001 level (2-tailed).

Novelty effect. In our selection of the participants' video recordings, we randomly selected a session for each participant. To check whether there was a novelty effect for the interaction types, we calculated correlations between the session number and the interaction types. Given the large amount of tied ranks, skewed distribution, and outliers, we again use Kendall's Tau-b correlation. The session number showed a positive association with exploratory ($\tau = .09$, 95% CI [-.23, .38], $p = .566$), relational ($\tau = .13$, 95% CI [-.04, .33], $p = .429$), functional ($\tau = .22$, 95% CI [-.07, .49], $p = .136$), social overtures ($\tau = .26$, 95% CI [-.09, .55], $p = .083$), and total number of spontaneous interactions ($\tau = .19$, 95% CI [-.10, .43], $p = .174$). None of these correlations were statistically significant.

Intercorrelations. The correlations between the individual characteristics can be seen in Table 4.4. Given the differences in type of data between the measures, we report either Kendall's Tau-b, due to small sample size and non-normality, or point-biserial correlations for sex. The likelihood ratio between sex and the ADOS-2 Module showed no significant difference ($\chi^2(2) = 0.14$, $p = .934$).

4.3.4 Manifestations of the interaction types

Below follows a qualitative analysis describing how the interaction types were manifested during the robot-assisted intervention and how they differed between the children. The frequencies of the manifestations for each interaction type can be seen in Table 4.5. The frequencies in this table are the frequencies for an average session

Table 4.4: Correlations between the independent variables of age, sex, ADOS-2 Module, ADOS-2 Calibrated Severity Score (CSS), CARS-2 score, and the VABS-2 communication (CD) and socialisation domain (SD). The correlations are Kendall's Tau-b correlations, except for correlations with sex, which are point-biserial correlations.

	<i>Age</i>	<i>Sex</i>	<i>ADOS-2 Module</i>	<i>ADOS-2 CSS</i>	<i>CARS-2</i>	<i>VABS-2 CD</i>	<i>VABS-2 SD</i>
Age	1	-.09	.09	.09	.10	-.21	-.30
Sex		1	n/a	-.20	-.02	-.22	.03
ADOS-2 Module			1	-.11	-.34*	.60**	.45*
ADOS-2 CSS				1	.40**	-.40*	-.45*
CARS-2					1	-.30	-.23
VABS-2 CD						1	.58**
VABS-2 SD							1

* Correlation is significant at the .05 level (2-tailed).

** Correlation is significant at the .01 level (2-tailed).

length. We grouped the children by the ADOS-2 Module that was used, given that the children's language ability, as measured by the ADOS-2 Module, showed the strongest association with the interaction types (see previous section).

Manifestations of exploratory interactions. The exploratory interactions that we observed involved explicitly inspecting the robot and touching various parts of it. These two manifestations of exploratory interactions were observed for children who saw the robot for the first time as well as for children who had interacted with the robot up to six times before. Sometimes the visual inspection would precede reaching for the robot, but more often the children would immediately reach for the robot. The robot stood at the rear of a table, which prevented the children from touching the robot when they were seated. Of the children, eight stood up to get a closer look at the robot, either by leaning forward over the table, or by walking to the rear of the table and standing next to the robot. Two of them persisted in visually inspecting the robot and accounted for six and ten annotations. Eleven children touched or attempted to touch the robot. Most of these children did so several times during a session. When the children reached for the robot, the adult and school staff member intervened, as it was often not possible to tell beforehand whether the child would gently touch the robot or would grab the robot to explore its properties (e.g. through licking, spinning or banging the robot on the ground), which could damage the robot.

Manifestations of relational interactions. The number of other objects that could have been used for relational interactions was limited to the learning materials and any items the children brought with them. Two children placed the laminated emotion cards in the robot's mouth after it had opened the mouth for a certain animation. When the robot closed its mouth, the card would stay clenched in the robot's mouth. One child in particular found this type of interaction interesting and accounted for five of the six observations.

Table 4.5: Frequency and range for the interaction types and the observed manifestations, indicating how frequently certain interactions and manifestations were observed. Given that the ADOS-2 Module showed the strongest correlation with the interaction types, this table also shows the number and percentage of autistic children with one or more observations for the interaction types and manifestations within each of the ADOS-2 Modules.

<i>Interaction manifestations</i>	<i>Frequency</i>		<i>Number of children showing the manifestation for each ADOS-2 Module*</i>		
	<i>Count</i>	<i>Range</i>	<i>Module 1, FNW</i>	<i>Module 1, SW</i>	<i>Module 2</i>
Exploratory interactions	57	0 - 15	4 (40%)	7 (50%)	2 (33%)
Touching robot	33	0 - 7	4 (40%)	6 (43%)	1 (17%)
Explicit visual inspection	24	0 - 10	2 (20%)	4 (29%)	2 (33%)
Relational interactions	6	0 - 5	0 (0%)	2 (14%)	0 (0%)
Using additional objects	6	0 - 5	0 (0%)	2 (14%)	0 (0%)
Functional interactions	71	0 - 17	4 (40%)	8 (57%)	6 (100%)
Controlling the robot	31	0 - 8	1 (10%)	4 (29%)	4 (67%)
Talking to the robot	15	0 - 11	0 (0%)	2 (14%)	2 (33%)
Imitation of the robot	12	0 - 5	2 (20%)	2 (14%)	2 (33%)
Physical behaviour	8	0 - 3	2 (20%)	3 (21%)	0 (0%)
Joint attention	4	0 - 1	1 (10%)	1 (7%)	2 (33%)
Social overtures	91	0 - 16	1 (10%)	9 (64%)	6 (100%)
Conversing with adult	27	0 - 10	0 (0%)	1 (7%)	6 (100%)
Requesting	25	0 - 7	1 (10%)	5 (36%)	3 (50%)
Shared enjoyment	20	0 - 4	1 (10%)	1 (7%)	5 (83%)
Social reference	11	0 - 4	0 (0%)	5 (36%)	4 (67%)
Directing attention	5	0 - 3	0 (0%)	1 (7%)	2 (33%)
Helping	4	0 - 2	1 (10%)	2 (14%)	0 (0%)
Interactions total	225	0 - 48	6 (60%)	11 (79%)	6 (100%)

* *Note.* One autistic child was not assessed with the ADOS-2 and is missing from these columns. We observed no spontaneous interactions for this child, who would likely have been assessed with Module 1, FNW.

Manifestations of functional interactions. The manifestations of the functional interactions were grouped into five behavioural categories, namely controlling the robot through the Wizard-of-Oz keypad, imitating the robot, talking to the robot, looking at the spot where the robot pointed towards when it tried to establish joint attention, and any physical behaviours directed at the robot, such as dancing. The majority of the functional interactions with the robot involved controlling the robot through the Wizard-of-Oz keypad used by the adult. While the keypad was hidden underneath the table, nine children had found out that the robot would respond to presses on the keypad. These children would reach for the keypad several times during the session, and sometimes were successful at pressing a button which resulted in the robot performing an action.

Of the children who used semantic speech, four children spoke to the robot. This was mostly confined to saying “hello” and “goodbye” to the robot. For the three children, we had annotated their third or fourth session with the robot. One child – for whom it was the second session – accounted for most of the annotations (11 out of 15). He introduced himself to the robot, asked it several questions, and asked the robot to keep its chest light on later during the session.

The functional interaction manifestations of imitation, physical behaviours, and joint attention, are more related to the content of the DE-ENIGMA intervention. Six children spontaneously imitated the robot. These were mostly imitations of the robot’s gestures and its speech (excluding children who used echolalia). One child initiated imitations of the robot’s emotions during stages of the learning content where the children were not specifically asked to imitate the robot. The physical behaviours with the robot included dancing together with the robot. One child waved at the robot in response to the robot saying hello. Another child blew air at the robot’s face. Notably, this child was listening to a social story where the robot was being pushed and started enacting the story by pushing the robot backwards himself. Four children followed the robot’s gaze when it pointed at an empty area in the room (joint attention). One child noticed that the robot was pointing at nothing in particular, and concluded that “its in his head”, i.e. the robot imagined something.

Manifestations of the social overtures. For the interactions with the adult, which were related to the robot in some form, we observed six types of interactions, namely conversing with the adult about the robot, making requests for a certain robot action, the sharing of enjoyment after the robot did something, social references after the robot did something, directing the attention of the adult to the robot, and helping the adult with the robot. Of these six types, the majority of the interactions involved requesting robot actions, the sharing of enjoyment, and for the autistic children with more expressive language, conversing with the adult about the robot.

Seven children initiated a conversation with the adult about the robot for a total of 27 observations. Prior to seeing the robot for the first time, one child seemed somewhat anxious and commented several times that the robot is not a “real robot” but a toy, and did not want to see the robot. He accounted for 10 of the 27 annotations. However, after the child calmed down and was familiarised with the robot’s actions, the child showed many signs of positive affect and continued conversing and initiating conversations with the adult. For the other six children who conversed with the adult about the robot, we had annotated their second, third, or fourth session. Four children came up with a rationale to explain why the robot did a certain action. During the social stories with the robot as the main character, two children spontaneously explained why they thought the robot would feel a certain emotion. Four children had questions for the adult regarding the robot. These questions had either to do with the appearance of the robot, if the robot could do certain things, such as having dinner, or how to control the robot through the keypad. Most of the requests were verbal, although some were non-verbal, where the child would look at the adult and imitate the robot’s action they wanted to see. Nine children requested to see specific robot behaviours. These requests include 12 requests for emotional facial expressions, seven

to see the robot dancing, five for other gestures, and for one request it was unsure what the child was requesting.

Some of the robot's actions elicited a social response from the child in the form of sharing their enjoyment or a social reference. Of the 20 observations of shared enjoyment half are related to the child making the robot do something through the Wizard-of-Oz keypad, or doing something to the robot like putting the laminated emotion cards in the robot's mouth. The other half of the shared enjoyment observations occurred after the robot had performed one of its gestures or expressions of emotion. Social references were elicited by the robot when it did something unexpected, such as when the robot started moving for the first time or when the robot got stuck halfway through an animation, or when the child did not know how to interpret the robot's behaviour. The child would then look at the adult for an explanation. The latter happened for two children when the robot tried to initiate joint attention. Rather than looking where the robot was pointing to, the children looked at the adult unsure what the robot was trying to communicate. For three children, unexpected actions of the robot also lead to them children directing the attention of the adult to the robot. Two other annotations of directing the adult's attention to the robot occurred when a child was talking with the adult about the robot, and the other when a child made the robot say "hi" using the keypad.

Lastly, we observed that four children spontaneously helped the adult to remove the blanket covering the robot at the start of the session. One child helped the adult to cover the robot with the blanket at the end of a session.

4.4 Discussion

4.4.1 Types of spontaneous interactions

We observed that autistic children spontaneously engage in a wide variety of interactions with the humanoid robot or the adult. In their interaction with the robot, autistic children most frequently initiated exploratory or functional interactions. Relational interactions were only observed for two children, which can be explained by the lack of other objects in the room. This made it difficult for children to engage in relational interactions.

Looking at the manifestations of the children's spontaneous interactions with the robot, it stands out that some children initiated the same spontaneous interaction many times during a session. Some children had a strong desire to touch or inspect the robot to the extent that they were preoccupied with engaging or trying to engage in such exploratory interactions. The soft, malleable material of the robot's face and hands were particularly interesting to the children. For the DE-ENIGMA robot-assisted intervention, the preoccupation with engaging in exploratory interactions was problematic, as the intervention was not designed to support learning through exploratory interaction, nor was the robot sufficiently robust that it could withstand tactile exploratory interactions for longer periods of time. The adult therefore dissuaded the children from touching the robot. Rather than preventing autistic children from engaging in exploratory interaction, it would be more motivating for them when they can learn about the targeted behaviour through this type of interaction. An example of

this are the exploratory play interactions such as those developed by [Boccanfuso et al. \(2016\)](#) or [Robins and Dautenhahn \(2014\)](#) to promote imaginary play and understand contingency and causality. In both studies, the robot facilitated touch interaction and reacted with a (affective) response. In addition to being able to learn through tactile exploratory interactions, the robot itself should also be designed to accommodate such interactions. Robots such as Probo ([Saldien et al., 2008](#)) and KASPAR ([Dautenhahn et al., 2009](#); [Robins and Dautenhahn, 2014](#)) have been specifically designed to accommodate tactile interaction. Such robot designs may be particularly suitable to facilitate learning for autistic children through tactile interaction. For robots that cannot withstand frequent tactile interaction, a more extensive familiarisation phase that includes exploratory interaction guided by an adult might reduce the desire to explore the robot's materials in some cases.

Next to the desire to explore the robot's materials, the children were also interested in making the robot do a certain behaviour. This is reflected by frequent observations of (trying to) controlling the robot through the keypad, or by requesting the behaviour to the adult, as well as some children becoming preoccupied with these interactions by repeatedly initiating them. The intervention used in the DE-ENIGMA database was set up to be led by the adult, which meant that those children who repeatedly wanted to control the robot's behaviour was distracting from progressing through the learning material and mastering the targeted skill. Other studies also report frequent observations of autistic children making requests, for example, during naturalistic observations of classroom activities ([Keen et al., 2002](#)), as well as in interactions with robots ([Tapus et al., 2012](#)). [Tapus et al. \(2012\)](#) used the humanoid NAO robot and reported that when the child made a request and the robot conformed, the child shared his enjoyment with the experimenter. We observed similar responses to requests or when the child was allowed to control the robot. Having child-led interactions would designate the robot to a more reactive role in the interaction. This notion was also highlighted by autism experts as an important design consideration for autistic children ([Robins et al., 2007](#)). The challenge then becomes to set boundaries and provide the context that the child can explore that aid in the learning of a targeted behaviour.

In terms of learning gains, previous research in human-delivered interventions show that having an autistic child lead the interaction has mixed results. [Kim and Mahoney \(2004\)](#) found that the engagement improves when an adult is more responsive and less directive in a human-delivered intervention. Conversely, some autistic children may learn better with a very structured, adult-directed intervention ([Kishida and Kemp, 2006](#)). For having the child-led interactions to lead to learning depends on the child's ability to initiate and engage with the robot on their own accord, and the robot's ability to elicit such initiations. Autistic children may not necessarily be inclined to engage in social interaction, and generally have lower rates of initiation, which in turn may reduce the amount of learning opportunities ([Corsello, 2005](#)). For child-led interaction in a robot-assisted intervention, it is then pivotal to identify what factors determine whether or not an autistic child may benefit from this approach.

We also frequently observed children conversing with the robot, or conversing about the robot with the adult. However, the number of children who initiated such

conversations was limited and was restricted to those that used language themselves, and one child in particular accounted for most of the talking to the robot. The observation that only a few children talked to the robot is noteworthy. The robot mostly used non-semantic speech (with the exception of the greeting and goodbye) and had an anthropomorphic design, yet the children's speech was primarily directed at the adult. While typically developing children readily make anthropomorphic inferences when interacting with a robot (Beran et al., 2011), and autistic children categorise robots very similarly to typically developing children (Peca et al., 2014), the resulting behaviour of verbal autistic children is different. Possibly, our finding may be better explained by a reduced tendency of autistic children (Chaminade et al., 2015) and autistic adults (Bird et al., 2007) to attribute human-like characteristics to artificial agents; the spontaneous interactions with the robot were more akin to interactions with an object.

For several children, the robot successfully elicited social overtures of the child towards the adult, which included the sharing enjoyment with the adult, making a social reference after the robot did something that seemingly was unexpected, directing the attention of the adult to the robot, or prosocial behaviour such as helping the adult with the robot. Such interactions are not indicative of interests of the child, but instead are learning goals for the development of certain social skills that are challenging for some autistic children. An adult could exploit these interactions as an opportunity for the child to further develop this social skill.

In the current study, we observed that six children did not engage with the robot or learning task. Similar to other studies that report aversive reactions of autistic children towards robots (e.g. Bekele et al., 2014; Short et al., 2017), we also observed aversive reactions of two of the six children. For one child, it was the first session with the robot. After lifting the blanket that initially covered the robot and showing the first robot behaviours, the child immediately showed signs of stress and left the room shortly after. The child showed a similar reaction in the second session, after which it was decided that there would be no third attempt. Possibly, the robot was too unfamiliar to the child, which triggered the stress response. For the other child who showed an aversive reaction, it was the second session that we annotated. Looking at the first session, the child did engage with the learning task and the robot. However, at some point, the robot's arms got stuck during one of its behaviours, which put the arms in an awkward position. The child showed a similar aversive reaction shortly after this happened. Possibly, this effect carried over to the next session. In the third (final) session, the child did not show an aversive reaction. The other four children were simply not drawn to the robot and showed little to no interaction towards the robot or the adult in the session that we annotated. After viewing their other sessions, we observed similar behaviour of these children, where they also did not show any interest in the task or the robot. The severity of their autistic features was similar to the other children, however they had in common that they were all non-verbal. Possibly, autistic children with limited language ability not only initiate fewer spontaneous interactions, as aforementioned, but are also more difficult to engage in a robot-assisted intervention in general. The DE-ENIGMA robot-assisted intervention may have been too complex for these children, where a simpler interaction type may

be better suited for engaging them. While we agree the commonly made claim that autistic children react positively towards robots by and large, robots are not inherently interesting to them. This highlights the need to specifically design the robot behaviour and learning task to accommodate the interests and needs of the autistic child.

4.4.2 Individual differences in the types of spontaneous interactions

For the functional interactions with the robot, we found associations with the children's language ability, severity of autistic features, and social functioning. Children with higher language ability, higher social functioning, or lower autistic features severity initiated more spontaneous functional interactions with the robot. In like manner, those children initiated social overtures directed at the adult more frequently. We found no evidence of an association between any of the children's individual characteristic and the frequency of spontaneous exploratory or relational interactions. For the latter, the observed frequency is too small to meaningfully interpret the correlations. On the other hand, many of the children engaged in exploratory interactions, but the correlations with the individual characteristics were low. This could indicate that children were equally interested in these types of interactions, or possibly some other individual characteristic influences an interest in this particular type of interaction.

Language ability and autism spectrum severity were measured through two different measurement tools, but we caution for making any inferences regarding one measure being a stronger predictor than the other. Even though in our sample one measurement tool had higher correlations than the other, the measures all follow a similar trend in their association with the interaction types. Given that the number of children included in the analysis is relatively small for the type of analysis, the confidence intervals of the correlations is large. For the children's language ability and severity of autistic features, small to large correlations with functional interactions and social overtures are also reasonably compatible with our data. Therefore, we warn that one of the two measures for measuring the same construct should not be interpreted as evidence for it more strongly associated with the interaction types than the other.

In our sample, the module that was used for the ADOS-2 assessment showed the highest correlations with the interaction types. The choice of module is primarily based on the child's expressive language ability. However, note that the ADOS-2 module is not the most reliable or valid measurement of an autistic child's expressive language ability, and should therefore be carefully interpreted — more factors may, unintentionally, have been considered for assigning the modules. In studies with autistic toddlers and children, positive associations between the complexity of object play and language ability have often been reported (Mundy et al., 1987; Jarrold et al., 1993; Toth et al., 2006; Thiemann-Bourque et al., 2012). However, these positive associations are not always found (Lewis, 2003; Kang et al., 2016). Kang et al. (2016) argue that the influence of language ability on symbolic play could possibly diminish with age. In our sample, age was not associated with the spontaneous interactions, nor did it meaningfully influence the correlations for the ADOS-2 Module

or the VABS-2 CD on the interaction types when we partialled out age. Play and language are believed to follow similar developmental trajectories and build on shared skills, such as representational skills (Lifter et al., 2011). Our finding that the spontaneous interactions of autistic children with a robot is associated with their language ability is therefore in line with this belief. Note that most of the observed manifestations classified as functional interactions did not require the children to be able to use expressive language, and therefore does not explain the association with language ability.

The autistic feature severity and the child's social functioning followed a similar trend to language ability in their association with the interaction types, but correlated less strongly in our sample. The moderate to strong correlations between the child's language ability, autistic feature severity, and social functioning, may be one explanation for finding a trend similar to that of language ability, as it suggests that they measure something similar, such as their developmental level. Indeed, as the developmental level increases, autistic children start engaging in more complex types of object play (Vig, 2007; Naber et al., 2008; Thiemann-Bourque et al., 2012). Even though robots may be seen as social actors rather than objects, it may be that autistic children similarly engage in more complex interactions with the robot as their developmental level increases. This could also provide an explanation why we found no evidence of a relation between the child's chronological age and the interaction types, as chronological age is not a good indicator for the developmental level of an autistic child due to the developmental nature of autism. Unfortunately, there was no measure available for our sample for assessing the relation of developmental levels on the spontaneous interactions in a robot-assisted intervention setting.

We argued previously that the ability to initiate may lead to more learning opportunities. In our sample, we found positive associations with the total number of spontaneous interactions with language ability, and social functioning. In like manner, a case study conducted by Duquette et al. (2008) found that non-verbal autistic children seemed less interested and engaged in human-mediated or robot-mediated sessions than pre-verbal autistic children. For the purpose of improving social initiation skills, it may be that it could be particularly beneficial to autistic children with limited language ability, but may also more difficult to achieve, as our results seem to indicate. Adding technology, such as a robot, may potentially serve as a scaffolding tool by providing an interesting, yet less complex, manner of social interaction. However, robot-assisted interventions that target different skills and are designed to capitalise on the robot's ability to elicit social interactions to another person may be less successful for these children as the robot may often fail to elicit such interactions.

4.4.3 Limitations

In the current study, we interpreted the autistic children's spontaneous interactions as interactions for which they were motivated. One of the criteria for the annotations was that the child's initiation was visibly unprompted. However, with this method it is not possible to exclude interactions that were prompted in previous interactions in the robot-assisted intervention. Additionally, while the children were motivated to initiate the unprompted interactions, the nature of the motivations may differ. The children

may have initiated an interaction for the sole purpose of having that interaction (they were intrinsically motivated), or to achieve another purpose (they were extrinsically motivated). For example, we observed that some of the school's staff members who were present during the session would encourage the child to say hello and goodbye. Such instances were not annotated, as they are prompted, but it may be that other children spontaneously said hello or goodbye because of similar rote learning. Their motivation may have been to adhere to a social norm, or to avoid a reminder to say hello and goodbye.

As a descriptive study, we did not select participants to answer our research question, and instead used an existing database of autistic child-robot interactions. The children featured in the DE-ENIGMA database are autistic children with no expressive speech up to the use simple phrases, many of whom had additional intellectual disabilities and language challenges. This limits our results to this specific subset of the autism spectrum, and does not necessarily generalise to autistic toddlers or autistic children with fluent expressive speech, who would be assessed with the ADOS-2 Module 3. While we looked at one of the largest samples of autistic children interacting with a robot, there may be other factors that influence whether and what type of spontaneous interactions they engage in, such as cultural differences (Rudovic et al., 2017). Additionally, not all interaction types were supported through the design of the intervention used in the DE-ENIGMA database, or through the protocol that was used by the adults. This affected exploratory and relational interactions in particular, as the exploratory interactions were often discouraged, and the relational interactions require the presence of additional objects. Also, the learning content of each session was different, which could have influenced the types of interaction that were observed. Future studies with different subsets of children on the autism spectrum are required to further investigate what individual factors influence the type of interaction autistic children spontaneously engage in, so that we may better tailor the robot's behaviour to the children's needs and interests.

Some children noticed that there was a keypad, and that pressing buttons on that keypad would result in the robot performing certain actions. From such actions, it is possible that they subsequently derived that the robot was in fact being controlled by the adult, which then could have influenced the agency they attributed to the robot. Given that most of the children had additional intellectual disabilities, it is uncertain whether they would be able to infer that the robot was being controlled by the adult, nor whether they considered this to be the most plausible explanation for the keypad's function. Moreover, as we mentioned in Section 2.3.1, it unclear to what extent autistic children consider a robot as a social actor to begin with.

Lastly, the autistic children in the DE-ENIGMA database interacted with a humanoid robot, with specific morphological and behavioural features. Robots with different morphology's and behaviours afford different types of interaction (Feil-Seifer and Matarić, 2011a). Therefore, the type of interactions initiated by the autistic children, and frequency thereof, may vary with other robots of a different morphology.

4.5 Conclusion

In this descriptive study, we investigated what types of interaction autistic children spontaneously engage in within a robot-assisted intervention setting, and how these types of interaction relate to the autistic traits of the autistic children. We frequently observed autistic children spontaneously engaging in exploratory and functional interactions with the robot, and robot-elicited interaction between the child and adult. In particular, autistic children with stronger language ability, social functioning, and lower in autistic features severity, initiated more functional interactions with the robot and robot-elicited interactions with the adult. None of the autistic traits were associated with the initiations of exploratory interaction with the robot.

To promote the engagement of autistic children to a robot-assisted intervention, we conclude that certain types of interaction may work better than other interaction types depending on the child's autism spectrum-specific characteristics. Facilitating learning through a specific interaction type, coupled with providing more autonomy over the robot's behaviour to the autistic children, may enable them to stay engaged longer, facilitate more learning opportunities, and ultimately improve the effectiveness of a robot-assisted intervention. Our results indicate that the child's language ability may prove a useful heuristic in predicting what type of interaction with the robot can be motivating to the child. To this end, other autism diagnostic assessments may also be insightful, but were so to a lesser degree in our sample. Facilitating a certain type of interaction will also affect the choice for a robot platform as it should support the facilitation of certain interaction types. In particular, exploratory interactions through touch are currently problematic for many robot platforms, as it can easily damage the robot. Our results indicate that such interactions with the robot are likely for autistic children, and should therefore be facilitated by the robot for it to become usable in practice. The differences among autistic children in their interaction within a robot-assisted intervention also underline the importance of reporting on autism-specific child characteristics to be able to generalise to other autistic children, which is currently not always the case (Begum et al., 2016).

The results of our study provide promising avenues for the design of deliberate robot behaviour to keep autistic children engaged in a robot-assisted intervention and to account for the heterogeneity of these children. More research is, however, required to translate our finding into interaction design for certain robot platform and assess whether they elicit and maintain the desired interaction. In the next chapter, where we discuss the evolution of the DE-ENIGMA intervention, we will discuss two of such studies, which were conducted to specifically assess iterations of the DE-ENIGMA intervention. Additionally, other researchers will need to look into conducting experimental research to draw more firm conclusions whether designing for certain interaction types, or having child-led interactions in a robot-assisted intervention actually improves engagement and provide autistic children with more learning opportunities. Moreover, more research is needed to assess how the differences in autistic traits influence the children's interaction. Research where the children are not only to interact with robots, but where they can also interact with people or objects. This is, however, beyond the scope of this dissertation.

5

The evolution of DE-ENIGMA intervention

To inform the design of the DE-ENIGMA intervention, we explained in Chapter 3 what user needs a robot-assisted intervention could address and what users require from such an intervention. Next, in Chapter 4, we analysed how autistic children spontaneously interact in a robot-assisted intervention setting, which provides us insight in how we might give shape to the robot's design to meet the user requirements. In the current chapter, we will show how we (the DE-ENIGMA consortium) developed the DE-ENIGMA intervention, based on the insights reported in the previous chapters.

5.1 Introduction

FOR the development of the DE-ENIGMA intervention, we adopted an iterative design approach, where we develop a prototype, test it, analyse the results, and update our design accordingly. This process resulted in five versions of our intervention. The initial intervention, which we refer to as v0, was developed for the DE-ENIGMA data collection study in the first months of the project. As we already discussed our analysis of this prototype in Chapter 4, we will not discuss this version further. Over the course of the DE-ENIGMA project, the prototypes have evolved from a highly structured, didactic intervention that leaned heavily on the (non-robot assisted) intervention outlined by [Howlin et al. \(1999\)](#), to a more modular and interactive type of intervention where the robot has a more prominent role. To this end, we have conducted nearly a dozen small exploratory studies in either the Netherlands, England, or Serbia.

In the next sections, I will discuss the exploratory studies that we conducted for v1 and v2 of the DE-ENIGMA intervention, as these have had a large impact on the design of the intervention. As I alluded on in Section 1.5, the focus of the DE-ENIGMA project shifted after we developed v2 of the intervention. The focus changed from developing a robot-assisted intervention for learning the basics of recognising emotions, and assessing whether this intervention leads to learning, to investigating the effect of the robot's predictability on the autistic children. This also influenced further proto-

types of the intervention. The goal of v3 was to keep the children engaged for multiple sessions and facilitate an experimental study on the effect of robot predictability on autistic children. This turned the intervention as a means to an end, rather than the end goal. Therefore, we only briefly touch on v3, but do not discuss the exploratory study conducted with v3, given that the goal of this study was not to improve the intervention itself. Instead, that study assessed whether the v3 intervention and the envisioned manipulation were suitable for a subsequent experimental study (reported in Chapter 9). Note that the work described in this chapter is highly collaborative work, and not solely my own.

5.2 Exploratory study, v1

5.2.1 Introduction

One possible way for autistic children to learn to recognise facial expressions may be to observe their face being mirrored by another party. In studies that did not involve a robot, mirroring has been reported to improve the social responsiveness of nonverbal autistic children (Field et al., 2001; Escalona et al., 2002). Deriso et al. (2012) designed a game called “The Emotion Mirror”, in which children can make facial expressions of emotion that are mirrored by a cartoony virtual agent. The role can also be reversed, where the agent makes a facial expression, and the child is rewarded when they mirror the agent’s expression. These studies suggest that mirroring activities might be an appropriate and effective way for autistic children to learn about emotions.

In v0 of the DE-ENIGMA intervention, autistic children were taught about facial expressions of emotion using direct instruction via repetitive rote-based learning of emotional expressions with reinforcement. While many of the autistic children were able to progress through the steps outlined in Howlin et al. (1999), autistic children with very limited productive language, and potentially weak receptive language as well, had difficulty progressing beyond the earliest steps. In our interviews with special education schools in England, we also found out that they generally do not report using rote instruction programs like in Howlin et al. (1999). Instead, they teach about emotions through personalised, contextual teaching, because it helps children understand what is being explained as specifically relevant to them. A generic social-story will not be as meaningful — and therefore may be more difficult to generalise to other settings — as using a real, just-concluded incident and real school settings as the basis of a social story. Other research groups have also sought to teach socio-emotional skills within more true-to-life contexts in order to promote generalisability of learning effects (e.g. “The Transporters”, Golan et al., 2010).

In the games designed for the first prototype, we seek to address the issue of generalisability by reducing the amount of direct instructions, and by using context and familiarity to promote children’s learning and understanding. For the first version of the DE-ENIGMA intervention (v1), we developed our own mirroring game using the robot Zeno, and a context game (explained in the next subsection). Before presenting the games to autistic children with limited language and cognitive ability, we conducted a small study with children from a mainstream school to assess the usability of the

DE-ENIGMA intervention v1. As such, it should be considered a fairly large exploratory study that enabled us to observe the feasibility of the setup and make design revisions for v2 of the DE-ENIGMA intervention, which was to be assessed with our target group of autistic children.

5.2.2 The DE-ENIGMA intervention, v1

Mirroring game. The mirroring game had two parts: mirroring and imitation. In the first part, the robot would mirror the child's expressions. In the second part, the robot would express a facial expression (happiness, sadness, fear, or anger) which the child then had to imitate. These facial expressions were well recognised by typically developing children (see [Schadenberg et al. \(2018\)](#)). The mirroring works as follows. A webcam captures the child's face and sends this to a laptop for analysis. Our partners at Imperial College London had developed an algorithm through which they can extract facial landmarks from video ([Asthana et al., 2015](#)). These facial landmarks are mapped onto the parameters of the robot's facial servo motors to create a facial configuration that matches the child's (see Figure 5.1), which is then sent to the robot for execution. The robot's resulting facial expression is a simplification of the the child's facial expression, as the robot only has five degrees of freedom in its face. The role of the adult using the robot is to guide the interaction and help the child figure out what is happening — understanding the cause-effect relationship between moving your own facial features and seeing the robot move its in similar fashion. With this game, we aim to help autistic children to attend to and play with emotional facial expressions, as a building block in recognising those expressions and grasping emotional concepts. We envision the mirroring game to be used as the first game the child will play with the robot. Through playful interactions, where the child has control over the robot's behaviour, the child can get familiar with the robot's movements and sounds. Next to familiarising, the game can be used to prepare the child to pay attention to the robot's facial features.

Context game. The goal of the context game is to evoke a physical emotional feeling (affect) and to teach that these feelings can be labelled with an emotion. To do so, we use images of things that the child likes or dislikes, which are provided by the child's parents or guardians. Those images are displayed on a tablet lying in front of the robot, and Zeno will ask if child can describe the item or situation depicted in the image. Next, Zeno will comment how the child is likely to feel, depending on whether the child's is supposed to like or dislike what is depicted on the image.

5.2.3 Materials and methods

Participants. The study was conducted at an international school in Enschede, the Netherlands, where 16 children participated, between the age of 4 to 6. Of the 16 children, one child was diagnosed with autism; the others were typically developing children. The children all spoke English at the school, but it may not have been their mother's tongue.

Figure 5.1: The mirroring game: a webcam captures the child's face from which the facial landmarks are extracted (the red dots, connected with the green lines). These are then converted and send as parameters to the robot, who then mirrors the child's facial expression.

Apparatus. In v1 of the DE-ENIGMA intervention, the Zeno robot was controlled by a researcher who was overseeing the intervention through a tablet interface. The images for the context game were displayed on another tablet for the child (referred to as the 'child's tablet'). Both tablets were Samsung Galaxy Tab S2 9.7 inch tablets with Google Android 6.0. Through the researcher's tablet interface, the researcher could choose between starting the mirroring game, the context game, having the robot say generic answers ("Well done", or "Listen and look"), and turn on/off the auditory sensitivity mode. In the latter, the robot would stop moving any of its servo motors below the neck, as we believe they are making quite a lot of noise and are unpleasant to listen to; much more so than the servo's in the robot's face. When idle, the robot would show 'life-like behaviour', which entails turning its head up, down, or sideways every few seconds, and blinking.

Experimental setup. The test took place in a hallway that was closed-off with a door (see Figure 5.2). Two camera's recorded the interaction, where one recorded the interaction over the shoulder of the child, and the other was placed on to the side of the child to capture the child's face. Four microphones were placed in front, above, and to the sides of the child, to record the audio. The child's tablet was put in a protective case, preventing the children from accessing the home button and closing the application.

Figure 5.2: One of the children interacting with the robot in the mirroring game of the DE-ENIGMA intervention v1. The woman is the researcher overseeing the intervention.

Procedure. The researcher would go to class and ask the child who scheduled to participate in the experiment to come and play with the robot. The child sat opposite to the robot and the researcher sat at the side of the table between the child and robot. The robot was initially covered by a blanket. After revealing the robot, the robot would come alive and start showing ‘life-like’ behaviour. The children would start with playing the mirroring game for a couple of minutes. Afterwards, they would proceed to the context game. After finishing this game, the children could return to class.

Analysis. The video recordings were analysed by three researchers (including the researcher who was overseeing the intervention). They paid attention to what went well and what needs further improvement. In specific, the researchers assessed whether children were able to complete and understand the activity, and whether the robot functioned correctly.

5.2.4 Results

The mirroring game. All 16 children understood that they were mirrored by the robot and seemed to enjoy it. For instance, some children spontaneously remarked that “Zeno is funny” or laughed frequently during the experiment. One child additionally moved his arms to see if the robot was also mirroring body gestures. The researcher then had to explain that the robot can only see the child’s face. The imitation part of the game was difficult for two of the youngest participants. The transition to a different task was not clear to the children, as they continued producing facial expressions

for the robot to mirror. The task was introduced by the researcher saying “now, you do the same as Zeno”, and the robot would say “Look at my face”, follow by an expression, and then “you do it”. The simple, non-ambiguous phrases were specifically chosen to be understandable for autistic children with limited receptive language, and who may also take phrases literally. However, for some of the participating children in this study these phrases were not clear, as they would imitate all of the robot’s behaviour, not just the facial expressions.

The mirroring algorithm and mapping onto the robot, on the other hand, did cause some issues. The mirroring was too sensitive. The over-sensitivity caused the servo motors to constantly make an unpleasant noise, which potentially could trigger auditory sensitivities. Furthermore, the algorithm had difficulty tracking the child’s neck position, which caused the robot to often look sideways, rather than towards the child.

The context game. There were some practical issues with the context game. First, while all parents provided two images for the study, they often did not provide the context of why the images depicts something that the child liked or disliked. Furthermore, the children did not know that the robot would show the images, and some children were surprised. The children also did not always agree with their parents on whether the images depicts something they like/dislike, which made it difficult to turn the interaction back to the targeted emotion.

5.2.5 Discussion and conclusion

By and large, both the mirroring game and context game were understood by typically developing children, aged 4 to 6 years, during this short interaction. The one autistic child also understood the games and in general, did not interact visibly different from the typically developing children. To improve the mirroring game, we turned off the neck tracking and toned down the sensitivity of the mirroring to generate more fluent motions and less noise. The context game could be improved by clearer instruction to the parents to provide some context for each image. Furthermore, the children would need to be made aware that their parents provided the images, and that the robot will use them in the game. For instance, a child could show the images to the robot, which would then start the conversation. Given the various practical issues that we encountered with the context game, and that the pedagogical design became clearer in the months after this study, we discontinued the context game for the DE-ENIGMA intervention.

5.3 Exploratory study, v2

5.3.1 Introduction

In addition to our assessment of v1, new sources of information became available to us for improving the second iteration of the DE-ENIGMA intervention — v2. These include our analysis of the user needs and requirements analyses (Chapter 3), the DE-ENIGMA database recordings (Chapter 4), and we further explored the autism and

developmental research literature. In relation to v1, v2 has a new pedagogical design and explores new ways for the child to interact with the robot. We will discuss both parts below.

The v2 games/activities were substantially redesigned, based on key findings in the autism and developmental research literature about face processing and acquisition of emotional skills and concepts. Recognising facial expressions is developmentally a relatively advanced skill that continues to develop during adolescence (Thomas et al., 2007). It builds on other skills that are prerequisites for learning about facial expressions. For instance, children first need to pay attention to faces and recognise that a face consists out of various facial features. In v2, we followed a *feature-based teaching strategy* — rather than the holistic faces from v1 — where the children first learn about individual facial features, before learning about specific facial configurations, which can be labelled with an emotion. The reason for this is that autistic children may have different processing strategies than typically developing children (see Weigelt et al., 2012, for a review) that privileges individual features (Dawson et al., 2005; Gauthier et al., 2009). Feature-based teaching of emotions may therefore be better aligned with the processing strategies of autistic children. Ryan and Charragáin (2010) developed a intervention that uses feature-based teaching of emotions, which was successful in teaching autistic children about emotions. In this programme, emotions such as surprise were first broken down in features including and “O” shaped mouth, raised eyebrows, and raised eyelids. In v2, we used a similar approach to that of Ryan and Charragáin (2010).

For v2, we also experimented with new ways for autistic children to interact with the robot. One of our goals for v2 was to provide children with more control over the interaction. This could allow the child to shape the interaction to their choosing — one of the user requirements from Chapter 3 — turning them in active participants. In our descriptive study on the DE-ENIGMA database (see Chapter 4), we also observed children frequently trying to control the robot through the keypad, and providing a degree of control over the robot was also suggested by educators (Robins et al., 2010b).

The most straightforward way to provide the children with a degree of control over the robot would be to further develop our existing tablet-interface. However, this may not be the most engaging and understandable way to communicate with a robot. Moreover, the children may associate the tablet with different activities, such as with playing their favourite app, rather than using the tablet to shape the interaction with the robot. In our descriptive study on the DE-ENIGMA database, we also observed that autistic children — in particular those with limited to no expressive language — seemed interested in tactile interaction with the robot. Unfortunately, the Zeno robot does not support the rough tactile interactions, as its face can be torn off or damaged relatively easily. Possibly, both the desire for tactile interactions and need for an interface for the child could be addressed by providing the children with a user interface that also provides tactile interactions — a *tangible user interface* (TUI). Such interfaces have been found to solicit more parallel and collaborative play than solitary play when compared to toys (Farr et al., 2010), and elicit a host of social interactions between autistic children or with an adult (Nonnis and Bryan-Kinns, 2019). However,

Figure 5.3: Example tablet screens of the child's tablet to select the whole face (left tablet screen) or specific facial features (right tablet screen).

we are concerned that a TUI may draw the attention away from the robot and the task. If the TUI is too interesting, and remains so, the children may have difficulty switching their attention back to the robot, and instead may focus on the tactile stimulation provided by the TUI.

The goal of the current exploratory study was twofold. First, we assessed how well autistic children understood and enjoyed each of the games of the DE-ENIGMA intervention v2. And second, we wanted to get a better understanding of how these children may interact with tangibles. For this, we used low-fi prototypes that could resemble a TUI, but which were not connected to the robot. We therefore refer to them simply as *tangibles*. Each set of tangibles could be used as a puzzle, but differed in the materials that were used and the complexity of the puzzlelike interaction they provided.

5.3.2 The DE-ENIGMA Intervention, v2

The second iteration of the DE-ENIGMA intervention consisted of five activities. These included the mirroring game and the emotion learning game, which features three steps. Two of those steps revolved around facial features, and one was about emotional facial expressions. The last new activity was one where the children interacted with tangibles. I will discuss each new activity below.

Step one: facial features. The goal of step one of the emotion learning game is to teach children the individual features that constitute a face, following a feature-based teaching strategy. In this case, we included the eyebrows, eyes, and mouth, as features, as the Zeno robot can also manipulate these features on its own face. The game starts with a picture of the robot's face being shown on the child's tablet (see the left tablet screen in Figure 5.3). The robot then prompts the child to touch the robot's face by saying "touch my picture". When they do, the robot says "This is my face" and moves each of its three facial features. After having pressed three times on

the robot's face on the tablet, the picture of the robot's face will change and outline the three facial features (see the right tablet screen in Figure 5.3). When one of the facial features is touched, the robot will label the feature (e.g. by saying "This is my mouth"), and moving the feature accordingly (e.g. opening and closing its mouth, and moving the mouth corners up and down). The interaction during this step is "child-led" in that the children explore each feature at their own pace and can repeat as often as they want.

Step two: discriminating facial features. After exploring and learning about the three facial features, the next step is about learning to discriminate those features. The interaction during this step is "robot-led", where it prompts the child to find an individual feature (e.g. "find my eyes"). A specific feature is chosen by the researcher through their own tablet interface. The child's tablet then shows the robot's face with the three facial features outlined, similar as in step one. When the child selects the correct feature, the robot responds positively (e.g. by saying "Well done, you found my eyes!") after which it moves the facial feature (e.g. opens and closes the eyes). When the child selects an incorrect feature, or does not respond for a set duration of time, the robot first repeats the prompt twice, after which it suggests to try a different feature.

Step three: emotional facial expressions. In the final step of the emotion learning game, the goal is about learning that certain configurations of the facial features express an emotion. For this study, the robot could express the happiness emotion, which was modelled after the expression designed by Salvador et al. (2015). We tested to what extent this expression was recognised by typically developing children and found that it was well recognised (Schadenberg et al., 2018). The third step is similar to step one in that the child can explore the features, and the robot will respond to the child's actions. This activity starts with a picture of the robot's "happy" face on the tablet, similar to the left tablet screen shown in Figure 5.3 (without the facial features outlined). The robot points towards the tablet and prompts the child by saying "touch my picture". When doing so, the robot displays the facial expression for happiness and says "Happy. This is my happy face.". The robot will show its happiness emotion three times, after which the game moves to feature-based teaching by highlighting the three facial features in the picture of the robot displayed on the child's tablet. When the child selects a facial feature, the robot responds by showing the expression for happiness, while explicitly stating the configuration of the chosen feature (e.g. "Happy. My mouth is up. I am smiling.").

Tangibles. For this study, we bought three tangibles that could be used for a puzzle-like interaction (with some customisation). The first tangibles were simple colourful blocks that we bought and were safe to use for children ages 0+ (see Figure 5.4). One set of blocks were coloured triangles, and another set were Lego-like bricks. The blocks were used for a free-play activity together with the robot, during which the robot would prompt the child to find certain colours and stack them together.

For the second tangible, we printed the robot's face — displaying a neutral facial

Figure 5.4: Colourful blocks made of squishy material in the form of triangles and Lego-like bricks.

Figure 5.5: Blocks that each display part of the robot's face.

expression — on canvas and divided it in four pieces: the three facial features that were taught in the emotion learning game plus one block for the robot's hair. The canvas was then sewed over the colourful blocks of the first tangible. The result can be seen in Figure 5.5. The children could use the blocks with the robot's individual facial features to construct a full face, during which the robot would provide encouragement and instructions.

The third tangible was a couple of plastic puzzle pieces with the robot's face printed on them (see Figure 5.6). In total, there were four emotions displayed by the robot (happiness, sadness, fear, and anger), and each emotion was divided into two puzzle pieces. To make the game easier, and to allow the robot to give instructions through its limited behaviour repertoire, the pieces for each emotion were outlined with a specific colour.

5.3.3 Materials and methods

Participants. Eight autistic children, of which two were girls, from a special education school in England participated in this study. The children were recruited from one of the schools who had participated in the DE-ENIGMA data collection study for building the DE-ENIGMA database, where half of the participants engaged in a robot-

Figure 5.6: Plastic puzzle pieces with the robot's facial expressions for respectively the emotions of happiness, sadness, fear, and anger.

assisted intervention. Four of the children had interacted with Zeno during this study, three had seen him once, but had not interacted with Zeno, and one child had never met Zeno before. Sessions took around 15 minutes and children completed up to three sessions over the course of the two-week study.

Apparatus. For v2, the researcher's tablet interface was extended to accommodate the new games, and removal of the context game. Furthermore, we added more generic robot behaviours, such as saying "Well done", so that the researcher could better adapt to the child's spontaneous behaviour.

Experimental setup. The robot was set up on a table in a small room in the children's school. The overall setup was similar to the previous study where we assessed v1 (see Subsection 5.2.3). This time, three cameras recorded the interaction: one provided a wide-angle, bird's-eye view of the complete setup, one recorded the interaction over the shoulder of the child and the third a close-up of the child's face. An external microphone was used to record audio.

Procedure. At the start of each session the robot was covered with a blanket. When a child entered, the adult would first show them which activity they were going to do using pictographic descriptions. Then the adult revealed the robot, which would then introduce itself and the activity. The first game the children would play was the mirroring game. Next, they would play the newly designed facial recognition game. The session ended with playing with the tangibles.

5.3.4 Results

The mirroring game. Out of the eight children, five played the mirroring game — two had technical difficulties, and one child was preoccupied with touching the robot's materials. Of those five children, one child made unprompted facial expressions, showed a lot of positive affect, and played for a lengthy amount of time. Another child played with the mirroring game for a bit, but then asked the robot to stop. The remaining three participants did not produce enough facial expressions for them to

learn from observation. Our impression was that they did not understand the cause-effect relationship of the mirroring game, where their facial expressions resulted in changes in the robot's facial expressions.

Emotion learning game, step one to three. Step one and step two generally seemed well understood by the children. The children understood the instructions and had little difficulty in exploring and discriminating facial features. On the other hand, step three was more difficult. The principle that facial features make up a facial expressions and that a change in those features also changes the expression was often not understood. Regarding the children's engagement in the emotion learning game, the steps did not seem to be encouraging engagement. We often observed children behaving disinterested and passive.

Seven out of eight children used the tablet correctly with varying degrees of difficulty (touch being too light or misplaced due to imprecise motor control). This caused frustration or disengagement in some children, who needed to press several times for the system to recognise it. One child appeared afraid of the robot and did not touch the tablet. Perhaps because he was aware it would make the robot move. Since there was an actual Zeno robot present and an image of Zeno on the tablet, some children seemed (initially) not to understand to which Zeno they needed to respond. The difference between tablet Zeno and actual Zeno may be difficult to understand for some (e.g. a child pointed to actual Zeno's face when prompted to find certain features on the tablet). For one child, the opportunity of interacting with the tablet was so distracting as to preclude most engagement with the content — the child was more interested in opening his favourite app.

Tangibles. The coloured blocks invite free play and most children immediately start playing with it, including those that do not show much engagement in the emotion learning game. Often, the children would use the blocks to build towers or other constructions, but did not demonstrate the required fine motor skills to stack the Lego-like bricks. Both the squishy and plastic blocks of Zeno's facial features were often intuitively combined into a full face, although some needed extra help to place individual features in the correct position. Some children labelled the individual features spontaneously or upon request. Most children were able to redirect their attention back to the robot when prompted by the researcher or robot. Only one child needed additional prompting from the researcher before redirecting the attention back to the robot.

5.3.5 Discussion and conclusion

The use of tangibles seems promising as children appeared to engage spontaneously with them. They also seemed more engaged playing with the tangibles, compared to their engagement in the emotion learning game¹¹. The actual task of connecting the blocks to form a face appears to be difficult in terms of the fine motor skills that are required, as well as figuring out what blocks belong together. The material of

¹¹Note that this is our impression, but that we did not carry out a formal analysis of the children's engagement.

the tangibles did not seem to influence the children, as they were using the blocks functionally, rather than exploring the materials themselves. Children will likely be able to disengage from the materials, but sometimes need some prompting (from adult or robot).

Overall, we recommend to further explore tangibles and turn them in a TUI. For example, a ‘smart puzzle’ featuring the robot’s face could be utilised as an interface for the emotion learning game. The facial features are then puzzle pieces that can be placed on the ‘puzzle board’ that is shaped like the robot’s face, but is missing facial features. By using multiple versions of each puzzle piece that represents a facial feature, the TUI could also be used for learning about emotions (e.g. having a piece where the mouth is smiling for happiness, and a piece where the mouth is in an ‘O’ shape for surprise). Through the use of sensors, the TUI can be connected to the robot and send feedback to the robot when the child places a particular puzzle piece on a certain position of the puzzle board.

Regarding the feature-based teaching strategy that we adopted for v2, the children were relatively good at recognising isolated features, but found putting them together (either spatially, or integrated into an expression) more challenging. For future iterations, we recommend to explicitly demonstrate and allow children to practice the various relationships between facial features. This should relate to both the spatial relationship of the features on the face and the correct configuration of features for making a complete face, as well as the relationship between facial features and facial expressions, where certain configurations of the individual features can be integrated to form an emotional expression. We also recommend that future iterations focus on developing a more engaging interaction in the emotion learning game. While the children carried out the tasks when prompted, we saw few spontaneous interactions and few instances of children showing positive affect. Based on the interactions we observed during the interaction with the tangibles, replacing the children’s tablet with a TUI might improve the children’s engagement. Additionally, the steps themselves are missing a game element that motivates the children, and adding one might also improve engagement.

With one exception, children did not appear to enjoy the mirroring game and the majority of children did not appear to recognise that Zeno was imitating their face, spontaneously or after explanation. During the mirroring task, some children also appeared to find the quick movements of the robot frightening, which may have contributed to disengagement or anxiety. We therefore discontinued the mirroring game.

5.4 From v3 to the final version of the DE-ENIGMA intervention

The goal of the DE-ENIGMA intervention was changed in v3 from “teaching children about facial expressions” to “keeping children engaged in a robot-assisted activity for at least four sessions of 15 minutes”. We judged that the emotion learning game was sufficiently engaging and well understood by autistic children (with high support needs) that it could meet our goal, provided that more content was added. Therefore, we added another step to v3, where the children would learn to discriminate facial

expressions of emotion.

Given the change in focus of the DE-ENIGMA project, we did not further explore the use of a tangible user interface, or further improved the emotion learning game to be more engaging. In v3, we instead focused on understanding how we could manipulate the predictability of the robot's behaviour, which is discussed in Chapters 6 to 9. Another effect of the change in focus was that we did not further integrate the perception modules (perception from 2D and 3D video, and audio), developed by Imperial College London, University of Augsburg, and the Institute of Mathematics of the Romanian Academy — into the intervention to make the robot more autonomous. As such, the final version remains semi-autonomous, and requires input from an adult user through a tablet interface. Rather than integrating the perception modules into the intervention, we developed a reporting tool that provides the adult user with a report of how the session went, based on the output of the perception modules. In the next section, I will elaborate on what the final version of the DE-ENIGMA intervention looks like.

5.5 The DE-ENIGMA intervention

5.5.1 Interaction design

Role of the robot, adult, and tablet. The robot's overall role is that of an *aide* to the adult (a parent, autism professional, or in case for our own studies, a researcher), who sits beside the child, oversees and guides the activities with the child. This is in alignment with the advice and needs of autism professionals outlined in Chapter 3. The robot's function is to display behaviours to inform or engage the child, to provide a safe and highly predictable environment, and reduce the complexity of looking and interpreting faces as well as learning to recognise emotions. Based on past work (Hourcade et al., 2012) and our exploratory studies reported in the previous sections, we use a tablet as a way for the child to communicate with the robot, as children are familiar with tablets. Using a tablet also allows us an alternate modality of communication for autistic children who may have comparably weak receptive language, or who better process visual information. The robot is semi-autonomous in that it can autonomously form a response to the child's input through the child's tablet and execute that response without intervention from the adult. The adult remains in control over what and when activities can be played, as well as handling the robot's responses to the unpredictable dynamics of the interaction (e.g. saying "I do not know", in response to a question from the child). Control over the robot can be enacted through a tablet interface that is used by either the adult who is overseeing the activity, or another adult. The other adult can be present in the room itself, or hidden from view to turn the interaction into a Wizard-of-Oz scenario (Riek, 2012).

Interaction design principles. In designing the robot's behaviour and the script, we applied several strategies: *Scaffolding* was applied to the intervention in such a way that we first introduce fundamental concepts before more complex concepts are introduced. Furthermore, the robot first attempted to interact with the child

in playing the games. When this proved too difficult, the adult offered additional explanation and demonstration. To create a safe environment where mistakes can be made, *positive feedback* was included throughout the intervention, where the robot said “Well done!” and then performed a cheering gesture.

To avoid distractions, robot *voice and motion were not concurrent*. We also removed the life-like behaviour that the robot displayed in v1 and v2, and removed all the lip synchronisation. The robot used *simple, non-ambiguous language* to match the children’s vocabulary. Rhetorical questions (e.g. “Can you find eyebrows?”) were avoided, and instead we used imperative phrases delivered in an encouraging tone (“Find eyebrows”), as autistic children can answer such questions literally (i.e., saying “yes”). The robot generally only said one sentence before pausing, where each sentence contained a maximum of five words. To make the interaction more *predictable*, we designed robot behaviours that did not show variations within sessions, nor over sessions. The robot behavioural repertoire was therefore limited, and fixed, as it would always use the same phrases or motions for its actions. Lastly, to *minimise the distraction caused by the child’s tablet*, we used no auditory or visual feedback (e.g. blinking lights). Furthermore, the tablet would go blank when the child did not need to interact with it.

5.5.2 DE-ENIGMA games

The final version of the DE-ENIGMA intervention consists of four different games. There are two types of games: *exploratory games* and *search games*. The former relates to games where the child can freely explore facial features or expressions at their own pace and where there are no incorrect answers. The latter relates to games where the child receives an instruction to find a specific facial feature or expression from among several options. During development and usability testing we found that it could be difficult for some children to discriminate between games (in particular those children with limited language skills), as they are similar in their structure and visual elements. For instance, the child’s tablet only displays images of the robot’s face. We therefore identified each of the games by a colour: the pink, blue, green, and yellow game (see Figure 5.7). These colours were selected considering possible colour blindness. The games have a consistent structure and duration to improve the predictability of the games. Each game is composed of four trials for the child to make a selection, (excluding instances where the child selected an incorrect answer). After four trials, the game ends and the tablet displays a blank screen.

The pink game. This exploratory game is similar to step *one* of the emotion learning game, described Subsection 5.3.2. The children can explore the individual facial features that constitute a face (eyebrows, eyes and mouth) at their own pace. When the children select a feature outlined on the child’s tablet, the robot will label and move the corresponding feature on his face.

The blue game. This search game is similar to step *two* of the emotion learning game, described Subsection 5.3.2. The game tests whether children can discriminate between facial features and link the correct verbal label to each feature. The robot

Figure 5.7: The child's tablet interface for each of the DE-ENIGMA games. Respectively the pink, blue, green, and yellow game.

prompts the child to find a feature, and demonstrates that feature when the correct answer is selected.

The green game. This exploratory game is similar to step *three* of the emotion learning game, described Subsection 5.3.2. The children can explore the facial features configurations of for four different emotions (happiness, sadness, anger, and fear) at their own pace. When a facial feature is selected, the robot responds by labelling the configuration of that feature for the respective emotion and then demonstrate the configuration for the selected feature.

The yellow game. This search game was newly developed for v4 of the DE-ENIGMA intervention. The goal of this game is to test whether children can discriminate between the emotional facial expressions they explored in the green game. The robot prompts the child to find an emotional facial expression (e.g. by saying “Find my happy face”). The child's tablet displays four images of different emotional expressions of the robot. When the child selects the correct image, the robot responds positively (e.g by saying “You found happy!”) and displays the corresponding emotion. When the child does not respond to the robot's prompt, or selects the incorrect image, the robot will encourage the child to try again. For children who find this game particularly challenging, the adult may also choose to reduce the number of images to choose to two images.

Generic robot actions. Next to the games related to recognising facial features and emotions, the adult also has several options of robot behaviours that can be used for personalising the interaction and meeting the unpredictable dynamics of interactions. These include behaviours where the robot displays an emotion (happiness, sadness, anger, or fear), performs a dance, plays a nursery rhyme, moves its torso to the each side respectively, says hello and introduces itself, says “yes”, “no”, “I do not know”, or “goodbye”, gives praise, or displays a cheering motion. These actions are used as a first activity of a session, where the adult can select any of these robot behaviours to see if it attracts the child's attention, and to let the child get more familiar with the various motions and sounds the robot makes.

Figure 5.8: The choice board through which the child can select certain games or robot behaviours.

Choice points. One of the results of the study described in Chapter 4 was that autistic children often make requests to the robot. As denying the child’s request could disrupt the interaction, we introduced “choice points” during the intervention. These choice points allowed the adult to defer any requests for one of the games or certain behaviours to a choice point, rather than denying the child’s request or interrupting the ongoing game. Moreover, by structuring the choice points and limiting them to two minutes, it allowed us to control it as much as possible with respect to following the same programme of content and the same order for each child. During the choice point, the child could choose from any content that they had already experienced (i.e. both games and generic robot actions). The child could make the choice through a “choice board”, which contained icons of the available options (see Figure 5.8). The icons can be added or removed from the choice board, as each was stuck on the board with velcro. The board and the icons on it were managed by the adult.

5.5.3 DE-ENIGMA system architecture

The DE-ENIGMA system that is used to run the DE-ENIGMA intervention consists of the following hardware: a Robokind R25 robot (see Figure 1.1 in Chapter 1.3), an Ubuntu 16.04 LTS and two Windows 10 OS processing machines, a Microsoft Lifecam HD-3000, a Zoom H2N microphone, a set of speakers (used instead of the robot’s build-in speakers), a router, and two Samsung Galaxy S2 tablets running Android 7.0. For a schematic overview of the system architecture for the DE-ENIGMA intervention v4, see Figure 5.9. The DE-ENIGMA system follows a modular setup and features four modules: a child-robot interaction module, and three perception modules. The communication between the modules is handled by Robot Operating System (Quigley et al., 2009) and Apache Apollo middleware.

Figure 5.9: System architecture diagram for the DE-ENIGMA system that runs the DE-ENIGMA intervention, v4.

Child-robot interaction module. The architecture of the child-robot interaction module of the DE-ENIGMA system uses the Flipper 2.0 Dialogue Manager (van Waterschoot et al., 2018), the Asap Behaviour Realiser (Reidsma and van Welbergen, 2013), and an Agent Control Engine that serves as an application programming interface for the robot. The dialogue manager’s information state is a representation of its beliefs about the world and is used to trigger certain robot behaviours. These behaviours are specified in Behavioural Markup Language — a description of nonverbal and verbal robot behaviour that is independent of a particular robot platform. When triggered, the corresponding behaviour is sent to the behaviour realiser that then translates the triggered behaviour to control primitives for the Zeno robot. These control primitives are sent to the robot through the agent control engine. The behaviour realiser is also responsible for synchronising, queuing, and monitoring the execution of the robot behaviour. For the robot’s speech, we recorded audio files in English and Serbian. For the former, a Dutch colleague (female) acted out the verbal scripts that we wrote. For the Serbian speech, one of our Serbian project members acted out the same scripts. All audio files were normalised in Audacity to ensure a consistent loudness.

Perception modules. The DE-ENIGMA system contains three perception modules. First, there is a module that recognises body posture, several gestures, and can distinguish individuals (Marinoiu et al., 2018). This module uses the 3D video feed from a Microsoft Kinect 2.0, and was trained to specifically recognise the body posture and gestures of autistic children. Second, a module for tracking valence and arousal through analysing the child’s speech, and real-time laughter and speech detection (Hagerer et al., 2017). To this end, the modules required an external microphone. The third perception module uses a video feed from a webcam as input and extracts the child’s facial features (Asthana et al., 2015). The output then feeds into a fa-

Figure 5.10: Example page of the automatically generated session reports that each cover a total of three pages.

cial expression recogniser, which is also used to keep track of the child's valence and arousal.

Tablets. The DE-ENIGMA system uses two tablets — one for the child (the child's tablet, and one for the autism professional (the adult's tablet). The tablets feed into the Flipper Dialogue Manager and were connected to the DE-ENIGMA system through wi-fi. The child's tablet was encased in a tablet case to prevent the child from pressing the home button in an attempt to access other apps. It displayed the learning content and let the child provide answers to the robot's questions, or select actions for the robot to display. The adult's tablet displayed buttons that could be used to directly trigger robot behaviours and initiate game activities.

Automatically generated session reports. Given that the goal of the DE-ENIGMA changed after v2 of the intervention, it was decided not to integrate the perception

modules with the robot. The modules did not always provide accurate and robust output, and integrating their output with the robot would mean that the robot would sometimes behave based on an incorrect analysis of the child. In turn, this could cause the robot to behave unpredictably from the child's perspective, which would influence the experiment that we planned on doing using v4 of the DE-ENIGMA intervention. Instead, the perception modules were used to automatically generate a report of a session which includes a list of activities that were played, time spent playing, the correctness of the answers, a heatmap of the child's valence and arousal during the session, the time the child spent talking, a summary of the detected gestures, and the child's overall physical activity during the session (see Figure 5.10 for an example page). The content of the report was based on the needfinding study (Schadenberg et al., 2020a) that we discussed in Chapter 3.3, additional interviews with six educators from England, and what the perception modules could actually detect.

This is the end of the first part of this dissertation, which described the background and development of the DE-ENIGMA intervention. Our work provides various leads for future research. In particular, we are excited about the use of a tangible user interface to improve the children's engagement to the intervention, provide control to the children to shape the interaction to their choosing, and may address the need of some for tactile stimulation. We have also seen that autistic children are highly idiosyncratic and that this can have a large influence on the interactions they have with robots. It will be essential for robot-assisted intervention to address these idiosyncrasies in some form if they are to be effective.

The final version of the DE-ENIGMA intervention (v4) was used in the last study that was conducted by the DE-ENIGMA consortium. In Chapter 9, I report on our analyses of this study, which aims to better understand how the robot's predictability influences the engagement of autistic children within a robot-assisted intervention. Speaking of robot predictability, the second part of this dissertation is all about this topic. As we have seen in Chapter 2, the highly predictable nature of robots is a commonly used argument for explaining why robots are promising for enhancing autism interventions. Not only is this argument used by researchers, but also by autism professionals (Alcorn et al. (2019), discussed in Chapter 3). This is not without cause, as predictability plays a major part in autism literature. In fact, difficulties in dealing or using predictions may be at the core of autism (Pellicano and Burr, 2012; Van de Cruys et al., 2014; Lawson et al., 2014; Sinha et al., 2014). In the second part of this dissertation, I will focus on the research that we conducted for getting a better understanding of this concept of predictability and how it relates to human-robot interaction in general.

Part III

PREDICTABILITY AND ITS EFFECTS ON HUMAN-ROBOT INTERACTION

“Laughter is an affect resulting from the sudden transformation of a heightened expectation into nothing.”

Immanuel Kant, Critique of Judgment

6

Autism, robots, and the role of predictability

This chapter is based on the following articles:

Schadenberg, B. R., Reidsma, D., Evers, V., Li, J. J., Davison, D. P., Heylen, D. K. J., Neves, C., Alvito, P., Shen, J., Pantić, M., Cummings, N., Schuller, B. W., Olaru, V., Sminchisescu, C., Babović Dimitrijević, S., Petrović, S., Williams, A., Alcorn, A. M., & Pellicano, E. (under revision). Predictable Robots for Autistic Children — Variance in Robot Behaviour, Idiosyncrasies in Autistic Children's Characteristics, and Child-Robot Engagement. *ACM Transactions on Computer-Human Interaction*. 39 pages. ([Schadenberg et al., under revision](#))

Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., & Evers, V. (in press). I See What You Did There: Understanding People's Social Perception of a Robot and its Predictability. *ACM Transactions on Human-Robot Interaction*. 27 pages. ([Schadenberg et al., in press](#))

In this chapter, we will discuss the concept of predictability and how it relates to humans, and autistic children in specific, interacting with robots. Is it a singular concept, or is it multi-faceted? In what way is autism related to predictability? And how does predictability shape human perception (or a robot)? After discussing literature that can provide answers to these questions, we synthesise a novel conceptualisation and operationalisation or predictability as it relates to humans interacting with robots.

6.1 Introduction

IN the previous chapters, we have seen that the predictability of a robot is a common argument used by researchers (e.g. [Dautenhahn and Werry, 2004](#); [Duquette et al., 2008](#); [Thill et al., 2012](#); [Huskens et al., 2013](#); [Sartorato et al., 2017](#); [David](#)

et al., 2020) and autism professionals (Alcorn et al., 2019) for why robots may be promising tools for those working with autistic children. Robots, as programmable systems, could be designed to provide such a highly predictable environment, which may be easier for autistic children to understand and to interact with, than with (relatively unpredictable) humans. The possibility to address the need for predictability through a robot is interesting to us, as it may also positively impact their engagement in learning.

The origin of the predictability argument for using robots for autistic children can be traced back to the importance that is placed on the concept of predictability for autistic individuals in general, where they are believed to favour environments that are predictable. As Verneti et al. (2018, p. 22) put it:

“... children with ASD [Autism Spectrum Disorder] exhibit more frequent social behaviours and social drive when interacting with familiar, therefore more predictable, social partners (e.g., caregivers) (Goldberg et al., 2017; Sigman et al., 1986). The decreased motivation towards social stimuli could also be reframed in terms of an aversion for more unpredictable stimulation.”

Predictability has been part of the definitions of autism since Leo Kanner first termed autism, based on his observations of eleven children (Kanner, 1943), albeit under a different name. During his observations of eleven children, he observed that the children had a “need for sameness” or a “resistance to (unexpected) change”. Nowadays, the insistence for sameness is a diagnostic feature of autism (American Psychiatric Association, 2013), where it is part of the non-social features of autism (the presence of rigid and repetitive patterns of behaviours and limited personal interests). A line of contemporary theories that try to explain autism further emphasise the importance of predictability, and hypothesise that difficulties in generating or using predictions lies at the heart of the condition (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014). The predictability of robots is not only of relevance to autistic individuals, but also to typically developing individuals. Predictability in human-robot interactions is considered to be essential for humans to understand robots (Sciutti et al., 2018), for coordinating actions with robots (Sebanz et al., 2006; Sebanz and Knoblich, 2009), for improving task performance (Koppenborg et al., 2017), safety (Alami et al., 2006; Heinzmann and Zelinsky, 2003), and trust in the robot (Hancock et al., 2011; Lewis et al., 2018), and can support decision-making processes for human-in-the-loop robotic systems (Noorman and Johnson, 2014).

While there seems to be a general agreement in literature regarding the importance of predictability to autistic individuals in general, and also for typically developing individuals interacting with robots, a conceptual definition that gives a shared meaning to predictability is missing. In the field of HRI, predictability is used in terms of it being a *property of robotic motion* (Lee et al., 2011; Basili et al., 2012; Lichtenthäler et al., 2012b; Dragan et al., 2013; Szafir et al., 2014; Dragan et al., 2015). In this context, the predictability of a motion relates to how well it matches the person’s prediction regarding the robot’s motion trajectory, given that the goal of the robot is known (Dragan et al., 2015). Studies also differ in *what about the robot’s*

behaviour a person is predicting, where a robot can be predictable in one aspect, but unpredictable in another. For instance, Dragan et al. (2015) differentiate between “predictable motions”, which are optimised to allow the prediction of the robot’s motion trajectory, and “legible motions”, which are optimised to allow the prediction of the goal the robot is trying to achieve with its motion. While these motions can result in the same motion in certain situations, in highly ambiguous situations, legible motions allow a person to predict its goal, but its trajectory is unpredictable, and vice versa for predictable motions (Dragan et al., 2013). Following on from predicting the end-state of a robot action, other studies are about predicting what action the robot will take after the on-going action (Takayama et al., 2011; Lichtenthäler et al., 2012b; Mubin and Bartneck, 2015; Driggs-Campbell and Bajcsy, 2016; Walker et al., 2016). And finally, people can also *attribute predictability* as a quality of a robot (Eyssel et al., 2011; Eyssel and Kuchenbrandt, 2011; Dragan and Srinivasa, 2014a; Dragan et al., 2015; Mubin and Bartneck, 2015; Driggs-Campbell and Bajcsy, 2016).

The lack of a conceptual definition that gives a shared meaning to predictability is problematic for us if we are to design a robot that is highly predictable in its behaviour. Because without it, we do not know what factors can make a robot more or less predictable. The goal of this chapter is therefore to better understand the concept of predictability and how it relates to predicting robot behaviour. In Section 6.2, we will first succinctly describe the role of predictability, drawing upon a line of contemporary accounts of autism. In Section 6.3, we further explore the concept of predictability and how it relates to predicting robot behaviour. We do this from the perspective of contemporary accounts of human cognition that explain human perception in terms of predictive processing (Friston, 2005, 2010; Clark, 2013; Hohwy, 2013). In this section, we will explain how predictive processing works, and apply insights from these accounts of human cognition to predicting robot behaviour in particular. Next, in Section 6.4, we provide a conceptualisation of predictability within the context of HRI. We then operationalise predictability in terms of variance in the robot’s behaviour, which is described in Section 6.5. We summarise and conclude on this chapter in Section 6.7.

6.2 Autism, predictability, and individual differences

6.2.1 Autism as atypical predictive processing

Our senses are constantly dealing with ambiguous sensory information, trying to make sense of it all. According to Bayesian accounts of perception, the human brain is constantly generating predictions on what sensory information is expected in order to resolve this ambiguity and attempts to minimise the error between the incoming sensory information and the prediction (Mumford, 1992; Friston, 2005; Bar, 2007; Clark, 2013; Hohwy, 2013). This is referred to as PREDICTIVE PROCESSING (explained in more detail in the next section). When a prediction does not match with the observed sensory information, additional cognitive effort is required to resolve the mismatch and to identify its cause. When these mismatches happen too often, it can lead to a person being overwhelmed or to anxiety (Paulus and Stein, 2006). Autistic individuals are believed to frequently experience these mismatches between predicted and

observed sensory information due to a decreased influence of predictions on perception, which is hypothesised in current Bayesian accounts of autism to be at the heart of their condition (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014)(see Palmer et al., 2017, for a review). According to the Bayesian accounts of autism, the non-social features are the result of trying to maintain a predictable environment through either self-generated sensory information — which can easily be predicted — or through enacting control over their environment. The social features of autism, on the other hand, are explained by the highly unpredictable nature of social environments, which can therefore be difficult to deal with. The Bayesian accounts of autism are not the only theories that relate prediction difficulties with autistic features. For instance, the Empathizing-Systemizing theory (Baron-Cohen, 2002, 2009) posits that the social features of autism arise from a below-average empathy, whereas the non-social features stem from above-average systemizing — identifying lawful patterns in the information. When faced with information that does not strictly adhere to discernible laws (i.e. unpredictable environments), such as in social interactions, autistic individuals struggle to make sense of it.

A decreased influence of predictions on the perception in autistic individuals is supported by various lab studies (Pellicano et al., 2007; Ewing et al., 2013; Turi et al., 2015; Balsters et al., 2017; Lawson et al., 2017; Goris et al., 2018), although support is not always found (Pell et al., 2016; Manning et al., 2017; Braukmann et al., 2018; Tewolde et al., 2018; Lieder et al., 2019). Nevertheless, literature has shown that, in practice, offering a more predictable environment is invaluable to facilitating learning by providing an environment that puts the child at ease by not having to deal with the discomfort resulting from unpredictability and requires fewer cognitive resources to be processed. Current educational practices therefore emphasise the need for structure at schools (e.g. through the TEACCH approach (Mesibov and Shea, 2010)), so that autistic children know what to expect during the day, increasing their engagement in learning (MacDuff et al., 1993; Bryan and Gast, 2000; O'Reilly et al., 2005). All things considered, there is compelling evidence that predictability is important to autistic children.

6.2.2 Intolerance of uncertainty

A construct that may be similar to a preference for predictability — or the intolerance of unpredictability — of autistic children is that of INTOLERANCE OF UNCERTAINTY (IU). IU is defined as “an individual’s dispositional incapacity to endure the aversive response triggered by the perceived absence of salient, key, or sufficient information, and sustained by the associated perception of uncertainty” (Carleton, 2016), and comprises two components, namely a “desire of predictability” and “uncertainty paralysis” (Berenbaum et al., 2008). Knowing what will happen in the future allows us to increase the odds of a desirable outcome, or brace for future adversity. However, this requires knowledge on the probability, timing, and nature of the future event, which is often not available as the future is intrinsically uncertain. Anticipating the future can therefore induce anxiety, as the uncertainty reduces our ability to efficiently and effectively prepare for the future. Individuals with heightened levels of anxiety are be-

lieved to have an excessive response to uncertainty (Grupe and Nitschke, 2013), and IU is considered to be a cognitive vulnerability factor in a broad range of emotional disorders like general anxiety disorder, social anxiety, and depression (Einstein, 2014; Hong and Cheung, 2015; Carleton, 2016). IU has also been studied with autistic individuals, who obtain higher IU scores than the general population (Chamberlain et al., 2013; Boulter et al., 2014; Neil et al., 2016a). Furthermore, the higher levels of IU might explain the high levels of anxiety that are consistently reported in autistic children (Simonoff et al., 2008; White et al., 2009), as IU was found to mediate the relation between autism and anxiety (Boulter et al., 2014; Wigham et al., 2015). However, the authors caution that a different causal story cannot be ruled out, given the non-experimental data used for the mediation analysis (Boulter et al., 2014).

6.2.3 Individual differences in the need for predictability

As we've seen in the previous chapters, autism is manifested very differently between autistic children. To explain such differences, and in light of the Bayesian accounts of autism, some autistic children should be better equipped to deal with unpredictability than others. Indeed, the mechanism for generating and using predictions appears to differ per individual. Next to being linked to explaining autism (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014), deviations in the prediction mechanism in the brain have been linked to delusions and hallucination in schizophrenia (Fletcher and Frith, 2009; Corlett et al., 2010) and depression (Huys et al., 2015). An individual's ability to deal with unpredictability is also a key component in explaining anxiety and is linked to various anxiety pathologies (for a review, see Carleton, 2016). While the above relates to an individual's ability to generate and use predictions, what someone considered to be predictable may also differ per person, as generating predictions is ultimately an individual process that is influenced by social factors such as the person's desires, affective states, and stereotypes (see Otten et al., 2017). Thus, what is considered to be predictable by one person, may be considered unpredictable by another.

As we are looking to balance the robot's predictability, what constitutes "sufficiently predictable" is therefore likely to vary between autistic children. Indeed, Goris et al. (2020) reported that, when autistic traits are measured in typically developing adults, these correlate significantly with preferences for predictability. Furthermore, autistic individuals' IU is also known to vary between autistic people (Chamberlain et al., 2013; Boulter et al., 2014).

6.3 The role of predicting in human perception and implications for robot predictability

The Bayesian accounts of autism (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014) provide an explanation of the underlying mechanisms for autism, where it is viewed as a special form of predictive processing. The latter is a general framework that provides an explanation of the mechanisms for human cognition. To better understand the concept of predictability and how it

relates to predicting robot behaviour, we will therefore first explore how typically developing individuals may predict robot behaviour, before we consider the special case of autistic individuals again (in Chapter 9). We also noted in the Introduction of this chapter that the predictability of robots may not only be a desirable feature for robots for autistic children, but may also be desirable for robots for typically developing individuals. Thus, the coming sections and chapters are more broadly applicable to people interacting with robots in general.

Depending on the perspective that is taken, various terms are being used to describe (the various stages in) how people understand robot behaviour. This process itself is often referred to as explainability or interpretability (Miller, 2019). In the field of Artificial Intelligence, this process is described in terms of legibility (Busch et al., 2017; Dragan and Srinivasa, 2014b), predictability (Koppenborg et al., 2017), explanation (de Graaf and Malle, 2019; Stange and Kopp, 2020), transparency (Felzmann et al., 2019), readability (Takayama et al., 2011), and anticipation (Gielniak and Thomaz, 2011). In this section, we look at the explainability of robots from a cognitive science point of view, as in the field of cognitive science, there already exists many decades worth of research on how people understand the world around them. Insights from this field may also help us to interpret how people understand robots in specific (Miller, 2019). In cognitive science, perception is generally viewed as a problem of causal inference (e.g. Helmholtz, 1860; Gregory, 1968; Ballard et al., 1983; Dayan et al., 1995). That is, the sensory information received by the brain is caused by objects and processes in the environment. Based on this sensory information, the brain needs to infer what caused it. Predictions, and thus the predictability of objects and processes in the environment, are believed to play a central role in how the brain solves the problem of causal inference (*predictive processing*, see (Rao and Ballard, 1999; Lee and Mumford, 2003; Friston, 2005, 2010; Clark, 2013; Hohwy, 2013)). This is based on the premise that the brain continually generates predictions on what input comes next based on current input and learned associations (Helmholtz, 1860; Gregory, 1968). When describing how people understand robots, we will therefore talk in terms of predictability. In the remainder of this section, we will provide a conceptualisation of predictability and how it relates to people predicting robot behaviour, based on insights from the predictive processing theory.

Before we apply the insights of predictive processing to predicting robot behaviour, we first provide a brief overview of the basics of this theory. For the human brain, inferring what caused sensory information is problematic, as the same cause can give different sensory information (e.g. seeing a robot from different angles), and different causes can provide similar sensory information (e.g. hearing one robot speak using the voice of another robot). To resolve the ambiguities of the various cause and effect possibilities of sensory information, the brain can utilise prior knowledge, making one cause more likely than another. In predictive processing, the brain is believed to do so by constantly attempting to match incoming bottom-up sensory information it receives from the environment with top-down predictions on what information is expected (Mumford, 1992). This is achieved using a hierarchical generative model that aims to minimise surprise (the PREDICTION ERROR), which is the mismatch between the prediction and the sensory input (Friston, 2005, 2010). To infer the most likely

causes of the sensory information (the POSTERIOR), the generative model combines two sources of information, namely the LIKELIHOOD, which is the internal representation of the sensory information, and the PRIOR, which reflects the initial predictions about the likely causes of the sensory information. By combining the prior and likelihood estimates — through an approximation of Bayes' theorem (Knill and Pouget, 2004) — an estimate of the posterior is generated.

What we eventually perceive (the posterior) further depends on the PRECISION associated with the prediction and with the sensory information. Precision is a function of reliability, where low precision sensory information means high ambiguity, which could result from bad lighting conditions, or seeing the robot's behaviour with peripheral vision. Similarly, low precision predictions may reflect that we are not sure what to expect from the robot, as is generally the case when we interact with a robot for the first time. The posterior would then rely more strongly on the internal representation of the sensory information, rather than on expectations encoded in the prior (Wyart et al., 2012). In other words, when people interact with a certain robot (or robot type) for the first time, the precision of their predictions will be low, and their perception of the robot is therefore less influenced by their expectations, and instead relies more strongly on the sensory information coming from the robot (Feldman and Friston, 2010; Wyart et al., 2012). Thus, a robot is inherently less predictable when interacting with it for the first time, but people's predictions also influence their perception to a lesser degree. The predictability of a robot may therefore be less impactful for understanding the robot during such initial interactions than for longer-term interactions.

Whereas predictions shape the perception in the here-and-now, the resulting prediction error shapes future perception. Minimising the prediction error can be done through one of two ways. By using the prediction error to adjust the structure of the generative model, it is possible to improve the precision of future priors and reduce future prediction errors (perceptual inference) (Rao and Ballard, 1999). Prediction errors can also be minimised by performing actions to change the sensory information causing these errors (active inference) (Friston, 2003, 2010; Brown et al., 2011). For instance, when a robot's speech is not sufficiently audible, the resulting prediction errors can be minimised by stepping closer to the robot to hear what it is saying. When predictions are incorrect, the internal models in the brain that generate these predictions are updated accordingly in order to generate more accurate predictions in the future (Yu and Dayan, 2005; Payzan-LeNestour and Bossaerts, 2011). Whether the predictions will become more accurate through this process of learning depends, in part, on the extent in which the robot's past behaviours are indicative of its future behaviour. In other words, the need to be structural regularities in the robot's behaviour. When people are able to discern these regularities, they can be used to update the internal models. The predictability of a robot therefore relates to how easily people can perceive the structural regularities in its behaviour so that they may learn to accurately predict its behaviour — predictability relates to facilitating this learning process. Note that this process of evaluating predictions shows similarities with the concept of explanation (i.e. how people explain robot behaviour).

In predictive processing, the generative models are assumed to be utilised by the

brain throughout the cortex — from low-level to high-level perceptual processing (Wacongne et al., 2011). Sensory information enters the cortex in the lower cortical hierarchies, which concern low-level processing such as perception of colour, slanted lines, or curves. Each higher hierarchical level captures a higher order regularity by building on the previous level. This implies that people also generate predictions on different aspects of the robot's behaviour (Kilner et al., 2007), or in different levels of detail (Kwisthout et al., 2017). Critically, the predictions generated at each of the hierarchical levels cascade downwards along the cortical hierarchies (Wacongne et al., 2011). Thus, when people may predict the robot to take a certain action (a higher-level prediction), they simultaneously may also predict certain robot motions required for this action — predictions are interrelated. For instance, a person may predict certain visual, auditory, or tactile sensory information coming from the robot (low-level predictions), as well as how this sensory information will be combined (e.g. in a grasping movement). Cascading up the hierarchical levels of the brain, predictions may eventually relate to future intentions, social conventions, or world knowledge (Kilner et al., 2007). As these predictions are interdependent, if we expect a robot to perform action 'X', we may also expect it to make sound 'A', motion 'B', and intention 'C'.

Taken together, predictive processing theory posit that the predictability of a robot shapes how people perceive and interact with it. When interacting with a robot for the first time, it will be less predictable, and predictability of a robot is dependent on the expectations one holds. As people interact with the robot, they will learn to predict its behaviour, and predictions will be based on previous interactions with the robot. The predictions relate to all aspects of the robot — from low-level motion primitives to high-level beliefs, desires, or intentions.

6.4 Conceptualising predictability within the context of HRI

6.4.1 Predictability and prediction levels

From the perspective of predictive processing, a relevant aspect of predictability is the ability to quickly and accurately learn to predict the behaviour of a robot. It is the speed and accuracy with which the priors can be learned (high-precision) that distinguishes between robots that are more or less predictable in their behaviour. In this definition of robot predictability, the accuracy of predictions relates to the *predictive value of the robot's behaviour*: the extent to which information about the robot's previous behaviours predicts its future behaviour. When the robot behaviour is highly structured — for instance, when it has little to no variation in its behaviour, or follows a short, predetermined script — previous behaviour will have high predictive value for predicting future behaviour. In addition to the predictive value of priors, predictability can differ in how easy it is to extract structural regularities from the robot's behaviour to form high-precision priors. Thus, a highly predictable robot is one whose behaviour follows a highly structured pattern that is easy to perceive. The more difficult it is to detect these patterns, the longer it will take to learn to accurately and quickly predict the robot's behaviour, where it ultimately becomes impossible to (fully) predict the robot's behaviour.

Regarding *what about* the robot's behaviour is predicted, we distinguish three "prediction levels" that we believe are relevant when designing human-robot interactions. These are based on the assumption that predictions are different in the level of detailed they encode and predictions are about different aspects of a robot's behaviour (Kilner et al., 2007; Kwisthout et al., 2017). Firstly, the "action level", where the goal is to predict the end-state of an on-going action. This level relates to predicting *how a robot will perform its on-going action* in terms of predicting the sensory information the robot will generate (i.e. robot motions, sounds, smells). At this point, people will have some sensory information of the on-going action which can be utilised to generate predictions. It is about predicting what is it going to say, or how will it move. For instance, as is done in research on legibility (e.g. Dragan et al., 2013; Dragan and Srinivasa, 2014b). Secondly, the "interaction level", where the goal is to predict the future action of the robot. Predictions on this level relate to how single robot actions are woven together to form the interaction. For instance, the robot is getting a mug from the kitchen for Filips; thus he may expect its next action will be bringing that mug to him (disregarding the predicted trajectory of this action, which is a prediction on the action level). If the robot were to go to its recharging station to recharge — mug in hand — that would likely be an unexpected action from Filips' perspective, as he may be unaware of the robot's internal state. Lastly, the "activity level", where the goal is to predict the gist of the activity. These are very vague predictions in terms of the details they encode and are related to how interactions are structured in terms of the "tasks", or "topics" within the activity. For instance, a companion robot may start the interaction with an introduction, before proceeding with task one, and then task two. While people may not always be able to predict how the robot will behave on action and interaction levels, they may still be able to predict that whatever actions the robot will perform, and however these actions will be performed, the future actions will be related to the current activity. For instance, during the introduction, one may expect that a social robot greets the person interacting with it, introduces itself, or ask who the person is, as these actions all relate to an (human-human) introduction. When the robot's behaviour is unstructured on an activity level, its future actions may be related to any of a myriad of topics, making it very difficult to predict the gist of the robot's future actions.

As with the hierarchical nature of processing sensory information in predictive processing, we assume that the prediction levels are interdependent. That is, correctly predicting the end-state of an on-going action will aid in predicting the robot's future actions and the structure of the interaction, and vice versa. A robot can also be predictable on one level, but unpredictable on another. When the robot shows highly complex behaviour, it may be difficult to predict the robot's behaviour on action and interaction levels, but it may be consistent in that its actions are related to the task at hand, which makes it predictable on an activity level.

6.4.2 Behavioural predictability and attributed predictability

We might expect that more predictable robots are also *perceived* as such by the people interacting with them. While this may often be the case, it may not always be so. To clearly distinguish between the predictability of a robot and the perception thereof

by people, we will refer to these two concepts as respectively *behavioural predictability* and *attributed predictability*. Results from [Driggs-Campbell and Bajcsy \(2016\)](#) illustrate the difference between attributed and behavioural predictability. In their study, participants were asked to predict when an autonomous car would change lane, where the car's motions were either mechanical/robot-like, or human-inspired. Even though participants were slower to predict the lane-changing behaviour of the car with human-inspired motions, they considered this car to be more predictable than the car with mechanical/robot-like motions. Critically, the results from this study imply that optimising for behavioural predictability may lead to different designs from optimising for attributed predictability.

A key question is how the concepts of behavioural and attributed predictability relate to each other. In general, people are good at dealing with unpredictability, thus some aspects of behavioural predictability (e.g. lower-level predictions on motions versus higher-level predictions on actions or intents) may be more important for the attributed predictability. The context in which a robot performs an unexpected action may also matter for the relationship between behavioural and attributed predictability. For instance, when a person is very familiar with a robot, any unexpected actions will have a larger impact than when a person interact with the robot for the first time. In the latter case, the person does not know what to expect from the robot from prior experiences. This lack of information on which to base predictions results in inaccurate predictions, as well as a high degree of uncertainty in the predictions ([Friston, 2003, 2010](#); [Clark, 2013](#); [Hohwy, 2013](#)). Due to this uncertainty, the predictions will have a reduced influence on perception, and *less attention is given any unexpected sensory information* ([Feldman and Friston, 2010](#)). Considering that these predictive processes are also believed to play a role in people's social perception ([Brown and Brüne, 2012](#); [Tamir and Thornton, 2018](#)), this weighting of the impact of predictions in proportion to their uncertainty on attention may similarly influence the relationship between behavioural and attributed predictability. That is, when predictions regarding the robot's behaviour are known to be uncertain, people pay less attention to any unexpected actions, which may reduce their impact on (social) perception (e.g. attributed predictability) is reduced. We will come back to this question in Chapter 8.

6.4.3 Predictability definitions

To conclude, there are many aspects to robot predictability. Our definitions of predictability, as it relates to HRI, are as follows. Behavioural predictability relates to:

“THE ABILITY TO QUICKLY AND ACCURATELY LEARN TO PREDICT THE BEHAVIOUR OF A ROBOT”

Whereas attributed predictability relates to:

“THE EXTENT TO WHICH A PERSON ATTRIBUTES PREDICTABILITY TO THE ROBOT AS ONE OF ITS QUALITIES”

Thus, predictability is an interaction between the person and the robot, where each person has his/her own initial expectations and learns to predict robot behaviour. In

turn, the robot can facilitate predictability by meeting initial expectations and by allowing the person to quickly and accurately learn to predict its behaviour through adhering to clear structural patterns in its behaviour. Possibly, such increases in behavioural predictability may also lead to increases in attributed predictability. But how do we exactly design “clear structural patterns of robot behaviour”? And when is a pattern of robot behaviour better at signalling its structure? To answer these questions, we will operationalise behavioural predictability in the next section.

6.5 Operationalisation of behavioural predictability as variance in a robot

6.5.1 Behavioural predictability as variance in robot behaviour

Through interacting with the robot, people become more familiar with its behaviour, improving their ability to predict its behaviour (Dragan and Srinivasa, 2014a). However, this requires one to perceive and learn the *structural regularities* in the robot’s behaviour — also referred to as invariance detection (Gibson and Pick, 2000; Gogate and Hollich, 2010). These structural regularities hold predictive value in that they can be used to generate predictions regarding the robot’s behaviour in the future. VARIANCE in robot behaviour can therefore lead to interactions with a robot that are too complex to be learned, and therefore predicted, or it can take the child longer to learn and predict. Designing a robot to be highly predictable then might entail programming robot behaviour with low variance such that it has high structural regularities that can be easily inferred from its behaviour.

6.5.2 Types of robot variance

Variance in robot behaviour can stem from any of the robot’s expressive modalities, as well as how those modalities are used within the interaction dynamics. There are several types of variance that we consider in this dissertation, namely variance in speech, motion, topic, and time:

SPEECH VARIANCE refers to variability in the words that the robot uses (diction) and how those words are spoken (prosody). Diction variance can be operationalised as using different sentence grammars and/or different words that have similar meaning. Prosody variance relates to using different prosody and/or different voice actors to change intonation, stress, rhythm and pitch of the sound. While the use of prerecorded speech can minimise speech variance, using models for natural language generation to generate dynamic speech may influence the variance in the robot’s speech significantly.

MOTION VARIANCE refers to variability in the robot’s motions, such as facial expressions and gestures. A robot could consistently use the same static animation to communicate a certain intent, which would keep the motion variance low. In contrast, when using models to generate dynamic motions, such as inverse kinematics, each robot motion can be unique in its trajectory, increasing variance in the robot’s motions.

TEMPORAL VARIANCE refers to variability in the timing of robot behaviour, both within actions (e.g. timing of a motion trajectory) as well as the timing between actions. While this type of variance is not a type of variance that is actively designed for, temporal delays in robot behaviour are likely to occur as the result of a lack of computational power, or due to the physics of the robot's motors, increasing temporal variance. Such temporal delays can disrupt the interaction of an autistic child (Ferrara and Hill, 1980).

TOPIC VARIANCE refers to variability in the use of different actions that address different topics at a particular point in an interaction. For example, at the first step of an interaction, increased topic variance could involve a robot saying either a greeting or a comment about a bird, as opposed to consistently saying a greeting. Types of topic variance related to a robot can include: (1) when a robot responds to its internal state (e.g. battery level, or its (faulty) perception) which the observer doesn't know about; (2) when a robot responds to an external event (e.g. a person walking in on the ongoing session) or item (e.g., a clock) in the environment; and (3) when a robot responds with an action that is not legible (i.e., understandable) to the observer (e.g., randomly saying "beep beep").

This is not an exhaustive list of variance types. For instance, the robot could also show variance in its appearance through wearing different clothing each day. However, the listed types of variance are the ones that we used to manipulate the behavioural predictability of the robot in the study that we will discuss in Chapter 9.

Finally, in addition to these types of variance, a robot can combine modalities into one action. For example, a robot capable of facial expressions of emotion can more clearly communicate these expressions when they are combined with short, emotional non-speech expressions (Schadenberg et al., 2018). This can also lead to variance when different combinations, intended to express the same multimodal intent, are perceived as different robot actions. For autistic children, processing multimodal information may be more difficult than for typically developing children (Happé and Frith, 2006; Collignon et al., 2013; Ostrolenk et al., 2019). Furthermore, autism professionals reported that presenting multimodal information might cause an information overload (Huijnen et al., 2017).

6.6 Reducible and irreducible unpredictability in robot behaviour

While prioritising the minimisation of the types of robot variance mentioned above could lead to highly predictable robots, doing so severely restricts the types of applications for robots for autism. For instance, by programming it to do the same thing, in exactly the same manner, again and again (e.g. by having a very limited repertoire of robot behaviours). For such robots, autistic children may quickly learn to predict their behaviour because there are only a few behaviours to predict. They can also offer endless repetitions of certain behaviours, which can be useful in certain use-cases. While such robots may play a role in the broader scope of robots for autism (Diep

et al., 2015), a *degree* of behavioural unpredictability is unavoidable for most of the robot roles envisioned by autism professionals (Huijnen et al., 2019; Alcorn et al., 2019) and researchers (Diehl et al., 2012; Scassellati et al., 2012; Cabibihan et al., 2013; Diehl et al., 2014). For instance, when the robot's role is that of a trainer or educator, where its task is to model, teach, and/or practice a targeted skill and to provide feedback to the child. When the robot is positioned as a social interaction partner, it will have to perform novel behaviour and in general have a larger range of behaviours, in order to deal with the complexity, dynamics, and unpredictability of interacting with a human. The resulting unpredictability of the robot's behaviour is *irreducible*, as it is required for the robot to perform its primary task. That is, even when the robot is carefully designed in terms of its predictability, irreducible unpredictability is unavoidable and does not lead to design considerations. For a robot that is positioned as a tutor, an example of this would be the behavioural unpredictability that can stem from introducing (novel) learning content. This results in novel actions that will be unpredictable initially, but without such actions, a robot cannot present any learning content and thus be considered a tutor.

There are also technical reasons why a robot is not always behaving predictably, in particular for robots that behave autonomously, as the technical systems that allow the robot to perceive, reason, and/or act, can each introduce variance. The robot may respond to what it is perceiving, which can be faulty, or the observer can be unaware of what the robot is responding to. The robot's reasoning may be too complex, or faulty, to be understood by the observer. Or, the intent the robot tries to convey may simply not be understood by the observer from the robot's behaviour. In all of these examples, the resulting robot behaviour is unexpected and was not predicted. As the resulting unpredictability is (to a certain extent) inherent to using technical systems, we also consider this to be irreducible unpredictability. In part, this can be solved by utilising a Wizard-of-Oz paradigm, where a person remotely controls the robot and creates the illusion of interacting with an autonomous robot, as this reduces the need for certain technical systems. However, for large-scale and long-term use of robots, the use of a Wizard-of-Oz system is not practical and requires that the robot can operate autonomously (Scassellati et al., 2012; Clabaugh and Matarić, 2019). The autism professional using the robot is too busy paying attention and interacting with the child to also control the robot, and we deem it unlikely that robot-assisted interventions provide enough value to warrant paying an additional person to control the robot.

In contrast to irreducible unpredictability, there are also aspects of robot interaction design that we consider to be *reducible*. These are instances where adding one feature may improve a certain aspect of the interaction, but may decrease the robot's behavioural predictability. They are design choices, rather than a necessity for the robot to perform its primary task. For instance, common design principles for achieving long-term interactions for typically developing children include using intelligent tutoring systems to provide the optimal challenge and keep children motivated (Ramachandran et al., 2017; Schadenberg et al., 2017), adding variance to speech to reduce boredom due to endless repetition of a robot's behaviour (Kidd and Breazeal, 2008; Kory and Breazeal, 2014; Coninx et al., 2015), or personalising speech (Lubold

et al., 2016; van Straten et al., 2018). For a tutoring robot, these features are desirable, but are not necessarily required for the robot to perform its primary function. In the case of robots for autistic children, choosing whether to implement such non-essential features that increase unpredictability is a balancing act. When the robot becomes too unpredictable it may have negative effects on the interaction and could even negate the intended effect of feature.

6.7 Conclusion

To summarise, the predictability of robots may make it easier for autistic children to engage in learning in a robot-assisted intervention and maintain this engagement. What constitutes “sufficiently predictable” is likely to differ between autistic children, due to how highly idiosyncratic autistic children are, and how this affects their interactions within a robot-assisted intervention. This also means that there is no one-size fits all solution to keeping them engaged in robot-assisted interventions whilst providing meaningful learning.

While predictability is an important concept in autism and human-robot interaction, there is no common understanding about the concept of predictability and there are various ways to approach this concept — it is an umbrella concept. To clearly distinguish between predicting robot behaviour, and considering a robot to be predictable, we refer to the former as behavioural predictability and the latter as attributed predictability. We applied insights from predictive processing to the concept on (behavioural) predictability and concluded that (1) there is inherent unpredictability during initial interactions, (2) predictability relates to a learning process, (3) predictability has a larger impact on perception when someone is already familiar with the robot, and (4) different types of predictions are made by the brain. From these insights we extracted our definition of behavioural predictability as it relates to robots — “the ability to quickly and accurately learn to predict the behaviour of a robot”. We distinguish three levels of what about the robot’s behaviour is predicted, namely predictions on an action, interaction, or topic levels. To learn to predict the behaviour of a robot, a person will need to perceive and learn the structural regularities in the robot’s behaviour. We therefore concluded that adding variance to a robot’s behaviour can make it more difficult for a person to learn the structural regularities, increasing the robot’s behavioural unpredictability. There are four types of variance that are relevant for this dissertation, namely speech, motion, temporal, and topic variance. The latter type of variance relates to generating predictions on a topic level, whereas the other types are on an action or interaction level.

Finally, we briefly discussed the difference between reducible and irreducible behavioural unpredictability. Robots cannot behave fully predictably as well as provide meaningful learning, as the learning gains may not generalise to the less predictable world of people (Alcorn et al., 2019). Moreover, a fully predictable robot would perpetuate in repetitive behaviour (Dautenhahn, 2007), limiting its long-term usefulness, and would also be unable to autonomously respond to the (unpredictable) dynamics of real-world settings (Clabaugh and Matarić, 2019). This leads to the question of what length we should go to for designing robot’s that are highly predictable.

The remainder of this part of the dissertation is structured as follows. In Chapter 8, we will discuss an experimental study that we conducted with typically developing individuals. The goal of this study was to get a better understanding of how behavioural and attributed predictability relate to each other, and how predictability influences people's social perception of a robot. The study was conducted with typically developing individuals, as our measurement tools for behavioural and attributed predictability were too complex for autistic children with high support needs. Given the similarities between uncertainty and predictability, we looked further into these two concepts and included a measure for Intolerance of Uncertainty. This measure could possibly help us to identify to what extent a person is sensitive to unpredictability, and account for the likely differences therein between autistic children. In Chapter 9, we report on our analyses of the final study that was conducted by the DE-ENIGMA consortium, using the DE-ENIGMA intervention. The goal our analyses was to investigate how the robot's predictability influences the engagement of autistic children in a robot-assisted intervention. But first, we will discuss the development of a new scale for measuring attributed unpredictability in the next chapter. This measurement was used in the study that we will discuss in Chapter 8.

7

Measuring robot predictability

This chapter is based on the following article:

Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., & Evers, V. (in press). I See What You Did There: Understanding People’s Social Perception of a Robot and its Predictability. *ACM Transactions on Human-Robot Interaction*. 27 pages. ([Schadenberg et al., in press](#))

To study the effects of a robot’s predictability, we need to be able to measure how well participants could predict the robot’s behaviour, as well as the extent to which a person associates the attribute of “(un)predictability” with the robot. In this chapter, we describe how we developed a new scale for measuring the attributed unpredictability of a robot.

7.1 Measuring the predictability of a robot

BOTH the behavioural and attributed predictability of robots have been studied and reported. In such studies, behavioural predictability has been measured through behavioural observations, for example by asking participants to make a prediction (e.g. [Driggs-Campbell and Bajcsy, 2016](#); [Takayama et al., 2011](#); [Lichtenthäler et al., 2012a,b](#); [Walker et al., 2016](#)), or subjectively through self report, by asking participants how predictable they judge the robot to be (e.g. [Dragan and Srinivasa, 2014a](#); [Dragan et al., 2015](#); [Mubin and Bartneck, 2015](#)). Attributed predictability either through a single item ([Mubin and Bartneck, 2015](#)), or multiple items specific to predictability of robot motion ([Dragan and Srinivasa, 2014a](#); [Dragan et al., 2015](#)). Single-items are more vulnerable to random measurement errors and unknown biases in the meaning and interpretation of that item. With multiple-item scales, the random measurement error is more likely to be cancelled out. Moreover, they cover a broader range of meanings of a construct, which can reduce the effect of

differently interpreting an item. We therefore developed a new multi-item scale that measures to what extent unpredictability is attributed to a robot, and which is not restricted to the predictability of robot motions.

This chapter briefly describes how we developed the new multi-item scale for measuring attributed unpredictability, and the final scale.

7.2 Materials and methods

To populate our multi-item scale of attributes that may indicate a robot's predictability, we first looked at English synonyms, antonyms, and related words for “predictable” and “unpredictable” using Merriam-Webster's thesaurus ([Merriam-Webster, 2018](#)), resulting in 22 items. We then asked participants to rate these according to how well they fit their intuitive concept of predictability. We analysed their answers to arrive at a reduced set of items to be used in our main study.

7.2.1 Participants

Participants were recruited through Amazon's Mechanical Turk (mTurk), which resulted in 99 adults (42 men, 57 women) participating. Three male participants were excluded from the analysis, because they indicated that they did not know what “predictable” or “unpredictable” means. Participants received \$0.50 for participating. To meet the recruitment criteria, participants had to be from the United States of America, have completed 500 or more assignments on mTurk, and have an approval rate of 99% or higher on mTurk.

7.2.2 Procedure

We informed participants that we are developing a measurement tool for measuring the “predictability” of a robot. To that end, we wanted to know to what extent they associated certain words with the general concept of a robot that is “predictable” and a robot that is “unpredictable”. For items related to “predictable”, participants were asked “Using the scale provided, how closely are the words below associated with a robot that is predictable”. The same question was used for items related to “unpredictable”, but with the different phrasing emphasising that the robot is unpredictable. The provided scale was a 7-point Likert scale, from 1 = *definitely not associated* to 7 = *definitely associated*. In case the participant did not know the word, they could choose *I do not know what this word means* instead. The items were presented in a randomised order.

We intentionally did not provide any description or visualisation of robots, as the questionnaire should serve as a general measure of the attributed predictability of robots. Moreover, we did not provide any definition of predictability, taking it as a type of *projective latent content* where it is assumed that most people understand, and to a certain degree share, a common meaning of the construct and one wants to approach that intuitive meaning rather than limit the user to restrictive definitions ([Potter and Levine-Donnerstein, 1999](#)).

Table 7.1: The component loadings for the items on unpredictability and the mean scores for how closely each items is related with an unpredictable robot.

Item	<i>M (SD)</i>	Component 1	Component 2
Unpredictable	6.55 (1.27)	.98	-.24
Irregular	6.09 (1.36)	.85	.04
Inconsistent	6.35 (1.13)	.84	-.06
Random	6.35 (1.11)	.78	.05
Variable	6.29 (0.96)	.65	.12
Erratic	6.28 (1.14)	.64	.35
Unreliable	5.98 (1.37)	.50	.41
Inconstant	5.93 (1.37)	.47	.53
Skittish	4.32 (1.87)	-.07	.87
Fickle	4.32 (1.74)	.02	.84
Eigenvalue		5.63	1.24
% of variance		56.31	12.38
Cronbach's α		.90	.72

7.3 Results

Prior to performing the analysis, we removed 4 items (mutable, immutable, mercurial, and capricious) for which more than 5% of the participants answered that they did not know what the words meant. We then conducted a Principle Component Analysis (PCA) on 18 items with varimax rotation. This resulted in four components, of which two were related to the items for predictability, and the other two were related to unpredictability. Given that unpredictability is the inverse of predictability, we expected one component on which both predictability and unpredictability items would load. Instead, it is likely that they loaded on different components because of asking both positively and negatively worded questions. This can produce an artefactual two component structure in the absence of two constructs (Marsh, 1996; Spector et al., 1997; Rodebaugh et al., 2004). Therefore, we proceeded with only the items related to unpredictability, for which we had more items than for predictability.

We performed a PCA on 10 unpredictability items (see Table 7.1) with varimax rotation. The KMO measure verified the sampling adequacy for the analysis (KMO = .87), and for the individual items the KMO values were $> .82$. Bartlett's test of sphericity indicated that correlations between items were sufficiently large for PCA ($\chi^2(45) = 577.90, p < .001$). Two components had eigenvalues above Kaiser's criterion of 1. These two components explained 68.67% of the variance. Correspondingly, the scree plot (see Figure 7.1) shows a point of inflexion at the third component. Table 7.1 shows the mean scores on how closely each item is associated with an unpredictable robot, and the component loadings for each item, where we highlighted items with loadings of .6 or higher and with a cross loading difference greater than .2. Component one seems to represent unpredictability associated with robots, as the mean association scores are high for these items. Component two contains items that had

Figure 7.1: Scree plot of the PCA on the unpredictability items.

relatively low association scores. It may therefore be indicative of a type of unpredictability that is less suited to ascribe to a robot, and was discarded from our set of items.

7.4 Final multi-item scale for attributed unpredictability and its limitations

The final scale for measuring attributed unpredictability of a robots asks participants how closely they consider each of the six items to be associated with the robot. The six items include “unpredictability”, “irregular”, “inconsistent”, “random”, “variable”, and “erratic”. In the study reported in the next chapter, we used our new scale to measure *attributed unpredictability* by taking the average of those items as the score for attributed unpredictability.

The scale would benefit from being further validated. Our goal was to develop a multi-item scale for measuring the attributed predictability that we could use in our future study (described in the next chapter), as such scales are less prone to random measurement errors and biases in the meaning and interpretation than is a single item. However, we did not assess whether our new scale is more sensitive and specific than measuring attributed predictability through a single-item, and thus cannot claim that our scale has better validity. Nevertheless, given our assessment of the scale, we felt confident that the scale has sufficient validity and reliability for us to use it in other studies.

Note that, as this is a dictionary study, conducted with participants from the United States, these six words may have different meanings to people from other countries. When using this new scale in other countries, we need to be mindful of this. Further-

more, as our scale relies on the participant's ability to read, understand the question, and know what the six items means, it is too complex to use for our target group of autistic children; those with limited language and high support needs. Although we do not assess this in this dissertation, it is likely also too complex for (autistic) children in general, which is another thing to be mindful of when using this new scale.

8

Understanding people's perception of a robot and its predictability

This chapter is based on the following article:

Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., & Evers, V. (in press). I See What You Did There: Understanding People's Social Perception of a Robot and its Predictability. *ACM Transactions on Human-Robot Interaction*. 27 pages. ([Schadenberg et al., in press](#))

As we have discussed in Chapter 6, generating predictions plays a central role in human perception, according to Bayesian accounts of human cognition. Furthermore, the concept of predictability has been argued to influence our ability to understand and coordinate our actions with robots, task performance, and people's trust in the robot. Even so, unpredictable actions are bound to happen at some point when interacting with a robot. But what happens when a robot performs an unpredictable action during social interaction with a person? How does that influence their social perception of the robot? This is the topic of the experimental study described in this chapter. Note that this study was with (typically developing) adults, rather than with autistic children.

8.1 Introduction

THE distinction between reducible and irreducible unpredictability (see Chapter 6.5) leads to the questions to what length we should go to for reducing the reducible unpredictability. For some robot applications, unpredicted actions will have clear (severe) negative effects, such as leading to dangerous situations, or reducing trust in the robot. For such cases, it is then vital to improve the robot's predictability and minimise reducible unpredictability. However, for robots used for

social interaction with typically developing individuals, results reported in literature are not as clear cut.

With robots used for social interaction, a *degree* of behavioural unpredictability may have positive effects. [Salem et al. \(2013\)](#) asked their participants to help a robot with unpacking boxes. The robot assisted the participant by providing information on where to put the items, where it showed either unimodal, congruent multimodal, or incongruent multimodal behaviour. The results indicated that despite decreasing task performance, the robot which showed incongruent multimodal behaviour was preferred by participants, in that it was considered more humanlike, more likeable, and the participants experienced greater shared reality with the robot. The authors argue that the incongruency of the robot's multimodal behaviour led to a degree of (behavioural) unpredictability, which was explained by participants attributing mental states to the robot, such as that the robot was "cheeky" or tried to fool them.

In a similar vein, robots that sometimes behave unpredictably have been found to increase behavioural engagement of adults ([Short et al., 2010](#)) and children ([Salter et al., 2008](#); [Lemaignan et al., 2015](#)). In the study conducted by [Short et al. \(2010\)](#), participants had to interact with a robot and play a rock-paper-scissors game. At a certain point in the interaction, the robot would verbally conclude that it had won the game, while the rules of the game dictated that it had lost. It did so three out of twenty times the game was played, making it inconsistent in its behaviour and therefore more difficult to predict its behaviour. The results indicated that while the inconsistent robot was judged as being less fair than when it showed consistent behaviour, it also increased behavioural engagement.

The positive effects of behavioural unpredictability are explained by arguing that it makes robots more humanlike, which in turn increases the attribution of mental states to the robot ([Salem et al., 2013](#)). Intuitively this explanation makes sense, as humans are unpredictable in their behaviour. Highly predictable behaviour may also be considered as boring, whereas unpredicted actions could spark interest ([Dautenhahn, 1999](#)). However, also in social interaction, unpredictable robot behaviour can lead to attribute negative attributes to the robot. [Salem et al. \(2015\)](#) compared a non-erroneous robot with a robot that showed erratic behaviour. The latter was operationalised by navigating in an erratic manner and by playing a song that was not selected by the participant. Participants considered the erratic robot to be less reliable, less technically competent, less understandable, and less trustworthy. Interestingly, the erratic robot was not perceived as more humanlike in this study.

One issue with studies on predictability and robots is that predictability is operationalised very differently. As a result, the unpredicted actions have different consequences for the person interacting with the robot. For example, the person may not be able to smoothly carry out their task because the robot does not comply with his/her commands ([Lemaignan et al., 2015](#); [Mubin and Bartneck, 2015](#)), is making mistakes ([Lemaignan et al., 2015](#); [Mubin and Bartneck, 2015](#)), or is illegible in its actions ([Salem et al., 2013](#)). In case of robots showing erratic navigation, while unpredictable, this is also not behaviour that would naturally occur in humans. This might explain why these behaviours were not perceived as being more human-like, which is argued to be the reason why a degree of unpredictability may lead to pos-

itive effects. Therefore, concluding that it is the robot's *predictability* that influences people's social perception of the robot, seems premature.

While the studies above made the robot differ in its behavioural predictability, this was not the main purpose of the studies — it was not assessed whether the behavioural predictability was successfully manipulated, nor was the attributed predictability assessed. It is also unclear to what extent the behavioural and attributed predictability led to the reported findings. All things considered, our current conceptual understanding of robot predictability — as it relates to HRI — cannot explain how the predictability of a robot affects people's social perception of a robot, nor how the behavioural and attributed predictability are related to this. This limits us in effectively taking robot predictability into account in the design of robot behaviour.

In this chapter, we report on a study where we investigated how the visibility of the cause of a robot's responsive actions influences the following: behavioural predictability, attributed unpredictability, an observer's social perception of the robot and whether this is mediated by the behavioural and attributed predictability, and how an observer's Intolerance of Uncertainty (IU) influences the observer's perception of the robot. The study was an online video-HRI study, where participants watched a video of an interaction between a person and a robot unfold. According to the predictive processing framework (see Chapter 6.3), people constantly generate predictions in order to make sense of sensory information. Thus this also applies to watching a video of a human-robot interaction. Moreover, the video-HRI methodology has been used extensively to study predictability and legibility (e.g. [Takayama et al., 2011](#); [Dragan and Srinivasa, 2013](#); [Dragan et al., 2013](#); [Dragan and Srinivasa, 2014a](#); [Lichtenthäler and Kirsch, 2014](#); [Nikolaidis et al., 2016](#)).

8.2 Research questions and hypotheses

8.2.1 Predictability operationalisation and manipulation

In the current study, participants watched a video of an interaction between a person and a robot unfold. While the participants will not be able to predict the robot's behaviour initially, the interaction follows a highly structured format. The structural regularities in the robot's behaviour should therefore be relatively easy to discern. For participants who are able to do so, this would presumably improve the accuracy of their predictions related to the robot's behaviour. However, for some participants, depending on the condition they are in, the robot will perform actions that cannot be predicted. These unexpected actions might make it more difficult to discern the structural regularities in the robot's behaviour, resulting in less accurate predictions.

The unexpected actions are operationalised through responsive actions (a type of topic variance). These could either be predicted based on the situational context, or could not be predicted because of a lack thereof. The situational context was manipulated by whether the participants could see what event the robot was responding to (the visibility of the cause; our independent variable). The events — the cause of the robot's responsive action — were either visible, obstructed from view, or there were no events. We consider this as a natural and realistic way to introduce unpredictability in the robot's behaviour, as in real interactions with a robot it may not always be

apparent to what a robot is responding to.

When the cause is visible, participants can explain and may even predict the responsive actions. In turn, this should make it easier to derive that the responsive actions are an exception (i.e. they hold no predictive value for future behaviour) and make it easier to discern the structural regularities in the robot's behaviour. When the cause is not visible, either because the event was obstructed or there was no event visible, the robot's action cannot be predicted, nor is there evidence regarding an explanation as to why the robot performed this action. This should make it more difficult to derive the structural regularities in the robot's behaviour. In case of the cause being obstructed from view, the participant has perceptual information that the information of the scene is incomplete, and it is therefore ambiguous whether the robot is responding to an event or not. It is therefore made explicit that any predictions are based on limited information. This might increase the uncertainty of any predictions and thereby reduce the impact of predictions on the (social) perception of the robot. For this condition, we are interested in understanding how knowing that you cannot fully understand the reasons behind the robot's behaviour influences a participants' social perception of the robot and the attributed predictability.

8.2.2 Research questions

The first set of research questions relates to the robot's predictability in relation to the visibility of the cause of responsive robot actions. Inspired by the findings of [Driggs-Campbell and Bajcsy \(2016\)](#), we are interested in how unpredictable robot actions influence the behavioural and attributed unpredictability depending on whether the participant see a cause for the robot's unexpected behaviour.

Topic 1: The first questions concern the relation between the *visibility of the cause of responsive robot actions* and the *robot's predictability*.

RQ 1a To what extent does the visibility of the cause of responsive robot actions influence the *behavioural predictability* of the robot?

RQ 1b To what extent does the visibility of the cause of responsive robot actions influence the *attributed unpredictability* of the robot?

We assume that when the cause of the responsive actions is not visible, the actions are unpredictable and make it more difficult to assess the structural regularities in the robot's behaviour. That is, the unpredictable actions do not have predictive value for predicting future actions, and hence should be discarded. We therefore expect that participants are less accurate and less confident in their prediction when the event was obstructed or not there, compared to when the event is visible — the behavioural predictability is reduced. Our hypotheses related to the attributed unpredictability are based on the notion of weighting the impact of predictions according to their uncertainty ([Friston, 2003, 2010](#); [Clark, 2013](#); [Hohwy, 2013](#)), and how this directs people's attention ([Feldman and Friston, 2010](#)) (see Section 6.3). We expect that the participants will associate unpredictability with the robot more strongly when the cause of the responsive actions is not visible (because the event is not there or is obstructed), compared to when it is. Presumably, uncertainty regarding predictions

reduces their impact on (social) perception. We therefore expect that in the ambiguous condition, where participants know they do not have full perceptual information of the scene, such ambiguity will mitigate their attribution of unpredictability. Regarding the relation between behavioural predictability and attributed unpredictability, we do expect that they are correlated with each other.

Our second set of research questions relates to the participants' social perception of the robot. Social perception was measured in terms of warmth, competence, and discomfort. While discomfort is a dimension unique to robots (Carpinella et al., 2017), warmth and competence are considered universal dimensions of human social cognition and are used to judge whether another person is a threat or a friend (Fiske et al., 2007). Together, these dimensions are believed to explain how people characterise others (Wojciszke et al., 1998). This leads to our second research question and sub-research questions:

Topic 2: The second set of research questions concerns *the impact of the visibility of the cause of responsive robot actions on the participants' social perception of the robot, and the extent to which this is mediated by the robot's predictability*

RQ 2a To what extent does the visibility of the cause of responsive robot actions influence a participant's social perception of the robot (i.e. what social attributes are attributed to the robot)?

RQ 2b To what extent does the robot's *behavioural predictability* mediate the social perception of the robot?

RQ 2c To what extent does the *attributed unpredictability* mediate the social perception of the robot?

We expect that when the cause of the responsive actions is not visible, the robot will negatively influence the participants' social perception of the robot, compared to when the cause is visible. Moreover, we expect that when participants know that they do not have full perceptual information, this will mitigate a reduction in warmth and competence and an increase in discomfort. Again, these hypotheses are based on the notion of weighting the impact of predictions according to their uncertainty (Friston, 2003, 2010; Clark, 2013; Hohwy, 2013), and how this directs people's attention (Feldman and Friston, 2010). Based on the results from Driggs-Campbell and Bajcsy (2016), we expect that the participants' social perception of the robot is mediated more strongly by the participants' attributed unpredictability of the robot. Given that we expect that behavioural predictability and attributed unpredictability are correlated, we also expect such a mediation effect for behavioural predictability, although one that is less strong than for attributed unpredictability.

The last of our research questions relate to a possible moderation effect of the participants' IU, given the relation of this construct with predictability (Grupe and Nitschke, 2013) and populations that are known to have difficulty dealing with unpredictability, such as autistic children (Chamberlain et al., 2013; Boulter et al., 2014; Neil et al., 2016a). This leads to our third research question and sub-research questions:

Topic 3: The final set of research questions concern the impact of the *participants' intolerance of uncertainty* on the (possible) *relations between the visibility of the cause of responsive robot actions and the participants' social perception of the robot or the robot's predictability*.

RQ 3a To what extent does the participants' intolerance of uncertainty moderate the *attributed unpredictability* of the robot?

RQ 3b To what extent does the participants' intolerance of uncertainty moderate the participants' social perception of the robot?

Given that those higher in IU respond more negatively to instances where they have insufficient information (Carleton, 2016), we expect that they will also respond more strongly to the robot's responsive actions when the cause is not visible in terms of their social perception of the robot and attributed unpredictability.

8.3 Materials and methods

8.3.1 Participants

Participants were recruited through Amazon's Mechanical Turk platform. In total, 169 adults (74 men, 94 women, 1 non-binary) participated through mTurk. They received \$1.20 for participating. To meet the recruitment criteria, participants had to be from the United States of America, have completed 100 or more assignments on mTurk, have an approval rate of 99% or higher on mTurk, and not have participated in our study for the questionnaire development reported in the previous section. Participants were asked how well they could hear the robot on a 7-point Likert scale, from 1 = *I didn't understand anything the robot said* to 7 = *I understood every word the robot said*. Overall, participants could understand the robot's speech well ($M = 6.48$, $SD = 0.70$). Eight participants who rated the robot's speech with a 4 or lower were excluded from the analysis, because it was essential for this study that they could clearly hear what the robot was saying. Another four participants were excluded from the analysis because their completion time was too fast to realistically assume they watched the video until the end. Their answers also corroborated this assumption. We also asked participants to what extent they had prior experience with robots, and included this factor as a covariate in our analysis. However, this factor did not influence any of the outcome variables, and thus we included participants regardless of their prior experience with robots. As a result, the analyses reported in this paper are based on the responses of 153 adults (62 men, 90 women, 1 non-binary).

This study was reviewed and approved by the ethics committee of the faculty of Electrical Engineering, Mathematics, and Computer Science of the University of Twente, and is registered under reference number "RP 2018-52". Prior to participating in the study, participants first had to sign an informed consent form.

8.3.2 Materials

For this study, we again used Robokind's R25 humanoid robot called "Zeno", or "Milo" (see Chapter 1.3). However, the software that we used was different from the DE-

ENIGMA system. Rather than using prerecorded human speech for the robot's speech, we generated it by the integrated Acapela Text-To-Speech engine with the American English children's voice "Josh". We also included lip synchronisation. The interaction was recorded with a 1080p camera and lasted 1 minute and 15 seconds. To improve the audio quality, we recorded it separately for the robot and for the adult using two microphones, and combined the audio feeds with the video feed afterwards.

8.3.3 Interaction design

Table 8.1: The interaction and structure of the storytelling game. The adult's speech is in italics, and non-verbal actions are in enclosed by square brackets.

Structure	Content		
	INTRODUCTION		
Introduction	“Hi, my name is Zeno.” [Robot waves] “Shall we start?” “Okay.” *Mask is applied smoothly*		
	STORYTELLING GAME		
	STORY PART 1	STORY PART 2	STORY PART 3
Story intro	“You stand in front of a small door.” *First event here*	“There is a distinct smell in this room.”	“When you look around, you see a pizza lying on the table.”
Story question	“How tall are you?”	“What is your favourite food?” *Second event here*	“Are you hungry?”
Adult’s answer	“Six foot.”	“Pizza.”	“Yes, I am.”
Story resolution	“You open the door, but have to crouch to get through it. You enter a kitchen.”	“The smell reminds you of what a pizza smells like.”	*Mask is removed, revealing the adult* and *Video ends*
	EVENTS AND CORRESPONDING RESPONSIVE ACTIONS		
First event	[Adult receives text message and views it on his mobile phone.] “Could you please put that away.” [Adult complies.] “Thank you.”		
Second event	[Another person taps the adult on shoulder. The robot turns head to the new person.] “Please go away, we are busy here.” [The other person leaves. The robot turns head back to the adult.] “Thank you. Let’s continue. What is your favourite food?”		

The participants watched a short interaction of a person interacting with a hu-

manoid robot. This *interaction* was the same for *all* participants. In the video, the robot is playing a storytelling game with an adult, where the robot tells a story in which the adult is the main character. This interaction was unrelated to the DE-ENIGMA project, and we developed it solely for this study. The game proceeds in a very structured manner: the robot tells something, asks input from the adult, and then uses his answer to proceed with the story in an obvious direction. This basic pattern is repeated 3 times, and forms the main structural regularity in the interaction that can be used to predict the robot's actions. The participants viewing the video received no information telling them what the interaction was about in order to prevent priming them to pay too close attention to regularities in the interaction.

The interaction pattern (what happens in the video) can be seen in Table 8.1. The interaction consists of an introduction, the storytelling game, and two events to which the robot responds. At the start of the video, the robot introduces itself and asks whether the person would like to start. The robot then starts telling the story, which is about the adult entering a kitchen and finding his favourite food. The story is divided into three parts that each follow the same fixed interaction pattern. First, the story proceeds. Second, the robot asks the adult a question related to the subject of the previous sentence about the story. Finally, the robot uses the adult's answer to resolve what happens with the adult and the subject. The video ends prior to the robot's response to the adult's answer to the third question ("Are you hungry?").

Twice during the interaction, the robot would display a responsive action to an event (the cause) that occurred within the situational context. The events are (a) the adult receives a text message on his cellphone and starts reading it, and (b) another person walks in and distracts the adult from the interaction with the robot.

8.3.4 Experimental design

The study was set up as a 3 by 1 between-participants design. The independent variable in our study is the visibility of the cause of the robot's responsive actions (which we will reference to as *visibility of the cause* for brevity), which had three levels. We measured four dependent variables, namely behavioural predictability, attributed predictability, robot social attributes, and the intolerance to uncertainty.

Screenshots of the video for each condition can be seen in Figure 8.1. In one condition, the events were visible, and therefore the responsive actions could be predicted. We refer to this condition as the *visible cause* condition. In the other two conditions, the events were not visible. This was either because there were no events, or because we limited the view of the participant shortly after the start of the interaction. We refer to the former as the *no cause* condition. In this condition, the responsive actions cannot be predicted. Moreover, the adult in the video who is interacting with the robot shows no response to the robot's responsive behaviour. This could indicate to the participant that there is simply no event that explains the robot's responsive actions. In the condition where we limited the view of the participant, the responsive actions are not predictable, but participants know that they do not have full perceptual information of the interaction. We limited the view by applying a mask that obstructed the view of the adult interacting with the robot. The mask was applied smoothly just after the introduction. This way, the participant knew that there was a person inter-

(a) VISIBLE CAUSE: perceptual information regarding the cause of the robot's responsive actions are visible.

(b) AMBIGUOUS: no perceptual information regarding the cause of the robot's responsive actions is visible because the event is obstructed by a mask.

(c) NO CAUSE: no perceptual information regarding the cause of the robot's responsive actions is visible. Instead, the adult provides perceptual information that there is no cause as he does not respond to the responsive actions.

Figure 8.1: The three conditions of the study, which differ in whether the participant can see the cause of the robot's responsive actions. In this case, another person distracting the adult from the interaction with the robot.

acting with the robot. After the mask was applied, participants would still hear the adult reply to the questions of the robot. We refer to this condition as *ambiguous*, as the robot may or may not be responding to actual events. Note that the interaction followed the exact same script for each condition with the same two events. We only manipulated the perceptual information available to participants.

8.3.5 Procedure

For the study itself, participants were instructed that they would watch a video of a person interacting with a robot and that afterwards they would need to answer a couple of questions regarding the content of the video, their impression of the robot, and some background questions. Prior to viewing the video, participants saw a test video where the robot would speak for 2 minutes, so that they could adjust their volume settings if this was required to hear the robot's speech clearly. Next, participants viewed one of three videos (one for each condition), which was selected randomly. After viewing the video, we first asked the participants to write down the last thing the robot said. This was followed by questions on behavioural predictability,

how well the participants understood what the robot was saying, the Robotic Social Attribute Scale (RoSAS, [Carpinella et al., 2017](#)), our new attributed unpredictability scale, and finally some background questions. The order in which the items of the RoSAS and attributed predictability questionnaires appeared were randomised between participants.

8.3.6 Measurements

Behavioural predictability. Similar to [Takayama et al. \(2011\)](#), we used an open question format to measure behavioural predictability, and asked how confident participants were in their prediction on a 7-point Likert scale. The open question (“What do you expect the robot will say next?”) was coded in terms of the presence of a keyword related to “eating” and to “pizza”. Answering this question requires the participant to connect the robot’s question (“whether the adult is *hungry*”) to the latest story development (“a *pizza* lying on the table”). When none of the keywords were present, the question was marked with a 0. When one keyword was present, related to either eating or pizza, behavioural predictability was marked with a 1, or with a 2 when keywords relating to both eating or pizza were present.

Attributed unpredictability. To measure to what extent participants associate unpredictability with the robot, we used the self report questionnaire with the 6 unpredictability items described in Chapter 7.

Robot social attributes. Participants’ social perception of the robot was measured through the RoSAS ([Carpinella et al., 2017](#)). This questionnaire measures the robot’s warmth, competence, and caused discomfort.

Intolerance of uncertainty. The participants’ IU was measured through the Intolerance of Uncertainty Scale - Short form (IUS-12, ([Carleton et al., 2007](#))). This is a 12-item questionnaire, for which we report the mean of the items.

8.4 Results

8.4.1 Behavioural predictability

There was a ceiling effect on the accuracy of the predictions, as most participants were able to make an accurate ($n = 127$), or partially accurate prediction ($n = 19$), and only 11 participants did not make a correct prediction. As a result, we used Fisher’s exact test to account for the three cells that had values less than 5, and found no significant difference between the accuracy of the prediction and the visibility of the cause ($p = .16$).

Figure 8.2 shows the participants’ confidence in their prediction on what the robot would say next. A one-way ANOVA indicated no significant difference between the participants’ confidence in their prediction for the three conditions ($F(2, 154) = 0.18$, $p = .83$). The average confidence scores when the cause was visible, ambiguous, or no cause condition were respectively 3.93 ($SD = 1.60$), 4.10 ($SD = 1.60$), and

Figure 8.2: Raincloud plots (scatter, box, and density plot) that show the participants' confidence in their prediction (behavioural predictability) for each of the three conditions.

3.94 ($SD = 1.65$). Given that there is no statistically significant relation between the visibility of the cause and behavioural predictability, there is also no mediation effect of behavioural predictability on the relation between visibility of the cause and the participants' social perception of the robot.

8.4.2 Attributed unpredictability

The reliability analysis of the attributed unpredictability questionnaire indicated that the internal consistency was excellent (Cronbach's $\alpha = .90$). The scores for the attributed unpredictability of the robot under each of the three conditions can be seen in Figure 8.3. Levene's test indicated that the variances for attributed unpredictability were significantly different for the cause visible condition ($F(2, 150) = 11.37, p < .001$). As the assumption of homogeneity of variance was not met, we used Welch's adjusted F ratio and the Games-Howell post hoc procedure. There was a significant effect of the visibility of the cause on the attributed unpredictability as determined by one-way ANOVA (Welch's $F(2, 90.65) = 22.20, p < .001, \hat{\omega}^2 = .22$). The participant's intolerance of uncertainty (IU) did not appear to moderate the relation between visibility of the cause and attributed unpredictability, as adding an interaction effect resulted in a non-significant result ($F(2, 152) = 0.14, p = .874$). Post hoc comparisons revealed that participants rated the robot as being less unpredictable when the cause was visible ($M = 12.65, SD = 5.54$) compared to the ambiguous condition ($M = 19.45, SD = 10.01, p < .001, d = 0.86$), or when no cause was visible ($M = 20.53, SD = 7.40, p < .001, d = 1.21$). There was no significant difference between the ambiguous and no cause conditions ($p = .769$).

Figure 8.3: Raincloud plots (scatter, box, and density plot) that show the scores on the attributed unpredictability of the robot for the three conditions.

8.4.3 Robotic Social Attributes

We expected that the participants' IU may influence their social perception of the robot. We therefore included IU as a covariate in the analysis of the three RoSAS subscales, given that the IU was not significantly related to the conditions, as determined by one-way ANOVA ($F(2, 150) = 0.49, p = .612$). IU was measured using the IUS-12 questionnaire, which had an excellent internal consistency (Cronbach's $\alpha = .90$). The assumption of homogeneity of variance was met for all three of the one-way ANCOVA's on the RoSAS subscales reported below. The assumption of homogeneity of regression slopes was also met, as IU showed no significant interaction effects with visibility of the cause. To assess whether IU moderated the relationship between visibility of the cause and the participants' social perception of the robot, we also ran the ANCOVA's with an interaction effect. The RoSAS subscales all had high internal consistency scores (Cronbach's α for warmth = .85, competence = .90, discomfort = .79). The mean score and confidence intervals of the RoSAS subscales can be seen in Figure 8.4.

Warmth. We found no significant effect of the visibility of the cause on warmth ($F(2, 149) = 0.22, p = .801$). Additionally, IU was not significantly related to warmth ($F(1, 149) = 0.24, p = .625$), nor did it interact with the visibility of the cause ($F(2, 149) = 2.04, p = .134$).

Competence. For competence, there was a significant difference between the visibility of the cause ($F(2, 149) = 7.78, p = .001, \omega^2 = .08$) after controlling for IU. Note that although IU was neither significantly related to competence ($F(1, 149) = 1.24, p = .267$) nor did it interact with the visibility of the cause ($F(2, 149) = 0.97, p =$

(a) Warmth subscale of the RoSAS.

(b) Competence subscale of the RoSAS.

(c) Discomfort subscale of the RoSAS.

Figure 8.4: Raincloud plots (scatter, box, and density plot) that show the scores on the three RoSAS subscales for each condition.

.381), we nevertheless kept it as a covariate. Post hoc comparisons with a Bonferroni correction revealed that the robot was considered to be more competent when cause was visible ($M = 6.80$, $SD = 1.48$), compared to when it was ambiguous ($M = 5.76$, $SD = 1.62$, $p = 0.003$, $d = 0.67$), or when there was no cause ($M = 5.78$, $SD = 1.51$, $p = 0.003$, $d = 0.68$). The difference between the ambiguous and no cause condition was not significant ($p = 1.000$).

As there was a significant difference in competence attributed to the robot under the condition of the cause being visible, we proceeded with a mediation analysis to investigate whether this difference is mediated by attributed unpredictability. Note that we cannot carry out a mediation analysis with behavioural predictability, because we did not find a significant difference between visibility of the cause and behavioural predictability (through quality or confidence). The mediation analysis was conducted in SPSS using the PROCESS macro (Hayes, 2017). To interpret the b coefficients, note that for visibility of the cause, we coded the condition “cause visible” as 0, “ambiguous” as 1, and “no cause” as 2. Competence and attributed unpredictability are both continuous variables with a value between 1 and 9.

The direct effect of visibility of the cause on competence was significant ($b = -0.32$, $t(149) = 2.04$, $p = .044$). To assess whether there can be a mediation effect (an indirect effect) of attributed unpredictability between the visibility of the cause and competence, the b coefficient between visibility of the cause and attributed unpredictability needs to be significant, as well as the b coefficient for attributed unpredictability and competence. This is the case, as the regression of visibility of the cause on attributed unpredictability was significant ($b = 0.65$, $t(150) = 5.11$, $p < .001$), as well as the regression of attributed unpredictability and competence ($b = -0.29$, $t(149) = -3.02$, $p = .003$). To calculate the indirect effect, we used bootstrapping with 5000 samples to generate the confidence interval for the b coefficient, which was -0.19 (95% CI $[-0.35, -0.05]$). This is significant at an α of .05 as the confidence interval does not include 0. As both the direct and indirect effect of visibility of the cause on competence are significant, attributed unpredictability only partially mediates this relationship. The mediation model with the regression coefficients can be seen in Figure 8.5. For all these steps, we controlled for IU. However, this variable was not significant in any of the steps.

Discomfort. We found no significant effect of the visibility of the cause on discomfort ($F(2, 149) = 1.94$, $p = .147$). However, IU is significantly related to discomfort ($F(1, 149) = 17.15$, $p < .001$, $\omega^2 = .10$), but did not interact with the visibility of the cause ($F(2, 149) = 1.51$, $p = .224$). Participants with higher scores on IU associated more discomfort with the robot.

8.5 Discussion

In this study, we investigated the relationship between the visibility of the cause of a robot’s responsive actions and action predictability, attributed unpredictability, the social attributes associated with the robot and the participants’ intolerance of uncertainty. Our results indicate that not seeing the cause of the robot’s responsive actions

Figure 8.5: Mediation model for predicting competence based on visibility of the cause with attributed unpredictability as mediator. The model includes the b coefficients and standard error for each relationship.

results in a stronger association of unpredictability with the robot, than when the cause was visible. We find no evidence that ambiguity in whether the robot responded to an actual event mitigated the effect of not seeing to what external event the robot responded. Interestingly, the participants' attributed unpredictability mediated the relation between the visibility of the cause and the attribution of competence to the robot. That is, participants rated the robot as more unpredictable, which also made them judge the robot as less competent. Considering a robot as competent is important, as it has been shown that robots that are considered less competent are also considered less trustworthy (Freedy et al., 2007). The mediation effect of attributed unpredictability was a partial mediation effect, as not seeing the cause of the robot's responsive actions also directly decreased the participants' attribution of competence to the robot. Even so, this direct effect was much smaller than the mediation effect of the attributed predictability. A possible explanation of the direct effect of seeing the cause of the robot's responsive actions is that they may have been perceived as mistakes. In turn, this could further decrease the robot's competence, as mistakes have been found to decrease the competence of a robot (Ragni et al., 2016). Unlike Dragan and Srinivasa (2014a), who found that higher familiarisation (leading to higher behavioural predictability) leads to more comfort, we found no effects between the visibility of the cause, attributed unpredictability and discomfort, nor with warmth.

We found no evidence that knowing that you cannot fully understand the reasons behind the robot's behaviour (in the ambiguous condition) mitigated the effect of not seeing the cause of the responsive robot actions on the attribution of competence. We hypothesised that this would mitigate the effect of unpredictable responsive actions, due to a reduced influence of predictions on social perception given the increased uncertainty regarding the predictions. Possibly, the explicit knowledge that limited information was available upon which to base predictions did not further decrease the already uncertain predictions due to seeing the robot for the first time. Alternatively, this may have to do with the fact that this was a video study in which participants did not interact with the robot. Prior research found that people try to make sense of a

robot more strongly when they expect to interact with the robot in the future (Eysselet al., 2011; Eysselet al. and Kuchenbrandt, 2011). Possibly, participants in our study were less inclined to make sense of the robot's behaviour as they did not interact with it nor would they in the future.

For behavioural predictability, we expected that not seeing the cause of the robot's responsive actions would make it more difficult for participants to infer the structural regularities of the interaction, and therefore make it more difficult to predict the robot's actions. We found no evidence that participants had difficulty predicting the robot's behaviour. Most participants were able to accurately predict the robot's next action. Neither did the visibility of the cause influence the participants' confidence in their prediction. This means that after watching just over 1 minute of interaction, they were able to discard the unpredictable actions as noise and use the structural regularities in the interaction to generate an accurate prediction. Unlike in the study of Takayama et al. (2011), our participants did not know what questions they would need to answer prior to watching the interaction, and were only given the instruction that they would view an adult interacting with a robot after which they would need to answer a couple of questions. Combined with a brief interaction, participants may not have had enough information upon which to make a reasoned prediction and went with their "gut feeling".

Reflecting back on how we manipulated the robot's behavioural predictability in our study, we wonder whether the manipulation was strong enough to find a difference, even though our pilot indicated that it was. The goal of our manipulation was to make the robot differ in its behavioural predictability in a realistic interaction, to limit confounding factors such as behaviour that is clearly erroneous or detrimental to the task at hand. Thus we opted for responsive robot actions, which are actions that are at risk of being perceived as unpredictable when the event the robot is responding to is not clear — a scenario that may well occur in real HRI's. In previous studies on predictability in HRI, behavioural predictability was also measured by asking participants to make a prediction (Takayama et al., 2011; Lichtenthäler et al., 2012a,b; Driggs-Campbell and Bajcsy, 2016; Walker et al., 2016) and indicate their certainty of their prediction (Takayama et al., 2011). However, in these studies, the robot's behavioural predictability was manipulated through limiting the information available to participants by having the robot perform fewer communicative actions (Takayama et al., 2011) or by presenting different and less information in a graphical user interface (Walker et al., 2016), or through how the robot moved (Lichtenthäler et al., 2012a,b; Driggs-Campbell and Bajcsy, 2016; Walker et al., 2016). Unlike these studies, we manipulated the structural regularities in a conversation, where the participants had to predict robot speech. The downside of this is that conversations need to be fairly structured to still make sense, limiting us in how unpredictable we could make the interaction whilst keeping it realistic. In that sense, manipulating robot behaviour through robot motions may be better suited for manipulating behavioural predictability without making the scenario unrealistic or unnatural.

Regarding the individual differences between people in their intolerance of uncertainty, we had expected that this would moderate the effects of the visibility of the cause and the participants' social perception of the robot. We hypothesised that

those high in IU would respond more strongly to not seeing the cause of the robot's responsive actions, which would negatively influence their social perception of the robot. We found no evidence that this is the case. IU only influenced the discomfort associated with the robot. Robots are a novel technology and are likely to induce uncertainty, which in turn may cause more discomfort for those who have a higher intolerance of uncertainty. Autistic individuals score relatively high on their intolerance of uncertainty compared to typically developing individuals (Chamberlain et al., 2013; Boulter et al., 2014; Neil et al., 2016b). Studies often report that the initial interactions with a robot can be stressful for some of these individuals and lead to an aversive reaction to the robot (e.g. Warren et al., 2015a; Short et al., 2017). However, even autistic individuals find a certain degree of unpredictability permissible, or even desirable, as a low-dose of novelty and surprise can be motivating for autistic children, rather than cause distress (Alcorn et al., 2014). We also expected that intolerance of uncertainty may be related to the warmth and competence associated with the robot, but found no evidence for this relation. Unlike discomfort, which is a dimension unique to people's social perception of robots (Carpinella et al., 2017), warmth and competence are argued to be universal dimensions of human social cognition and are used to judge whether another person is a threat or a friend (Fiske et al., 2007). This means that we found no evidence of a bias of people against our robot in terms of whether the robot intends good or ill.

8.5.1 Implications

Our findings imply that we have to be careful when designing responsive behaviours for a robot, as the attributed unpredictability may be negatively affected when the cause of the responsive action is not clear. This can hold even when the event is visible, as the person interacting with the robot may not perceive it, or perceive it differently. In turn, the attributed unpredictability can lead to a negative view of the robot's competence, as is indicated by our results. In human-human interaction, people actively improve the behavioural predictability of their actions by using additional signalling actions expressed through a gesture, deictics, facial expressions, body language and posture (Clark, 1996; Klein et al., 2005). Similarly, implementing signalling actions (Takayama et al., 2011; Szafir et al., 2015) or adding transparency (Chai et al., 2014) in robot behaviour may be beneficial. These additional actions take time and effort from the robot as well as cognitive effort from the person interacting with it, and therefore should not be implemented *en masse*. But selectively adding them for responsive robot actions may help in communicating the robot's intent and future responsive actions and thereby maintaining attributed predictability. Additionally, endowing robots with the ability to take the perspective of the person interacting with it into account (e.g. Trafton et al., 2005; Nikolaidis et al., 2016) may help identify when a robot's responsive action may be unpredictable and requires additional actions to make it predictable. For example, when a robot is responding to certain events that it has detected, it may be advisable that the robot first communicates its perception to the user, prior to taking the responsive action, or makes the reason for the response transparent in its responsive action. However, more research is needed to explore whether these designs can mitigate or prevent the negative ef-

fect of not seeing the cause of a robot's responsive actions on attributed predictability. We investigated one scenario where a robot can behave unpredictably and threaten the attributed predictability and perceived competence of the robot. Identifying other scenarios that can threaten the attributed predictability requires more research, as well as what design solutions can resolve those situations. Note that findings from [Driggs-Campbell and Bajcsy \(2016\)](#) indicate that reduced behavioural predictability does not always lead to reduced perceived competence, thus the results from such studies are also likely to be nuanced.

The above applies to the design of robot actions that may be unpredictable, and how they can then negatively influence the robot's attributed predictability. For unpredictable robot motions, the design considerations are different. While we did not experimentally assess the predictability of robot motion, the framework for robot predictability we presented in Chapter 6.4, highlights that in initial interactions the person's expectations play a large role in generating predictions, as the person does not yet have experience interacting with the specific robot (model). Studies from [Driggs-Campbell and Bajcsy \(2016\)](#) and [Dragan and Srinivasa \(2014a\)](#) indicate that conforming to such expectations is important for the attributed predictability. An autonomous car that drove with human-inspired motions was found to increase attributed predictability, compared to a car that drove with more mechanical/robot-like motions ([Driggs-Campbell and Bajcsy, 2016](#)). As people are already familiar with cars driven by humans, they may also expect similar motions from robot-driven cars. Similarly, a humanoid robot that moves more naturally (i.e. more like a human would) was also found to increase attributed predictability ([Dragan and Srinivasa, 2014a](#)). We expect there to be a correlation between behavioural and attributed predictability, and thus improving the behavioural predictability may also improve the attributed predictability. In human-human interaction, people also actively improve the predictability of their motions by improving the legibility ([Pezzulo et al., 2013](#)) or by being less variable in their motions ([Vesper et al., 2011](#)). Likewise, making robot motions better at expressing intent can also improve their behavioural predictability (e.g. [Gielniak and Thomaz, 2011](#); [Dragan and Srinivasa, 2014b](#); [Szafir et al., 2014](#)).

8.5.2 Generalisation to autistic individuals

The current study was conducted in the context of the DE-ENIGMA project, where it was a first step in trying to understand what predictability is about in the context of autistic people interacting with robots. From an ethical point of view, autistic children are a vulnerable population, which should only be recruited by researcher if there is no other way to conduct the study. For certain research questions, it makes sense to first conduct research on typically developing individuals, before validating the results with autistic individuals. Many studies have shown that by measuring autistic traits and correlating these measures to the outcome variables, insights can be gained that are also relevant to autistic populations (e.g. [Grinter et al., 2009](#); [Robertson and Simmons, 2013](#); [Goris et al., 2021](#)). To get an idea of how the results of the current study could be generalised to autistic populations, we included IU as a measure. On average, autistic individuals score higher on IU than typically developing individuals ([Chamberlain et al., 2013](#); [Boulter et al., 2014](#); [Neil et al., 2016a](#)). As we mentioned

in the previous section, IU was only related to the discomfort attributed to the robot, and we hypothesise that autistic individuals may also experience higher levels of discomfort when seeing a robot for the first time. Regarding our other outcome measures, we unfortunately cannot generate insights on how they may generalise to autistic populations.

8.5.3 Limitations

The participants in this study had to watch a video showing an interaction of another person interacting with a robot. While this allowed us to carefully control what perceptual information was available to participants in order to make sense of the robot's behaviour, we need to be careful with generalising the reported findings to individuals who are actually engaging with a robot. In particular, the explored conditions may have different effects on how the robot is perceived when the participant is actually engaged with the robot. Furthermore, while we asked participants to rate the *robot's* unpredictability, and their social perception of the *robot*, we cannot be entirely certain that the participants based their attributed unpredictability solely on the robot's behaviour. The scene itself may also have induced a feeling of unpredictability that was then attributed to the robot.

8.6 Conclusion

The predictability of a robot is often mentioned in HRI as an important quality of the robot (e.g. [Heinzmann and Zelinsky, 2003](#); [Alami et al., 2006](#); [Hancock et al., 2011](#); [Noorman and Johnson, 2014](#); [Koppenborg et al., 2017](#); [Lewis et al., 2018](#); [Sciutti et al., 2018](#)), yet the current conceptual understanding of robot predictability is inadequate as the concept is multi-faceted, rather than a singular concept. This limits us in effectively taking robot predictability into account in the design of robot behaviour. The aim of the study reported in this article was to improve our conceptual understanding of predictability and robots, how robot predictability relates to people's social perception of a robot, and to what extent this is influenced by a person's intolerance of uncertainty. To that end, we presented a framework of what we believe robot predictability is, and we carried out an experimental study. In this study, participants would watch a human interacting with a robot, but we limited the information available to the participants in relation to the robot's responsive actions, to make it more difficult to learn to predict the robot's behaviour. We conclude that not seeing what a robot responds to, negatively influences the attributed predictability and the competence attributed to the robot. The negative effect of not seeing the cause of the robot's responsive actions on competence was largely mediated by the attributed predictability.

We consider our results as preliminary evidence that the predictability attributed to a robot is an important factor to take into account when designing robot behaviour. We have presented our view on the concept of predictability and how it relates to HRI, which may be helpful in understanding and improving robot predictability, and laid out design several potential options on how this could be achieved. Key is to take the attributed predictability of a robot as an important outcome measure that can be used

to evaluate different designs (e.g. as in [Walker et al. \(2016\)](#)), and can lead to human-robot interactions that are easier to understand. While we cannot draw conclusions regarding the behavioural predictability of a robot — due to a ceiling effect — we do believe that both behavioural and attributed predictability are important for effective HRI, given the central role of predictability in human perception.

In future work, we should look at further investigating the relationship between behavioural and attributed predictability. Furthermore, as predicting the robot's behaviour can be learned, investigating this relationship as it evolves over time will be interesting. While people can learn to predict unnatural robot behaviours (at least to a certain degree ([Dragan and Srinivasa, 2014a](#))), or initially seemingly random robot behaviour, this does not mean that the robot is also considered to be more predictable.

As this study was conducted in the broader context of the DE-ENIGMA project, we are interested in the effect of robot predictability on the engagement of autistic children. We will look at this in the next chapter.

9

Robot predictability and the engagement of autistic children

This chapter is based on the following article:

Schadenberg, B. R., Reidsma, D., Evers, V., Li, J. J., Davison, D. P., Heylen, D. K. J., Neves, C., Alvito, P., Shen, J., Pantić, M., Cummings, N., Schuller, B. W., Olaru, V., Sminchisescu, C., Babović Dimitrijević, S., Petrović, S., Williams, A., Alcorn, A. M., & Pellicano, E. (under revision). Predictable Robots for Autistic Children — Variance in Robot Behaviour, Idiosyncrasies in Autistic Children's Characteristics, and Child-Robot Engagement. *ACM Transactions on Computer-Human Interaction*. 39 pages. ([Schadenberg et al., under revision](#))

As we have seen in the previous chapter, unpredictable actions can influence people's social perception of the robot. In Chapter 6, we also discussed how predictability is an important concept in autism and how it shapes current practices of autism professionals. However, what is not clear is whether we should account for the robot's predictability when designing a robot-assisted intervention, in order to reap the benefits of providing a highly predictable interaction — one of the main arguments for using a robot to enhance interventions. After all, robots are generally more predictable than humans, given the repetitive nature of robot actions of contemporary robots. Nevertheless, robots can still be unpredictable (although less frequently so than humans). Are contemporary robots then sufficiently predictable already? Or should we account for a robot's predictability during the design phase, and try to maximise it?

9.1 Introduction

IN the context of social skill learning, experiencing discomfort due to dealing with unpredictability is problematic as it prevents children from being in a state where they are ready to learn. Incorporating a robot in social skill learning might be

helpful in that it can provide a highly predictable manner of learning social skills, as we can systematically control the predictability of the robot's behaviour (e.g. see our study in the previous chapter). Indeed, the predictability of a robot is a commonly used argument for why robots may be promising tools for autism professionals working with autistic children (e.g. [Dautenhahn and Werry, 2004](#); [Duquette et al., 2008](#); [Thill et al., 2012](#); [Huskens et al., 2013](#); [Sartorato et al., 2017](#); [David et al., 2020](#)).

The predictability of robots may make it easier for autistic children to engage in learning in a robot-assisted intervention and maintain this engagement. However, robots cannot behave fully predictably as well as provide meaningful learning, as the learning gains may not generalise to the less predictable world of people ([Alcorn et al., 2019](#)). Moreover, a fully predictable robot would perpetuate in repetitive behaviour ([Dautenhahn, 2007](#)), limiting its long-term usefulness, and would also be unable to autonomously respond to the (unpredictable) dynamics of real-world settings ([Clabaugh and Matarić, 2019](#)). Both make large scale deployment of such a robot difficult. Thus, there is a trade-off between making the robot more predictable and providing the child with meaningful learning content with a tractable robot-assisted intervention, which requires a degree of unpredictability. This unpredictability stems from (a) adding novel content, which has not yet been learned, (b) making the learned skill's approximate — at least in part — what happens in real-world settings with humans, to facilitate generalisability to such highly unpredictable environments which autistic children need to learn to deal with, and (c) implementing complex robot behaviours (e.g. responsive actions to the environment), which may shroud the structural regularities of its behaviour and thereby prevent children from learning to predict the robot's behaviour. Thus, on the one hand, we want autistic children to be in a state where learning can occur, where the children show high levels of engagement with the learning material. This may require very predictable interactions. On the other hand, we want them to learn skills that are meaningful in the human world, which comes with a degree of unpredictability. A balance needs to be struck then, where the robot is sufficiently predictable to maintain engagement while still providing learning content that is representative of human-human interaction.

What constitutes “sufficiently predictable” is likely to differ between autistic children as they are a notoriously heterogeneous group ([Happé et al., 2006](#)). Some children may be better equipped to deal with unpredictability than others. For these children, predictability of the robot's behaviour may not be as important, and the focus can lie on optimising learning. Furthermore, children's reactions to the robot's unpredictability may be very different, and they may find different aspects of unpredictability problematic ([Goris et al., 2020](#)). This makes it difficult to generalise findings on predictability and autism to autistic children working with robots.

The effectiveness of robot-assisted interventions designed for social skill learning presumably depends — in part — on the *interplay between robot predictability, engagement in learning, and the individual differences between different autistic children*. To better understand this interplay, we report on a study where autistic children participated in a robot-assisted activity, where we manipulated the variance in the robot's behaviour as a way to operationalise predictability (see Chapter 6.5), and measured the children's individual characteristics as well as their engagement related behaviours.

9.2 Research questions and hypotheses

To make an informed decision on how to position and design the robot's behaviour in terms of its predictability, we need to better understand the interplay between robot predictability, engagement, and the individual differences between different autistic children. The aim of our study was to investigate this interplay, where we specifically looked at *behavioural engagement* and *visual attention* (as a proxy for cognitive engagement) in relation to the robot's predictability. We measured these two facets of engagement through manual annotations of observable child behaviours (see Section 9.3). In our study, autistic children engaged in a robot-assisted activity that was about the basics of recognising emotional facial expressions. Over the course of four sessions, the children interacted with a robot that was either low or high in variance of its behaviour, which was how we operationalised robot predictability (see Chapter 6.5). We aimed to address the following research questions:

Research question 1: How does variance in the robot's behaviour affect the autistic child's engagement?

- To what extent does variance in the robot's behaviour affect the autistic child's *behavioural* engagement?
- To what extent does variance in the robot's behaviour affect the autistic child's *visual* attention?

We hypothesised that initially, in the first session, there should be no difference between the low and high variance conditions, as the robot's behaviour is novel and still has to be learned. However, based on the importance of predictability to autistic children, we expected that over sessions the behavioural engagement and visual attention of autistic children should increasingly diverge between the low-variance robot compared to the high-variance robot in favour of the low-variance condition. That is, we expected an interaction effect between robot predictability on behavioural engagement and on visual attention over sessions, but no main effects. For visual attention, this means that we expected the children to increasingly look less towards the robot (the source of the unpredictability) and more to non-activity related locations that provided little unpredicted sensory input, such as the walls.

Our second research question is related to the individual differences between autistic children.

Research question 2: How do individual differences between autistic children influence their engagement in the activity in relation to the variance in the robot's behaviour?

- To what extent do autistic children's *autistic features* moderate the relation between the two facets of engagement and robot predictability?
- To what extent do autistic children's *expressive language ability* moderate the relation between the two facets of engagement and robot predictability?

- To what extent do autistic children's *intolerance of uncertainty* moderate the relation between the two facets of engagement and robot predictability?

Based on the preliminary findings of [Goris et al. \(2020\)](#), [Rudovic et al. \(2017\)](#), and our findings from the descriptive study reported in Chapter 4, we hypothesised that autistic children with higher autistic features should be less behaviourally engaged and pay less visually attention to activity-relevant locations (main effects), and their behavioural engagement should be more strongly and negatively affected than those with lower autistic features (interaction effect). Our hypothesis was similar for autistic children with lower expressive language ability. For autistic children who were more sensitive to unpredictability, as measured through their intolerance of uncertainty, we expected that they should respond more strongly to more variance in the robot's behaviour (interaction effect).

This study is complementary to another study which will be published elsewhere (hereafter referred to as the *complementary study*) and is the first study that assesses the claim on the benefits of robots being highly predictable. That complementary study is about micro behavioural analysis of autistic children in light of robot predictability. In the study reported in the current article, we used the same study design and data collection as the other study, but focused on different research questions and conducted different analyses to address those questions.

9.3 Materials and methods

9.3.1 Participants

Autistic children from the United Kingdom were recruited from a special education institution in the Greater London area. In total, 27 children were recruited of whom 24 (8 girls) were included in the analysis. These were autistic children with limited spoken communication and high support needs. For the three children who were excluded from the analysis, participation was discontinued due to administrative error (1 girl), or due to elevated anxiety during the session (2 boys)¹². All included participants had previously received an independent clinical diagnosis of autism according to the ICD-10 ([World Health Organization, 1992](#)), DSM-IV-TR ([American Psychiatric Association, 2000](#)), or DSM-V ([American Psychiatric Association, 2013](#)). All children were assessed by the Autism Diagnostic Observation Scale - second edition (ADOS-2, [Lord et al., 2012](#)), the Childhood Autism Rating Scale - second edition (CARS-2, [Schopler et al., 2010](#)), the Social Communication Questionnaire (SCQ, [Rutter et al., 2003](#)), and a bespoke scale of expressive language ability. In addition to receiving a clinical diagnosis of autism, all children scored above the autism cutoff on ADOS-2 (4 or higher). The participants' characteristics can be viewed in Table 9.1.

This study was reviewed and approved by the ethics committee of University College London, Institute of Education (REC 1175). For all children, parental consent

¹²Both were in the high-variance condition (see Section 9.3.3). For one boy, the activity was inducing too much anxiety — he did not get past the introduction in session one. The other boy repeatedly needed to be calmed in the second session, and eventually refused to go on. It was then decided to stop the experiment for this child.

Table 9.1: Participant characteristics.

	Condition	
	<i>Low-variance</i>	<i>High-variance</i>
<i>n</i> (sex)	12 (5 female)	12 (3 female)
Age (years:months)		
<i>M</i> (<i>SD</i>)	8:8.92 (1:8.70)	8:4.42 (1:8.49)
Range	6:10 - 11:7	6:10 - 11:4
ADOS-2 ^a Calibrated Severity Score		
<i>M</i> (<i>SD</i>)	6.08 (1.78)	6.25 (1.60)
Range	4 - 10	4 - 10
CARS-2 ^b		
<i>M</i> (<i>SD</i>)	27.83 (4.09)	29.17 (6.56)
Range	20.5 - 33.0	21.5 - 38.5
SCQ ^c		
<i>M</i> (<i>SD</i>)	22.25 (8.72)	25.50 (5.79)
Range	8 - 37	17 - 33
Bespoke Scale of Expressive Language		
<i>M</i> (<i>SD</i>)	2.92 (0.79)	2.42 (1.24)
Range	2 - 4	0 - 4
IUSC-S ^d		
<i>M</i> (<i>SD</i>)	2.64 (1.04)	3.19 (1.16)
Range	1.42 - 4.00	1.00 - 4.92

^aAutism Diagnostic Observation Scale - Second Edition (Lord et al., 2012)

^bChildhood Autism Rating Scale - Second Edition (Schopler et al., 2010)

^cSocial Communication Questionnaire (Rutter et al., 2003)

^dIntolerance of Uncertainty Scale for Children - Simplified

was obtained prior to their participation in our study. Our experimental protocols followed the ethical standards laid down in the 1964 Declaration of Helsinki.

9.3.2 Materials

For the experiment, we used the DE-ENIGMA system and DE-ENIGMA intervention v4, described in Chapter 5.5. Note that we only used the data from the webcams for our analyses.

9.3.3 Experiment design

We used the data collected in the complementary study. We will briefly summarise the experiment design here. The between-participants study involved two conditions, where the independent variable was the robot's behavioural variance which was either low or high. The children were randomly assigned to one of the two condition and remained in this condition throughout the experiment. In the *low-variance condition*, the robot's behaviour showed minimal variance. This means that the verbal

Figure 9.1: The experimental setup. The school's staff member who accompanied the child would sit behind the child in the corner of the room. The red circles outline the camera's that were used for annotating our engagement measures, which also included the camera from which this image was taken.

and physical behaviour for a certain action was always the same. For example, the robot would always say “Hi, my name is Zeno” and wave with its right arm as a way of greeting the child. In contrast, in the *high-variance condition*, we implemented speech, motion, temporal, and topic variance, to increase the behavioural variance of *all* of the robot's actions. Thus, the robot displayed behavioural variance throughout the interaction. In this condition, each of the robot's actions had four variations that differed the aforementioned types of variance, excluding the additional actions that were designed to introduce topic variance. This was done as follows:

- **SPEECH VARIANCE:** variability in the words that the robot uses (diction) and how those words are spoken (prosody).
- **MOTION VARIANCE:** variability in patterns in the movements of the robot's face and arms during emotional facial expressions and robot actions.
- **TEMPORAL VARIANCE:** variability in time offsets between the child pressing a button and the robot responding to the button press.
- **TOPIC VARIANCE:** the use of different actions that address different topics at a particular point in an interaction. Note that in the low-variance condition, the robot actions for topic variance were responsive actions in that the robot responded to an event (e.g. hearing a noise). In the high-variance condition, there would be no observable event that explained the robot's action.

The variant actions have been designed to display *unimodal* variance compared to their invariant counterparts. That is, for each robot action, the action can show vari-

ance on either speech or motion — not both modalities. An intent may translate to a combination of both speech and motion actions, but these will always be shown sequentially as two unimodal actions, and will not be combined to create a multimodal stimulus. For instance, when greeting the child, the robot would first say “Hello, my name is Zeno”, which was followed by a wave. We consider this as a single robot action.

The choice for which variant to display for an action was determined through an algorithm. The action variant was chosen semi-randomly, where it would pick one of the four variants, excluding the antecedent variant for that action (if any). The latter was to prevent variants to be shown twice in a row. For actions with topic variance, a different selection mechanism needed to be used, as the variance for these actions related to whether the action is contingent or non-contingent on (internal or external) events in the environment. We therefore opted to use the Wizard-of-Oz paradigm, where another experimenter was controlling the robot (the wizard), without being visible to the child. The wizard was responsible for selecting the topic variant actions. To standardise the number of topic variance actions, the wizard would be notified through the wizard’s control interface when it was time for the robot to perform a topic variant/invariant action. At this time, the wizard would look for one of the topic variant actions that was congruent with the condition the child was in. In the low-variance condition, the wizard would select topic variant actions that were a response to an observable event. For instance, the wizard could have the robot say “what was that noise?” in response to noise outside of the experiment room. For the high-variance condition, the wizard would ensure that the event the robot would respond to was not perceivable. In this case, the wizard would have the robot respond to noise when there was no noise to be heard.

9.3.4 Experimental setup

The study took place at the children’s school, in one of the offices. This room was converted into an experiment room. A picture of the experimental setup can be seen in Figure 9.1. The robot stood on a table facing the child and acted partly autonomously and was partly controlled by the wizard. This person was sitting in the same room behind a room divider and could view the interaction through a webcam. The wizard was responsible for selecting the correct task within the activity, selecting the topic variance actions, and for responding to any unscripted interactions through using a preset of robot actions such as having the robot say “no”, “yes”, “I don’t know”. Within each task, the robot behaved autonomously, although the wizard could interrupt these behaviour at any time when the situation demanded it. The child would sit in front of the robot and next to the adult who was presenting the DE-ENIGMA activity. The recording equipment was placed behind and next to the table on which the robot stood. All computers and laptops were also placed behind the room divider and were managed by another researcher.

9.3.5 Experiment procedure

The autistic children would engage in the robot-assisted activity individually, once per day, for four to five sessions. The fifth session was only applicable to children who did not finish the activity in four sessions. Each session lasted around 15 to 20 minutes. The sessions were scheduled on consecutive school days as much as possible. Due to weekends and the children's schedule some variability was inevitable. For the low-variance condition there was an average of 0.36 days in between sessions ($SD = 0.68$), and for the high-variance condition this was 0.64 days ($SD = 1.22$).

Each child was assigned to one of three adults who would lead the sessions. The adult assigned to the child would also remain with that child for each of the child's sessions. The adult was tasked with augmenting the robot's instructions, supporting the children in using the tablet, giving feedback, and responding to their communicative overtures. Theirs was a supportive role, as the robot delivered the majority of the instructions and feedback. The children were often accompanied by a school staff member, who would sit in the back of the experiment room. They were asked not to participate in the activity unless they thought there was an issue such as when the child showed anxiety.

The content of each day of participation was scheduled to be delivered in a specific order and was identical across both conditions. When the children met the robot for the first time, the robot was covered by a blanket when they walked in the experiment room. The session would start with the robot being uncovered, introducing itself, and showing what movements it could do, some facial expressions, and what it sounded like. The goal here was to let the child get comfortable with the robot, what it looked like, and how it behaved. If the child liked any of the actions, they could be repeated by the wizard. When the child appeared comfortable, the adult would suggest to transition to the main activity — which the wizard then started — finishing the introduction. In subsequent sessions, the introduction was similar but shorter. After the introduction, the DE-ENIGMA activity would start. During the activity, the children would engage in the games that were set for that day, as well as in “choice points”, where the child could choose one of the games or specific robot behaviours they liked. As described in Chapter 4, autistic children often make requests to the robot and denying the child's request could disrupt the interaction. These choice points allowed the adult to defer any requests for one of the games or certain behaviours to a choice point, rather than denying the child's request or interrupting the ongoing game. Moreover, by structuring the choice points and limiting them to two minutes, it allowed us to control it as much as possible with respect to following the same programme of content and the same order for each child. During the choice point, the child could choose from any content that they had already experienced. The child could make the choice through a “choice board”, which contained icons of the available options. This board and the icons on it were managed by the adult, and the wizard executed the child's requests. For children who would or could not choose, the adult would first prompt the child by suggesting an activity or a robot action. If the child still did not express any preference, the adult ask the robot what to do next, upon which the wizard would select an activity. When the session was over, the robot would say goodbye, and the child went back to class.

9.3.6 Measures

Engagement measures. We chose to operationalise and measure the autistic children’s engagement in terms of *behavioural engagement* and *visual attention* as these can be annotated as *patterns of manifest content* — child behaviours that are directly observable (Potter and Levine-Donnerstein, 1999). This is in contrast to scoring engagement holistically as *projective latent content* (Potter and Levine-Donnerstein, 1999), as this would require the coder to employ subjective interpretations of the meaning of the behaviour, which is difficult without being familiar with how the participating children generally behave and understanding the meaning of their sometimes atypical and idiosyncratic behaviours.

BEHAVIOURAL ENGAGEMENT. We measured behavioural engagement through expert annotations, using the coding scheme described in Table 9.2 on segments of 5 seconds. With this coding scheme, we make the distinction between autistic children being behaviourally engaged or disengaged with the HRI using a 5-point ordinal scale that denotes the amount of behavioural engagement.

While we believe that stimming behaviour — a self-stimulatory behaviour marked by a repetitive action or movement of the body — or fidgeting can be indicative of behavioural disengagement, it was problematic to annotate. Stimming behaviour provides sensory input for one of the senses, preventing the child from using this modality for engaging in the activity. However, whilst stimming, the child can still engage in the activity — and potentially learn — through using other modalities. For example, a child may be rubbing their hands while speaking to the adult or robot about the activity. Because we do not distinguish between modalities for annotating behavioural engagement, we decided to code all stimming behaviours that did not prevent the child from engaging in the activity and interaction were coded as passive. When the stimming was all consuming and prevented the child from engaging with the activity, we coded it as disengagement.

VISUAL ATTENTION. The experiment procedure and the activity did not allow us to directly measure cognitive engagement, which is why we opted to measure this through annotating the children’s visual attention — a proxy for cognitive engagement. Visual attention relates to the extent the participant paid overt visual attention to the ongoing activity and interaction, and tells us more on what or whom the children were engaged with. This was also measured through expert annotations by coding the children’s gaze direction, but on segments of 2.5 seconds. From these annotations, we calculated the percentage of time spend looking at a certain direction for each session.

The coding scheme that we used included the following gaze directions: “robot”, “tablet”, “activity materials”, “adult”, “school staff member”, “elsewhere”, or “mixed”. The latter was used to annotate instances where the child did not focus their gaze at any one point. The activity materials refers to any materials that were used in the activity, which primarily was the visual choice board for selecting the next task. And elsewhere refers to any gaze direction that was not in any of the other categories

Table 9.2: Coding scheme for annotating the observed behavioural engagement of an autistic child.

<i>Level</i>	<i>Meaning</i>	<i>Description</i>	<i>Example</i>
-2	Fully behaviourally disengaged	The child exhibits non-task related behaviour for most or all of the segment.	Child is playing around with objects that does not involve the task or interaction with the robot or adult. Child is engaging in stimming behaviours that prevent the child from partaking in the task. Child indicates wanting to stop with the interaction, e.g. through asking whether the game or session is finished, or turning the tablet to the adult prior to the game's conclusion. Talking with the adult about something unrelated to the activity or the ongoing triadic interaction, such as children saying they are hungry.
-1	Partly behaviourally disengaged	The child exhibits non-task related behaviour for some of the segment.	
0	Passive	Child does not behaviourally engage with the task, nor does the child show engagement in other activities.	Child is seemingly listening to the adult or robot, but does not communicate back. Child is looking at the task material, but does not physically interact with it. Also, echolalic and undirected vocalisations were included on this level, as well as stimming behaviours that do not prevent the child from partaking in the task. Covering ears due to auditory sensitivity.
1	Partly behaviourally engaged	The child behaviourally engages with the task through interacting with the adult, with the robot directly or through the tablet, or other task materials for some of the segment.	

Table 9.2: Table continued.

<i>Level</i>	<i>Meaning</i>	<i>Description</i>	<i>Example</i>
2	Fully behaviourally engaged	The child behaviourally engages with the task through interacting with the adult, with the robot directly or through the tablet, or other task materials for most or all of the segment.	Pressing a button on the tablet, requesting facial expressions of the robot, choosing an activity, talking with the adult about the robot, dancing with the robot. Also includes non-verbal communication with adult, such as sharing enjoyment, or social references after the robot did something.

and thus *was not related to the activity*. Gaze direction such as the walls, the floor, the cameras, or the room divider. In contrast, the locations “robot”, “tablet”, “activity materials”, and “adult” can be considered “activity-relevant locations”. For our analysis of visual attention, we specifically look at the gaze directions towards the robot, and elsewhere, as our hypotheses relate to these locations. The other locations were annotated to enable better interpretation of the data.

The annotations were done at two levels. The first (primary) level represents the gaze direction where the child spent looking for the majority of the segment (1.25 seconds or more). The second (secondary) level was optional and could be used to annotate a secondary gaze direction during the segment which lasted at least 0.75 seconds, but no more than 1.25 seconds. This excludes brief glances, where the child’s gaze would not stay on one direction.

Individual factors. The complementary study also measured several individual differences between the autistic children. In the current study, we use a subset of those measures, namely the CARS-2 for measuring autistic features, a bespoke scale of expressive language, and an adapted version of the Intolerance of Uncertainty Scale for Children - parent report form (IUSC, [Comer et al., 2009](#)).

The IUSC questionnaire was adapted from the complementary study to accommodate to children with limited spoken language, as we believed for many of them the IUSC questions were too difficult for parents to answer about their child. In the remainder of this article, we refer to this adapted version of the parent report of the IUSC as the Intolerance of Uncertainty Scale for Children - Simplified (*IUSC-S*). The adapted questions of the IUSC-S can be seen in the next subsection, as well as principle component analysis of this adapted questionnaire. Based on this analysis, we excluded three questions for computing the IUSC-S scores.

The CARS-2 ([Schopler et al., 2010](#)) is a 15-item autism screening and diagnostic tool and was administered to obtain a general measure of characteristics of autism. It was completed based on direct behaviour observation by a professional as well as reports from parents, teachers, or caretakers. The measure was completed by the adult who worked with the specific child. The total score on the CARS-2 reflects the

severity of autistic features with scores of 15.0-29.5 indicating minimal-to-no evidence, 30.0-36.5 is mild-to-moderate severity, and 37.0 and higher is severe autistic features.

The bespoke scale of expressive language measures the spoken language ability of the child. The adult who gave the sessions rated the child's expressive language after the last session. This scale ranged from 0 to 4, where 0 means "no words", 1 is "some vocalisations or word approximations", 2 is "single words", 3 is "simple sentences (2 to 3 words)", and 4 is "more complex speech, including complex sentences. The bespoke score reflects the level of expressive language that was generally used by the child during the sessions.

9.3.7 Intolerance of Uncertainty Scale - Simplified

The IUSC-S questionnaire can be seen in Table 9.3. To assess how many dimensions of intolerance of uncertainty the IUSC-S measured, we conducted a Principle Component Analysis (PCA) on 15 items with varimax rotation. The Kaiser-Meyer-Olkin (KMO) measure verified the sampling adequacy for the analysis (KMO = .65), and all KMO values for individual items were > .55. Given our sample size of 24, this is sufficient, but mediocre. Bartlett's test of sphericity indicated that correlations between items were sufficiently large for PCA ($\chi^2(105) = 312.61, p < .001$). Two components had eigenvalues above Kaiser's criterion of 1 and together explained 70.85% of the variance. Table 9.4 shows the component loadings for each item, where we highlighted items with loadings of .72 or higher and with a cross loading difference greater than .2, based on recommendations from Stevens (2002). Based on these component loadings, we exclude item 2, 7, and 11 for computing the IUSC-S scores.

9.3.8 Annotation procedures

As only a few children participated in five sessions, we only annotated the videos from the first four sessions. These sessions were annotated in terms of behavioural engagement levels and visual attention locations using the ELAN transcription software¹³, developed by the Max Planck Institute for Psycholinguistics in Nijmegen, the Netherlands.

Behavioural engagement annotation. To determine whether a child is engaged or disengaged requires taking the situational context into account. For example, the child looking away may be part of the current activity, a response to the adult, or indicative of disengaging from the interaction. To preserve the situational context, we annotated 1 minute out of every 2 minutes, excluding only instances where there were technical problems with the system. We started the annotations from the moment the adult said hello to the robot and ended when the robot had said goodbye. This resulted in the annotation of 1660 minutes of video recording, of which 848 minutes were in the low-variance condition and 812 minutes were in the high-variance condition. During initial testing of the coding scheme for behavioural engagement, we noted that the autistic children sometimes had brief behavioural disengagement episodes of

¹³<https://tla.mpi.nl/tools/tla-tools/elan/>

Table 9.3: The Intolerance of Uncertainty Scale for Children - Simplified (IUSC-S), Parent-report form, adapted from the complementary study.

<i>Questions</i>
1. Uncertainty makes my child's life intolerable.
2. My child's mind can't be relaxed if he/she doesn't know what will happen tomorrow.
3. Uncertainty makes my child uneasy, anxious, or stressed.
4. Unforeseen events upset my child greatly.
5. It frustrates my child to not have all the information he/she needs in a situation.
6. Uncertainty keeps my child from living a full life.
7. When it's time to act, uncertainty paralyzes my child.
8. When my child is uncertain he/she can't function very well.
9. Other children seem to be more certain than my child.
10. Uncertainty makes my child unhappy or sad.
11. My child always wants to know what the future has in store for him/her.
12. My child can't stand being taken by surprise.
13. Uncertainty keeps my child from sleeping soundly.
14. My child tries to get away from all uncertain situations.
15. The ambiguities of life stress my child.

Table 9.4: The component loadings for the items on the intolerance of uncertainty as measured by the IUSC-S.

<i>Item</i>	<i>Component 1</i>	<i>Component 2</i>
Item 6	.90	-.12
Item 4	.87	.20
Item 15	.87	-.09
Item 3	.86	-.25
Item 1	.86	-.12
Item 12	.84	.02
Item 10	.84	-.23
Item 5	.79	-.27
Item 14	.78	.27
Item 13	.77	.17
Item 9	.74	-.17
Item 8	.73	.11
Item 7	.70	.55
Item 11	.58	.58
Item 2	.48	-.63
Eigenvalue	9.16	1.47
% of variance	61.08	9.77
Cronbach's α	.96	.34

Table 9.5: Confusion matrix of the annotations for behavioural engagement between the main coder and secondary coder.

Main coder	Secondary coder					Total
	-2	-1	0	1	2	
-2	188	47	28	14	4	281
-1	6	105	53	15	1	180
0	7	20	706	60	2	795
1	9	17	79	796	57	958
2	2	4	3	60	198	267
Total	212	193	869	945	262	2481

a couple of seconds. We therefore divided each minute into segments of 5 seconds, similar to [Kim et al. \(2012\)](#) and [Simpson et al. \(2013\)](#).

A single main coder annotated all the recordings, amounting to 9947 segments. To calculate the inter-rater reliability of these annotations, a second coder annotated a randomly selected session for each participant. This amounted to 25% of the recordings being dual-coded, which also contained 25% of all the segments. To determine the agreement between the two coders, Cohen's κ statistic was used. There was good agreement between the two coders for behavioural engagement (Cohen's $\kappa = .72$, 95% CI [.70, .75], $p < .001$).

To gain insight into the nature of the coder disagreements, we inspected the confusion matrix for behavioural engagement. As can be seen in Table 9.5, the main coder annotated more instances of disengagement (approximately 14%), which were coded as passive by the secondary coder. Note that the main coder was not blind to conditions, but the secondary coder was. Given the agreement between the coders, we do not think it likely that this influenced our results. All in all, we deem these values high enough to continue our analysis on the basis of the main coder.

Visual attention annotation. For visual attention, we annotated the same minutes as for behavioural engagement. However, annotating segments of 5 seconds proved too long, as there were often more than three gaze directions. This made it difficult to code with our annotation scheme, which accounted for two gaze directions. The segments for visual attention therefore lasted 2.5 seconds to ensure that there were generally fewer than three gaze directions per segment.

Again, a single main coder annotated all the recordings and second coder annotated a randomly selected session for each participant. The main and secondary coder were the same annotators as for the annotation of behavioural engagement. For visual attention, this amounted to a total of 19855 segments. The dual-coding resulted into 25% of the segments being dual-coded. There was very good agreement for the primary gaze direction annotations (Cohen's $\kappa = .87$, 95% CI [.86, .89], $p < .001$) (see Table 9.6) and good agreement for the additional, secondary gaze direction annotations on the (Cohen's $\kappa = .69$, 95% CI [.66, .73], $p < .001$) (see Table 9.7).

Table 9.6: Confusion matrix of the *primary* annotations for visual attention between the main coder and secondary coder.

Main coder	Secondary coder							Total
	Robot	Tablet	Teaching Mats.	Adult	Assis- tant	Else- where	Mixed	
Robot	1599	23	10	16	0	60	27	1735
Tablet	19	1070	5	7	0	24	12	1137
Teaching Mats.	7	2	534	5	0	13	4	565
Adult	12	8	6	373	0	28	19	446
Assistant	1	0	0	0	10	2	4	17
Elsewhere	28	40	11	16	0	904	19	1018
Mixed	11	7	5	9	0	25	55	112
Total	1677	1150	571	426	10	1056	140	5030

Table 9.7: Confusion matrix of the *secondary* annotations for visual attention between the main coder and secondary coder.

Main coder	Secondary coder							Total
	Robot	Tablet	Teaching Mats.	Adult	Assis- tant	Else- where	Mixed	
Robot	305	16	7	13	1	23	16	381
Tablet	19	122	1	6	0	7	6	161
Teaching Mats.	10	1	60	4	0	5	3	83
Adult	13	7	3	118	0	7	9	157
Assistant	0	0	0	0	7	0	0	7
Elsewhere	17	8	4	11	0	119	10	169
Mixed	6	4	0	6	1	3	52	72
Total	370	158	75	158	9	164	96	1030

We deem these values high enough to continue our analysis on the basis of the main coder. Furthermore, inspection of the confusion matrices showed no biases on the coder disagreements for one of the two coders.

9.3.9 Data Analysis

For the analysis of both behavioural engagement and visual attention, we used growth models with a maximum likelihood estimation method. These were modelled using the statistical program “R” (R Core Team, 2020), version 3.6.3, with the “nmlr” package (Pinheiro et al., 2020), and analysed with two-tailed tests and a 95% confidence level. Growth models allow us to take the multi-level nature into account, where the behavioural engagement scores/visual attention locations for each of the sessions are level 1 variables, the child is a level 2 variable, and the adult leading the session a level 3 variable. The adults were randomly assigned to the children regardless of

the condition that was assigned to the child. The adult is therefore a crossed effect. For the visual attention analysis, we removed annotations where the child looked at the school staff member, as they were not always present in each session. Additionally, we removed annotations coded as “mixed”, due to the uncertainty regarding the direction of the child’s gaze.

As recommended by Raudenbush and Bryk (2002), we first make a “basic” model and then add in variables as appropriate. In this article, we consider the model that includes random coefficients, the condition, and the session (time) as the basic (growth) model. We expected that the children will have different intercepts, as some children will be behaviourally more engaged than other children, or prefer looking at a certain gaze location more. Furthermore, we also expected that the slopes will be different between children, where one child loses interest faster than other children, resulting in difference in behavioural engagement and visual attention. The random coefficients account for these random effects in our models. Next, we further explored to what extent the individual factors improved the model fit of the basic conditional growth model. Finally, we investigated whether the different adults leading the sessions influenced the two facets of engagement of the children by adding the adult as a factor.

9.3.10 Manipulation check

The study protocol had a flexible activity selection and session duration so as to support each autistic child’s individual preferences. This means that the variability introduced by the robot differed per session and per participant. In turn, this means that we cannot be certain that the robot displayed high amount of variance in the high-variance condition, and vice versa for the low-variance condition. To check to what extent our manipulation of robot variance succeeded, we calculated the following variables:

1. AVERAGE NUMBER OF ROBOT ACTIONS PER MINUTE. These should be similar between conditions, and serves as a baseline to put item (2) and (3) into perspective.
2. AVERAGE NUMBER OF UNIQUE ROBOT ACTIONS PER MINUTE. These are unique *within* a session. Of the total number of robot actions, the high-variance conditions should have a more *unique* robot actions per minute. This reflects the larger pool of unique actions that were implemented in the high-variance condition.
3. AVERAGE NUMBER OF NOVEL ROBOT ACTIONS PER MINUTE. These are actions that the child has not seen before in either the current or previous sessions. While the content of the interaction differed per session, there are robot actions that are used throughout the sessions, such as giving praise. Therefore, this number should drop over sessions in both conditions, but the high-variance condition should introduce more novel actions per minute than the low-variance condition in all sessions.

4. THE CUMULATIVE AVERAGE NUMBER OF REPETITIONS PER ROBOT ACTION. This number reflects how often the robot displayed a certain robot action, either in the current session or previous sessions. The higher the number, the more opportunities the child had to learn to predict this action. In the high-variance condition, this large amount of unique actions should result in fewer repetitions per action than in the low-variance condition.

For the manipulation check, we consider all *variants* of robot actions as being unique actions. For the variables that were calculated per minute, we excluded the time when there was a technical difficulty. To assess to what extent there is a difference between the two conditions on each of these four items, we conducted four mixed ANOVAs, where the condition is a between-subject variable and session a within-subject variable. Furthermore, we assessed each outlier and consider to remove them from the analysis. This was done by placing the outlier in context of what happened during the session as well as relating the value of the outlier to the values of the other condition.

9.4 Results

9.4.1 Manipulation check

On average the children engaged in the DE-ENIGMA activity for 16 minutes and 13 seconds ($SD = 3\text{min}, 1\text{s}$) in the low-variance condition. For the high-variance condition, the average was 15 minutes and 39 seconds ($SD = 1\text{min}, 23\text{s}$). There was no significant difference in the time spent in the activity between the conditions ($F(1, 22) = 3.30, p = .08, \eta_p^2 = .13, \eta_p^2$ 90% CI[.00, .34]), over sessions ($F(3, 66) = 2.55, p = .06, \eta_p^2 = .10, \eta_p^2$ 90% CI[.00, .20]), nor an interaction effect ($F(3, 66) = 0.22, p = .88, \eta_p^2 = .01, \eta_p^2$ 90% CI[.00, .03]). Note that the variables reported below are all normalised to account for differences between children and sessions in the time that the session lasted.

The extent to which there was variance in the robot's behaviour per session can be seen in Figure 9.2 for both conditions. The average number of robot actions per minute for the low-variance condition was 5.43 ($SD = 0.55$) and 5.36 ($SD = 0.82$) for the high-variance condition. There was no significant difference in the average number of robot action per minute over sessions ($F(3, 66) = 0.58, p = .632, \eta_p^2 = .03, \eta_p^2$ 90% CI[.00, .07]), between conditions ($F(1, 22) = 0.09, p = .763, \eta_p^2 < .01, \eta_p^2$ 90% CI[.00, .12]), or an interaction effect ($F(3, 66) = 0.91, p = .441, \eta_p^2 = .04, \eta_p^2$ 90% CI[.00, .10]). As expected, there was a significant difference between conditions in the average number of *unique* robot actions per minute ($F(1, 22) = 136.78, p < .001, \eta_p^2 = .86, \eta_p^2$ 90% CI[.74, .90]), where this number was significantly higher in the high-variance condition ($M = 3.80, SD = 0.63$) than in the low-variance condition ($M = 1.80, SD = 0.30$). Similarly, the difference between conditions for the average number of *novel* robot action per minute was significant ($F(1, 22) = 363.44, p < .001, \eta_p^2 = .94, \eta_p^2$ 90% CI[.89, .96]). The robot in the high-variance condition displayed a higher number of novel actions per minute than in the low-variance condition for each of the four sessions. For the average number of repetitions per action there was

Figure 9.2: Boxplots that show to what extent variance in the robot's behaviour was achieved. **(A)** shows the average number of actions the robot performed per minute, **(B)** shows the average number of unique actions performed by the robot actions per minute within a session, **(C)** shows the average number of actions that the robot performed for the first time per minute, **(D)** shows the cumulative average number of repetitions per robot action.

also a significant difference between the conditions ($F(1, 22) = 406.46, p < .001, \eta_p^2 = .95, \eta_p^2$ 90% CI[.90, .96]). For each session, the high-variance condition had fewer repetitions per action than in the low-variance condition. Based on these tests, we conclude that the robot's variance was significantly different between the conditions, as well as sufficiently large given the reported effect sizes.

While we consider the manipulation of the extent of variance in the robot's behaviour successful, there were four sessions of three children where the variance of the robot's behaviour was more akin to that of the other condition. In the low-variance condition, the first session of one child was only partially recorded in our logfiles due to technical issues. The data loss resulted in a shorter session recording, which in

turn resulted in a higher number of unique and novel actions per minute, as well as fewer repetitions per action. Nonetheless, it was a regular session where the child was engaged for most of the time. It is therefore likely that the system performed as intended and produced low variance in the robot's behaviour. In the high-variance condition, two children had sessions where the robot showed few unique actions. For one child, this was the case in session 3 and 4, while for the other it was only for session 3. Note that for session 3, these instances are not statistical outliers and therefore are not shown as such in Figure 9.2. For all three sessions, the child was disengaged for most of the session, limiting the progression through the content. In turn, the robot only displayed a limited repertoire of its behaviours, which increased the chance of displaying actions that had already been displayed before. Given that this happened in the last two sessions, we cannot exclude that this happened due to their experiences in the first and second session. Moreover, the mean number of novel actions per minute and repetitions per action are in line with the high-variance condition. Therefore, we include these sessions in the main analysis.

9.4.2 Behavioural engagement

The parameter estimates for each of the multi-level models are presented in Table 9.8. First, we fitted a basic conditional growth model (MODEL A) with random slopes and intercept. As fixed effects, this model contains the session, condition, and an interaction effect between the two. For the covariance structure, we used a first-order autoregressive covariance structure. To further explore the trend of behavioural engagement over sessions, we fitted a quadratic and cubic trend instead of the linear trend. The quadratic trend best fitted the change in behavioural engagement over sessions (MODEL B). Next, we investigated whether accounting for the adult who was leading the session improved the model (MODEL C). This required a three-level model, where the adult is a crossed effect. Compared to MODEL B, the three-level model did not significantly improve the model fit ($\chi^2(1) = 0.01, p = .999$). Thus, while the children received their sessions from one of three adults, this does not explain the variance in their behavioural engagement.

Next, we investigated whether individual differences in characteristics of the children could explain the variance in behavioural engagement. We took MODEL B, as it had the best model fit, and investigated whether the CARS-2 score, expressive language score, or IUSC-S score, improved the model fit. Adding the CARS-2 score as covariate significantly improved MODEL B ($\chi^2(1) = 9.60, p = .002$). Similarly, the expressive language score for expressive language significantly improved MODEL B ($\chi^2(1) = 8.59, p = .003$). Adding the IUSC-S did not significantly improve MODEL B ($\chi^2(1) = 0.03, p = .867$).

Based on the model fit of the models reported above, and whether they significantly improved the model fit, we consider MODEL D with the CARS-2 score as covariate, and MODEL E, which has the expressive language score as covariate, as the models that best explain the variance in behavioural engagement. However, to understand to what extent the CARS-2 score and the expressive language score explain the same variance in behavioural engagement, we fitted a model using both covariates (MODEL G). This did not significantly improve the model fit compared to the best model with

Table 9.8: Model fit measures for each model as they increase in complexity, as well as chi-square statistic with 1 degree of freedom and statistical significance. Each model is compared to the first model with lower complexity as reflected by the model's degrees of freedom. For the models with one or two covariates, the models are compared respectively to MODEL B and MODEL D. The models with an additional interaction effect between the covariate and the condition are compared with MODEL G.

<i>Model</i>	<i>df</i>	<i>AIC</i>	<i>BIC</i>	<i>Log-likelihood χ^2</i>		<i>p</i>
<i>Basic growth models</i>						
(A) Conditional	9	26.46	49.54	-4.23	-	-
(B) Conditional with quadratic trend	9	21.46	44.54	-1.73	-	-
(C) Three-level	12	27.45	58.22	-1.72	0.01	.999
<i>MODEL B with covariate</i>						
(D) CARS-2 score	10	13.86	39.50	3.07	9.60	.002
(E) Exp. lang. score	10	14.87	40.52	2.56	8.59	.003
(F) IUSC-S score	10	23.43	49.08	-1.72	0.03	.867
<i>MODEL D with another covariate</i>						
(G) CARS-2 score and exp. lang. score	11	12.05	40.26	4.98	3.81	.051
<i>MODEL G with interaction between covariate and condition</i>						
(H) CARS-2 score * condition	12	14.05	44.82	4.98	<0.01	.961
(I) Exp. lang. score * condition	12	11.08	41.85	6.46	2.97	.085
(J) IUSC-S score * condition	13	15.14	48.48	5.43	0.91	.634

only one covariate (MODEL D). Furthermore, while this model has a better fit than MODEL D, it is also more complex, as reflected by a higher BIC.

To investigate whether the covariates actually moderated the effect of the condition (robot predictability) on behavioural engagement, we further added an interaction effect between the covariate and the condition to MODEL G. For the CARS-2 score, this did not significantly improve the model fit of MODEL G ($\chi^2(1) < 0.01$, $p = .961$). Nor was the model fit improved when adding an interaction effect for the expressive language score ($\chi^2(1) = 2.97$, $p = .085$) or the IUSC-S score ($\chi^2(1) = 0.91$, $p = .634$).

The model parameters of all MODEL D, E, and G can be seen in Table 9.9. For MODEL D, the CARS-2 score was significant ($t(21) = -3.52$, $p = .002$). The condition was not significant ($t(21) = -0.64$, $p = .529$), nor was session² ($t(70) = -1.78$, $p = .079$), or the interaction between the condition and session² ($t(70) = -1.18$, $p = .243$). The relationship between the two conditions and behavioural engagement showed significant variance in intercepts across the children. In addition, the slopes significantly varied across children, and the slopes and intercepts were negatively and

Table 9.9: Parameter estimates for the growth model on behavioural engagement with the CARS-2 score (MODEL D), bespoke expressive language (Exp. Lang.) score (MODEL E), or both CARS-2 and bespoke expressive language score as covariate (MODEL G). The statistically significant parameter estimates for the fixed effects are in bold, excluding the intercept.

Variable	MODEL D: CARS-2		MODEL E: Exp. Lang.		MODEL G: CARS-2 + Exp. Lang.	
	<i>b</i> (SE)	95% CI	<i>b</i> (SE)	95% CI	<i>b</i> (SE)	95% CI
<i>Fixed effects</i>						
Intercept	1.59 (0.33)	0.95, 1.59	-0.09 (0.20)	-0.47, 0.29	0.89 (0.41)	0.09, 1.69
Session ²	-0.01 (0.01)	-0.03, 0.00	-0.01 (0.01)	-0.03, 0.00	-0.01 (0.01)	-0.03, 0.00
Condition	-0.09 (0.14)	-0.37, 0.19	-0.08 (0.13)	-0.34, 0.18	-0.05 (0.12)	-0.31, 0.20
Ses ² :Cond	-0.01 (0.01)	-0.03, 0.01	-0.01 (0.01)	-0.03, 0.01	-0.01 (0.01)	-0.03, 0.01
CARS-2	-0.04 (0.01)	-0.06, -0.02	-	-	-0.03 (0.01)	-0.05, -0.01
Exp. Lang.	-	-	0.20 (0.06)	0.07, 0.32	0.13 (0.06)	0.01, 0.25
	<i>SD</i>	<i>95% CI</i>	<i>SD</i>	<i>95% CI</i>	<i>SD</i>	<i>95% CI</i>
<i>Random effects</i>						
Intercept	0.41	0.30, 0.56	0.33	0.24, 0.46	0.35	0.25, 0.50
Session	0.12	0.09, 0.17	0.12	0.09, 0.17	0.12	0.09, 0.17

MODEL D: *Behavioural engagement* ~ *Session*² * *Condition* + *CARS-2* + (*Session* | *Participant*).

MODEL E: *Behavioural engagement* ~ *Session*² * *Condition* + *Exp. Lang.* + (*Session* | *Participant*).

MODEL G: *Behavioural engagement* ~ *Session*² * *Condition* + *CARS-2* + *Exp. Lang.* + (*Session* | *Participant*).

significantly correlated ($r = -.73$, 95% CI[-.88, -.45]).

The parameter estimates for MODEL E, including the expressive language score, show a similar trend to MODEL D. The covariate, expressive language score, was significant ($t(21) = 3.25$, $p = .004$). But neither the condition ($t(21) = -0.64$, $p = .530$), session² ($t(70) = -1.78$, $p = .079$), nor the interaction effect between the condition and session² was significant ($t(70) = -1.18$, $p = .243$). The random intercept and slopes showed a significant and negative correlation ($r = -.52$, 95% CI[-.78, -.12]).

The marginal means for behavioural engagement, estimated by MODEL G, can be

Figure 9.3: The predicted values (marginal effects) for behavioural engagement of growth MODEL G, presented in Table 9.9.

seen in Figure 9.3. MODEL G, which includes the CARS-2 and expressive language score as covariates, shows that both covariates significantly contribute to predicting the variance in the behavioural engagement of the children. For the CARS-2 the parameter estimate is -0.03 ($t(20) = -2.51$, $p = .021$). This means that the model estimates that autistic children who scored higher on the CARS-2 were less behaviourally engaged. For the expressive language score, the parameter estimate is 0.13 ($t(20) = 2.18$, $p = .042$), which means that autistic children with more complex expressive language, with scores ranging from 0 to 4, were also more behaviourally engaged than those with less complex expressive language. The model shows that the condition was not significant, nor did the condition interact with the sessions.

9.4.3 Visual attention

The visual attention over the sessions can be seen in Figure 9.4. As our hypotheses only relate to the “robot” and “elsewhere” gaze direction, we will not report on the other annotated gaze directions. Again, we fitted multi-level models, using random slopes, random intercepts, and first-order autoregressive covariance structure, to model the participants’ visual attention with the robot and their visual attention elsewhere. The parameter estimates for the conditional growth model on visual attention *towards the robot* can be seen Table 9.10. The relationship between the robot’s variance and visual attention with the robot showed significant variance in intercepts across the children, but the slopes were non-significant. We investigated whether accounting for the differences between children, as measured by the individual factors, improved the model. In contrast to the previous results on behavioural engagement, this was not the case for any of the measures. The model that best fits the data is therefore a conditional growth model with random intercepts. The visual attention towards the robot significantly decreased over sessions ($t(70) = -2.91$, $p = .005$).

Figure 9.4: The (unmodelled) percentage of time a participant spent looking at each of the annotated directions for each of the sessions (visual attention) and each of the conditions.

Table 9.10: Parameter estimates for the conditional growth models for visual attention towards the robot and elsewhere. The statistically significant parameter estimates for the fixed effects are in bold, excluding the intercept.

Variable	TOWARDS THE ROBOT		ELSEWHERE	
	<i>b</i> (SE)	95% CI	<i>b</i> (SE)	95% CI
<i>Fixed effects</i>				
Intercept	0.46 (0.04)	0.38, 0.55	0.22 (0.04)	0.14, 0.29
Session	-0.03 (0.01)	-0.05, -0.01	-0.00 (0.01)	-0.02, 0.02
Condition	0.03 (0.06)	-0.09, 0.15	-0.10 (0.05)	-0.21, 0.00
Session:Condition	-0.02 (0.01)	-0.05, 0.01	0.05 (0.01)	0.03, 0.08
	<i>SD</i>	95% CI	<i>SD</i>	95% CI
<i>Random effects</i>				
Intercept	0.10	0.07, 0.14	0.10	0.08, 0.15

VISUAL ATTENTION TOWARDS THE ROBOT $\sim$ Session * Condition + (1 | Participant).

VISUAL ATTENTION ELSEWHERE $\sim$ Session * Condition + (1 | Participant).

Figure 9.5: The predicted values (marginal effects) for the visual attention towards the robot (A) and elsewhere (B) of the growth models presented in Table 9.10.

This could be due to other factors that influence visual attention, such as a novelty effect that wears off, or boredom. There was no effect of condition ($t(22) = 0.56$, $p = .579$) nor was there an interaction effect between condition and session ($t(70) = -1.24$, $p = .219$).

The model parameters for the children's visual attention *elsewhere* can also be seen in Table 9.10. We found no significant variance in the slopes across children. The variance in intercepts did significantly differ. Again, accounting for differences on the autism-specific measures did not improve the model. Therefore, the data was best described by a conditional growth model with random intercept. The visual attention elsewhere did not significantly decrease over sessions ($t(70) = -0.12$, $p = .908$), nor was there a significant difference between conditions ($t(22) = -1.75$, $p = .095$). There was, however, a significant interaction effect between the condition

and session ($t(70) = 3.90, p < .001$). In the high-variance condition, children looked increasingly “elsewhere” over sessions compared to the low-variance condition. Thus, in contrast to the results for behavioural engagement, this shows an impact of robot predictability on the children’s visual attention — which is indicative of engagement — to the robot-assisted activity.

9.5 Discussion

The goal of our study was to investigate the interplay between robot predictability, the behavioural engagement and visual attention to the activity, and the idiosyncrasies therein between autistic children. To that end, we manipulated the variance in the robot’s behaviour as a way to operationalise predictability, and measured the children’s behavioural engagement and visual attention in a robot-assisted activity, as well as individual factors.

We found that for the less predictable robot, autistic children paid less *visual attention* to activity-relevant locations as the sessions progressed. Rather, the children started to pay more attention to locations where nothing moves (annotated as “elsewhere”), such as the walls, the cameras, or the room divider. In contrast to our predictions, however, the robot’s predictability did not impact on *behavioural engagement*. The children continued to engage with the robot-assisted activity regardless of the robot’s predictability. We did find, though, that *individual differences* in children’s background characteristics were related to their behavioural engagement. Higher autistic features were related to less behavioural engagement of autistic children, while the children’s expressive language ability was related to greater behavioural engagement. Visual attention, on the other hand, was not influenced by any of the individual differences measured. Finally, we found no evidence for a relation between the autistic children’s intolerance of uncertainty and their response to the robot’s predictability in terms of their behavioural engagement and visual attention.

In conclusion, it appears that the children continued engaging in the robot-assisted activity, but started to pay less visual attention over time to the activity-relevant locations when the robot was less predictable. Instead, the children started to pay more attention to locations that visually do not change (i.e. no changes in sensory information), such as the walls, or the floor — they look away from the activity. We believe this might be a coping strategy to minimise sensory input and deal with anxiety resulting from the inability to learn to accurately predict the robot’s actions. Similarly to how stimming can be used by autistic children to generate predictable sensory information to deal with an overload of unpredictability (Sinha et al., 2014). For learning, paying less visual attention to the activity-relevant locations is problematic as it indicates that the child is less strongly engaged with learning tasks. In particular, this may impact long-term use of robot-assisted interventions, as in our study this effect got stronger over sessions. We may also expect that eventually the children may start to be less *behaviourally* engaged, when the encouragement from the adult starts to fail in motivating the child to engage with the robot and deal with the resulting unpredictability. However, as we did not measure the children’s learning, we cannot draw a firm conclusion about whether the diminished visual attention also impacts

the learning.

The result that individual differences influence the behavioural engagement of autistic children is in line with our findings from the descriptive study reported in Chapter 4, as well as findings from [Rudovic et al. \(2017\)](#), who also found correlations between autistic features — measured through the CARS-2 — and behavioural engagement. Similarly, [Kostrubiec and Kruck \(2020\)](#) found correlations between autistic features — measured via the SCQ — and the proportion of prosocial behaviours in a robot-assisted intervention. Not only do individual differences seem to influence to what degree autistic children are behaviourally engaged in activities with robots, they are also correlated with how they behaviourally engage ([Schadenberg et al., 2020b](#)). While researchers often report increased engagement, increased levels of attention, and novel social behaviours when incorporating a robot in the interaction ([Scassellati, 2007](#); [Ricks and Colton, 2010](#); [Diehl et al., 2012](#); [Scassellati et al., 2012](#)), the finding that individual differences influence the behavioural engagement of autistic children raises the question whether the children’s individual differences only influences their behavioural engagement, or whether they *moderate* the effect of a robot on behavioural engagement. Understanding this relationship would allow us to better determine which autistic children may benefit in particular from such robot activities or interventions. Lastly, these findings also suggest that the way autistic children interact with, or in the presence of, robots is indicative of their autistic features. For robots that are used for diagnosing autism, this finding could indicate that the way children interact with a robot could tell us something about the severity of the autistic features.

Whether certain individual factors also moderate the effect of robot predictability on engagement remains unanswered by this study, as we found no evidence that this was the case. In particular, we had expected that those higher in IU would also be more strongly affected by the robot’s predictability, given the similarity between IU and unpredictability. While other studies did find relationships between IU and autistic features, such as the presence of sensory sensitivities ([Neil et al., 2016b](#)), repetitive motor behaviours and insistence on sameness ([Wigham et al., 2015](#)), we found no results in our study that support a relationship between IU and predictability. Similarly, to our study with typically developing adults reported in the previous chapter, we found no relationship either between IU and the social perception of the robot in terms of warmth, competence, and discomfort. While unpredictability and uncertainty are often used interchangeably, suggesting conceptual similarity ([Grupe and Nitschke, 2013](#)), our results suggest that greater caution is warranted when doing so. In the current study, we defined and manipulated robot predictability in terms of making it more or less difficult to learn to predict the robot’s behaviour. However, the questionnaires on IU are more about uncertainty regarding events further into the future than the robot’s next action(s) would be. The children knew that they were going to interact with “Zeno the robot” at a certain point during the day, that they would then play several games with the robot, and then go back to class — a fixed structure. In that sense, the children could predict the robot’s behaviour on a more abstract level (i.e. the robot will perform actions related to the games), but not on precisely when, where, and how those robot actions would appear. Possibly, IU relates

to an intolerance for situations without a clear structure on what is going to happen in a longer time span, resulting in the unpredictability of future events. Predictability, as we defined and manipulated it, related more to predicting sensory information in the immediate future. This distinction may also be relevant when considering the role of predictability and uncertainty in the non-social features of ASC. In particular in insistence of sameness (e.g. inflexible adherence to routines, or ritualised patterns) given that both IU and insistence of sameness refer to liking things to be predictable and a dislike of change (Chamberlain et al., 2013).

To our knowledge, this is the first study that has operationalised and manipulated predictability in a real-world setting and showed how the extent of unpredictability can be quantified. The metrics in our manipulation check can be used to compare the degree of unpredictability between studies using robots. In general, robotic technology is uniquely positioned for investigating predictability in autistic children, as they allow us to carefully manipulate its predictability (unlike with humans), but they also elicit social interactions in autistic children. There is some preliminary work in trying to teach autistic children to deal with unpredictability (e.g. Rodgers et al., 2017; Hallett et al., 2020). To this end, robots may be particularly useful in that its unpredictability can be carefully increased in both intensity as well as in predicting different aspects of the environment (e.g. predicting robot actions, or predicting future events).

9.5.1 Limitations

In our study, we manipulated the predictability of the robot's behaviour through its variance. This ought to have made it more difficult to learn to predict the robot's behaviour. Through our manipulation check, we concluded that this manipulation was successful. However, possibly even in the high-variance condition, the robot's behaviour was not *problematically* unpredictable, as is the unpredictability of human behaviour for example. In practice, robots can be more unpredictable than we could manipulate in our study, as we needed to keep the conditions comparable and avoid confounding factors. In the future, robots will become more sophisticated and capable of levels closer to humanlike behaviour. In turn, so too does their ability to be unpredictable. Thus, robot predictability could affect the engagement of autistic children more strongly when they become more sophisticated and humanlike. In our study, we looked at behavioural predictability, which is not the same as attributed predictability (see Chapter 6.4). In the previous chapter, and in the study by Driggs-Campbell and Bajcsy (2016), results indicated that even though a robot's behaviour is more difficult to predict, it can still be considered to be more predictable. Current measures for attributed predictability are too complex to be used for autistic children in our study, as they rely heavily on language and intellectual ability. When new measures become available, it would be interesting to assess the autistic children's attributed predictability in relation to different levels of variance in the robot's behaviour.

Note that there are several sources of variance that could not be controlled fully in our study. First, there is inherent variation in robot motion due to limits in reproducibility of motion by the robot's stepper motors, which may differ slightly with each operation. Second, we were not able to control for the affective quality and the at-

tractiveness for engagement of each specific variant we chose, despite the fact that we selected ones with similar verbal and/or motion qualities. For example, children may have experienced more negative affect when the robot used the specific term “Time for dancing”, which was present only in the high-variance condition. Additionally, for some sessions, technical difficulties resulted in unintended behaviours and form of behaviours, introducing some variability. These instances were not annotated, but the effects could possibly carry over to later in the session.

Finally, we also note that, as with most studies that concern robots and autistic children (see (Diehl et al., 2012; Begum et al., 2016)), our sample size was relatively low. This can negatively influence the generalisability of our results and prevents us from drawing strong conclusions. It is therefore important to take the reported margins of error into account when interpreting our results.

9.6 Conclusion

Predictability is important to autistic individuals, and robots have been suggested to meet this need as they can be programmed to be predictable, as well as elicit social interaction. However, little was known about the interplay between robot predictability, engagement in learning, and the individual differences between autistic children. Here, we systematically manipulated the robot’s predictability, and measured the behavioural and visual attention of the autistic children. Additionally, we also measured several individual factors, including the children’s autistic features, expressive language ability, and intolerance of uncertainty. We found that the children will continue engaging in the activity behaviourally, but start to pay less visual attention over time to activity-relevant locations when the robot is less predictable. Instead, they increasingly start to look away from the activity. Ultimately, this could negatively influence learning. In particular for tasks with a visual component, where paying less visual attention leads to fewer opportunities for learning. Furthermore, we found that the severity of autistic features and expressive language ability had a significant impact on behavioural engagement.

We consider our results as preliminary evidence that robot predictability is an important factor for keeping children in a state where learning can occur. In particular, in long-term interactions with many sessions, our results indicate that the trend of paying less visual attention increases over time. As individual differences between autistic children were shown to have a significant impact on behavioural engagement, future studies should therefore be careful to account for these differences. Finally, our study indicates that predictability can be studied in real-life scenarios with real stimuli, rather than artificial stimuli in lab settings. We also showed how the degree of predictability can be quantified in a way that it can be used as a manipulation check to display the degree of unpredictability between conditions in a real-world setting.

Once the engagement of the children with the robot has been further clarified, future research should consider looking into the affective component of engagement to investigate whether a more predictable robot is more enjoyable to the children. Additionally, future research is needed to determine whether “higher quality” engagement also leads to more and faster learning. After all, increased engagement alone is in-

sufficient to justify robots to assist in interventions for autistic children when it does not also lead to increased learning. In our study, we opted for a holistic approach to increasing the robot's unpredictability by implementing several types of variance. Future research should examine to what extent each types of variance influences autistic children. This would allow us to more carefully take the robot's predictability into account when designing its behaviour. Lastly, the goal of our study was to investigate how we should design robots to best aid autistic children in learning, specifically focusing on their need for predictable environments whilst taking the children's idiosyncrasies into account. Future research could also investigate the role of a robot's predictability in engaging typically-developing children in learning, in order to assess whether or not predictability is uniquely important to autistic children.

Part IV

DELIBERATION

“What’s past is prologue.”

William Shakespeare, *The Tempest*

10

Discussion and conclusion

In this chapter, I will reflect on the work we carried out and summarise the findings and contribution of this dissertation by answering the research questions. Next, I will reflect back on the work that we carried out and touch upon new research questions. I will discuss these in the future work section. And finally, the last section of this dissertation — some closing remarks.

10.1 Main findings and contributions

IN this section, I will briefly summarise the main findings from this dissertation, and conclude on the research questions that I formulated at the beginning of this dissertation in Chapter 1.5. These were:

Research question 1: How can a robot-assisted intervention be designed to engage autistic children in learning?

Research question 2: What is predictability in relation to people interacting with robots?

Research question 3: How does robot predictability influence human-robot interaction?

10.1.1 RQ1: on designing robot-assisted interventions that engage in learning

To answer the first research question, we started by investigating the theoretical background on robot-assisted interventions in Chapter 2. There, we looked at what current robot-assisted interventions look like, what the concept of engagement relates to, how children interact in robot-assisted intervention settings, and in what ways these children differ from each other in terms of their autism. We saw that the current state of the art robot-assisted interventions can maintain engagement for a month and operate autonomously, but that the role of the robot in this is unclear. Furthermore,

stronger experimental designs are needed to assess whether robot-assisted interventions can lead to learning that generalises to other settings and to humans. What also became apparent is that autistic children can interact with a robot very differently from each other, where the robot may not be perceived as a social entity. These differences can, in part, be explained by the idiosyncrasies of how each child is affected by autism.

During our assessment of the theoretical background on robot-assisted interventions, we noticed that several papers mentioned specific user requirements that are considered important for the design of a robot-assisted intervention. However, we found no papers that provided a more comprehensive list of user requirements. To this end, we conducted a systematic literature search to identify such user requirements, and also included two studies that we ourselves conducted on this topic (Chapter 3). This resulted in a list of user requirements that can serve as a basis for the development of any robot-assisted intervention. Again, we saw that the idiosyncrasies of autistic children play a large role in *how* the user requirements can be addressed, and thus, whether a child can successfully engage in a robot-assisted intervention.

To quantify the role of the autistic children's idiosyncrasies, we conducted a large descriptive study to assess the role of autistic traits in how autistic children spontaneously interact within a robot-assisted intervention (Chapter 4). In this study, we found large differences between autistic children in how the type of interactions they engage in, where some children were predominantly engaged in spontaneous exploratory interactions, and others engaged in functional interactions, or did not initiate in any spontaneous interactions. The outcome of this study can be used to tailor the autistic child-robot interaction to specific subgroups of autistic children who share similar autistic traits. The study also provided several avenues for designing interactions that may be particularly engaging for autistic children, as we observed that they spontaneously engaged in such interactions.

In Chapter 5, we explored several of the potentially highly engaging interactions, and assessed how we could address several of the user requirements that we had identified earlier. We consider the use of tangibles to be particularly promising as they can provide tactile interaction as well as serve as a user interface for the child to communicate with the robot. Importantly, such tangible user interfaces do not depend on language understanding, which many autistic children have difficulty with. Furthermore, from observing the interactions of the children with the tangibles, they seemed to be highly engaged with the tangibles, while still switching their attention to the robot when required.

To conclude, the results of our studies indicate that for designing robot-assisted interventions which can engage autistic children in learning, sustain such engagement for longer periods of time, and lead to learning that could generalise to other settings and to humans, the idiosyncrasies of autistic children will need to be taken into account. The solution to this is likely to be two-fold. First, there is no one-size-fits-all robot that can do this for any autistic child — the differences between the children are simply too large. A robot-assisted intervention will therefore have to be designed for a specific subgroup of autistic children who share certain commonalities. For instance, a robot for autistic children with high support needs (e.g. high autism severity, and

low language ability) may need to use very simple and short speech, and teach about basic social skills. Such a robot is unlikely to keep the interest of an autistic child who can function at a regular school, but who has difficulty with more advanced social skills. For this child, a robot would not have to adjust its speech much, and the interaction itself can be more advanced in terms of complexity. Our results suggest that autistic traits, such as those measured by the ADOS-2 or the CARS-2, as well as the children's language ability, can be useful for identifying such subgroups of autistic children. Other studies have shown that the children's sensory profiles, measured through the Adolescent/Adult Sensory Profile (Brown and Dunn, 2002), is indicative of how they interact with a robot (Chevalier et al., 2017c,b). The children's sensory profile may be particularly useful for addressing any hypo- or hyper-sensitivities.

Second, individual differences can also (at least to a certain extent) be addressed by personalising the interaction. For instance, by drawing the child's attention by utilising the child's specific interests (Putnam et al., 2019). But also its appearance could be personalised (Huijnen et al., 2017), or the difficulty of the learning content (Scasellati et al., 2018; Clabaugh et al., 2019). With advances in robot perception, the interests or preferences of a child may be autonomously detected through the robot's sensors, and can be used to personalise its behaviour. A different approach would be to design the intervention in such a way that the children can decide themselves on how to interact with the robot. For instance, by giving the child control over the interaction. The child could then decide to interact with the robot through tactile interaction, or other ways of interaction. This way, the child can personalise the interaction themselves. A key challenge would then be to still have the child engage with the learning material, as this material may not always be as interesting as other things the child can think of.

10.1.2 RQ2: explication of predictability as a concept

We answer the second research question by providing an explication of predictability, in the context of HRI, gained from applying insights from the predictive processing account of human cognition. When interacting with a robot for the first time, it will be less predictable, and predictability of a robot is dependent on the expectations one holds. As people interact with the robot, they will learn to predict its behaviour, and predictions will be based on previous interactions with the robot. The predictions relate to all aspects of the robot — from low-level motion primitives to high-level beliefs, desires, or intentions. We therefore defined predictability for HRI as: “the ability to quickly and accurately learn to predict the behaviour of a robot”. Robots are more predictable when their behaviour has higher predictive value: the extent to which information about the robot's previous behaviours predicts its future behaviour. Regarding what about the robot's behaviour is predicted, we distinguish three prediction levels, namely the action level (predicting the end-state of an on-going action), the interaction level, (predicting the future action of the robot), and the activity level (predicting the gist of the robot's future actions). The description of predictability above relates to what we call behavioural predictability. This is different from the predictability that people can attribute to a robot, which we refer to as attributed predictability.

In order to conduct an experiment where we manipulate the robot's behavioural predictability, further operationalisation was required. To this end, we operationalised behavioural predictability in terms of variance in the robot's behaviour. The more the robot shows variance in its behaviour, the more difficult it becomes to perceive and learn the structural regularities in the robot's behaviour. While behavioural variance can be reduced, a degree of variance (and therefore risk of unpredictable behaviour) is required for most robot applications in order for the robot to perform its primary task. After all, a fully predictable robot would perpetuate in repetitive behaviour, limiting its long-term usefulness, and would also be unable to autonomously respond to the (unpredictable) dynamics of real-world settings. We call the resulting unpredictability from such sources of variance as irreducible unpredictability. On the other hand, some sources of behavioural variance are not required for the robot to perform its primary task, and are thus reducible. For these sources, designers can choose whether the resulting behavioural unpredictability outweighs some other benefit that can be gained by keeping the source that (also) adds behavioural variance.

To conclude on the second research question, predictability is a broad concept that plays a central role in how we perceive and think about that robot. We believe that the explication of predictability we provided is fundamental for setting up future studies that explore this concept in relation to HRI. In general, our explication, and distinction between attributed and behavioural predictability, improves our ability to communicate more precisely what one means with predictability when referring to this concept. Our explication can also provide a theoretical basis from which possible designs for improving a robot's predictability can be derived. In particular, robots could be made more behaviourally predictable by providing the user with more *control* over the robot's behaviour, ensuring the robot's behaviour is *consistent*, and ensuring that the (external or internal) *contingencies* between the robot's actions and events are apparent to the user.

10.1.3 RQ3: influence of a robot's predictability on HRI

To address the third research question, we conducted two experimental studies. One with typically developing individuals, and another with autistic children. In Chapter 8, I reported on the study with the typically developing individuals, where we investigated how a robot's behavioural predictability influenced the social perception of people and the attributed predictability of the robot. In this study, we operationalised the robot's predictability through adding topic variance. This was done by manipulating whether participants could either see the cause of a robot's responsive action, or could not see this, because there was no cause, or because we obstructed the visual cues. Our results indicate that when the cause of the robot's responsive actions was not visible, participants rated the robot as more unpredictable and less competent, compared to when it was visible. The relationship between seeing the contingent relation between an event and a resulting robot's action and the attribution of competence was partially mediated by the attribution of unpredictability to the robot. The results provide evidence that it is important that a person is aware of the contingent relation between an event and a resulting robot's action upon which the robot acted, as the attribution of competence plays a large role in how a person

characterises others (Wojciszke et al., 1998), and is believed to play a similar role in characterising robots (Carpinella et al., 2017). We argued that the effects of unpredictability may be mitigated when the robot identifies when a person may not be aware of what the robot wants to respond to, and uses additional actions to make its response predictable.

In the final joint-study of the DE-ENIGMA consortium (Chapter 9), we investigated how the robot's predictability influenced the engagement of autistic children within a robot-assisted intervention. To this end, in one condition, the robot showed behaviour that was inconsistent with its previous behaviours in terms of its motions, speech, and timing of its actions, as well as performing non-contingent actions (i.e. actions that had no apparent cause). In the other condition, the robot was consistent and only showed behaviours that were contingent on a visible event. We found that the children will continue engaging in the activity behaviourally, but may start to pay less visual attention over time to activity-relevant locations when the robot is less predictable. Instead, they increasingly start to look away from the activity. Ultimately, this could negatively influence learning, in particular for tasks with a visual component. Furthermore, the severity of autistic features and expressive language ability had a significant impact on behavioural engagement. We consider our results as preliminary evidence that robot predictability is an important factor for keeping children in a state where learning can occur. Furthermore, our study, and analysis thereof, shows that predictability can be studied in real-life scenarios with real stimuli, rather than through artificial stimuli in lab settings, and thus improve the external validity of studies that investigate predictability.

To conclude on the third research question, a robot's predictability influences our social perception of it and to what extent we consider the robot to be predictable. For autistic children, a less predictable robot negatively influenced the children's visual attention to activity-relevant locations. These are the first studies that directly manipulated a robot's predictability and studied its effects. In combination with the theoretical accounts of autism (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014), the importance of predictability for autistic children in practice (e.g. MacDuff et al., 1993; Bryan and Gast, 2000; O'Reilly et al., 2005; Mesibov and Shea, 2010), we conclude that predictability is an important concept to take into account for robots that are to interact with autistic children. Similarly, for typically developing individuals, predictability has been argued to be an important factor for people interacting with robots (e.g. Sciutti et al., 2018; Sebanz et al., 2006; Sebanz and Knoblich, 2009; Hancock et al., 2011; Lewis et al., 2018; Noorman and Johnson, 2014). Our results provide experimental evidence that this is indeed the case for the social perception of the robot.

10.2 Reflection and future work

The conclusions reported in the previous section provide answers that are a first step in understanding the concept of predictability and its effects on HRI, as well as how we can design robot-assisted intervention for autistic children that may sustain engagement. As with all research, answers lead to more questions. Questions that will

need to be addressed in future work. Some of these new questions are tied to the limitations of our own work. Whereas other new questions that we have stem from our increased understanding on the topics addressed in this dissertation. In this section, I will reflect back on the limitations of our work and discuss three research directions that can advance the research reported in this dissertation.

10.2.1 Utilising and understanding (robot) predictability

Teaching autistic children to deal with unpredictability. With our conceptualisation and operationalisation of robot predictability, we can now carefully manipulate the robot's predictability. This enables us to start developing robot-assisted interventions aimed at teaching autistic children to deal with unpredictability. The robot then serves as a scaffolding tool that can adjust its predictability. If difficulties of dealing with unpredictability are indeed at the core of autism (Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014), *such an intervention could then address the core of the condition*, rather than addressing symptoms, and by doing so possibly reducing autism symptoms that are unwanted by the autistic individual. Critically, robots are uniquely positioned to take on this role, as they also elicit social interactions, turning them into a predictability-controlled surrogate. Note that humans cannot fulfil this role, as we cannot systematically control the predictability of our behaviour. Even actors, who are adept at controlling their behaviour, will show uncontrolled variance.

The relation between predictable environments, learning, and comfort. The results of the study reported in Chapter 9 indicate that autistic children start to pay less visual attention when the robot is less predictable. This could negatively influence the child's learning, as well as their comfort during the intervention. These are two critical factors for any robot-assisted intervention to be successful. It is therefore important that we continue this line of research and start investigating experimentally whether a more predictable environment, engendered through the use of a robot, improves learning and/or comfort.

Understanding the prediction mechanism of autistic individuals. The explication of robot predictability that we provided is based on insights from the predictive processing account of human cognition. This account, and similar Bayesian accounts, have been used to explain autism and how the cognition of autistic individuals differs from typically developing individuals (e.g. Pellicano and Burr, 2012; Lawson et al., 2014; Sinha et al., 2014; Van de Cruys et al., 2014). Each of these accounts of autism posits a different element of predictive/Bayesian processing that is atypical. Whether it is attenuated priors that cause atypicalities in the prediction mechanism (Pellicano and Burr, 2012), aberrant encoding of precision (Lawson et al., 2014), inaccuracies in estimating the predictive relationship between two events from an observed temporal sequence (Sinha et al., 2014), or the inability to flexibly process prediction errors (Van de Cruys et al., 2014), each account states that the mechanism is atypical in some way. Until we know more about whether and how the prediction mechanism is different in autistic individuals, and between autistic individuals, we cannot be

certain that the explication of predictability that we provided generalises to autistic processing of sensory information. Without this knowledge, we could be adjusting the robot's predictability that improves its behavioural and attributed predictability for typically developing individuals, but fails to do so for autistic individuals, simply because their brain generates or uses predictions differently. If we are to effectively utilise the ability for us to control a robot's predictability, understanding the prediction mechanism of autistic individuals will be critical.

10.2.2 Towards designing effective robot-assisted interventions

Understanding the idiosyncrasies of autistic children. Throughout this dissertation, we have stressed the importance of taking the idiosyncrasies of autistic children in account. Not only in statistical analyses, but also in the design of robot-assisted interventions. A key challenge is to figure out how we can design for a spectrum condition, such as autism. When you design a robot for a subgroup of children that is too large, the effectiveness of the intervention will be reduced for certain children, whereas designing for a too small of a subgroup limits the impact the intervention could have. In our work, we mostly looked at autistic traits and language ability for categorising subgroups of autistic children, as there are excellent tools for measuring these, such as the ADOS-2 (Lord et al., 2012) or CARS-2 (Schopler et al., 2010). The cognitive ability also play a large role in how a robot-assisted intervention should look like, as the children with lower cognitive abilities require more support. In future work, it will be important to assess different ways of personalising content to individual autistic children. Personalising not only in terms of game difficulty (Scassellati et al., 2018; Clabaugh et al., 2019) and feedback (Clabaugh et al., 2019), but also in the robot's behaviour (e.g. whether and how it uses speech), learning content, and adapting to sensory issues. Furthermore, automatic detection for the need for personalising may help alleviate the adult from manually inputting this, or assist the adult in making a choice on how to personalise the intervention.

Assessing engagement and predictability in long-term studies. Through our research, we hope to be able to better design robots so that they can engage autistic children in learning a targeted skill. Current interventions for autistic children generally take many sessions before a targeted skill is internalised by the child. As such, a robot will also need to be able to sustain long-term interactions with an autistic child. The results that we found were all from studies where children only interacted with a robot for a handful of sessions. What is important during the first couple of sessions may change as children continue to interact with a robot. For example, the influence of predictability is presumably different during initial interactions than for long-term interactions (Wyart et al., 2012). While people's (social) perception is initially less dependent on predictions, it should become more dependent on predictions later as their precision increases. And with higher precision predictions, more attention will be paid when such a prediction turns out to be inaccurate (Feldman and Friston, 2010). Moreover, our own results reported in Chapter 9 indicated that autistic children start to pay less visual attention to a robot-assisted activity when it is less predictable in its behaviour. In long-term interactions, this effect may be ex-

acerbated. While the insights from the research presented in this dissertation can be used to design more engaging interactions, it is yet unclear how these insights relate to longer-term interactions. This is left for future research.

Designing and assessing tangible user interfaces for autistic children. We believe that providing tangible interactions, through which an autistic child can communicate with a robot, is a promising avenue for providing an accessible user interface for autistic children. Providing the children with an interface to communicate with the robot can give them more control to shape the interaction to their choosing, which in turn can improve the predictability of the robot (if you command the robot to do “A”, you also predict it to do so). A *tangible* user interface may have several other benefits as well. First, a TUI does not rely on the child’s ability to understand and use language, which many autistic children have difficulties with. In turn, the robot does not necessarily require robust speech recognition and natural language processing, which are still unsolved problems in artificial intelligence for typically developing individuals, and in particular for children (Kennedy et al., 2017) (let alone the atypical use of language by autistic children). Second, TUIs can be designed in a way that it does not look like a known device, such as a tablet. In our experiments, some children wanted to play with the tablet, rather than with the robot, and were fully absorbed to find the home button and access their favourite app. This obviously was not the goal of our robot-assisted intervention and would prevent them from engaging with any of the provided learning content. Last, a TUI could provide tactile stimulation for those children who actively seek out such stimulation — so-called “sensory seeker” (Hazen et al., 2014). In our descriptive study, reported in Chapter 4, we saw that quite a few children spontaneously touched (or tried to touch) the robot, and persisted in such behaviour. In Serbia, the occupational therapists from the DE-ENIGMA project used a pin-art frame to calm down such children. In general, the pin-art frame was a great success for many autistic children during free-play. However, there is the risk that the TUI might be too interesting and absorb the child’s attention. Future research will need to determine whether this is so.

Note that the evidence that we provide in this dissertation regarding tangibles is limited to exploratory research and descriptive observations. The seemingly engaging interactions that we observed when the children played with the tangibles could have different causes. For instance, the interaction could have been more playful, or self-directed, than the activities with the robot, and explain why we observed seemingly high engagement.

Designing robot-assisted interventions *with* autistic children. In our systematic literature search, reported in Chapter 3, we noted that none of the papers that we found involved the autistic children themselves in their research. That is, we are designing *for* autistic children, but not *with* them. In the DE-ENIGMA project, and my own work in this, we also did not the children in the design process itself. Thus, while we consider our design direction to be suitable to autistic children, they may view it differently, and might prefer a different direction. To our knowledge, there is little participatory design research with autistic children within the field of HRI.

In particular with those children with high support needs and limited language and cognitive ability. For future work, it is therefore important to try and include autistic children in the design of robot-assisted interventions. For instance, by using a combination of ethnography and structured observations to understand their experiences (e.g. [Pellicano et al., 2014](#)). Or, alternatively, autistic children who are further in their development could be involved in a study — for whom there are many methods to engage them in research (see [Frauenberger et al., 2017](#); [Spiel et al., 2017](#)) — and represent autistic children with more difficulty in their communication and cognition.

10.3 Closing remarks

The use of robots for aiding autistic children in some way has come a long way since the pioneering work from [Emanuel and Weir \(1976\)](#). Robots are no longer restricted to a dozen or so actions that it can perform, but instead they can utilise powerful perception, reasoning, action cycles, allowing for complex and diverse robot behaviour. As we have seen in this dissertation, a very limited set of robot behaviours may, however, not necessarily be a bad thing when it comes to interacting with an autistic child. The robot's behaviour may be more easily learned by the child, improving its predictability. This can allow the child to focus on what's important in a task, without having to continuously resolve unpredicted perceptual events. Thus, robots for autistic children may not have to be very sophisticated in order to provide value to them, or to those who are to work with the robot. This could be important for robotics companies who are developing a robot for autistic children, as one reason why robotics companies are having difficulty to succeed is that it takes too much capital to develop a robot that can solve actual user needs [Fresh Consulting \(2020\)](#). A robot for autistic children does not necessarily need to be expensive in terms of hardware, nor very rich in its behaviour, to provide value and address actual needs of these children. I therefore hope that such robots could be developed and improved to see a larger uptake and improve the lives of the children, their families, and the people who work with them. Hopefully, the work described in this dissertation furthers this goal. Only time will tell...

About the author

Bob Schadenberg is post-doctoral researcher at the University of Twente's Human Media Interaction group. He received a bachelor's degree in Psychology and master's degree in Human Machine Communication at the University of Groningen, the Netherlands. His research interest lies on the intersection of Psychology and Artificial Intelligence (AI), and focuses on understanding how people make sense of autonomous systems. To this end, he specifically likes to apply insights from cognitive science to people interacting with autonomous systems. As our understanding of how people make sense of autonomous systems grows, so too does our ability to support this process in the development of such systems.

After earning his master's degree, Bob worked as a product researcher at Philips, where he conducted consumer and usability research for the development of personal care products. And later, he worked at the IBM CIC as a data scientist and application developer. In 2015, just before Christmas, Bob embarked on his Ph.D journey, under the supervision of dr. ir. Dennis Reidsma, prof. dr. Dirk Heylen, and prof. dr. Vanessa Evers, on a topic where he could combine his knowledge of psychology and AI: autistic child-robot interaction.

Besides the EU-funded project DE-ENIGMA, Bob was briefly involved in the EU-project ALIZ-E, while conducting research at TNO for his master's thesis. This project set out to build the artificial intelligence for small social robots, which was to sustain long-term engagement with children with diabetes. In his research for this project, he was one of the first researchers who applied models from Intelligent Tutoring Systems to Human-Robot Interaction. Currently, Bob is working on the EU-funded project HARMONY. The goal of this project is to develop assistive robotic mobile manipulation technologies for use in hospital environments. In this project, Bob continues his line of research on understanding how people make sense of autonomous systems, and with this knowledge, develops explainable behaviour for such systems.

When Bob is not at work, he enjoys being in the outdoors: long-distance hiking, running, or speeding through the woods on his cyclocross bike. His motto during such activities is that a degree of physical and mental hardship puts life in perspective. Bob also enjoys a glass of whisky, a good fiction book, or playing Dungeon & Dragons with friends. But most of all, he loves spending time with friends, family, and his girlfriend.

List of publications

The following publications were a core part of this dissertation:

- **Schadenberg, B. R.**, Reidsma, D., Heylen, D. K. J., & Evers, V. (2020). Differences in Spontaneous Interactions of Autistic Children in an Interaction With an Adult and Humanoid Robot. *Frontiers in Robotics and AI*, 7, 19 pages. doi:10.3389/frobt.2020.00028 ([Schadenberg et al., 2020b](#))
- **Schadenberg, B. R.**, Li, J. J., Petrović, S., Reidsma, D., Heylen, D. K. J., & Evers, V. (2020). Helping Educators Monitor Autistic Children's Progress Across Sessions. In the *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (New York, NY, USA: ACM), 436–438. doi:10.1145/3371382.3378378 ([Schadenberg et al., 2020a](#))
- **Schadenberg, B. R.**, Reidsma, D., Heylen, D. K. J., & Evers, V. (in press). I See What You Did There: Understanding People's Social Perception of a Robot and its Predictability. *ACM Transactions on Human-Robot Interaction*. 27 pages. ([Schadenberg et al., in press](#))
- **Schadenberg, B. R.**, Reidsma, D., Evers, V., Li, J. J., Davison, D. P., Heylen, D. K. J., Neves, C., Alvito, P., Shen, J., Pantić, M., Cummings, N., Schuller, B. W., Olaru, V., Sminchisescu, C., Babović Dimitrijević, S., Petrović, S., Williams, A., Alcorn, A. M., & Pellicano, E. (under revision). Predictable Robots for Autistic Children — Variance in Robot Behaviour, Idiosyncrasies in Autistic Children's Characteristics, and Child-Robot Engagement. *ACM Transactions on Computer-Human Interaction*. 39 pages. ([Schadenberg et al., under revision](#))

Other publications where Bob was a (co-)author:

JOURNAL ARTICLES:

- **Schadenberg, B. R.**, Neerincx, M. A., Cnossen, F., & Looije, R. (2017). Personalising game difficulty to keep children motivated to play with a social robot: A Bayesian approach. *Cognitive Systems Research* 43, 222–231. doi:10.1016/j.cogsys.2016.08.003 ([Schadenberg et al., 2017](#))
- Alcorn, A. M., Ainger, E., Charisi, V., Mantinioti, S., Petrović, S., **Schadenberg, B. R.**, Tavassoli, T., & Pellicano, E. (2019). Educators' Views on Using Humanoid Robots With Autistic Learners in Special Education Settings in England. *Frontiers in Robotics and AI* 6, 1–15. doi:10.3389/frobt.2019.00107 ([Alcorn et al., 2019](#))

CONFERENCE PAPERS:

- Chevalier, P., Li, J. J., Ainger, E., Alcorn, A. M., Babovic, S., Charisi, V., Petrović, S., **Schadenberg, B. R.**, Pellicano, E., & Evers, V. (2017a). Dialogue Design for a Robot-Based Face-Mirroring Game to Engage Autistic Children with Emotional Expressions. In *Social Robotics: 9th International Conference, ICSR*

2017, Tsukuba, Japan, November 22-24, 2017, Proceedings, eds. A. Kheddar, E. Yoshida, S. S. Ge, K. Suzuki, J.-J. Cabibihan, F. Eysse, and H. He (Cham: Springer International Publishing), vol. 10652 of Lecture Notes in Computer Science. 546–555. doi:10.1007/978-3-319-70022-9_54 ([Chevalier et al., 2017a](#))

- Li, J. J., Davison, D., Alcorn, A. M., Williams, A., Babović Dimitrijević, S., Petrović, S., Chevalier, P., **Schadenberg, B. R.**, Ainger, E., Pellicano, E., & Evers, V. (2020). Non-Participatory User-Centered Design of Accessible Teacher-Teleoperated Robot and Tablets for Minimally Verbal Autistic Children. In *Proceedings of the 13th ACM International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '20)* (Corfu, Greece), 51–59. doi: 10.1145/3389189.3393738 ([Li et al., 2020](#))

WORKSHOP PAPERS, EXTENDED ABSTRACTS, AND DOCTORAL CONSORTIA:

- Davison, D., Gorer, B., Kolkmeier, J., Linssen, J., **Schadenberg, B. R.**, van de Vijver, B., Campbell, N., Dertien, E., & Reidsma, D. (2017). Things that Make Robots Go HMMM: Heterogeneous Multilevel Multimodal Mixing to Realise Fluent, Multiparty, Human-Robot Interaction. In *Proceedings of the 12th Summer Workshop on Multimodal Interfaces (eNTERFACE '16)*, eds. K. P. Truong and D. Reidsma (Enschede, The Netherlands: CTIT), 6–20 ([Davison et al., 2017](#))
- Zaga, C., Charisi, V., **Schadenberg, B. R.**, Reidsma, D., Neerincx, M. A., Prescott, T., Zillich, M., Verschure, P. F. M. J., & Evers, V. (2017). Growing-Up Hand in Hand with Robots. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17* (New York, New York, USA: ACM Press), 429–430. doi:10.1145/3029798.3029804 ([Zaga et al., 2017](#))
- **Schadenberg, B. R.**, Heylen, D. K. J., & Evers, V. (2018). Affect bursts to constrain the meaning of the facial expressions of the humanoid robot Zeno. In *Proceedings of the 1st Workshop on Social Interaction and Multimodal Expression for Socially Intelligent Robots (RO-MAN 2017)* (Lisbon, Portugal: CEUR Workshop Proceedings), 30–39 ([Schadenberg et al., 2018](#))
- **Schadenberg, B. R.** (2019). Predictability in Human-Robot Interactions for Autistic Children. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI) (Daegu, South Korea: IEEE), 748–750. doi:10.1109/HRI.2019.8673135 ([Schadenberg, 2019](#))

Bibliography

- Alami, R., Albu-Schaeffer, A., Bicchi, A., Bischoff, R., Chatila, R., De Luca, A., et al. (2006). Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Beijing, China: IEEE), 1–16. doi:10.1109/IROS.2006.6936985
- Alcorn, A. M., Ainger, E., Charisi, V., Mantinioti, S., Petrović, S., Schadenberg, B. R., et al. (2019). Educators' Views on Using Humanoid Robots With Autistic Learners in Special Education Settings in England. *Frontiers in Robotics and AI* 6, 1–15. doi:10.3389/frobt.2019.00107
- Alcorn, A. M., Pain, H., and Good, J. (2014). Motivating children's initiations with novelty and surprise. In *Proceedings of the 2014 conference on Interaction design and children - IDC '14* (Aarhus, Denmark: ACM Press), 225–228. doi:10.1145/2593968.2610458
- American Psychiatric Association (2000). *Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition, Text Revision (DSM-IV-TR)* (Washington, DC: Author), 4th edn. doi:10.1176/appi.books.9780890423349
- American Psychiatric Association (2013). *Diagnostic and Statistical Manual of Mental Disorders* (Washington, DC: Author), 5th edn. doi:10.1176/appi.books.9780890425596
- Anagnostou, E., Jones, N., Huerta, M., Halladay, A. K., Wang, P., Scahill, L., et al. (2015). Measuring social communication behaviors as a treatment endpoint in individuals with autism spectrum disorder. *Autism* 19, 622–636. doi:10.1177/1362361314542955
- Anzalone, S. M., Boucenna, S., Ivaldi, S., and Chetouani, M. (2015). Evaluating the Engagement with Social Robots. *International Journal of Social Robotics* 7, 465–478. doi:10.1007/s12369-015-0298-7
- Asthana, A., Zafeiriou, S., Tzimiropoulos, G., Cheng, S., and Pantic, M. (2015). From Pixels to Response Maps: Discriminative Image Filtering for Face Alignment in the Wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37, 1312–1320. doi:10.1109/TPAMI.2014.2362142
- Azevedo, R. (2015). Defining and Measuring Engagement and Learning in Science: Conceptual, Theoretical, Methodological, and Analytical Issues. *Educational Psychologist* 50, 84–94. doi:10.1080/00461520.2015.1004069
- Baer, D. M., Wolf, M. M., and Risley, T. R. (1968). Some current dimensions of applied behavior analysis. *Journal of Applied Behavior Analysis* 1, 91–97. doi:10.1901/jaba.1968.1-91
- Bal, V. H., Katz, T., Bishop, S. L., and Krasileva, K. (2016). Understanding definitions of minimally verbal across instruments: evidence for subgroups within minimally verbal children and adolescents with autism spectrum disorder. *Journal of Child Psychology and Psychiatry* 57, 1424–1433. doi:10.1111/jcpp.12609
- Ballard, D. H., Hinton, G. E., and Sejnowski, T. J. (1983). Parallel visual computation. *Nature* 306, 21–26. doi:10.1038/306021a0
- Balsters, J. H., Apps, M. A. J., Bolis, D., Lehner, R., Gallagher, L., and Wenderoth, N. (2017). Disrupted prediction errors index social deficits in autism spectrum disorder. *Brain* 140, 235–246. doi:10.1093/

- brain/aww287
- Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Sciences* 11, 280–289. doi:10.1016/j.tics.2007.05.005
- Barakova, E. I., Gillesen, J. C. C., Huskens, B. E. B. M., and Lourens, T. (2013). End-user programming architecture facilitates the uptake of robots in social therapies. *Robotics and Autonomous Systems* 61, 704–713. doi:10.1016/j.robot.2012.08.001
- Baron-Cohen, S. (2002). The extreme male brain theory of autism. *Trends in Cognitive Sciences* 6, 248–254. doi:10.1016/S1364-6613(02)01904-6
- Baron-Cohen, S. (2009). Autism: The Empathizing-Systemizing (E-S) Theory. *Annals of the New York Academy of Sciences* 1156, 68–80. doi:10.1111/j.1749-6632.2009.04467.x
- Baron-Cohen, S., Leslie, A. M., and Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition* 21, 37–46. doi:10.1016/0010-0277(85)90022-8
- Basili, P., Huber, M., Kourakos, O., Lorenz, T., Brandt, T., Hirche, S., et al. (2012). Inferring the goal of an approaching agent: A human-robot study. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication* (Paris, France: IEEE), 527–532. doi:10.1109/ROMAN.2012.6343805
- Beatty, J. (1982). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin* 91, 276–292. doi:10.1037/0033-2909.91.2.276
- Beatty, J. and Lucero-Wagoner, B. (2000). The pupillary system. In *Handbook of psychophysiology*, eds. J. T. Cacioppo, L. G. Tassinary, and G. G. Berntson (New York, NY, US: Cambridge University Press). 2nd edn., 142–162
- Begum, M., Serna, R. W., Kontak, D., Allspaw, J., Kuczynski, J., Yanco, H. A., et al. (2015). Measuring the Efficacy of Robots in Autism Therapy. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15* (New York, New York, USA: ACM Press), 335–342. doi:10.1145/2696454.2696480
- Begum, M., Serna, R. W., and Yanco, H. A. (2016). Are Robots Ready to Deliver Autism Interventions? A Comprehensive Review. *International Journal of Social Robotics* 8, 157–181. doi:10.1007/s12369-016-0346-y
- Bekele, E. T., Crittendon, J. A., Swanson, A. R., Sarkar, N., and Warren, Z. E. (2014). Pilot clinical application of an adaptive robotic system for young children with autism. *Autism* 18, 598–608. doi:10.1177/1362361313479454
- Beran, T. N., Ramirez-Serrano, A., Kuzyk, R., Fior, M., and Nugent, S. (2011). Understanding how children understand robots: Perceived animism in child–robot interaction. *International Journal of Human Computer Studies* 69, 539–550. doi:10.1016/j.ijhcs.2011.04.003
- Berenbaum, H., Bredemeier, K., and Thompson, R. J. (2008). Intolerance of uncertainty: Exploring its dimensionality and associations with need for cognitive closure, psychopathology, and personality. *Journal of Anxiety Disorders* 22, 117–125. doi:10.1016/j.janxdis.2007.01.004
- Bernard-Opitz, V., Sriram, N., and Nakhoda-Sapuan, S. (2001). Enhancing Social Problem Solving in Children with Autism and Normal Children Through Computer-Assisted Instruction. *Journal of Autism and Developmental Disorders* 31, 377–384. doi:10.1023/A:1010660502130
- Bickmore, T. W., Schulman, D., and Yin, L. (2010). Maintaining engagement in long-term interventions with relational agents. *Applied Artificial Intelligence* 24, 648–666. doi:10.1080/08839514.2010.492259
- Bird, G., Leighton, J., Press, C., and Heyes, C. (2007). Intact automatic imitation of human and robot actions in autism spectrum disorders. *Proceedings of the Royal Society B: Biological Sciences* 274, 3027–3031. doi:10.1098/rspb.2007.1019

- Boccanfuso, L., Barney, E., Foster, C., Ahn, Y. A., Chawarska, K., Scassellati, B., et al. (2016). Emotional robot to examine different play patterns and affective responses of children with and without ASD. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Christchurch, New Zealand), 19–26. doi:10.1109/HRI.2016.7451729
- Boucenna, S., Narzisi, A., Tilmont, E., Muratori, F., Pioggia, G., Cohen, D., et al. (2014). Interactive Technologies for Autistic Children: A Review. *Cognitive Computation* 6, 722–740. doi:10.1007/s12559-014-9276-x
- Boulter, C., Freeston, M. H., South, M., and Rodgers, J. (2014). Intolerance of Uncertainty as a Framework for Understanding Anxiety in Children and Adolescents with Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 44, 1391–1402. doi:10.1007/s10803-013-2001-x
- Bradley, M. M., Miccoli, L., Escrig, M. A., and Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology* 45, 602–607. doi:10.1111/j.1469-8986.2008.00654.x
- Braukmann, R., Ward, E., Hessels, R. S., Bekkering, H., Buitelaar, J. K., and Hunnius, S. (2018). Action prediction in 10-month-old infants at high and low familial risk for Autism Spectrum Disorder. *Research in Autism Spectrum Disorders* 49, 34–46. doi:10.1016/j.rasd.2018.02.004
- Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 77–101. doi:10.1191/1478088706qp063oa
- Brown, C. and Dunn, W. (2002). *Adolescent/adult sensory profile* (San Antonio, Texas, USA: Psychological Corporation)
- Brown, E. C. and Brüne, M. (2012). The role of prediction in social neuroscience. *Frontiers in Human Neuroscience* 6, 1–19. doi:10.3389/fnhum.2012.00147
- Brown, H., Friston, K., and Bestmann, S. (2011). Active Inference, Attention, and Motor Preparation. *Frontiers in Psychology* 2, 1–10. doi:10.3389/fpsyg.2011.00218
- Brugha, T. S., McManus, S., Bankart, J., Scott, F., Purdon, S., Smith, J., et al. (2011). Epidemiology of Autism Spectrum Disorders in Adults in the Community in England. *Archives of General Psychiatry* 68, 459. doi:10.1001/archgenpsychiatry.2011.38
- Bryan, L. C. and Gast, D. L. (2000). Teaching on-task and on-schedule behaviors to high-functioning children with autism via picture activity schedules. *Journal of Autism and Developmental Disorders* 30, 553–567. doi:10.1023/A:1005687310346
- Busch, B., Grizou, J., Lopes, M., and Stulp, F. (2017). Learning Legible Motion from Human–Robot Interactions. *International Journal of Social Robotics* 9, 765–779. doi:10.1007/s12369-017-0400-4
- Cabibihan, J.-J., Javed, H., Ang, M., and Aljunied, S. M. (2013). Why Robots? A Survey on the Roles and Benefits of Social Robots in the Therapy of Children with Autism. *International Journal of Social Robotics* 5, 593–618. doi:10.1007/s12369-013-0202-2
- Carleton, R. N. (2016). Into the unknown: A review and synthesis of contemporary models involving uncertainty. *Journal of Anxiety Disorders* 39, 30–43. doi:10.1016/j.janxdis.2016.02.007
- Carleton, R. N., Norton, M. P. J., and Asmundson, G. J. (2007). Fearing the unknown: A short version of the Intolerance of Uncertainty Scale. *Journal of Anxiety Disorders* 21, 105–117. doi:10.1016/j.janxdis.2006.03.014
- Carpinella, C. M., Wyman, A. B., Perez, M. A., and Stroessner, S. J. (2017). The Robotic Social Attributes Scale (RoSAS). In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17* (New York, New York, USA: ACM Press), 254–262. doi:10.1145/2909824.3020208
- Casby, M. W. (2003). Developmental Assessment of Play. *Communication Disorders Quarterly* 24, 175–183. doi:10.1177/15257401030240040301
- Case-Smith, J. and Arbesman, M. (2008). Evidence-Based Review of Interventions for Autism Used in

- or of Relevance to Occupational Therapy. *American Journal of Occupational Therapy* 62, 416–429. doi:10.5014/ajot.62.4.416
- Castro-González, A., Admoni, H., and Scassellati, B. (2016). Effects of form and motion on judgments of social robots' animacy, likability, trustworthiness and unpleasantness. *International Journal of Human-Computer Studies* 90, 27–38. doi:10.1016/j.ijhcs.2016.02.004
- Chai, J. Y., She, L., Fang, R., Ottarson, S., Little, C., Liu, C., et al. (2014). Collaborative effort towards common ground in situated human-robot dialogue. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14* (Bielefeld, Germany: ACM), 33–40. doi:10.1145/2559636.2559677
- Chamberlain, P. D., Rodgers, J., Crowley, M. J., White, S. E., Freeston, M. H., and South, M. (2013). A potentiated startle study of uncertainty and contextual anxiety in adolescents diagnosed with autism spectrum disorder. *Molecular Autism* 4, 31. doi:10.1186/2040-2392-4-31
- Chaminade, T., Rosset, D., Da Fonseca, D., Hodgins, J. K., and Deruelle, C. (2015). Anthropomorphic bias found in typically developing children is not found in children with autistic spectrum disorder. *Autism* 19, 248–251. doi:10.1177/1362361313512425
- Chevalier, P., Li, J. J., Ainger, E., Alcorn, A. M., Babovic, S., Charisi, V., et al. (2017a). Dialogue Design for a Robot-Based Face-Mirroring Game to Engage Autistic Children with Emotional Expressions. In *Social Robotics: 9th International Conference, ICSR 2017, Tsukuba, Japan, November 22-24, 2017, Proceedings*, eds. A. Kheddar, E. Yoshida, S. S. Ge, K. Suzuki, J.-J. Cabibihan, F. Eyssel, and H. He (Cham: Springer International Publishing), vol. 10652 of *Lecture Notes in Computer Science*. 546–555. doi:10.1007/978-3-319-70022-9_54
- Chevalier, P., Martin, J.-C., Isableu, B., Bazile, C., and Tapus, A. (2017b). Impact of sensory preferences of individuals with autism on the recognition of emotions expressed by two robots, an avatar, and a human. *Autonomous Robots* 41, 613–635. doi:10.1007/s10514-016-9575-z
- Chevalier, P., Raiola, G., Martin, J.-c., Isableu, B., Bazile, C., and Tapus, A. (2017c). Do Sensory Preferences of Children with Autism Impact an Imitation Task with a Robot? In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17* (New York, New York, USA: ACM Press), 177–186. doi:10.1145/2909824.3020234
- Christensen, L., Hutman, T., Rozga, A., Young, G. S., Ozonoff, S., Rogers, S. J., et al. (2010). Play and Developmental Outcomes in Infant Siblings of Children with Autism. *Journal of Autism and Developmental Disorders* 40, 946–957. doi:10.1007/s10803-010-0941-y
- Clabaugh, C., Mahajan, K., Jain, S., Pakkar, R., Becerra, D., Shi, Z., et al. (2019). Long-Term Personalization of an In-Home Socially Assistive Robot for Children With Autism Spectrum Disorders. *Frontiers in Robotics and AI* 6, 1–18. doi:10.3389/frobt.2019.00110
- Clabaugh, C. and Matarić, M. J. (2019). Escaping Oz: Autonomy in Socially Assistive Robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 2, 33–61. doi:10.1146/annurev-control-060117-104911
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36, 181–204. doi:10.1017/S0140525X12000477
- Clark, H. H. (1996). *Using language* (Cambridge, UK: Cambridge University Press). doi:10.1017/CBO9780511620539
- Coeckelbergh, M., Pop, C. A., Simut, R. E., Peca, A., Pintea, S., David, D. O., et al. (2016). A Survey of Expectations About the Role of Robots in Robot-Assisted Therapy for Children with ASD: Ethical Acceptability, Trust, Sociability, Appearance, and Attachment. *Science and Engineering Ethics* 22, 47–65. doi:10.1007/s11948-015-9649-x
- Collignon, O., Charbonneau, G., Peters, F., Nassim, M., Lassonde, M., Lepore, F., et al. (2013). Reduced multisensory facilitation in persons with autism. *Cortex* 49, 1704–1710. doi:10.1016/j.cortex.2012.06.001

- Colton, M. B., Ricks, D. J., Goodrich, M. A., Dariush, B., Fujimura, K., and Fujiki, M. (2009). Toward therapist-in-the-loop assistive robotics for children with autism and specific language impairment. In *AISB New Frontiers in Human-Robot Interaction Symposium* (Edinburgh, Scotland), 1–5
- Comer, J. S., Roy, A. K., Furr, J. M., Gotimer, K., Beidas, R. S., Dugas, M. J., et al. (2009). The Intolerance of Uncertainty Scale for Children: A psychometric evaluation. *Psychological Assessment* 21, 402–411. doi:10.1037/a0016719
- Coninx, A., Baxter, P. E., Oleari, E., Bellini, S., Bierman, B. P., Blanson Henkemans, O. A., et al. (2015). Towards Long-Term Social Child-Robot Interaction: Using Multi-Activity Switching to Engage Young Users. *Journal of Human-Robot Interaction* 5, 32–64. doi:10.5898/JHRI.5.1.Coninx
- Connell, J. P. and Wellborn, J. G. (1991). Competence, autonomy, and relatedness: A motivational analysis of self-system processes. In *Self processes and development* (Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc). 43–77
- Corlett, P., Taylor, J., Wang, X.-J., Fletcher, P., and Krystal, J. (2010). Toward a neurobiology of delusions. *Progress in Neurobiology* 92, 345–369. doi:10.1016/j.pneurobio.2010.06.007
- Corrigan, L. J., Peters, C., Küster, D., and Castellano, G. (2016). Engagement Perception and Generation for Social Robots and Virtual Agents. In *Toward Robotic Socially Believable Behaving Systems - Volume I*, eds. A. Esposito and L. C. Jain (Cham: Springer International Publishing), vol. 105 of *Intelligent Systems Reference Library*. 29–51. doi:10.1007/978-3-319-31056-5_4
- Corsello, C. M. (2005). Early Intervention in Autism. *Infants & Young Children* 18, 74–85. doi:10.1097/00001163-200504000-00002
- Costa, S. C., Lehmann, H., Dautenhahn, K., Robins, B., and Soares, F. O. (2015). Using a Humanoid Robot to Elicit Body Awareness and Appropriate Physical Interaction in Children with Autism. *International Journal of Social Robotics* 7, 265–278. doi:10.1007/s12369-014-0250-2
- Costescu, C. A., Vanderborght, B., and David, D. O. (2015). Reversal Learning Task in Children with Autism Spectrum Disorder: A Robot-Based Approach. *Journal of Autism and Developmental Disorders* 45, 3715–3725. doi:10.1007/s10803-014-2319-z
- Dautenhahn, K. (1999). Robots as social actors: Aurora and the case of autism. In *Proceedings Third Cognitive Technology Conference CT'99* (San Francisco, CA, USA: M.I.N.D. Lab), vol. 359, 359–374
- Dautenhahn, K. (2007). Socially intelligent robots: dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362, 679–704. doi:10.1098/rstb.2006.2004
- Dautenhahn, K., Nehaniv, C. L., Walters, M. L., Robins, B., Kose-Bagci, H., Mirza, N. A., et al. (2009). KASPAR – a minimally expressive humanoid robot for human–robot interaction research. *Applied Bionics and Biomechanics* 6, 369–397. doi:10.1080/11762320903123567
- Dautenhahn, K. and Werry, I. (2004). Towards interactive robots in autism therapy: Background, motivation and challenges. *Pragmatics & Cognition* 12, 1–35. doi:10.1075/pc.12.1.03dau
- David, D. O., Costescu, C. A., Matu, S.-A., Szentagotai, A., and Dobrean, A. (2018). Developing Joint Attention for Children with Autism in Robot-Enhanced Therapy. *International Journal of Social Robotics* 10, 595–605. doi:10.1007/s12369-017-0457-0
- David, D. O., Costescu, C. A., Matu, S.-A., Szentagotai, A., and Dobrean, A. (2020). Effects of a Robot-Enhanced Intervention for Children With ASD on Teaching Turn-Taking Skills. *Journal of Educational Computing Research* 58, 29–62. doi:10.1177/0735633119830344
- Davison, D., Gorer, B., Kolkmeier, J., Linssen, J., Schadenberg, B. R., van de Vijver, B., et al. (2017). Things that Make Robots Go HMMM: Heterogeneous Multilevel Multimodal Mixing to Realise Fluent, Multiparty, Human-Robot Interaction. In *Proceedings of the 12th Summer Workshop on Multimodal Interfaces (eINTERFACE '16)*, eds. K. P. Truong and D. Reidsma (Enschede, The Netherlands: CTIT), 6–20

- Dawson, G., Webb, S. J., and McPartland, J. (2005). Understanding the Nature of Face Processing Impairment in Autism: Insights From Behavioral and Electrophysiological Studies. *Developmental Neuropsychology* 27, 403–424. doi:10.1207/s15326942dn2703_6
- Dayan, P., Hinton, G. E., Neal, R. M., and Zemel, R. S. (1995). The Helmholtz Machine. *Neural Computation* 7, 889–904. doi:10.1162/neco.1995.7.5.889
- de Graaf, M. M. and Malle, B. F. (2019). People’s Explanations of Robot Behavior Subtly Reveal Mental State Inferences. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Daegu, South Korea: IEEE), 239–248. doi:10.1109/HRI.2019.8673308
- Deci, E. L. and Ryan, R. M. (1985). *Intrinsic Motivation and Self-Determination in Human Behavior* (Boston, MA: Springer US). doi:10.1007/978-1-4899-2271-7
- Denham, S. A., Blair, K. A., DeMulder, E., Levitas, J., Sawyer, K., Auerbach-Major, S., et al. (2003). Preschool Emotional Competence: Pathway to Social Competence? *Child Development* 74, 238–256. doi:10.1111/1467-8624.00533
- Deriso, D., Susskind, J., Krieger, L., and Bartlett, M. (2012). Emotion Mirror: A Novel Intervention for Autism Based on Real-Time Expression Recognition. In *Computer Vision – ECCV 2012. Workshops and Demonstrations* (Florence, Italy: Springer Berlin Heidelberg), 671–674. doi:10.1007/978-3-642-33885-4_{_}79
- Desideri, L., Negrini, M., Malavasi, M., Tanzini, D., Rouame, A., Cutrone, M. C., et al. (2018). Using a Humanoid Robot as a Complement to Interventions for Children with Autism Spectrum Disorder: a Pilot Study. *Advances in Neurodevelopmental Disorders* 2, 273–285. doi:10.1007/s41252-018-0066-4
- Dickerson, P., Robins, B., and Dautenhahn, K. (2013). Where the action is: A conversation analytic perspective on interaction between a humanoid robot, a co-present adult and a child with an ASD. *Interaction Studies* 14, 297–316. doi:10.1075/is.14.2.07dic
- Diehl, J. J., Crowell, C. R., Villano, M., Wier, K., Tang, K., and Riek, L. D. (2014). Clinical Applications of Robots in Autism Spectrum Disorder Diagnosis and Treatment. In *Comprehensive Guide to Autism*, eds. V. B. Patel, V. R. Preedy, and C. R. Martin (New York, NY: Springer New York). 411–422. doi:10.1007/978-1-4614-4788-7_14
- Diehl, J. J., Schmitt, L. M., Villano, M., and Crowell, C. R. (2012). The clinical use of robots for individuals with Autism Spectrum Disorders: A critical review. *Research in Autism Spectrum Disorders* 6, 249–262. doi:10.1016/j.rasd.2011.05.006
- Diep, L., Cabibihan, J.-J., and Wolbring, G. (2015). Social Robots: Views of special education teachers. In *Proceedings of the 3rd 2015 Workshop on ICTs for improving Patients Rehabilitation Research Techniques - REHAB '15* (New York, New York, USA: ACM Press), 160–163. doi:10.1145/2838944.2838983
- Dragan, A. D., Bauman, S., Forlizzi, J., and Srinivasa, S. S. (2015). Effects of Robot Motion on Human-Robot Collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15* (New York, New York, USA: ACM Press), vol. 1, 51–58. doi:10.1145/2696454.2696473
- Dragan, A. D., Lee, K. C., and Srinivasa, S. S. (2013). Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Tokyo, Japan: IEEE), 301–308. doi:10.1109/HRI.2013.6483603
- Dragan, A. D. and Srinivasa, S. S. (2013). Generating Legible Motion. In *Robotics: Science and Systems IX* (Berlin, Germany: Robotics: Science and Systems Foundation), 92–103. doi:10.15607/RSS.2013.IX.024
- Dragan, A. D. and Srinivasa, S. S. (2014a). Familiarization to robot motion. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14* (New York, New York, USA: ACM Press), 366–373. doi:10.1145/2559636.2559674
- Dragan, A. D. and Srinivasa, S. S. (2014b). Integrating human observer inferences into robot motion

- planning. *Autonomous Robots* 37, 351–368. doi:10.1007/s10514-014-9408-x
- Driggs-Campbell, K. and Bajcsy, R. (2016). Communicating intent on the road through human-inspired control schemes. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Daejeon, South Korea: IEEE), 3042–3047. doi:10.1109/IROS.2016.7759471
- Duquette, A., Michaud, F., and Mercier, H. (2008). Exploring the use of a mobile robot as an imitation agent with children with low-functioning autism. *Autonomous Robots* 24, 147–157. doi:10.1007/s10514-007-9056-5
- Eaves, L. C. and Ho, H. H. (1997). School Placement and Academic Achievement in Children with Autistic Spectrum Disorders. *Journal of Developmental and Physical Disabilities* 9, 277–291. doi:10.1023/A:1024944226971
- Einstein, D. A. (2014). Extension of the Transdiagnostic Model to Focus on Intolerance of Uncertainty: A Review of the Literature and Implications for Treatment. *Clinical Psychology: Science and Practice* 21, 280–300. doi:10.1111/cpsp.12077
- Elsabbagh, M., Divan, G., Koh, Y.-J., Kim, Y. S., Kauchali, S., Marcín, C., et al. (2012). Global Prevalence of Autism and Other Pervasive Developmental Disorders. *Autism Research* 5, 160–179. doi:10.1002/aur.239
- Emanuel, R. and Weir, S. (1976). Catalysing communication in an autistic child in a LOGO-like learning environment. In *Proceedings of the 2nd Summer Conference on Artificial Intelligence and Simulation of Behaviour* (Edinburgh, UK: IOS Press), 118–129. doi:10.5555/3015508.3015520
- Escalona, A., Field, T., Nadel, J., and Lundy, B. (2002). Brief Report: Imitation Effects on Children with Autism. *Journal of Autism and Developmental Disorders* 32, 141–144. doi:10.1023/A:1014896707002
- Ewing, L., Pellicano, E., and Rhodes, G. (2013). Atypical updating of face representations with experience in children with autism. *Developmental Science* 16, 116–123. doi:10.1111/desc.12007
- Eyssel, F. and Kuchenbrandt, D. (2011). Manipulating anthropomorphic inferences about NAO: The role of situational and dispositional aspects of effectance motivation. In *2011 RO-MAN* (Atlanta, Georgia: IEEE), 467–472. doi:10.1109/ROMAN.2011.6005233
- Eyssel, F., Kuchenbrandt, D., and Bobinger, S. (2011). Effects of anticipated human-robot interaction and predictability of robot behavior on perceptions of anthropomorphism. In *Proceedings of the 6th International Conference on Human-Robot Interaction - HRI '11* (Lausanne, Switzerland: ACM Press), 61–67. doi:10.1145/1957656.1957673
- Farr, W., Yuill, N., and Raffle, H. (2010). Social benefits of a tangible user interface for children with Autistic Spectrum Conditions. *Autism* 14, 237–252. doi:10.1177/1362361310363280
- Feil-Seifer, D. J. and Matarić, M. J. (2009). Toward Socially Assistive Robotics for Augmenting Interventions for Children with Autism Spectrum Disorders. In *Experimental Robotics*, eds. O. Khatib, V. Kumar, and G. J. Pappas (Berlin Heidelberg: Springer), vol. 54. 201–210. doi:10.1007/978-3-642-00196-3_24
- Feil-Seifer, D. J. and Matarić, M. J. (2011a). Automated detection and classification of positive vs. negative robot interactions with children with autism using distance-based features. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11* (Lausanne, Switzerland: ACM Press), 323. doi:10.1145/1957656.1957785
- Feil-Seifer, D. J. and Matarić, M. J. (2011b). Socially Assistive Robotics. *IEEE Robotics & Automation Magazine* 18, 24–31. doi:10.1109/MRA.2010.940150
- Feinman, S. (1982). Social Referencing in Infancy. *Merrill-Palmer Quarterly* 28, 445–470
- Feldman, H. and Friston, K. J. (2010). Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience* 4, 1–23. doi:10.3389/fnhum.2010.00215
- Felzmann, H., Fosch-Villaronga, E., Lutz, C., and Tamo-Larrieux, A. (2019). Robots and Transparency:

- The Multiple Dimensions of Transparency in the Context of Robot Technologies. *IEEE Robotics & Automation Magazine* 26, 71–78. doi:10.1109/MRA.2019.2904644
- Ferrara, C. and Hill, S. D. (1980). The responsiveness of autistic children to the predictability of social and nonsocial toys. *Journal of Autism and Developmental Disorders* 10, 51–57. doi:10.1007/BF02408432
- Field, T., Field, T., Sanders, C., and Nadel, J. (2001). Children with Autism Display more Social Behaviors after Repeated Imitation Sessions. *Autism* 5, 317–323. doi:10.1177/1362361301005003008
- Fiske, S. T., Cuddy, A. J., and Glick, P. (2007). Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences* 11, 77–83. doi:10.1016/j.tics.2006.11.005
- Fletcher, P. C. and Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews Neuroscience* 10, 48–58. doi:10.1038/nrn2536
- Frauenberger, C., Good, J., and Keay-Bright, W. (2011). Designing technology for children with special needs: bridging perspectives through participatory design. *CoDesign* 7, 1–28. doi:10.1080/15710882.2011.587013
- Frauenberger, C., Makhaeva, J., and Spiel, K. (2017). Blending Methods: Developing Participatory Design Sessions for Autistic Children. In *Proceedings of the 2017 Conference on Interaction Design and Children* (New York, NY, USA: ACM), 39–49. doi:10.1145/3078072.3079727
- Fredricks, J. A., Blumenfeld, P. C., and Paris, A. H. (2004). School Engagement: Potential of the Concept, State of the Evidence. *Review of Educational Research* 74, 59–109. doi:10.3102/00346543074001059
- Freedy, A., DeVisser, E., Weltman, G., and Coeyman, N. (2007). Measurement of trust in human-robot collaboration. In *2007 International Symposium on Collaborative Technologies and Systems* (Orlando, Florida, USA: IEEE), 106–114. doi:10.1109/CTS.2007.4621745
- Fresh Consulting (2020). *Why Robotics Companies Fail*. Tech. rep., Fresh Consulting
- Friston, K. (2003). Learning and inference in the brain. *Neural Networks* 16, 1325–1352. doi:10.1016/j.neunet.2003.06.005
- Friston, K. J. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences* 360, 815–836. doi:10.1098/rstb.2005.1622
- Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* 11, 127–138. doi:10.1038/nrn2787
- Frith, U. and Happé, F. (1994). Autism: beyond “theory of mind”. *Cognition* 50, 115–132. doi:10.1016/0010-0277(94)90024-8
- Frost-Karlsson, M., Galazka, M. A., Gillberg, C., Gillberg, C., Miniscalco, C., Billstedt, E., et al. (2019). Social scene perception in autism spectrum disorder: An eye-tracking and pupillometric study. *Journal of Clinical and Experimental Neuropsychology* 41, 1024–1032. doi:10.1080/13803395.2019.1646214
- Gauthier, I., Klaiman, C., and Schultz, R. T. (2009). Face composite effects reveal abnormal face processing in Autism spectrum disorders. *Vision Research* 49, 470–478. doi:10.1016/j.visres.2008.12.007
- Gernsbacher, M. A. (2017). Editorial Perspective: The use of person-first language in scholarly writing may accentuate stigma. *Journal of Child Psychology and Psychiatry* 58, 859–861. doi:10.1111/jcpp.12706
- Gibson, E. J. and Pick, A. D. (2000). *Perceptual Learning and Development: An Ecological Approach* (New York, NY: Oxford University Press)
- Gielniak, M. J. and Thomaz, A. L. (2011). Generating anticipation in robot motion. In *2011 RO-MAN* (Atlanta, Georgia: IEEE), 449–454. doi:10.1109/ROMAN.2011.6005255

- Giullian, N., Ricks, D. J., Atherton, J. A., Colton, M. B., Goodrich, M. A., and Brinton, B. (2010). Detailed requirements for robots in autism therapy. In *2010 IEEE International Conference on Systems, Man and Cybernetics* (IEEE), 2595–2602. doi:10.1109/ICSMC.2010.5641908
- Gogate, L. J. and Hollich, G. (2010). Invariance detection within an interactive system: A perceptual gateway to language development. *Psychological Review* 117, 496–516. doi:10.1037/a0019049
- Golan, O., Ashwin, E., Granader, Y., McClintock, S., Day, K., Leggett, V., et al. (2010). Enhancing Emotion Recognition in Children with Autism Spectrum Conditions: An Intervention Using Animated Vehicles with Real Emotional Faces. *Journal of Autism and Developmental Disorders* 40, 269–279. doi:10.1007/s10803-009-0862-9
- Goldberg, M. C., Allman, M. J., Hagopian, L. P., Triggs, M. M., Frank-Crawford, M. A., Mostofsky, S. H., et al. (2017). Examining the reinforcing value of stimuli within social and non-social contexts in children with and without high-functioning autism. *Autism* 21, 881–895. doi:10.1177/1362361316655035
- Goodrich, M. A., Colton, M. B., Brinton, B., Fujiki, M., Atherton, J. A., Robinson, L., et al. (2012). Incorporating a robot into an autism therapy team. *IEEE Intelligent Systems* 27, 52–59. doi:10.1109/MIS.2012.40
- Goris, J., Braem, S., Nijhof, A. D., Rigoni, D., Deschrijver, E., Van de Cruys, S., et al. (2018). Sensory Prediction Errors Are Less Modulated by Global Context in Autism Spectrum Disorder. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* 3, 667–674. doi:10.1016/j.bpsc.2018.02.003
- Goris, J., Brass, M., Cambier, C., Delplanque, J., Wiersema, J. R., and Braem, S. (2020). The Relation Between Preference for Predictability and Autistic Traits. *Autism Research* 13, 1144–1154. doi:10.1002/aur.2244
- Goris, J., Silveti, M., Verguts, T., Wiersema, J. R., Brass, M., and Braem, S. (2021). Autistic traits are related to worse performance in a volatile reward learning task despite adaptive learning rates. *Autism* 25, 440–451. doi:10.1177/1362361320962237
- Gotham, K., Pickles, A., and Lord, C. (2009). Standardizing ADOS Scores for a Measure of Severity in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 39, 693–705. doi:10.1007/s10803-008-0674-3
- Gotham, K., Risi, S., Pickles, A., and Lord, C. (2007). The Autism Diagnostic Observation Schedule: Revised Algorithms for Improved Diagnostic Validity. *Journal of Autism and Developmental Disorders* 37, 613–627. doi:10.1007/s10803-006-0280-1
- Greenwood, C. R. (1991). Longitudinal Analysis of Time, Engagement, and Achievement in At-Risk versus Non-Risk Students. *Exceptional Children* 57, 521–535. doi:10.1177/001440299105700606
- Gregory, R. L. (1968). Perceptual illusions and brain models. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 171, 279–296. doi:10.1098/rspb.1968.0071
- Grinter, E. J., Maybery, M. T., Van Beek, P. L., Pellicano, E., Badcock, J. C., and Badcock, D. R. (2009). Global Visual Processing and Self-Rated Autistic-like Traits. *Journal of Autism and Developmental Disorders* 39, 1278–1290. doi:10.1007/s10803-009-0740-5
- Grove, R., Hoekstra, R. A., Wierda, M., and Begeer, S. (2018). Special interests and subjective wellbeing in autistic adults. *Autism Research* 11, 766–775. doi:10.1002/aur.1931
- Grupe, D. W. and Nitschke, J. B. (2013). Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience* 14, 488–501. doi:10.1038/nrn3524
- Grynszpan, O., Weiss, P. L., Perez-Diaz, F., and Gal, E. (2014). Innovative technology-based interventions for autism spectrum disorders: A meta-analysis. *Autism* 18, 346–361. doi:10.1177/1362361313476767

- Grzadzinski, R., Carr, T., Colombi, C., McGuire, K., Dufek, S., Pickles, A., et al. (2016). Measuring Changes in Social Communication Behaviors: Preliminary Development of the Brief Observation of Social Communication Change (BOSCC). *Journal of Autism and Developmental Disorders* 46, 2464–2479. doi:10.1007/s10803-016-2782-9
- Grzadzinski, R., Huerta, M., and Lord, C. (2013). DSM-5 and autism spectrum disorders (ASDs): an opportunity for identifying ASD subtypes. *Molecular Autism* 4, 12. doi:10.1186/2040-2392-4-12
- Hagerer, G., Cummins, N., Eyben, F., and Schuller, B. (2017). "Did You Laugh Enough Today?" -Deep Neural Networks for Mobile Andwearable Laughter Trackers. In *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*. 2044–2045. doi:10.21437/Interspeech.2017-2036
- Halberstadt, A. G., Denham, S. A., and Dunsmore, J. C. (2001). Affective Social Competence. *Social Development* 10, 79–119. doi:10.1111/1467-9507.00150
- Hallett, V., Mueller, J., Breese, L., Hollett, M., Beresford, B., Irvine, A., et al. (2020). Introducing 'Predictive Parenting': A Feasibility Study of a New Group Parenting Intervention Targeting Emotional and Behavioral Difficulties in Children with Autism Spectrum Disorder. *Journal of Autism and Developmental Disorders* doi:10.1007/s10803-020-04442-2
- Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., and Parasuraman, R. (2011). A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 53, 517–527. doi:10.1177/0018720811417254
- Happé, F. and Frith, U. (2006). The Weak Coherence Account: Detail-focused Cognitive Style in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 36, 5–25. doi:10.1007/s10803-005-0039-0
- Happé, F., Ronald, A., and Plomin, R. (2006). Time to give up on a single explanation for autism. *Nature Neuroscience* 9, 1218–1220. doi:10.1038/nn1770
- Hayes, A. F. (2017). *Introduction to Mediation, Moderation, and Conditional Process Analysis* (New York, New York, USA: The Guilford Press), 2nd edn.
- Hayes, G. R., Gardere, L. M., Abowd, G. D., and Truong, K. N. (2008). CareLog: A Selective Archiving Tool for Behavior Management in Schools. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)* (Florence, Italy: ACM Press), 685–694. doi:10.1145/1357054.1357164
- Hazen, E. P., Stornelli, J. L., O'Rourke, J. A., Koesterer, K., and McDougale, C. J. (2014). Sensory Symptoms in Autism Spectrum Disorders. *Harvard Review of Psychiatry* 22, 112–124. doi:10.1097/01.HRP.0000445143.08773.58
- Heaton, P. (2003). Pitch memory, labelling and disembedding in autism. *Journal of Child Psychology and Psychiatry* 44, 543–551. doi:10.1111/1469-7610.00143
- Heinzmann, J. and Zelinsky, A. (2003). Quantitative Safety Guarantees for Physical Human-Robot Interaction. *The International Journal of Robotics Research* 22, 479–504. doi:10.1177/02783649030227004
- Helmholz, H. v. (1860). *Handbuch der physiologischen optik* (New York: Dover: ed. & trans. J. P. C. Southall), vol. 3 edn.
- Hess, E. H. and Polt, J. M. (1964). Pupil Size in Relation to Mental Activity during Simple Problem-Solving. *Science* 143, 1190–1192. doi:10.1126/science.143.3611.1190
- Hill, E. L. (2004). Executive dysfunction in autism. *Trends in Cognitive Sciences* 8, 26–32. doi:10.1016/j.tics.2003.11.003
- Hohwy, J. (2013). *The Predictive Mind* (Oxford, UK: Oxford University Press). doi:10.1093/acprof:oso/9780199682737.001.0001

- Hong, R. Y. and Cheung, M. W. (2015). The Structure of Cognitive Vulnerabilities to Depression and Anxiety. *Clinical Psychological Science* 3, 892–912. doi:10.1177/2167702614553789
- Hourcade, J. P., Bullock-Rest, N. E., and Hansen, T. E. (2012). Multitouch tablet applications and activities to enhance the social skills of children with autism spectrum disorders. *Personal and Ubiquitous Computing* 16, 157–168. doi:10.1007/s00779-011-0383-3
- Howlin, P., Baron-Cohen, S., and Hadwin, J. (1999). *Teaching children with autism to mind-read: A practical guide for teachers and parents* (Chichester, UK: J. Wiley & Sons)
- Howlin, P., Goode, S., Hutton, J., and Rutter, M. (2004). Adult outcome for children with autism. *Journal of Child Psychology and Psychiatry* 45, 212–229. doi:10.1111/j.1469-7610.2004.00215.x
- Huijnen, C. A. G. J., Lexis, M. A. S., and de Witte, L. P. (2016). Matching Robot KASPAR to Autism Spectrum Disorder (ASD) Therapy and Educational Goals. *International Journal of Social Robotics* 8, 445–455. doi:10.1007/s12369-016-0369-4
- Huijnen, C. A. G. J., Lexis, M. A. S., Jansens, R., and de Witte, L. P. (2017). How to Implement Robots in Interventions for Children with Autism? A Co-creation Study Involving People with Autism, Parents and Professionals. *Journal of Autism and Developmental Disorders* 47, 3079–3096. doi:10.1007/s10803-017-3235-9
- Huijnen, C. A. G. J., Lexis, M. A. S., Jansens, R., and de Witte, L. P. (2019). Roles, Strengths and Challenges of Using Robots in Interventions for Children with Autism Spectrum Disorder (ASD). *Journal of Autism and Developmental Disorders* 49, 11–21. doi:10.1007/s10803-018-3683-x
- Huskens, B. E. B. M., Palmen, A., Van der Werff, M., Lourens, T., and Barakova, E. I. (2015). Improving Collaborative Play Between Children with Autism Spectrum Disorders and Their Siblings: The Effectiveness of a Robot-Mediated Intervention Based on Lego® Therapy. *Journal of Autism and Developmental Disorders* 45, 3746–3755. doi:10.1007/s10803-014-2326-0
- Huskens, B. E. B. M., Verschuur, R., Gillesen, J. C. C., Didden, R., and Barakova, E. I. (2013). Promoting question-asking in school-aged children with autism spectrum disorders: Effectiveness of a robot intervention compared to a human-trainer intervention. *Developmental Neurorehabilitation* 16, 345–356. doi:10.3109/17518423.2012.739212
- Huys, Q. J., Daw, N. D., and Dayan, P. (2015). Depression: A Decision-Theoretic Analysis. *Annual Review of Neuroscience* 38, 1–23. doi:10.1146/annurev-neuro-071714-033928
- Jain, S., Thiagarajan, B., Shi, Z., Clabaugh, C., and Matarić, M. J. (2020). Modeling engagement in long-term, in-home socially assistive robot interventions for children with autism spectrum disorders. *Science Robotics* 5, eaaz3791. doi:10.1126/scirobotics.aaz3791
- Jarrold, C. (2003). A Review of Research into Pretend Play in Autism. *Autism* 7, 379–390. doi:10.1177/1362361303007004004
- Jarrold, C., Boucher, J., and Smith, P. (1993). Symbolic play in autism: A review. *Journal of Autism and Developmental Disorders* 23, 281–307. doi:10.1007/BF01046221
- Javed, H., Burns, R., Jeon, M., Howard, A. M., and Park, C. H. (2019). A Robotic Framework to Facilitate Sensory Experiences for Children with Autism Spectrum Disorder. *ACM Transactions on Human-Robot Interaction* 9, 1–26. doi:10.1145/3359613
- Javed, H., Lee, W., and Park, C. H. (2020). Toward an Automated Measure of Social Engagement for Children With Autism Spectrum Disorder—A Personalized Computational Modeling Approach. *Frontiers in Robotics and AI* 7, 14. doi:10.3389/frobt.2020.00043
- Johnson, S. C. (2003). Detecting agents. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 358, 549–559. doi:10.1098/rstb.2002.1237
- Jones, A., Bull, S., and Castellano, G. (2018). “I Know That Now, I’m Going to Learn This Next” Promoting Self-regulated Learning with a Robotic Tutor. *International Journal of Social Robotics* 10, 439–454.

- doi:10.1007/s12369-017-0430-y
- Jones, A. and Castellano, G. (2018). Adaptive Robotic Tutors that Support Self-Regulated Learning: A Longer-Term Investigation with Primary School Children. *International Journal of Social Robotics* 10, 357–370. doi:10.1007/s12369-017-0458-z
- Kahneman, D. and Beatty, J. (1966). Pupil Diameter and Load on Memory. *Science* 154, 1583–1585. doi:10.1126/science.154.3756.1583
- Kanda, T., Hirano, T., Eaton, D., and Ishiguro, H. (2004). Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interaction* 19, 61–84. doi:10.1080/07370024.2004.9667340
- Kang, E., Klein, E. F., Lillard, A. S., and Lerner, M. D. (2016). Predictors and Moderators of Spontaneous Pretend Play in Children with and without Autism Spectrum Disorder. *Frontiers in Psychology* 7, 1–14. doi:10.3389/fpsyg.2016.01577
- Kanner, L. (1943). Autistic disturbances of affective contact. *Nervous Child* 2, 217–250
- Keen, D. (2009). Engagement of Children With Autism in Learning. *Australasian Journal of Special Education* 33, 130–140. doi:10.1375/ajse.33.2.130
- Keen, D., Woodyatt, G., and Sigafoos, J. (2002). Verifying Teacher Perceptions of the Potential Communicative Acts of Children with Autism. *Communication Disorders Quarterly* 23, 131–140. doi:10.1177/15257401020230030201
- Kennedy, J., Lemaignan, S., Montassier, C., Lavalade, P., Irfan, B., Papadopoulos, F., et al. (2017). Child Speech Recognition in Human-Robot Interaction: Evaluations and Recommendations. *ACM/IEEE International Conference on Human-Robot Interaction*, 82–90. doi:10.1145/2909824.3020229
- Kenny, L., Hattersley, C., Molins, B., Buckley, C., Povey, C., and Pellicano, E. (2016). Which terms should be used to describe autism? Perspectives from the UK autism community. *Autism* 20, 442–462. doi:10.1177/1362361315588200
- Kidd, C. and Breazeal, C. (2008). Robots at home: Understanding long-term human-robot interaction. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Nice, France: IEEE), 3230–3235. doi:10.1109/IROS.2008.4651113
- Kientz, J. A., Boring, S., Abowd, G. D., and Hayes, G. R. (2005). Abaris: Evaluating Automated Capture Applied to Structured Autism Interventions. In *UbiComp 2005: Ubiquitous Computing*, eds. M. Beigl, S. Intille, J. Rekimoto, and H. Tokuda (Tokyo, Japan: Springer), vol. 3660, 323–339. doi:10.1007/11551201_19
- Kientz, J. A., Hayes, G. R., Westeyn, T., Starner, T., and Abowd, G. D. (2007). Pervasive Computing and Autism: Assisting Caregivers of Children with Special Needs. *IEEE Pervasive Computing* 6, 28–35. doi:10.1109/MPRV.2007.18
- Kilner, J. M., Friston, K. J., and Frith, C. D. (2007). Predictive coding: an account of the mirror neuron system. *Cognitive Processing* 8, 159–166. doi:10.1007/s10339-007-0170-2
- Kim, E. S., Berkovits, L. D., Bernier, E. P., Leyzberg, D., Shic, F., Paul, R., et al. (2013). Social Robots as Embedded Reinforcers of Social Behavior in Children with Autism. *Journal of Autism and Developmental Disorders* 43, 1038–1049. doi:10.1007/s10803-012-1645-2
- Kim, E. S., Daniell, C. M., Makar, C., Elia, J., Scassellati, B., and Shic, F. (2015). Potential clinical impact of positive affect in robot interactions for autism intervention. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)* (Xi'an, China: IEEE), 8–13. doi:10.1109/ACII.2015.7344544
- Kim, E. S., Paul, R., Shic, F., and Scassellati, B. (2012). Bridging the Research Gap: Making HRI Useful to Individuals with Autism. *Journal of Human-Robot Interaction* 1, 26–54. doi:10.5898/JHRI.1.1.Kim
- Kim, J.-M. and Mahoney, G. (2004). The Effects of Mother's Style of Interaction on Children's Engage-

- ment. *Topics in Early Childhood Special Education* 24, 31–38. doi:10.1177/02711214040240010301
- Kim, M.-G., Park, M., Choi, Y. S., Kim, K.-H., Lee, N., and Sohn, D.-S. (2019). Understanding Work Environment of Therapists to Incorporate Robots into Education for Children with Developmental Disorders. In *2019 16th International Conference on Ubiquitous Robots (UR)* (IEEE), 603–608. doi:10.1109/URAI.2019.8768629
- King, G., Currie, M., and Petersen, P. (2014). Child and parent engagement in the mental health intervention process: a motivational framework. *Child and Adolescent Mental Health* 19, 2–8. doi:10.1111/camh.12015
- Kishida, Y. and Kemp, C. (2006). A measure of engagement for children with intellectual disabilities in early childhood settings: A preliminary study. *Journal of Intellectual & Developmental Disability* 31, 101–114. doi:10.1080/13668250600710823
- Klein, G., Feltovich, P. J., Bradshaw, J. M., and Woods, D. D. (2005). Common Ground and Coordination in Joint Activity. In *Organizational Simulation* (Hoboken, NJ, USA: John Wiley & Sons, Inc.). 139–184. doi:10.1002/0471739448.ch6
- Klin, A., Danovitch, J. H., Merz, A. B., and Volkmar, F. R. (2007). Circumscribed Interests in Higher Functioning Individuals with Autism Spectrum Disorders: An Exploratory Study. *Research and Practice for Persons with Severe Disabilities* 32, 89–100. doi:10.2511/rpsd.32.2.89
- Klin, A., Lin, D. J., Gorrindo, P., Ramsay, G., and Jones, W. (2009). Two-year-olds with autism orient to non-social contingencies rather than biological motion. *Nature* 459, 257–261. doi:10.1038/nature07868
- Knill, D. C. and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences* 27, 712–719. doi:10.1016/j.tins.2004.10.007
- Koegel, L. K., Carter, C. M., and Koegel, R. L. (2003). Teaching Children with Autism Self-Initiations as a Pivotal Response. *Topics in Language Disorders* 23, 134–145. doi:10.1097/00011363-200304000-00006
- Koegel, L. K., Koegel, R. L., Harrower, J. K., and Carter, C. M. (1999a). Pivotal Response Intervention I: Overview of Approach. *Journal of the Association for Persons with Severe Handicaps* 24, 174–185. doi:10.2511/rpsd.24.3.174
- Koegel, L. K., Koegel, R. L., Shoshan, Y., and McNeerney, E. (1999b). Pivotal Response Intervention II: Preliminary Long-Term Outcome Data. *Journal of the Association for Persons with Severe Handicaps* 24, 186–198. doi:10.2511/rpsd.24.3.186
- Koegel, R. L. and Koegel, L. K. (2006). *Pivotal response treatments for autism: Communication, social, & academic development* (Baltimore, MD, USA: Paul H. Brookes Publishing)
- Koppenborg, M., Nickel, P., Naber, B., Lungfiel, A., and Huelke, M. (2017). Effects of movement speed and predictability in human-robot collaboration. *Human Factors and Ergonomics in Manufacturing & Service Industries* 27, 197–209. doi:10.1002/hfm.20703
- Kory, J. and Breazeal, C. (2014). Storytelling with robots: Learning companions for preschool children's language development. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication* (Edinburgh, UK: IEEE), 643–648. doi:10.1109/ROMAN.2014.6926325
- Kostrubiec, V. and Kruck, J. (2020). Collaborative Research Project: Developing and Testing a Robot-Assisted Intervention for Children With Autism. *Frontiers in Robotics and AI* 7, 1–16. doi:10.3389/frobt.2020.00037
- Kozima, H., Michalowski, M. P., and Nakagawa, C. (2009). Keepon. *International Journal of Social Robotics* 1, 3–18. doi:10.1007/s12369-008-0009-8
- Kriz, S., Ferro, T. D., Damera, P., and Porter, J. R. (2010). Fictional robots as a data source in HRI research: Exploring the link between science fiction and interactional expectations. In *19th*

- International Symposium in Robot and Human Interactive Communication* (IEEE), 458–463. doi:10.1109/ROMAN.2010.5598620
- Kwisthout, J., Bekkering, H., and van Rooij, I. (2017). To be precise, the details don't matter: On predictive processing, precision, and level of detail of predictions. *Brain and Cognition* 112, 84–91. doi:10.1016/j.bandc.2016.02.008
- Lang, R., O'Reilly, M., Healy, O., Rispoli, M., Lydon, H., Streusand, W., et al. (2012). Sensory integration therapy for autism spectrum disorders: A systematic review. *Research in Autism Spectrum Disorders* 6, 1004–1018. doi:10.1016/j.rasd.2012.01.006
- Lawson, R. P., Mathys, C., and Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience* 20, 1293–1299. doi:10.1038/nn.4615
- Lawson, R. P., Rees, G., and Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience* 8, 302. doi:10.3389/fnhum.2014.00302
- Lee, M. K., Forlizzi, J., Kiesler, S., Cakmak, M., and Srinivasa, S. S. (2011). Predictability or adaptivity? In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11* (Lausanne, Switzerland: ACM Press), 179. doi:10.1145/1957656.1957720
- Lee, T. S. and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A* 20, 1434. doi:10.1364/JOSAA.20.001434
- Leite, I., Martinho, C., and Paiva, A. (2013). Social Robots for Long-Term Interaction: A Survey. *International Journal of Social Robotics* 5, 291–308. doi:10.1007/s12369-013-0178-y
- Lemaignan, S., Fink, J., Mondada, F., and Dillenbourg, P. (2015). You're Doing It Wrong! Studying Unexpected Behaviors in Child-Robot Interaction. In *International Conference on Social Robotics*, eds. A. Tapus, E. André, J.-C. Martin, F. Ferland, and M. Ammi (Cham: Springer International Publishing), vol. 9388 of *Lecture Notes in Computer Science*. 390–400. doi:10.1007/978-3-319-25554-5_39
- Leslie, A. M. (1987). Pretense and representation: The origins of "theory of mind.". *Psychological Review* 94, 412–426. doi:10.1037/0033-295X.94.4.412
- Lewis, M., Sycara, K., and Walker, P. (2018). The Role of Trust in Human-Robot Interaction. In *Studies in Systems, Decision and Control*, 117 (Cham: Springer International Publishing). 135–159. doi:10.1007/978-3-319-64816-3_{_}8
- Lewis, V. (2003). Play and Language in Children with Autism. *Autism* 7, 391–399. doi:10.1177/1362361303007004005
- Li, J. J., Davison, D., Alcorn, A. M., Williams, A., Dimitrijevic, S. B., Petrović, S., et al. (2020). Non-Participatory User-Centered Design of Accessible Teacher-Teleoperated Robot and Tablets for Minimally Verbal Autistic Children. In *Proceedings of the 13th ACM International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '20)* (Corfu, Greece), 51–59. doi:10.1145/3389189.3393738
- Libby, S., Powell, S., Messer, D., and Jordan, R. (1998). Spontaneous Play in Children with Autism: A Reappraisal. *Journal of Autism and Developmental Disorders* 28, 487–497. doi:10.1023/A:1026095910558
- Lichtenthäler, C. and Kirsch, A. (2014). Goal-predictability vs. trajectory-predictability. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14* (New York, New York, USA: ACM Press), 228–229. doi:10.1145/2559636.2559802
- Lichtenthäler, C., Lorenz, T., Karg, M., and Kirsch, A. (2012a). Increasing perceived value between human and robots - Measuring legibility in human aware navigation. In *2012 IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)* (Munich, Germany: IEEE), 89–94. doi:10.1109/ARSO.2012.6213405
- Lichtenthäler, C., Lorenz, T., and Kirsch, A. (2012b). Influence of legibility on perceived safety in

- a virtual human-robot path crossing task. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication* (Paris, France: IEEE), 676–681. doi:10.1109/ROMAN.2012.6343829
- Lieder, I., Adam, V., Frenkel, O., Jaffe-Dax, S., Sahani, M., and Ahissar, M. (2019). Perceptual bias reveals slow-updating in autism and fast-forgetting in dyslexia. *Nature Neuroscience* 22, 256–264. doi:10.1038/s41593-018-0308-9
- Lifter, K., Foster-Sanda, S., Arzamarski, C., Briesch, J., and McClure, E. (2011). Overview of play: Its uses and importance in early intervention/early childhood special education. *Infants & Young Children* 24, 225–245. doi:10.1097/IYC.0b013e31821e995c
- Lillard, A. S. (2001). Pretend Play as Twin Earth: A Social-Cognitive Analysis. *Developmental Review* 21, 495–531. doi:10.1006/drev.2001.0532
- Lord, C., Rutter, M., DiLavore, P. C., Risi, S., Gotham, K., and Bishop, S. L. (2012). *Autism diagnostic observation schedule: ADOS-2* (Los Angeles, CA, USA: Western Psychological Services)
- Lubold, N., Walker, E., and Pon-Barry, H. (2016). Effects of voice-adaptation and social dialogue on perceptions of a robotic learning companion. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Christchurch, New Zealand: IEEE), 255–262. doi:10.1109/HRI.2016.7451760
- Macdonald, H., Rutter, M., Howlin, P., Rios, P., Conteur, A. L., Evered, C., et al. (1989). Recognition and Expression of Emotional Cues by Autistic and Normal Adults. *Journal of Child Psychology and Psychiatry* 30, 865–877. doi:10.1111/j.1469-7610.1989.tb00288.x
- MacDuff, G. S., Krantz, P. J., and McClannahan, L. E. (1993). Teaching children with autism to use photographic activity schedules: Maintenance and generalization of complex response chains. *Journal of Applied Behavior Analysis* 26, 89–97. doi:10.1901/jaba.1993.26-89
- Manning, C., Kilner, J., Neil, L., Karaminis, T., and Pellicano, E. (2017). Children on the autism spectrum update their behaviour in response to a volatile environment. *Developmental Science* 20, e12435. doi:10.1111/desc.12435
- Marinoiu, E., Zanfir, M., Olaru, V., and Sminchisescu, C. (2018). 3D Human Sensing, Action and Emotion Recognition in Robot Assisted Therapy of Children with Autism. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE)*, 2158–2167. doi:10.1109/CVPR.2018.00230
- Marsh, H. W. (1996). Positive and negative global self-esteem: A substantively meaningful distinction or artifacts? *Journal of Personality and Social Psychology* 70, 810–819. doi:10.1037/0022-3514.70.4.810
- Martin-Ortiz, M., Kim, M.-G., and Barakova, E. I. (2017). Mobile Application for Executing Therapies with Robots. In *14th International Work-Conference on Artificial Neural Networks, IWANN 2017*, eds. I. Rojas, G. Joya, and A. Catala (Cham: Springer International Publishing), vol. 10306 of *Lecture Notes in Computer Science*. 82–92. doi:10.1007/978-3-319-59147-6_8
- Matson, J. L. and Nebel-Schwalm, M. S. (2007). Comorbid psychopathology with autism spectrum disorder in children: An overview. *Research in Developmental Disabilities* 28, 341–352. doi:10.1016/j.ridd.2005.12.004
- Maughan, L., Gutnikov, S., and Stevens, R. (2007). Like more, look more. Look more, like more: The evidence from eye-tracking. *Journal of Brand Management* 14, 335–342. doi:10.1057/palgrave.bm.2550074
- McCann, J. and Peppé, S. (2003). Prosody in autism spectrum disorders: a critical review. *International Journal of Language & Communication Disorders* 38, 325–350. doi:10.1080/1368282031000154204
- McConnell, S. R. (2002). Interventions to Facilitate Social Interaction for Young Children with Autism: Review of Available Research and Recommendations for Educational Intervention and Future Research. *Journal of Autism and Developmental Disorders* 32, 351–372. doi:10.1023/A:1020537805154

- McCormick, L., Noonan, M. J. O., and Heck, R. (1998). Variables Affecting Engagement in Inclusive Preschool Classrooms. *Journal of Early Intervention* 21, 160–176. doi:10.1177/105381519802100208
- [Dataset] Merriam-Webster (2018). Merriam-Webster Thesaurus, accessed September 10th, 2018
- Mesibov, G. B. and Shea, V. (2010). The TEACCH Program in the Era of Evidence-Based Practice. *Journal of Autism and Developmental Disorders* 40, 570–579. doi:10.1007/s10803-009-0901-6
- Miller, B. W. (2015). Using Reading Times and Eye-Movements to Measure Cognitive Engagement. *Educational Psychologist* 50, 31–42. doi:10.1080/00461520.2015.1004068
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267, 1–38. doi:10.1016/j.artint.2018.07.007
- Mubin, O. and Bartneck, C. (2015). Do as I say: Exploring human response to a predictable and unpredictable robot. In *Proceedings of the 2015 British HCI Conference on - British HCI '15* (Lincoln, United Kingdom: ACM Press), 110–116. doi:10.1145/2783446.2783582
- Mumford, D. (1992). On the computational architecture of the neocortex. *Biological Cybernetics* 66, 241–251. doi:10.1007/BF00198477
- Mundy, P., Sigman, M., Ungerer, J. A., and Sherman, T. (1987). Nonverbal communication and play correlates of language development in autistic children. *Journal of Autism and Developmental Disorders* 17, 349–364. doi:10.1007/BF01487065
- Naber, F. B. A., Bakermans-Kranenburg, M. J., van IJzendoorn, M. H., Swinkels, S. H. N., Buitelaar, J. K., Dietz, C., et al. (2008). Play Behavior and Attachment in Toddlers with Autism. *Journal of Autism and Developmental Disorders* 38, 857–866. doi:10.1007/s10803-007-0454-5
- Nazneen, N., Boujarwah, F. A., Sadler, S., Mogus, A., Abowd, G. D., and Arriaga, R. I. (2010). Understanding the challenges and opportunities for richer descriptions of stereotypical behaviors of children with asd. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility (ASSETS '10)* (New York, New York, USA: ACM Press), 67–74. doi:10.1145/1878803.1878817
- Nazneen, N., Rozga, A., Romero, M., Findley, A. J., Call, N. A., Abowd, G. D., et al. (2012). Supporting parents for in-home capture of problem behaviors of children with developmental disabilities. *Personal and Ubiquitous Computing* 16, 193–207. doi:10.1007/s00779-011-0385-1
- Neil, L., Cappagli, G., Karaminis, T., Jenkins, R., and Pellicano, E. (2016a). Recognizing the same face in different contexts: Testing within-person face recognition in typical development and in autism. *Journal of Experimental Child Psychology* 143, 139–153. doi:10.1016/j.jecp.2015.09.029
- Neil, L., Olsson, N. C., and Pellicano, E. (2016b). The Relationship Between Intolerance of Uncertainty, Sensory Sensitivities, and Anxiety in Autistic and Typically Developing Children. *Journal of Autism and Developmental Disorders* 46, 1962–1973. doi:10.1007/s10803-016-2721-9
- Nikolaidis, S., Dragan, A. D., and Srinivasa, S. S. (2016). Viewpoint-based legibility optimization. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Christchurch, New Zealand: IEEE), 271–278. doi:10.1109/HRI.2016.7451762
- Nonnis, A. and Bryan-Kinns, N. (2019). Mazi: Tangible Technologies as a Channel for Collaborative Play. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19* (New York, New York, USA: ACM Press), 1–13. doi:10.1145/3290605.3300670
- Noorman, M. and Johnson, D. G. (2014). Negotiating autonomy and responsibility in military robots. *Ethics and Information Technology* 16, 51–62. doi:10.1007/s10676-013-9335-0
- O'Brien, H. L. and Toms, E. G. (2008). What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the American Society for Information Science and Technology* 59, 938–955. doi:10.1002/asi.20801
- O'Neill, M. and Jones, R. S. (1997). Sensory-perceptual abnormalities in autism: A case for more re-

- search? *Journal of Autism and Developmental Disorders* 27, 283–293. doi:10.1023/A:1025850431170
- Ostrolenk, A., Bao, V. A., Mottron, L., Collignon, O., and Bertone, A. (2019). Reduced multisensory facilitation in adolescents and adults on the Autism Spectrum. *Scientific Reports* 9, 11965. doi:10.1038/s41598-019-48413-9
- Otten, M., Seth, A. K., and Pinto, Y. (2017). A social Bayesian brain: How social knowledge can shape visual perception. *Brain and Cognition* 112, 69–77. doi:10.1016/j.bandc.2016.05.002
- Ozonoff, S., Pennington, B. F., and Rogers, S. J. (1991). Executive Function Deficits in High-Functioning Autistic Individuals: Relationship to Theory of Mind. *Journal of Child Psychology and Psychiatry* 32, 1081–1105. doi:10.1111/j.1469-7610.1991.tb00351.x
- O'Reilly, M., Sigafoos, J., Lancioni, G., Edrisinha, C., and Andrews, A. (2005). An Examination of the Effects of a Classroom Activity Schedule on Levels of Self-Injury and Engagement for a Child with Severe Autism. *Journal of Autism and Developmental Disorders* 35, 305–311. doi:10.1007/s10803-005-3294-1
- Palmer, C. J., Lawson, R. P., and Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. *Psychological Bulletin* 143, 521–542. doi:10.1037/bul0000097
- Papert, S. (1973). Uses of technology to enhance education. *Artificial Intelligence Memo No. 298*
- Paul, R., Augustyn, A., Klin, A., and Volkmar, F. R. (2005). Perception and Production of Prosody by Speakers with Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders* 35, 205–220. doi:10.1007/s10803-004-1999-1
- Paulus, M. P. and Stein, M. B. (2006). An Insular View of Anxiety. *Biological Psychiatry* 60, 383–387. doi:10.1016/j.biopsych.2006.03.042
- Payzan-LeNestour, E. and Bossaerts, P. (2011). Risk, Unexpected Uncertainty, and Estimation Uncertainty: Bayesian Learning in Unstable Settings. *PLoS Computational Biology* 7, e1001048. doi:10.1371/journal.pcbi.1001048
- Peca, A., Simut, R. E., Pintea, S., Costescu, C. A., and Vanderborght, B. (2014). How do typically developing children and children with autism perceive different social robots? *Computers in Human Behavior* 41, 268–277. doi:10.1016/j.chb.2014.09.035
- Pell, P. J., Mareschal, I., Calder, A. J., von dem Hagen, E. A. H., Clifford, C. W., Baron-Cohen, S., et al. (2016). Intact priors for gaze direction in adults with high-functioning autism spectrum conditions. *Molecular Autism* 7, 25. doi:10.1186/s13229-016-0085-9
- Pellicano, E. and Burr, D. (2012). When the world becomes 'too real': A Bayesian explanation of autistic perception. *Trends in Cognitive Sciences* 16, 504–510. doi:10.1016/j.tics.2012.08.009
- Pellicano, E., Hill, V., Croydon, A., Greathead, S., Kenny, L., Yates, R., et al. (2014). *My Life at School: Understanding the experiences of children and young people with special educational needs in residential special schools* (London, UK: Office of the Children's Commissioner)
- Pellicano, E., Jeffery, L., Burr, D., and Rhodes, G. (2007). Abnormal Adaptive Face-Coding Mechanisms in Children with Autism Spectrum Disorder. *Current Biology* 17, 1508–1512. doi:10.1016/j.cub.2007.07.065
- Perlman, R. (1974). TORTIS – Toddler's own recursive turtle interpreter system. *Artificial Intelligence Memo No. 311*, 1–11
- Pezzulo, G., Donnarumma, F., and Dindo, H. (2013). Human Sensorimotor Communication: A Theory of Signaling in Online Social Interactions. *PLoS ONE* 8, e79876. doi:10.1371/journal.pone.0079876
- Pierucci, J. M., Barber, A. B., Gilpin, A. T., Crisler, M. E., and Klinger, L. G. (2015). Play Assessments and Developmental Skills in Young Children With Autism Spectrum Disorders. *Focus on Autism and Other Developmental Disabilities* 30, 35–43. doi:10.1177/1088357614539837

- Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., and R Core Team (2020). *nlme: Linear and Nonlinear Mixed Effects Models*. R package version 3.1-149
- Plaisted, K., O’Riordan, M., and Baron-Cohen, S. (1998). Enhanced Discrimination of Novel, Highly Similar Stimuli by Adults with Autism During a Perceptual Learning Task. *Journal of Child Psychology and Psychiatry* 39, S0021963098002601. doi:10.1017/S0021963098002601
- Pop, C. A., Pinteá, S., Vanderborght, B., and David, D. O. (2014). Enhancing play skills, engagement and social skills in a play task in ASD children by using robot-based interventions. A pilot study. *Interaction Studies* 15, 292–320. doi:10.1075/is.15.2.14pop
- Pop, C. A., Simut, R. E., Pinteá, S., Saldien, J., Rusu, A. S., Vanderfaeillie, J., et al. (2013). Social Robots vs. Computer Display: Does the Way Social Stories are Delivered Make a difference for Their Effectiveness on ASD Children? *Journal of Educational Computing Research* 49, 381–401. doi:10.2190/EC.49.3.f
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology* 32, 3–25. doi:10.1080/00335558008248231
- Potter, W. J. and Levine-Donnerstein, D. (1999). Rethinking validity and reliability in content analysis. *Journal of Applied Communication Research* 27, 258–284. doi:10.1080/00909889909365539
- Putnam, C., Hanschke, C., Todd, J., Gemmell, J., and Kollia, M. (2019). Interactive Technologies Designed for Children with Autism. *ACM Transactions on Accessible Computing* 12, 1–37. doi:10.1145/3342285
- Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., et al. (2009). ROS: an open-source Robot Operating System Morgan. In *ICRA Workshop on Open Source Software* (Kobe, Japan), 5
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. Vienna, Austria
- Ragni, M., Rudenko, A., Kuhnert, B., and Arras, K. O. (2016). Errare humanum est: Erroneous robots in human-robot interaction. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (New York, New York, USA: IEEE), 501–506. doi:10.1109/ROMAN.2016.7745164
- Ramachandran, A., Huang, C.-m., and Scassellati, B. (2017). Give Me a Break! In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI ’17* (New York, New York, USA: ACM Press), 146–155. doi:10.1145/2909824.3020209
- Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2, 79–87. doi:10.1038/4580
- Raudenbush, S. W. and Bryk, A. S. (2002). *Hierarchical Linear Models. Applications and Data Analysis Methods* (Thousand Oaks, CA: Sage Publications), 2nd edn.
- Reidsma, D. and van Welbergen, H. (2013). AsapRealizer in practice – A modular and extensible architecture for a BML Realizer. *Entertainment Computing* 4, 157–169. doi:10.1016/j.entcom.2013.05.001
- Rich, C., Ponsleur, B., Holroyd, A., and Sidner, C. L. (2010). Recognizing engagement in human-robot interaction. In *Proceeding of the 5th ACM/IEEE international conference on Human-robot interaction - HRI ’10* (New York, New York, USA: ACM Press), 375. doi:10.1145/1734454.1734580
- Ricks, D. J. and Colton, M. B. (2010). Trends and considerations in robot-assisted autism therapy. In *2010 IEEE International Conference on Robotics and Automation (IEEE)*, August, 4354–4359. doi:10.1109/ROBOT.2010.5509327
- Riek, L. D. (2012). Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *Journal of Human-Robot Interaction* 1, 119–136. doi:10.5898/JHRI.1.1.Riek
- Robertson, A. E. and Simmons, D. R. (2013). The Relationship between Sensory Sensitivity and Autistic Traits in the General Population. *Journal of Autism and Developmental Disorders* 43, 775–784. doi:10.1007/s10803-012-1608-7

- Robins, B., Amirabdollahian, F., Ji, Z., and Dautenhahn, K. (2010a). Tactile interaction with a humanoid robot for children with autism: A case study analysis involving user requirements and results of an initial implementation. In *19th International Symposium on Robot and Human Interactive Communication (IEEE)*, 704–711. doi:10.1109/ROMAN.2010.5598641
- Robins, B. and Dautenhahn, K. (2014). Tactile Interactions with a Humanoid Robot: Novel Play Scenario Implementations with Children with Autism. *International Journal of Social Robotics* 6, 397–415. doi:10.1007/s12369-014-0228-0
- Robins, B., Dautenhahn, K., and Dickerson, P. (2009). From Isolation to Communication: A Case Study Evaluation of Robot Assisted Play for Children with Autism with a Minimally Expressive Humanoid Robot. In *2009 Second International Conferences on Advances in Computer-Human Interactions (Cancun, Mexico: IEEE)*, 205–211. doi:10.1109/ACHI.2009.32
- Robins, B., Dautenhahn, K., te Boekhorst, R., and Billard, A. (2005). Robotic assistants in therapy and education of children with autism: can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society* 4, 105–120. doi:10.1007/s10209-005-0116-3
- Robins, B., Dickerson, P., Stribling, P., and Dautenhahn, K. (2004). Robot-mediated joint attention in children with autism: A case study in robot-human interaction. *Interaction Studies* 5, 161–198. doi:10.1075/is.5.2.02rob
- Robins, B., Ferrari, E., Dautenhahn, K., Kronreif, G., Prazak-Aram, B., Gelderblom, G.-J., et al. (2010b). Human-centred design methods: Developing scenarios for robot assisted play informed by user panels and field trials. *International Journal of Human-Computer Studies* 68, 873–898. doi:10.1016/j.ijhcs.2010.08.001
- Robins, B., Otero, N., Ferrari, E., and Dautenhahn, K. (2007). Eliciting Requirements for a Robotic Toy for Children with Autism - Results from User Panels. In *RO-MAN 2007 - The 16th IEEE International Symposium on Robot and Human Interactive Communication (Jeju, Korea: IEEE)*, 101–106. doi:10.1109/ROMAN.2007.4415061
- Rodebaugh, T. L., Woods, C. M., Thissen, D. M., Heimberg, R. G., Chambless, D. L., and Rapee, R. M. (2004). More Information From Fewer Questions: The Factor Structure and Item Properties of the Original and Brief Fear of Negative Evaluation Scale. *Psychological Assessment* 16, 169–181. doi:10.1037/1040-3590.16.2.169
- Rodgers, J., Hodgson, A., Shields, K., Wright, C., Honey, E., and Freeston, M. (2017). Towards a Treatment for Intolerance of Uncertainty in Young People with Autism Spectrum Disorder: Development of the Coping with Uncertainty in Everyday Situations (CUES©) Programme. *Journal of Autism and Developmental Disorders* 47, 3959–3966. doi:10.1007/s10803-016-2924-0
- Rudovic, O., Lee, J., Dai, M., Schuller, B. W., and Picard, R. W. (2018). Personalized machine learning for robot perception of affect and engagement in autism therapy. *Science Robotics* 3, eaao6760. doi:10.1126/scirobotics.aao6760
- Rudovic, O., Lee, J., Mascarell-Maricic, L., Schuller, B. W., and Picard, R. W. (2017). Measuring Engagement in Robot-Assisted Autism Therapy: A Cross-Cultural Study. *Frontiers in Robotics and AI* 4, 36. doi:10.3389/frobt.2017.00036
- Rutter, M., Bailey, A., and Lord, C. (2003). *The social communication questionnaire: Manual* (Los Angeles, CA, USA: Western Psychological Services)
- Ryan, C. and Charragáin, C. N. (2010). Teaching Emotion Recognition Skills to Children with Autism. *Journal of Autism and Developmental Disorders* 40, 1505–1511. doi:10.1007/s10803-010-1009-8
- Ryan, R. M. and Deci, E. L. (2000). Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology* 25, 54–67. doi:10.1006/ceps.1999.1020
- Saldaña, J. (2015). *The coding manual for qualitative researchers* (Thousand Oaks, CA: Sage Publications), 3rd edn.

- Saldien, J., Goris, K., Yilmazyildiz, S., Verhelst, W., and Lefeber, D. (2008). On the design of the huggable robot Probo. *Journal of Physical Agents* 2, 3–11. doi:10.14198/JoPha.2008.2.2.02
- Salem, M., Eyssel, F., Rohlfing, K., Kopp, S., and Joubin, F. (2013). To Err is Human(-like): Effects of Robot Gesture on Perceived Anthropomorphism and Likability. *International Journal of Social Robotics* 5, 313–323. doi:10.1007/s12369-013-0196-9
- Salem, M., Lakatos, G., Amirabdollahian, F., and Dautenhahn, K. (2015). Would You Trust a (Faulty) Robot? In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15* (New York, New York, USA: ACM Press), 141–148. doi:10.1145/2696454.2696497
- Salter, T., Werry, I., and Michaud, F. (2008). Going into the wild in child–robot interaction studies: issues in social robotic development. *Intelligent Service Robotics* 1, 93–108. doi:10.1007/s11370-007-0009-9
- Salvador, M. J., Silver, S., and Mahoor, M. H. (2015). An emotion recognition comparative study of autistic and typically-developing children using the zeno robot. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (Seattle, WA, USA: IEEE), 6128–6133. doi:10.1109/ICRA.2015.7140059
- Sartorato, F., Przybyłowski, L., and Sarko, D. K. (2017). Improving therapeutic outcomes in autism spectrum disorders: Enhancing social communication and sensory processing through the use of interactive robots. *Journal of Psychiatric Research* 90, 1–11. doi:10.1016/j.jpsychires.2017.02.004
- Scassellati, B. (2005). Using social robots to study abnormal social development. *Proceedings of the 5th international workshop on epigenetic robotics* , 11–14
- Scassellati, B. (2007). How Social Robots Will Help Us to Diagnose, Treat, and Understand Autism. In *Robotics Research*, eds. S. Thrun, R. Brooks, and H. Durrant-Whyte (Berlin, Heidelberg: Springer Berlin Heidelberg). 552–563. doi:10.1007/978-3-540-48113-3_47
- Scassellati, B., Admoni, H., and Matarić, M. J. (2012). Robots for Use in Autism Research. *Annual Review of Biomedical Engineering* 14, 275–294. doi:10.1146/annurev-bioeng-071811-150036
- Scassellati, B., Boccanfuso, L., Huang, C.-m., Mademtzi, M., Qin, M., Salomons, N., et al. (2018). Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics* 3, eaat7544. doi:10.1126/scirobotics.aat7544
- Schadenberg, B. R. (2019). Predictability in Human-Robot Interactions for Autistic Children. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Daegu, South Korea: IEEE), 748–750. doi:10.1109/HRI.2019.8673135
- Schadenberg, B. R., Heylen, D. K. J., and Evers, V. (2018). Affect bursts to constrain the meaning of the facial expressions of the humanoid robot Zeno. In *Proceedings of the 1st Workshop on Social Interaction and Multimodal Expression for Socially Intelligent Robots (RO-MAN 2017)* (Lisbon, Portugal: CEUR Workshop Proceedings), 30–39
- Schadenberg, B. R., Li, J. J., Petrović, S., Reidsma, D., Heylen, D. K. J., and Evers, V. (2020a). Helping Educators Monitor Autistic Children's Progress Across Sessions. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (New York, NY, USA: ACM), 436–438. doi:10.1145/3371382.3378378
- Schadenberg, B. R., Neerinx, M. A., Cnossen, F., and Looije, R. (2017). Personalising game difficulty to keep children motivated to play with a social robot: A Bayesian approach. *Cognitive Systems Research* 43, 222–231. doi:10.1016/j.cogsys.2016.08.003
- Schadenberg, B. R., Reidsma, D., Evers, V., Davison, D. P., Li, J. J., Heylen, D. K. J., et al. (under revision). Predictable Robots for Autistic Children — Variance in Robot Behaviour, Idiosyncrasies in Autistic Children's Characteristics, and Child-Robot Engagement. *ACM Transactions on Computer-Human Interaction* , 1–39
- Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., and Evers, V. (2020b). Differences in Spontaneous

- Interactions of Autistic Children in an Interaction With an Adult and Humanoid Robot. *Frontiers in Robotics and AI* 7, 19. doi:10.3389/frobt.2020.00028
- Schadenberg, B. R., Reidsma, D., Heylen, D. K. J., and Evers, V. (in press). “I See What You Did There”: Understanding People’s Social Perception of a Robot and its Predictability. *ACM Transactions on Human-Robot Interaction* 10, 1–27. doi:10.1145/3461534
- Schopler, E., Van Bourgondien, M. E., Wellman, G. J., and Love, S. R. (2010). *The Childhood Autism Rating Scale (CARS-2)* (Los Angeles, CA, USA: Western Psychological Services)
- Schröder, M. (2003). Experimental study of affect bursts. *Speech Communication* 40, 99–116. doi:10.1016/S0167-6393(02)00078-X
- Sciutti, A., Mara, M., Tagliasco, V., and Sandini, G. (2018). Humanizing Human-Robot Interaction: On the Importance of Mutual Understanding. *IEEE Technology and Society Magazine* 37, 22–29. doi:10.1109/MTS.2018.2795095
- Sebanz, N., Bekkering, H., and Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in Cognitive Sciences* 10, 70–76. doi:10.1016/j.tics.2005.12.009
- Sebanz, N. and Knoblich, G. (2009). Prediction in Joint Action: What, When, and Where. *Topics in Cognitive Science* 1, 353–367. doi:10.1111/j.1756-8765.2009.01024.x
- Shah, A. and Frith, U. (1983). An Islet of Ability in Autistic Children: A Research Note. *Journal of Child Psychology and Psychiatry* 24, 613–620. doi:10.1111/j.1469-7610.1983.tb00137.x
- Shane, H. C. and Albert, P. D. (2008). Electronic Screen Media for Persons with Autism Spectrum Disorders: Results of a Survey. *Journal of Autism and Developmental Disorders* 38, 1499–1508. doi:10.1007/s10803-007-0527-5
- Short, E. S., Deng, E. C., Feil-Seifer, D. J., and Matarić, M. J. (2017). Understanding Agency in Interactions Between Children With Autism and Socially Assistive Robots. *Journal of Human-Robot Interaction* 6, 21. doi:10.5898/JHRI.6.3.Short
- Short, E. S., Hart, J. W., Vu, M., and Scassellati, B. (2010). No fair!! An interaction with a cheating robot. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Osaka, Japan: IEEE), 219–226. doi:10.1109/HRI.2010.5453193
- Shvarts, A. and Bakker, A. (2019). The early history of the scaffolding metaphor: Bernstein, Luria, Vygotsky, and before. *Mind, Culture, and Activity* 26, 4–23. doi:10.1080/10749039.2019.1574306
- Sidner, C. L., Lee, C., Kidd, C. D., Lesh, N., and Rich, C. (2005). Explorations in engagement for humans and robots. *Artificial Intelligence* 166, 140–164. doi:10.1016/j.artint.2005.03.005
- Sigman, M., Mundy, P., Sherman, T., and Ungerer, J. (1986). Social interactions of autistic, mentally retarded and normal children and their caregivers. *Journal of Child Psychology and Psychiatry* 27, 647–656. doi:10.1111/j.1469-7610.1986.tb00189.x
- Silver, M. and Oakes, P. (2001). Evaluation of a New Computer Intervention to Teach People with Autism or Asperger Syndrome to Recognize and Predict Emotions in Others. *Autism* 5, 299–316. doi:10.1177/1362361301005003007
- Simonoff, E., Pickles, A., Charman, T., Chandler, S., Loucas, T., and Baird, G. (2008). Psychiatric Disorders in Children With Autism Spectrum Disorders: Prevalence, Comorbidity, and Associated Factors in a Population-Derived Sample. *Journal of the American Academy of Child & Adolescent Psychiatry* 47, 921–929. doi:10.1097/CHI.0b013e318179964f
- Simpson, K., Keen, D., and Lamb, J. (2013). The use of music to engage children with autism in a receptive labelling task. *Research in Autism Spectrum Disorders* 7, 1489–1496. doi:10.1016/j.rasd.2013.08.013
- Simut, R. E., Vanderfaeillie, J., Peca, A., Van de Perre, G., and Vanderborght, B. (2016). Children with Autism Spectrum Disorders Make a Fruit Salad with Probo, the Social Robot: An Interaction Study.

- Journal of Autism and Developmental Disorders* 46, 113–126. doi:10.1007/s10803-015-2556-9
- Sinatra, G. M., Heddy, B. C., and Lombardi, D. (2015). The Challenges of Defining and Measuring Student Engagement in Science. *Educational Psychologist* 50, 1–13. doi:10.1080/00461520.2014.1002924
- Sinha, P., Kjelgaard, M. M., Gandhi, T. K., Tsourides, K., Cardinaux, A. L., Pantazis, D., et al. (2014). Autism as a disorder of prediction. *Proceedings of the National Academy of Sciences* 111, 15220–15225. doi:10.1073/pnas.1416797111
- Sparrow, S. S., Cicchetti, D. V., and Balla, D. A. (2005). *Vineland Adaptive Behavior Scales 2 (Vineland-2)* (San Antonio, TX: Pearson), 2nd edn.
- Spector, P. E., Van Katwyk, P. T., Brannick, M. T., and Chen, P. Y. (1997). When Two Factors Don't Reflect Two Constructs: How Item Characteristics Can Produce Artifactual Factors. *Journal of Management* 23, 659–677. doi:10.1177/014920639702300503
- Spiel, K., Frauenberger, C., Keyes, O., and Fitzpatrick, G. (2019). Agency of Autistic Children in Technology Research—A Critical Literature Review. *ACM Transactions on Computer-Human Interaction* 26, 1–40. doi:10.1145/3344919
- Spiel, K., Malinverni, L., Good, J., and Frauenberger, C. (2017). Participatory Evaluation with Autistic Children. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA: ACM), vol. 2017-May, 5755–5766. doi:10.1145/3025453.3025851
- Srinivasan, S. and Bhat, A. (2013). The Effect of Robot-Child Interactions on Social Attention and Verbalization Patterns of Typically Developing Children and Children With Autism Between 4 and 8 Years. *Autism* 3, 111. doi:10.4172/2165-7890.1000111
- Srinivasan, S. M., Pescatello, L. S., and Bhat, A. N. (2014). Current perspectives on physical activity and exercise recommendations for children and adolescents with autism spectrum disorders. *Physical Therapy* 94, 875–889. doi:10.2522/ptj.20130157
- Stagnitti, K. (2004). Understanding play: The Implications for play assessment. *Australian Occupational Therapy Journal* 51, 3–12. doi:10.1046/j.1440-1630.2003.00387.x
- Stange, S. and Kopp, S. (2020). Effects of a Social Robot's Self-Explanations on How Humans Understand and Evaluate Its Behavior. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (New York, NY, USA: ACM), Section 2, 619–627. doi:10.1145/3319502.3374802
- Stevens, J. P. (2002). *Applied multivariate statistics for the social sciences* (New York, NY, USA: Taylor & Francis), 5th edn.
- Stone, W. L., Ousley, O. Y., Yoder, P. J., Hogan, K. L., and Hepburn, S. L. (1997). Nonverbal communication in two- and three-year-old children with autism. *Journal of Autism and Developmental Disorders* 27, 677–696. doi:10.1023/A:1025854816091
- Strand, P. S., Downs, A., and Barbosa-Leiker, C. (2016). Does facial expression recognition provide a toehold for the development of emotion understanding? *Developmental Psychology* 52, 1182–1191. doi:10.1037/dev0000144
- Syrdal, D. S., Dautenhahn, K., Robins, B., Karakosta, E., and Jones, N. C. (2020). Kaspar in the wild: Experiences from deploying a small humanoid robot in a nursery school for children with autism. *Paladyn, Journal of Behavioral Robotics* 11, 301–326. doi:10.1515/pjbr-2020-0019
- Szafir, D., Mutlu, B., and Fong, T. (2014). Communication of intent in assistive free flyers. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14* (Bielefeld, Germany: ACM Press), 358–365. doi:10.1145/2559636.2559672
- Szafir, D., Mutlu, B., and Fong, T. (2015). Communicating Directionality in Flying Robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15* (New

- York, New York, USA: ACM Press), 19–26. doi:10.1145/2696454.2696475
- Tager-Flusberg, H. and Kasari, C. (2013). Minimally Verbal School-Aged Children with Autism Spectrum Disorder: The Neglected End of the Spectrum. *Autism Research* 6, 468–478. doi:10.1002/aur.1329
- Takayama, L., Dooley, D., and Ju, W. (2011). Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11* (Lausanne, Switzerland: ACM Press), 69. doi:10.1145/1957656.1957674
- Tamir, D. I. and Thornton, M. A. (2018). Modeling the Predictive Social Mind. *Trends in Cognitive Sciences* 22, 201–212. doi:10.1016/j.tics.2017.12.005
- Tapus, A., Peca, A., Aly, A., Pop, C. A., Jisa, L., Pintea, S., et al. (2012). Children with autism social engagement in interaction with Nao, an imitative robot: A series of single case experiments. *Interaction Studies* 13, 315–347. doi:10.1075/is.13.3.01tap
- Tewolde, F. G., Bishop, D. V. M., and Manning, C. (2018). Visual Motion Prediction and Verbal False Memory Performance in Autistic Children. *Autism Research* 11, 509–518. doi:10.1002/aur.1915
- Thiemann-Bourque, K. S., Brady, N. C., and Fleming, K. K. (2012). Symbolic Play of Preschoolers with Severe Communication Impairments with Autism and Other Developmental Delays: More Similarities than Differences. *Journal of Autism and Developmental Disorders* 42, 863–873. doi:10.1007/s10803-011-1317-7
- Thill, S., Pop, C. A., Belpaeme, T., Ziemke, T., and Vanderborght, B. (2012). Robot-Assisted Therapy for Autism Spectrum Disorders with (Partially) Autonomous Control: Challenges and Outlook. *Paladyn, Journal of Behavioral Robotics* 3, 209–217. doi:10.2478/s13230-013-0107-7
- Thomas, L. A., De Bellis, M. D., Graham, R., and LaBar, K. S. (2007). Development of emotional facial recognition in late childhood and adolescence. *Developmental Science* 10, 547–558. doi:10.1111/j.1467-7687.2007.00614.x
- Toth, K., Munson, J., Meltzoff, A. N., and Dawson, G. (2006). Early Predictors of Communication Development in Young Children with Autism Spectrum Disorder: Joint Attention, Imitation, and Toy Play. *Journal of Autism and Developmental Disorders* 36, 993–1005. doi:10.1007/s10803-006-0137-7
- Trafton, J., Cassimatis, N., Bugajska, M., Brock, D., Mintz, F., and Schultz, A. C. (2005). Enabling Effective Human–Robot Interaction Using Perspective-Taking in Robots. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 35, 460–470. doi:10.1109/TSMCA.2005.850592
- Turi, M., Burr, D. C., Iglizzi, R., Aagten-Murphy, D., Muratori, F., and Pellicano, E. (2015). Children with autism spectrum disorder show reduced adaptation to number. *Proceedings of the National Academy of Sciences* 112, 7868–7872. doi:10.1073/pnas.1504099112
- Uljarević, M. and Hamilton, A. (2013). Recognition of Emotions in Autism: A Formal Meta-Analysis. *Journal of Autism and Developmental Disorders* 43, 1517–1526. doi:10.1007/s10803-012-1695-5
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., De-Wit, L., et al. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review* 121, 649–675. doi:10.1037/a0037665
- van Straten, C. L., Smeekens, I., Barakova, E. I., Glennon, J. C., Buitelaar, J. K., and Chen, A. (2018). Effects of robots' intonation and bodily appearance on robot-mediated communicative treatment outcomes for children with autism spectrum disorder. *Personal and Ubiquitous Computing* 22, 379–390. doi:10.1007/s00779-017-1060-y
- van Waterschoot, J., Bruijnes, M., Flokstra, J., Reidsma, D., Davison, D., Theune, M., et al. (2018). Flipper 2.0. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents - IVA '18* (New York, New York, USA: ACM Press), 43–50. doi:10.1145/3267851.3267882
- Vernetti, A., Senju, A., Charman, T., Johnson, M. H., and Gliga, T. (2018). Simulating interaction: Using gaze-contingent eye-tracking to measure the reward value of social signals in toddlers with and

- without autism. *Developmental Cognitive Neuroscience* 29, 21–29. doi:10.1016/j.dcn.2017.08.004
- Vesper, C., van der Wel, R. P. R. D., Knoblich, G., and Sebanz, N. (2011). Making oneself predictable: reduced temporal variability facilitates joint action coordination. *Experimental Brain Research* 211, 517–530. doi:10.1007/s00221-011-2706-z
- Vig, S. (2007). Young Children's Object Play: A Window on Development. *Journal of Developmental and Physical Disabilities* 19, 201–215. doi:10.1007/s10882-007-9048-6
- Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., and Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences* 108, 20754–20759. doi:10.1073/pnas.1117807108
- Wainer, J., Robins, B., Amirabdollahian, F., and Dautenhahn, K. (2014). Using the Humanoid Robot KASPAR to Autonomously Play Triadic Games and Facilitate Collaborative Play Among Children With Autism. *IEEE Transactions on Autonomous Mental Development* 6, 183–199. doi:10.1109/TAMD.2014.2303116
- Walker, P., Lewis, M., and Sycara, K. (2016). The effect of display type on operator prediction of future swarm states. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (Budapest, Hungary: IEEE), 002521–002526. doi:10.1109/SMC.2016.7844619
- Warren, Z. E., Zheng, Z., Das, S., Young, E. M., Swanson, A. R., Weitlauf, A. F., et al. (2015a). Brief Report: Development of a Robotic Intervention Platform for Young Children with ASD. *Journal of Autism and Developmental Disorders* 45, 3870–3876. doi:10.1007/s10803-014-2334-0
- Warren, Z. E., Zheng, Z., Swanson, A. R., Bekele, E. T., Zhang, L., Crittendon, J. A., et al. (2015b). Can Robotic Interaction Improve Joint Attention Skills? *Journal of Autism and Developmental Disorders* 45, 3726–3734. doi:10.1007/s10803-013-1918-4
- Weigelt, S., Koldewyn, K., and Kanwisher, N. (2012). Face identity recognition in autism spectrum disorders: A review of behavioral studies. *Neuroscience & Biobehavioral Reviews* 36, 1060–1084. doi:10.1016/j.neubiorev.2011.12.008
- White, S. W., Oswald, D., Ollendick, T., and Scahill, L. (2009). Anxiety in children and adolescents with autism spectrum disorders. *Clinical Psychology Review* 29, 216–229. doi:10.1016/j.cpr.2009.01.003
- Wigham, S., Rodgers, J., South, M., McConachie, H., and Freeston, M. H. (2015). The Interplay Between Sensory Processing Abnormalities, Intolerance of Uncertainty, Anxiety and Restricted and Repetitive Behaviours in Autism Spectrum Disorder. *Journal of Autism and Developmental Disorders* 45, 943–952. doi:10.1007/s10803-014-2248-x
- Williams, E. I. (2003). A Comparative Review of Early Forms of Object-Directed Play and Parent-Infant Play in Typical Infants and Young Children with Autism. *Autism* 7, 361–374. doi:10.1177/1362361303007004003
- Williams, E. I., Reddy, V., and Costall, A. (2001). Taking a Closer Look at Functional Play in Children with Autism. *Journal of Autism and Developmental Disorders* 31, 67–77. doi:10.1023/A:1005665714197
- Wilson, K. P., Carter, M. W., Wiener, H. L., DeRamus, M. L., Bulluck, J. C., Watson, L. R., et al. (2017). Object play in infants with autism spectrum disorder: A longitudinal retrospective video analysis. *Autism & Developmental Language Impairments* 2, 1–12. doi:10.1177/2396941517713186
- Wojciszke, B., Bazinska, R., and Jaworski, M. (1998). On the Dominance of Moral Categories in Impression Formation. *Personality and Social Psychology Bulletin* 24, 1251–1263. doi:10.1177/01461672982412001
- Wood, D. and Middleton, D. (1975). A Study of Assisted Problem-solving. *British Journal of Psychology* 66, 181–191. doi:10.1111/j.2044-8295.1975.tb01454.x
- Wood, L. J., Robins, B., Lakatos, G., Syrdal, D. S., Zaraki, A., and Dautenhahn, K. (2019). Developing a protocol and experimental setup for using a humanoid robot to assist children with autism to

- develop visual perspective taking skills. *Paladyn, Journal of Behavioral Robotics* 10, 167–179. doi:10.1515/pjbr-2019-0013
- World Health Organization (1992). *The ICD-10 classification of mental and behavioural disorders: clinical descriptions and diagnostic guidelines* (Geneva: World Health Organization)
- Wu, S.-C. and Remington, R. W. (2003). Characteristics of covert and overt visual orienting: Evidence from attentional and oculomotor capture. *Journal of Experimental Psychology: Human Perception and Performance* 29, 1050–1067. doi:10.1037/0096-1523.29.5.1050
- Wyart, V., Nobre, A. C., and Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences* 109, 3593–3598. doi:10.1073/pnas.1120118109
- Yirmiya, N., Kasari, C., Sigman, M., and Mundy, P. (1989). Facial Expressions of Affect in Autistic, Mentally Retarded and Normal Children. *Journal of Child Psychology and Psychiatry* 30, 725–735. doi:10.1111/j.1469-7610.1989.tb00785.x
- Yu, A. J. and Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron* 46, 681–692. doi:10.1016/j.neuron.2005.04.026
- Zaga, C., Charisi, V., Schadenberg, B. R., Reidsma, D., Neerincx, M. A., Prescott, T., et al. (2017). Growing-Up Hand in Hand with Robots. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17* (New York, New York, USA: ACM Press), 429–430. doi:10.1145/3029798.3029804
- Zheng, Z., Nie, G., Swanson, A. R., Weitlauf, A. F., Warren, Z. E., and Sarkar, N. (2020). A Randomized Controlled Trial of an Intelligent Robotic Response to Joint Attention Intervention System. *Journal of Autism and Developmental Disorders* 50, 2819–2831. doi:10.1007/s10803-020-04388-5
- Zubrycki, I. and Granosik, G. (2016). Understanding Therapists' Needs and Attitudes Towards Robotic Support. The Robotrapia Project. *International Journal of Social Robotics* 8, 553–563. doi:10.1007/s12369-016-0372-9

SIKS dissertation list

- 2011 01 Botond Cseke (RUN), Variational Algorithms for Bayesian Inference in Latent Gaussian Models
02 Nick Tinnemeier (UU), Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language
03 Jan Martijn van der Werf (TUE), Compositional Design and Verification of Component-Based Information Systems
04 Hado van Hasselt (UU), Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference
05 Bas van der Raadt (VU), Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.
06 Yiwen Wang (TUE), Semantically-Enhanced Recommendations in Cultural Heritage
07 Yujia Cao (UT), Multimodal Information Presentation for High Load Human Computer Interaction
08 Nieske Vergunst (UU), BDI-based Generation of Robust Task-Oriented Dialogues
09 Tim de Jong (OU), Contextualised Mobile Media for Learning
10 Bart Bogaert (UvT), Cloud Content Contention
11 Dhaval Vyas (UT), Designing for Awareness: An Experience-focused HCI Perspective
12 Carmen Bratosin (TUE), Grid Architecture for Distributed Process Mining
13 Xiaoyu Mao (UvT), Airport under Control. Multiagent Scheduling for Airport Ground Handling
14 Milan Lovric (EUR), Behavioral Finance and Agent-Based Artificial Markets
15 Marijn Koolen (UvA), The Meaning of Structure: the Value of Link Evidence for Information Retrieval
16 Maarten Schadd (UM), Selective Search in Games of Different Complexity
17 Jiyin He (UVA), Exploring Topic Structure: Coherence, Diversity and Relatedness
18 Mark Ponsen (UM), Strategic Decision-Making in complex games
19 Ellen Rusman (OU), The Mind's Eye on Personal Profiles
20 Qing Gu (VU), Guiding service-oriented software engineering - A view-based approach
21 Linda Terlouw (TUD), Modularization and Specification of Service-Oriented Systems
22 Junte Zhang (UVA), System Evaluation of Archival Description and Access
23 Wouter Weerkamp (UVA), Finding People and their Utterances in Social Media
24 Herwin van Welbergen (UT), Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior
25 Syed Waqar ul Qounain Jaffry (VU), Analysis and Validation of Models for Trust Dynamics
26 Matthijs Aart Pontier (VU), Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots
27 Aniel Bhulai (VU), Dynamic website optimization through autonomous management of design patterns
28 Rianne Kaptein (UVA), Effective Focused Retrieval by Exploiting Query Context and Document Structure
29 Faisal Kamiran (TUE), Discrimination-aware Classification
30 Egon van den Broek (UT), Affective Signal Processing (ASP): Unraveling the mystery of emotions
31 Ludo Waltman (EUR), Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
32 Nees-Jan van Eck (EUR), Methodological Advances in Bibliometric Mapping of Science
33 Tom van der Weide (UU), Arguing to Motivate Decisions
34 Paolo Turrini (UU), Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations
35 Maaike Harbers (UU), Explaining Agent Behavior in Virtual Training
36 Erik van der Spek (UU), Experiments in serious game design: a cognitive approach
37 Adriana Burlutiu (RUN), Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference
38 Nyree Lemmens (UM), Bee-inspired Distributed Optimization
39 Joost Westra (UU), Organizing Adaptation using Agents in Serious Games

- 40 Viktor Clerc (VU), Architectural Knowledge Management in Global Software Development
 41 Luan Ibraimi (UT), Cryptographically Enforced Distributed Data Access Control
 42 Michal Sindlar (UU), Explaining Behavior through Mental State Attribution
 43 Henk van der Schuur (UU), Process Improvement through Software Operation Knowledge
 44 Boris Reuderink (UT), Robust Brain-Computer Interfaces
 45 Herman Stehouwer (UvT), Statistical Language Models for Alternative Sequence Selection
 46 Beibei Hu (TUD), Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work
 47 Azizi Bin Ab Aziz (VU), Exploring Computational Models for Intelligent Support of Persons with Depression
 48 Mark Ter Maat (UT), Response Selection and Turn-taking for a Sensitive Artificial Listening Agent
 49 Andreea Niculescu (UT), Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality
-
- 2012 01 Terry Kakeeto (UvT), Relationship Marketing for SMEs in Uganda
 02 Muhammad Umair (VU), Adaptivity, emotion, and Rationality in Human and Ambient Agent Models
 03 Adam Vanya (VU), Supporting Architecture Evolution by Mining Software Repositories
 04 Jurriaan Souer (UU), Development of Content Management System-based Web Applications
 05 Marijn Plomp (UU), Maturing Interorganisational Information Systems
 06 Wolfgang Reinhardt (OU), Awareness Support for Knowledge Workers in Research Networks
 07 Rianne van Lambalgen (VU), When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions
 08 Gerben de Vries (UVA), Kernel Methods for Vessel Trajectories
 09 Ricardo Neisse (UT), Trust and Privacy Management Support for Context-Aware Service Platforms
 10 David Smits (TUE), Towards a Generic Distributed Adaptive Hypermedia Environment
 11 J.C.B. Rantham Prabhakara (TUE), Process Mining in the Large: Preprocessing, Discovery, and Diagnostics
 12 Kees van der Sluijs (TUE), Model Driven Design and Data Integration in Semantic Web Information Systems
 13 Suleman Shahid (UvT), Fun and Face: Exploring non-verbal expressions of emotion during playful interactions
 14 Evgeny Knutov (TUE), Generic Adaptation Framework for Unifying Adaptive Web-based Systems
 15 Natalie van der Wal (VU), Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.
 16 Fiemke Both (VU), Helping people by understanding them - Ambient Agents supporting task execution and depression treatment
 17 Amal Elgammal (UvT), Towards a Comprehensive Framework for Business Process Compliance
 18 Eltjo Poort (VU), Improving Solution Architecting Practices
 19 Helen Schonenberg (TUE), What's Next? Operational Support for Business Process Execution
 20 Ali Bahramisharif (RUN), Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing
 21 Roberto Cornacchia (TUD), Querying Sparse Matrices for Information Retrieval
 22 Thijs Vis (UvT), Intelligence, politie en veiligheidsdienst: verenigbare grootheden?
 23 Christian Muehl (UT), Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction
 24 Laurens van der Werff (UT), Evaluation of Noisy Transcripts for Spoken Document Retrieval
 25 Silja Eckartz (UT), Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application
 26 Emile de Maat (UVA), Making Sense of Legal Text
 27 Hayrettin Gurkok (UT), Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games
 28 Nancy Pascall (UvT), Engendering Technology Empowering Women
 29 Almer Tigelaar (UT), Peer-to-Peer Information Retrieval
 30 Alina Pommeranz (TUD), Designing Human-Centered Systems for Reflective Decision Making
 31 Emily Bagarukayo (RUN), A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure
 32 Wietske Visser (TUD), Qualitative multi-criteria preference representation and reasoning
 33 Rory Sie (OUN), Coalitions in Cooperation Networks (COCOON)
 34 Pavol Jancura (RUN), Evolutionary analysis in PPI networks and applications
 35 Evert Haasdijk (VU), Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics
 36 Denis Ssebugwawo (RUN), Analysis and Evaluation of Collaborative Modeling Processes
 37 Agnes Nakakawa (RUN), A Collaboration Process for Enterprise Architecture Creation
 38 Selmar Smit (VU), Parameter Tuning and Scientific Testing in Evolutionary Algorithms
 39 Hassan Fatemi (UT), Risk-aware design of value and coordination networks

- 40 Agus Gunawan (UvT), Information Access for SMEs in Indonesia
 41 Sebastian Kelle (OU), Game Design Patterns for Learning
 42 Dominique Verpoorten (OU), Reflection Amplifiers in self-regulated Learning
 43 Withdrawn
 44 Anna Tordai (VU), On Combining Alignment Techniques
 45 Benedikt Kratz (UvT), A Model and Language for Business-aware Transactions
 46 Simon Carter (UVA), Exploration and Exploitation of Multilingual Data for Statistical Machine Translation
 47 Manos Tsagkias (UVA), Mining Social Media: Tracking Content and Predicting Behavior
 48 Jorn Bakker (TUE), Handling Abrupt Changes in Evolving Time-series Data
 49 Michael Kaisers (UM), Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions
 50 Steven van Kervel (TUD), Ontology driven Enterprise Information Systems Engineering
 51 Jeroen de Jong (TUD), Heuristics in Dynamic Scheduling; a practical framework with a case study in elevator dispatching
-
- 2013 01 Viorel Milea (EUR), News Analytics for Financial Decision Support
 02 Erietta Liarou (CWI), MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing
 03 Szymon Klarman (VU), Reasoning with Contexts in Description Logics
 04 Chetan Yadati (TUD), Coordinating autonomous planning and scheduling
 05 Dulce Pumareja (UT), Groupware Requirements Evolutions Patterns
 06 Romulo Goncalves (CWI), The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience
 07 Giel van Lankveld (UvT), Quantifying Individual Player Differences
 08 Robbert-Jan Merk (VU), Making enemies: cognitive modeling for opponent agents in fighter pilot simulators
 09 Fabio Gori (RUN), Metagenomic Data Analysis: Computational Methods and Applications
 10 Jeewanie Jayasinghe Arachchige (UvT), A Unified Modeling Framework for Service Design.
 11 Evangelos Pournaras (TUD), Multi-level Reconfigurable Self-organization in Overlay Services
 12 Marian Razavian (VU), Knowledge-driven Migration to Services
 13 Mohammad Safiri (UT), Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly
 14 Jafar Tanha (UVA), Ensemble Approaches to Semi-Supervised Learning Learning
 15 Daniel Hennes (UM), Multiagent Learning - Dynamic Games and Applications
 16 Eric Kok (UU), Exploring the practical benefits of argumentation in multi-agent deliberation
 17 Koen Kok (VU), The PowerMatcher: Smart Coordination for the Smart Electricity Grid
 18 Jeroen Janssens (UvT), Outlier Selection and One-Class Classification
 19 Renze Steenhuisen (TUD), Coordinated Multi-Agent Planning and Scheduling
 20 Katja Hofmann (UvA), Fast and Reliable Online Learning to Rank for Information Retrieval
 21 Sander Wubben (UvT), Text-to-text generation by monolingual machine translation
 22 Tom Claassen (RUN), Causal Discovery and Logic
 23 Patricio de Alencar Silva (UvT), Value Activity Monitoring
 24 Haitham Bou Ammar (UM), Automated Transfer in Reinforcement Learning
 25 Agnieszka Anna Latoszek-Berendsen (UM), Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System
 26 Alireza Zarghami (UT), Architectural Support for Dynamic Homecare Service Provisioning
 27 Mohammad Huq (UT), Inference-based Framework Managing Data Provenance
 28 Frans van der Sluis (UT), When Complexity becomes Interesting: An Inquiry into the Information eXperience
 29 Iwan de Kok (UT), Listening Heads
 30 Joyce Nakatumba (TUE), Resource-Aware Business Process Management: Analysis and Support
 31 Dinh Khoa Nguyen (UvT), Blueprint Model and Language for Engineering Cloud Applications
 32 Kamakshi Rajagopal (OUN), Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development
 33 Qi Gao (TUD), User Modeling and Personalization in the Microblogging Sphere
 34 Kien Tjin-Kam-Jet (UT), Distributed Deep Web Search
 35 Abdallah El Ali (UvA), Minimal Mobile Human Computer Interaction
 36 Than Lam Hoang (TUE), Pattern Mining in Data Streams
 37 Dirk Börner (OUN), Ambient Learning Displays
 38 Eelco den Heijer (VU), Autonomous Evolutionary Art
 39 Joop de Jong (TUD), A Method for Enterprise Ontology based Design of Enterprise Information Systems
 40 Pim Nijssen (UM), Monte-Carlo Tree Search for Multi-Player Games

	41	Jochem Liem (UVA), Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning
	42	Léon Planken (TUD), Algorithms for Simple Temporal Reasoning
	43	Marc Bron (UVA), Exploration and Contextualization through Interaction and Concepts
2014	01	Nicola Barile (UU), Studies in Learning Monotone Models from Data
	02	Fiona Tulyano (RUN), Combining System Dynamics with a Domain Modeling Method
	03	Sergio Raul Duarte Torres (UT), Information Retrieval for Children: Search Behavior and Solutions
	04	Hanna Jochmann-Mannak (UT), Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation
	05	Jurriaan van Reijssen (UU), Knowledge Perspectives on Advancing Dynamic Capability
	06	Damian Tamburri (VU), Supporting Networked Software Development
	07	Arya Adriansyah (TUE), Aligning Observed and Modeled Behavior
	08	Samur Araujo (TUD), Data Integration over Distributed and Heterogeneous Data Endpoints
	09	Philip Jackson (UvT), Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language
	10	Ivan Salvador Razo Zapata (VU), Service Value Networks
	11	Janneke van der Zwaan (TUD), An Empathic Virtual Buddy for Social Support
	12	Willem van Willigen (VU), Look Ma, No Hands: Aspects of Autonomous Vehicle Control
	13	Arlette van Wissen (VU), Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains
	14	Yangyang Shi (TUD), Language Models With Meta-information
	15	Natalya Mogles (VU), Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare
	16	Krystyna Milian (VU), Supporting trial recruitment and design by automatically interpreting eligibility criteria
	17	Kathrin Dentler (VU), Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability
	18	Mattijs Ghijsen (UVA), Methods and Models for the Design and Study of Dynamic Agent Organizations
	19	Vinicius Ramos (TUE), Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support
	20	Mena Habib (UT), Named Entity Extraction and Disambiguation for Informal Text: The Missing Link
	21	Kassidy Clark (TUD), Negotiation and Monitoring in Open Environments
	22	Marieke Peeters (UU), Personalized Educational Games - Developing agent-supported scenario-based training
	23	Eleftherios Sidiourgos (UvA/CWI), Space Efficient Indexes for the Big Data Era
	24	Daive Ceolin (VU), Trusting Semi-structured Web Data
	25	Martijn Lappenschaar (RUN), New network models for the analysis of disease interaction
	26	Tim Baarslag (TUD), What to Bid and When to Stop
	27	Rui Jorge Almeida (EUR), Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty
	28	Anna Chmielowiec (VU), Decentralized k-Clique Matching
	29	Jaap Kabbedijk (UU), Variability in Multi-Tenant Enterprise Software
	30	Peter de Cock (UvT), Anticipating Criminal Behaviour
	31	Leo van Moergestel (UU), Agent Technology in Agile Multiparallel Manufacturing and Product Support
	32	Naser Ayat (UvA), On Entity Resolution in Probabilistic Data
	33	Tesfa Tegegne (RUN), Service Discovery in eHealth
	34	Christina Manteli (VU), The Effect of Governance in Global Software Development: Analyzing Trans-active Memory Systems.
	35	Joost van Ooijen (UU), Cognitive Agents in Virtual Worlds: A Middleware Design Approach
	36	Joos Buijs (TUE), Flexible Evolutionary Algorithms for Mining Structured Process Models
	37	Maral Dadvar (UT), Experts and Machines United Against Cyberbullying
	38	Danny Plass-Oude Bos (UT), Making brain-computer interfaces better: improving usability through post-processing.
	39	Jasmina Maric (UvT), Web Communities, Immigration, and Social Capital
	40	Walter Omona (RUN), A Framework for Knowledge Management Using ICT in Higher Education
	41	Frederic Hogenboom (EUR), Automated Detection of Financial Events in News Text
	42	Carsten Eijckhof (CWI/TUD), Contextual Multidimensional Relevance Models
	43	Kevin Vlaanderen (UU), Supporting Process Improvement using Method Increments
	44	Paulien Meesters (UvT), Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.
	45	Birgit Schmitz (OUN), Mobile Games for Learning: A Pattern-Based Approach
	46	Ke Tao (TUD), Social Web Data Analytics: Relevance, Redundancy, Diversity
	47	Shangsong Liang (UVA), Fusion and Diversification in Information Retrieval
2015	01	Niels Netten (UvA), Machine Learning for Relevance of Information in Crisis Response

- 02 Faiza Bukhsh (UvT), Smart auditing: Innovative Compliance Checking in Customs Controls
 - 03 Twan van Laarhoven (RUN), Machine learning for network data
 - 04 Howard Spoelstra (OUN), Collaborations in Open Learning Environments
 - 05 Christoph Bösch (UT), Cryptographically Enforced Search Pattern Hiding
 - 06 Farideh Heidari (TUD), Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes
 - 07 Maria-Hendrike Peetz (UvA), Time-Aware Online Reputation Analysis
 - 08 Jie Jiang (TUD), Organizational Compliance: An agent-based model for designing and evaluating organizational interactions
 - 09 Randy Klaassen (UT), HCI Perspectives on Behavior Change Support Systems
 - 10 Henry Hermans (OUN), OpenU: design of an integrated system to support lifelong learning
 - 11 Yongming Luo (TUE), Designing algorithms for big graph datasets: A study of computing bisimulation and joins
 - 12 Julie M. Birkholz (VU), Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks
 - 13 Giuseppe Procaccianti (VU), Energy-Efficient Software
 - 14 Bart van Straalen (UT), A cognitive approach to modeling bad news conversations
 - 15 Klaas Andries de Graaf (VU), Ontology-based Software Architecture Documentation
 - 16 Changyun Wei (UT), Cognitive Coordination for Cooperative Multi-Robot Teamwork
 - 17 André van Cleeff (UT), Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs
 - 18 Holger Pirk (CWI), Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories
 - 19 Bernardo Tabuenca (OUN), Ubiquitous Technology for Lifelong Learners
 - 20 Lois Vanhée (UU), Using Culture and Values to Support Flexible Coordination
 - 21 Sibren Fetter (OUN), Using Peer-Support to Expand and Stabilize Online Learning
 - 22 Zheming Zhu (UT), Co-occurrence Rate Networks
 - 23 Luit Gazendam (VU), Cataloguer Support in Cultural Heritage
 - 24 Richard Berendsen (UVA), Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation
 - 25 Steven Woudenberg (UU), Bayesian Tools for Early Disease Detection
 - 26 Alexander Hogenboom (EUR), Sentiment Analysis of Text Guided by Semantics and Structure
 - 27 Sándor Héman (CWI), Updating compressed column stores
 - 28 Janet Bagorogoza (TiU), Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO
 - 29 Hendrik Baier (UM), Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains
 - 30 Kiavash Bahreini (OU), Real-time Multimodal Emotion Recognition in E-Learning
 - 31 Yakup Koç (TUD), On the robustness of Power Grids
 - 32 Jerome Gard (UL), Corporate Venture Management in SMEs
 - 33 Frederik Schadd (TUD), Ontology Mapping with Auxiliary Resources
 - 34 Victor de Graaf (UT), Gesocial Recommender Systems
 - 35 Jungxao Xu (TUD), Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction
-
- 2016
- 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
 - 02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
 - 03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
 - 04 Laurens Rietveld (VU), Publishing and Consuming Linked Data
 - 05 Evgeny Sherkhonov (UVA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
 - 06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
 - 07 Jeroen de Man (VU), Measuring and modeling negative emotions for virtual training
 - 08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
 - 09 Archana Nottamkandath (VU), Trusting Crowdsourced Information on Cultural Artefacts
 - 10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
 - 11 Anne Schuth (UVA), Search Engines that Learn from Their Users
 - 12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
 - 13 Nana Baah Gyan (VU), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
 - 14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
 - 15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
 - 16 Guangliang Li (UVA), Socially Intelligent Autonomous Agents that Learn from Human Reward
 - 17 Berend Weel (VU), Towards Embodied Evolution of Robot Organisms
 - 18 Albert Meroño Peñuela (VU), Refining Statistical Data on the Web

- 19 Julia Efremova (Tu/e), Mining Social Structures from Genealogical Data
 - 20 Daan Odijk (UVA), Context & Semantics in News & Web Search
 - 21 Alejandro Moreno Céleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
 - 22 Grace Lewis (VU), Software Architecture Strategies for Cyber-Foraging Systems
 - 23 Fei Cai (UVA), Query Auto Completion in Information Retrieval
 - 24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
 - 25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior
 - 26 Dilhan Thilakarathne (VU), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
 - 27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
 - 28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
 - 29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning
 - 30 Ruud Mattheij (UvT), The Eyes Have It
 - 31 Mohammad Khelghati (UT), Deep web content monitoring
 - 32 Eelco Vriezolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
 - 33 Peter Bloem (UVA), Single Sample Statistics, exercises in learning from just one example
 - 34 Dennis Schunselaar (TUE), Configurable Process Trees: Elicitation, Analysis, and Enactment
 - 35 Zhaochun Ren (UVA), Monitoring Social Media: Summarization, Classification and Recommendation
 - 36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies
 - 37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
 - 38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
 - 39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
 - 40 Christian Detweiler (TUD), Accounting for Values in Design
 - 41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
 - 42 Spyros Martzoukos (UVA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
 - 43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
 - 44 Thibault Sellam (UVA), Automatic Assistants for Database Exploration
 - 45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
 - 46 Jorge Gallego Perez (UT), Robots to Make you Happy
 - 47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
 - 48 Tanja Buttler (TUD), Collecting Lessons Learned
 - 49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
 - 50 Yan Wang (UVT), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains
-
- 2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
 - 02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
 - 03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
 - 04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
 - 05 Mahdiah Shadi (UVA), Collaboration Behavior
 - 06 Damir Vandic (EUR), Intelligent Information Systems for Web Product Search
 - 07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
 - 08 Rob Konijn (VU), Detecting Interesting Differences: Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
 - 09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
 - 10 Robby van Delden (UT), (Steering) Interactive Play Behavior
 - 11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter #anticipointment
 - 12 Sander Leemans (TUE), Robust Process Mining with Guarantees
 - 13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
 - 14 Shoshannah Tekofsky (UvT), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
 - 15 Peter Berck (RUN), Memory-Based Text Correction
 - 16 Aleksandr Chuklin (UVA), Understanding and Modeling Users of Modern Search Engines
 - 17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution
 - 18 Ridho Reinanda (UVA), Entity Associations for Search

- 19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
 - 20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
 - 21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)
 - 22 Sara Magliacane (VU), Logics for causal inference under uncertainty
 - 23 David Graus (UVA), Entities of Interest — Discovery in Digital Traces
 - 24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning
 - 25 Veruska Zamborlini (VU), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
 - 26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
 - 27 Michiel Joosse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
 - 28 John Klein (VU), Architecture Practices for Complex Contexts
 - 29 Adel Alhuraibi (UvT), From IT-BusinessStrategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
 - 30 Wilma Latuny (UvT), The Power of Facial Expressions
 - 31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
 - 32 Thaer Samar (RUN), Access to and Retrievability of Content in Web Archives
 - 33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
 - 34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
 - 35 Martine de Vos (VU), Interpreting natural science spreadsheets
 - 36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
 - 37 Alejandro Montes Garcia (TUE), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
 - 38 Alex Kayal (TUD), Normative Social Applications
 - 39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR
 - 40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems
 - 41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle
 - 42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets
 - 43 Maaike de Boer (RUN), Semantic Mapping in Video Retrieval
 - 44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering
 - 45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement
 - 46 Jan Schneider (OU), Sensor-based Learning Support
 - 47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration
 - 48 Angel Suarez (OU), Collaborative inquiry-based learning
-
- 2018
- 01 Han van der Aa (VUA), Comparing and Aligning Process Representations
 - 02 Felix Mannhardt (TUE), Multi-perspective Process Mining
 - 03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction
 - 04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks
 - 05 Hugo Huurdeman (UVA), Supporting the Complex Dynamics of the Information Seeking Process
 - 06 Dan Ionita (UT), Model-Driven Information Security Risk Assessment of Socio-Technical Systems
 - 07 Jieting Luo (UU), A formal account of opportunism in multi-agent systems
 - 08 Rick Smetsers (RUN), Advances in Model Learning for Software Systems
 - 09 Xu Xie (TUD), Data Assimilation in Discrete Event Simulations
 - 10 Julienka Mollee (VUA), Moving forward: supporting physical activity behavior change through intelligent technology
 - 11 Mahdi Sargolzaei (UVA), Enabling Framework for Service-oriented Collaborative Networks
 - 12 Xixi Lu (TUE), Using behavioral context in process mining
 - 13 Seyed Amin Tabatabaei (VUA), Computing a Sustainable Future
 - 14 Bart Joosten (UvT), Detecting Social Signals with Spatiotemporal Gabor Filters
 - 15 Naser Davarzani (UM), Biomarker discovery in heart failure
 - 16 Jaebok Kim (UT), Automatic recognition of engagement and emotion in a group of children
 - 17 Jianpeng Zhang (TUE), On Graph Sample Clustering
 - 18 Henriette Nakad (UL), De Notaris en Private Rechtspraak
 - 19 Minh Duc Pham (VUA), Emergent relational schemas for RDF
 - 20 Manxia Liu (RUN), Time and Bayesian Networks

	21	Aad Sloomaker (OUN), EMERGO: a generic platform for authoring and playing scenario-based serious games
	22	Eric Fernandes de Mello Araujo (VUA), Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks
	23	Kim Schouten (EUR), Semantics-driven Aspect-Based Sentiment Analysis
	24	Jered Vroon (UT), Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots
	25	Riste Gligorov (VUA), Serious Games in Audio-Visual Collections
	26	Roelof Anne Jelle de Vries (UT), Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology
	27	Maikel Leemans (TUE), Hierarchical Process Mining for Scalable Software Analysis
	28	Christian Willemse (UT), Social Touch Technologies: How they feel and how they make you feel
2019	01	Rob van Eijk (UL), Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification
	02	Emmanuelle Beauxis Aussalet (CWI, UU), Statistics and Visualizations for Assessing Class Size Uncertainty
	03	Eduardo Gonzalez Lopez de Murillas (TUE), Process Mining on Databases: Extracting Event Data from Real Life Data Sources
	04	Ridho Rahmadi (RUN), Finding stable causal structures from clinical data
	05	Sebastiaan van Zelst (TUE), Process Mining with Streaming Data
	06	Chris Dijkshoorn (VU), Nichesourcing for Improving Access to Linked Cultural Heritage Datasets
	07	Soude Fazeli (TUD), Recommender Systems in Social Learning Platforms
	08	Frits de Nijs (TUD), Resource-constrained Multi-agent Markov Decision Processes
	09	Fahimeh Alizadeh Moghaddam (UVA), Self-adaptation for energy efficiency in software systems
	10	Qing Chuan Ye (EUR), Multi-objective Optimization Methods for Allocation and Prediction
	11	Yue Zhao (TUD), Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs
	12	Jacqueline Heinerman (VU), Better Together
	13	Guanliang Chen (TUD), MOOC Analytics: Learner Modeling and Content Generation
	14	Daniel Davis (TUD), Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses
	15	Erwin Walraven (TUD), Planning under Uncertainty in Constrained and Partially Observable Environments
	16	Guangming Li (TUE), Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models
	17	Ali Hurriyetoglu (RUN), Extracting actionable information from microtexts
	18	Gerard Wagenaar (UU), Artefacts in Agile Team Communication
	19	Vincent Koeman (TUD), Tools for Developing Cognitive Agents
	20	Chide Groenouwe (UU), Fostering technically augmented human collective intelligence
	21	Cong Liu (TUE), Software Data Analytics: Architectural Model Discovery and Design Pattern Detection
	22	Martin van den Berg (VU), Improving IT Decisions with Enterprise Architecture
	23	Qin Liu (TUD), Intelligent Control Systems: Learning, Interpreting, Verification
	24	Anca Dumitrache (VU), Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing
	25	Emiel van Miltenburg (VU), Pragmatic factors in (automatic) image description
	26	Prince Singh (UT), An Integration Platform for Synchromodal Transport
	27	Alessandra Antonaci (OUN), The Gamification Design Process applied to (Massive) Open Online Courses
	28	Esther Kuindersma (UL), Cleared for take-off: Game-based learning to prepare airline pilots for critical situations
	29	Daniel Formolo (VU), Using virtual agents for simulation and training of social skills in safety-critical circumstances
	30	Vahid Yazdanpanah (UT), Multiagent Industrial Symbiosis Systems
	31	Milan Jelisavcic (VU), Alive and Kicking: Baby Steps in Robotics
	32	Chiara Sironi (UM), Monte-Carlo Tree Search for Artificial General Intelligence in Games
	33	Anil Yaman (TUE), Evolution of Biologically Inspired Learning in Artificial Neural Networks
	34	Negar Ahmadi (TUE), EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES
	35	Lisa Facey-Shaw (OUN), Gamification with digital badges in learning programming
	36	Kevin Ackermans (OUN), Designing Video-Enhanced Rubrics to Master Complex Skills
	37	Jian Fang (TUD), Database Acceleration on FPGAs
	38	Akos Kadar (OUN), Learning visually grounded and multilingual representations
2020	01	Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
	02	Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
	03	Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding
	04	Maarten van Gompel (RUN), Context as Linguistic Bridges

05	Yulong Pei (TUE), On local and global structure mining
06	Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
07	Wim van der Vegt (OUN), Towards a software architecture for reusable game components
08	Ali Mirsoleimani (UL), Structured Parallel Programming for Monte Carlo Tree Search
09	Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
10	Alifah Syamsiyah (TUE), In-database Preprocessing for Process Mining
11	Sepideh Mesbah (TUD), Semantic-Enhanced Training Data Augmentation Methods for Long-Tail Entity Recognition Models
12	Ward van Breda (VU), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
13	Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
14	Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
15	Konstantinos Georgiadis (OUN), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
16	Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
17	Daniele Di Mitri (OUN), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences
18	Georgios Methenitis (TUD), Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems
19	Guido van Capelleveen (UT), Industrial Symbiosis Recommender Systems
20	Albert Hankel (VU), Embedding Green ICT Maturity in Organisations
21	Karine da Silva Miras de Araujo (VU), Where is the robot?: Life as it could be
22	Maryam Masoud Khamis (RUN), Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar
23	Rianne Conijn (UT), The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging
24	Lenin da Nobrega Medeiros (VUA/RUN), How are you feeling, human? Towards emotionally supportive chatbots
25	Xin Du (TUE), The Uncertainty in Exceptional Model Mining
26	Krzysztof Leszek Sadowski (UU), GAMBIT: Genetic Algorithm for Model-Based mixed-Integer optimization
27	Ekaterina Muravyeva (TUD), Personal data and informed consent in an educational context
28	Bibeg Limbu (TUD), Multimodal interaction for deliberate practice: Training complex skills with augmented reality
29	Ioan Gabriel Bucur (RUN), Being Bayesian about Causal Inference
30	Bob Zadok Blok (UL), Creatief, Creatieve, Creatiefst
31	Gongjin Lan (VU), Learning better – From Baby to Better
32	Jason Rhuggenaath (TUE), Revenue management in online markets: pricing and online advertising
33	Rick Gilsing (TUE), Supporting service-dominant business model evaluation in the context of business model innovation
34	Anna Bon (MU), Intervention or Collaboration? Redesigning Information and Communication Technologies for Development
35	Siamak Farshidi (UU), Multi-Criteria Decision-Making in Software Production
<hr/>	
2021	01 Francisco Xavier Dos Santos Fonseca (TUD), Location-based Games for Social Interaction in Public Space
	02 Rijk Mercuur (TUD), Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models
	03 Seyyed Hadi Hashemi (UVA), Modeling Users Interacting with Smart Devices
	04 Ioana Jivet (OU), The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning
	05 Davide Dell'Anna (UU), Data-Driven Supervision of Autonomous Systems
	06 Daniel Davison (UT), "Hey robot, what do you think?" How children learn with a social robot
	07 Armel Lefebvre (UU), Research data management for open science
	08 Nardie Fanchamps (OU), The Influence of Sense-Reason-Act Programming on Computational Thinking
	09 Cristina Zaga (UT), The Design of Robothings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play
	10 Quinten Meertens (UvA), Misclassification Bias in Statistical Learning
	11 Anne van Rossum (UL), Nonparametric Bayesian Methods in Robotic Vision
	12 Lei Pi (UL), External Knowledge Absorption in Chinese SMEs

This dissertation explores how we can design robot-assisted interventions for autistic children in such a way that it might facilitate engagement. One aspect to this is taking the autistic children's individual differences into account. Another aspect that needs to be considered is the predictability of a robot, which is a commonly used argument for why robots are promising tools for those working with autistic children in the first place.

The work described in this dissertation is a step towards better understanding the concept of predictability and its effects on (autistic) human-robot interaction, as well as how we can design robot-assisted intervention for autistic children that may sustain engagement in learning.