

Responsive Social Agents

Feedback-sensitive behavior generation for social interactions

Jered Vroon, Gwenn Englebienne, and Vanessa Evers

Human Media Interaction, University of Twente, 7500AE Enschede, the Netherlands
{j.h.vroon, g.englebienne, v.evers}@utwente.nl

Abstract. How can we generate appropriate behavior for social artificial agents? A common approach is to (1) establish with controlled experiments which action is most appropriate in which setting, and (2) select actions based on this knowledge and an estimate of the setting. This approach faces challenges, as it can be very hard to acquire and reason with all the required knowledge. Estimating the setting is challenging too, as many relevant aspects of the setting (e.g. personality of the interactee) can be unobservable. We formally describe an alternative approach that can handle these challenges; **responsiveness**. This is the idea that a social agent can utilize the many feedback cues given in social interactions to continuously adapt its behavior to something more appropriate. We theoretically discuss the relative advantages and disadvantages of these two approaches, which allows for more explicitly considering their application in social agents.

Keywords: Control architectures, Social robotics, Feedback

1 Introduction

Robotic and other artificial agents are increasingly often being deployed in settings where they have to interact with humans in a socially appropriate way. From telepresence robots to educational agents; they all interact with humans to serve their purpose.

How to generate socially appropriate behavior for such agents? This question involves all behaviors such agents can show, from how they position themselves to the sounds they use. This paper theoretically discusses approaches to answering this question, using social positioning (proxemics) for mobile agents as a running example.

Commonly, socially appropriate behavior for artificial agents is investigated with psychological experiments measuring the effect of particular conditions in interactions between social agents and participants. Ideally, this results in the generalized knowledge that within a particular setting, a particular behavior is more appropriate.

Behaviors can be generated based on this generalized knowledge, by first estimating the current setting and then using the generalized knowledge to select the appropriate behavior for that setting. We will refer to this as the **setting-specific approach**. For example, in social distancing for mobile agents, it is common to derive appropriate distances from a combination of factors, ranging from size and human-likeness of the agent [12] to experience with pets and robots of the interactee [9].

Such a setting-specific approach faces several practical challenges. Firstly, the generalized knowledge required to select the appropriate action can involve a complex

interplay of many different variables, which makes it hard to acquire and reason with. Secondly, many relevant aspects of the setting can be hard or impossible to observe, making estimating the current setting into a very challenging task. To continue with our social distancing example; hearing problems may well influence what is an appropriate interaction distance, but may be impossible to detect beforehand.

Fortunately, interactions with humans provide extra information that could help overcome these challenges: feedback. Feedback can be anything from asking someone not to speak too loud, or cupping a hand to your ear to indicate hearing problems, to taking a step back if someone gets too close (e.g. [4]).

In this paper we discuss the idea that agents can generate social behavior by being **responsive** to these feedback cues. Such agents could try a behavior to get started and then continuously adapt it to something more appropriate based on the feedback cues they recognize. Responsiveness would thus provide a pathway to finding the appropriate behavior that does not rely on or assume knowing all relevant aspects of the setting.

We will theoretically discuss the setting-specific and the responsive approach to generating social behavior, by formally defining both (Section 2), discussing the challenges faced by setting-specific approaches (Section 3), and how responsiveness can (partly) resolve these (Section 4). Though responsiveness may seem straightforward, it is not commonly used in social agents; for example, even though using responsiveness may well be suitable for doing social distancing, we are not aware of any existing artificial agents doing so (Section 5). With our specification of responsiveness, we aim to contribute to the development of social agents by allowing people to more explicitly consider its application, limitations, and opportunities (Section 6).

2 Terminology

In this section, we formally define the setting-specific approach and the responsive approach. We start with the basic building blocks (2.1) with which we define agents and interaction (2.2) and discuss what makes behavior “appropriate” (2.3). We then define the two approaches (2.4). Symbolic representations (building on our earlier work [10]) will be introduced solely to make the relations between the terms more explicit.

2.1 Variables, time spans, and value assignments

We will treat agents as entities, roughly separable from the setting in which they exist, that gather observations and produce actions based on those. The state of the setting causes these observations and can in turn be influenced by these actions, allowing for interactions (Fig. 1). Actions, observations and the state of the setting will all be formalized as *value assignments* over a particular *time span* to a set of *variables*.

Variables We will be talking about **variables** (denoted by v), which can be among others **input variables** (v^I), **output variables** (v^O) or **setting variables** (v^S). Each variable v^I has a **domain** (D_{v^I}), which is the set of values that variable v^I can take. A **variable set** (V) is a set of variables, each of which can have a different domain. A variable set containing only input variables, output variables or setting variables is, respectively, an **input set** (V^I), **output set** (V^O) or **setting** (V^S).

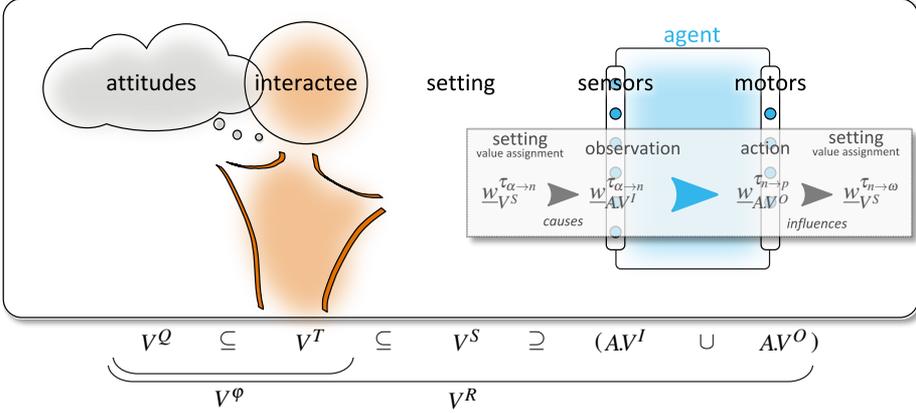


Fig. 1. Overview of the terminology involved in the relationship between an agent and the setting in which it exists. The goal of the agent is a mapping from observations ($w_{AV^I}^{\tau_{\alpha \rightarrow n}}$) to actions ($w_{AV^O}^{\tau_{n \rightarrow p}}$) such that the social appropriateness of those actions during the interaction ($\underline{SA}_{V^Q'}(w_{AV^O}^{\tau_{\alpha \rightarrow \omega}})$) is optimal, sufficient, or improving.

Time spans In addition to these, we will use the variable **time** (denoted by t). Its domain (D_t) is a totally ordered set of values, representing a series of successive moments in time. τ_α indicates the first moment of an interaction, τ_ω the last. Moments in between will be indicated with letters such that alphabetical ordering indicates succession, e.g. τ_q comes before τ_r . A **time span** ($\tau_{m \rightarrow n}$) between two moments ($\{\tau_m, \tau_n\} \in D_t, \tau_m \leq \tau_n$) is the complete subset of successive moments in time between them ($\tau_{m \rightarrow n} = \{x \mid x \in D_t, x \geq \tau_m \wedge x \leq \tau_n\}$). Implementations may rely, without loss of generality, on discretised time or event-based observation.

Value assignments A **single value assignment** for a variable v' and a moment in time τ_o (denoted by $w_{v'}^{\tau_o}$) is defined as a function that returns the value of that variable at that moment in time ($w_{v'}^{\tau_o} : v' \mapsto D_{v'}$). We also define a **value assignment** for a set of variables V' and a time span $\tau_{m \rightarrow n}$ (denoted by $w_{V'}^{\tau_{m \rightarrow n}}$), to give the single value assignments for all variables in that variable set and all moments in that time span. The **value assignment set** for a variable set V' and a time span $\tau_{m \rightarrow n}$ (denoted by $\underline{W}_{V'}^{\tau_{m \rightarrow n}}$) is the set of all possible value assignments for that V' and $\tau_{m \rightarrow n}$ ($\underline{W}_{V'}^{\tau_{m \rightarrow n}} = \{w_{V'}^{\tau_{m \rightarrow n}} \mid \cdot\}$).

2.2 Agents and interaction

An (artificial) **agent** (denoted by A) has **sensors** (an input set, AV^I), **actuators** (an output set, AV^O) and “inner workings” to connect those. It produces **actions** (value assignments for its actuators, $w_{AV^O}^{\tau_{m \rightarrow n}}$) that are affected by its **observations** (value assignments for its sensors, $w_{AV^I}^{\tau_{m \rightarrow n}}$). We use (partial specifications of) settings as theoretical constructs to discuss the environment in which an agent exists. The actions of an agent influence the setting to some extent. Likewise, to some extent, the observations of an

agent reflect the setting, based on which the agent can **estimate** it (a value o being estimated is denoted by o^E). The more reliably a value can be estimated by an agent in practice, the more **estimable** it will be said to be.

We will refer to the (human) other agents with which the agent is interacting in the setting as **interactees** (denoted by V^T). These are part of the setting ($V^T \subseteq V^S$).

2.3 Appropriate behavior

Central in deciding if the behavior of an agent in an interaction is socially appropriate, are the **attitudes** of the interactee(s) (denoted by V^Q), loosely defined as a subset of the variables used to express interactees and their properties ($V^Q \subseteq V^T \subseteq V^S$). Attitudes can range from, for example, comfort to perception of the agent as intelligent or sensitive.

The actions of an agent to some extent influence the setting, which can include the attitudes of the involved interactees. Depending on the goals of the agent, different attitudes can be more or less desirable; for example, an agent may want to avoid selecting actions that make the interactee more uncomfortable. We define the **social appropriateness** function (denoted by \underline{SA}_{V^Q}), for a set of attitudes V^Q and a setting during an interaction, that for all possible actions returns a numerical value, such that a higher value indicates that that action would lead to a more ‘desirable’ value for those attitudes ($\underline{SA}_{V^Q} : \underline{W}_{AVO}^{\tau_{m \rightarrow n}} \mapsto \mathbb{R}$). As with the setting, we use this function as a theoretical construct for discussion purposes; an agent can at best estimate it.

2.4 Approaches to finding socially appropriate behavior

How can an agent select actions such that their social appropriateness is optimal, sufficient, or at least improving? We here define two approaches, both of which focus on the strategy used to find socially appropriate behavior for an agent, not on the actual implementation of these steps. Different ways of generating behavior, e.g. static, scripted, dynamic, adaptive, might thus all be used to implement either of the two approaches.

Setting-specific approach The setting-specific approach depends on prior knowledge about how the social appropriateness of different actions is dependent on the values for particular variables in the setting. We therefore define the **knowledge** function (denoted by \underline{K}) that, for all value assignments to (a subset of) the setting $\underline{W}_{VS}^{\tau_{\alpha \rightarrow n}}$, it returns the most appropriate action ($\underline{K} : \underline{W}_{VS}^{\tau_{\alpha \rightarrow n}} \mapsto \underline{W}_{AVO}^{\tau_{n \rightarrow p}}$).

We define **relevant setting variables** (denoted by V^R) as a subset of the variables in the setting ($V^R \subseteq V^S$) such that their values contain all information required to distinguish between setting value assignments where \underline{K} should give different outcomes. A knowledge function that uses at least the relevant setting variables should thus have

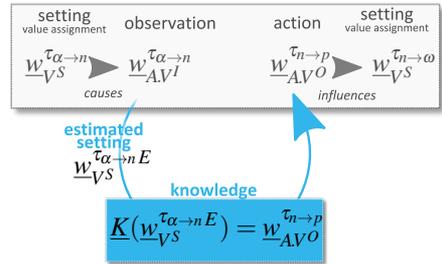


Fig. 2. Steps taken in a setting-specific approach.

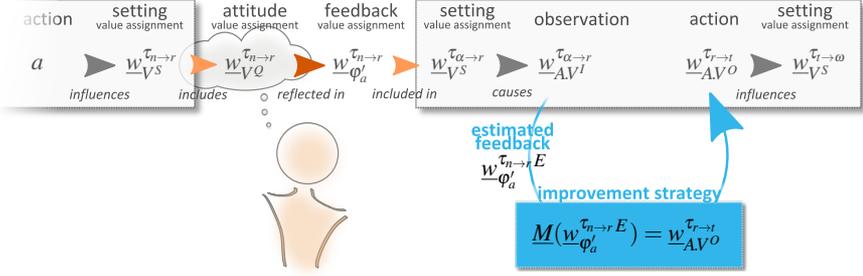


Fig. 3. Steps taken in a responsive approach. Includes an overview of how feedback variables can be assigned a value in response to an earlier action a of the agent.

enough information to select the most appropriate action. Such a knowledge function is the ideal, as in practice approximations often have to be used instead (Section 3).

From the knowledge, the setting-specific approach works in two steps to produce an action based on observations (Fig. 2). First, the available observations are used to estimate value assignments to (a subset of) the setting. Second, these estimates are used with the knowledge function to try and select the best action. If the knowledge is approximated, or if the relevant setting variables are not fully estimable, this may result in the best *known* action, rather than in the best action.

Responsive approach Central to the responsive approach is feedback; any action a of the agent influence the attitudes of the interactee, which in turn can be reflected by **feedback variables** (denoted by ϕ'_a). Feedback variables provide information about the underlying appropriateness of a previous action, e.g. if it was optimal/sufficient (**basic feedback**), how it compares to other earlier actions (**comparable feedback**), or even which actions would be more/less suitable (**directional feedback**). The **feedback set** (denoted by V^φ) is the set of all available feedback variables.

Feedback variables encode the information about the underlying appropriateness; different feedback variables can code (partially) overlapping information and, importantly, the encoding may be flawed. The greater the certainty with which the underlying appropriateness can be derived from a feedback variable, the more **legible** we will say it to be. Feedback variables can be less legible because they reflect things besides the underlying appropriateness, or because they differ between interactees.

The responsive approach works in two steps (Fig. 3). First, the feedback variables are estimated from the available observations, and interpreted as relating to particular previous actions of the agent. Second, this estimated feedback is used to adapt the subsequent actions of the agent. For this, we define an **improvement strategy** (denoted by M) as a function that, based on all available feedback on previous actions $w_{V\varphi}^{\tau_{\alpha \rightarrow n}}$ returns a suggested action $w_{AVO}^{\tau_{n \rightarrow p}}$, such that, possibly after several iterations, the actions will be sufficient and/or improving ($M : W_{V\varphi}^{\tau_{\alpha \rightarrow n}} \mapsto W_{AVO}^{\tau_{n \rightarrow p}}$). For example, an improvement strategy using comparable feedback could be to try and select actions that are more dissimilar to actions with lower social appropriateness.

3 Implications and challenges for a setting-specific approach

Assuming that from the observations an agent can perfectly derive the relevant setting variables, and assuming that the agent has full knowledge of the optimal action for all settings, it is trivial to prove that a setting-specific approach will yield the optimal action. However these assumptions will likely hold only in constrained settings, which social interactions usually are not. This presents several challenges to the approach.

3.1 Estimating the required setting variables

A setting-specific approach depends on estimating the relevant setting variables, which is a challenging task; not only because the set of relevant setting variables can still be very large, despite being a subset of all setting variables, but also because many of these setting variables will only partially be observable, if at all. They may be internal to the interactee (e.g. personality traits), include cases that are hard to classify (e.g. gender and cultural background), or algorithms that reliably detect them may not exist (yet). In our running example; it would be challenging to evaluate someone’s hearing, but it could well influence the appropriate interaction distance [11]. Though this may partly be resolved by using a reasonable approximation in a limited setting, there is no guarantee that the behaviors based on such approximations would be sufficiently appropriate.

These challenges can **not** be overcome within a purely setting-specific approach. Not taking into account these practical limitations in detecting and estimating the relevant setting variables severely challenges the implementation of autonomous agents that actually use knowledge which depends on these setting variables.

3.2 The knowledge to select the best action

The knowledge required for a setting-specific approach is in practice usually approximated by a combination of findings from scientific studies. This allows for the design of agents that can effectively select a reasonable action based on a well-chosen selection of relevant setting variables. It also introduces several challenges.

Establishing which setting variables are relevant Establishing which setting variables are relevant can be challenging, given the sheer amount of setting variables in real-world settings and because it is hard to predict which variables will be relevant. Though in controlled experiments one can try to focus on specific setting variables, every aspect of the world could be a relevant setting variable. Thus, even listing all setting variables would be challenging, let alone investigating their relevance with scientific rigor.

Combinatorial explosion As the number of setting variables that have to be considered increases, so does the complexity of the knowledge function. If the different variables are dependent on each other, all combinations of those factors have to be considered to reliably derive the appropriate behavior, resulting in exponential growth¹.

¹ Even when limiting ourselves to ‘just’ the relevant setting variables (V^R) this would already be $\prod_{v \in V^R} |D_v|$ combinations (since $|D_v| \geq 2$ for all meaningful variables, this is at least $2^{|V^R|}$).

An implementation of such a knowledge function would thus quickly become intractable. This can partly be avoided by instead using approximations, though this would necessarily introduce uncertainty about the appropriateness of the selected action. The complexity could also be reduced by explicitly establishing which setting variables are independent of each other – but that is a challenging task itself.

In addition, this combinatorial explosion also poses a significant challenge to acquiring the required (prior) knowledge in a scientifically sound way; given the sheer number of combinations, it would be infeasible to test all combinations against each other in a controlled experiment. While approximations may be acceptable for implementations, they are less appropriate for scientific experiments.

Stereotyping by using generalized findings The knowledge function of an agent is commonly acquired through controlled experiments, which investigate how the effects of particular setting variables on particular attitudes could be generalized to a population.

When individual differences play a role in establishing the appropriate behavior, this can pose a challenge to a setting-specific approach. For example, an agent may well need to adapt its behavior when interacting with people who had a negative prior experience with similar agents.

To some extent, these individual differences can be handled by introducing them as setting variables. However, this would pose its own challenges if it introduces (partly) unobservable variables or results in a large increase in the number of variables.

4 Implications and challenges for a responsive approach

The responsive and setting-specific approach are both aimed at the same goal, though they use different steps. In this section we will discuss how a responsive approach could circumvent some of the challenges faced by a setting-specific approach, and vice versa.

4.1 Estimating the required setting variables

The setting-specific approach needs to estimate all relevant setting variables, whereas the responsive approach depends on a legible set of feedback variables. The more legible the feedback variables are, the more information they provide about the social appropriateness of previous actions (on a set of attitudes), and the less feedback variables a responsive approach will need. If feedback variables are available that are legible and estimable, a responsive approach can thus use these to avoid the aforementioned combinatorial explosion faced by a setting-specific approach.

Such a combination of legible and estimable feedback variables may actually be common, since there is an incentive for the interactee to provide them. For if the interactee provides legible and estimable feedback variables, a responsive agent, artificial or not, can use these to try and improve its behavior – which would benefit both the agent *and* the interactee. Using a responsive approach could thus turn finding socially appropriate actions into a collaborative effort. Therefore, the interactee may actively provide legible and estimable feedback variables, be it consciously and/or subconsciously.

4.2 The improvement strategy to select better actions

Another important difference between the responsive and the setting-specific approach is that the former uses an improvement strategy function instead of a (prior) knowledge function. This gives the responsive approach a reduced dependency on knowledge for all setting variables and allows for individualized instead of stereotyped adaptation.

Reduced dependency on all setting variables Since a responsive approach does not use prior knowledge (but only feedback), it avoids many of the challenges faced by a setting-specific approach, such as the combinatorial explosion and the challenges of establishing the prior knowledge. Only if the feedback variables would not be legible could similar challenges also arise for a responsive approach.

Individualized instead of stereotyped adaptation A responsive approach per definition uses the feedback given by individual interactees, rather than working from knowledge generalized to the population of interactants. Since feedback variables are individual, a responsive approach can be used to adapt to the individual preferences of interactees. Some feedback variables may even encode a combination of different attitudes prioritized based on the preferences of individual interactees.

This circumvents the stereotyping challenge faced by a setting-specific approach. It also shows that a purely responsive approach could easily miss out on the advantages of such stereotyping. Herein, the two approaches can complement each other. A setting-specific approach could be used to select an initial ‘stereotyped’ action, that can then be refined into more ‘personalized’ actions using a responsive approach.

Defining an improvement strategy The responsive approach depends on suitable improvement strategy functions. In contrast to the knowledge function of the setting-specific approach, an improvement strategy can be defined to deliberately use various aspects of the interaction. For example, an improvement strategy could be to directly ask the interactees for the desired actions. Furthermore, interactees might even appreciate the attempts of a responsive agent to try and improve the interaction, regardless of the appropriateness of the selected actions. While introducing such interesting options, this flexibility could also make it a challenge to create suitable improvement strategies.

4.3 Quality of the selected action

Where a setting-specific approach can ideally aim for selecting the most appropriate action, a responsive approach instead aims for improvement. Consequently, a responsive approach will be most suitable if the cost of selecting an inappropriate action is not too high and/or if no systems exists that reliably deliver the most appropriate action. In some cases, showing responsive behavior may actually *be* the most appropriate action.

5 Applications of responsiveness in social agents

There is a variety of existing work in artificial agents that we feel aligns with our definition of the responsive approach. Our aim here is not to give a complete overview, but

instead to illustrate how solutions fitting within the framework of a responsive approach exist and have been shown to be effective.

Most of the work on responsiveness in human-agent interaction focuses on agents that deliberately *provide* legible feedback variables, rather than being responsive themselves. This includes prior work using the term ‘robot responsiveness’, which primarily investigated different (dynamic) non-verbal feedback behaviors a robot could use when listening to an interactee – showing various positive effects of giving the appropriate feedback behaviors [2,3]. Jung *et al.* looked at the effects of robots using backchanneling on human-robot teamwork and found both improved team functioning and decreased perceived competence [5]. In the field of our running example, it has been found that people adapt their proxemic preferences when interacting with an agent that provided (feedback) information on its effectiveness at different interaction distances [6].

To our knowledge, there is no work on social positioning, our running example, with artificial agents using a responsive approach, even though various feedback variables may be available (e.g. [4,7]). In fact, a large part of the work on social positioning in human-human interaction seems to be strongly in line with the responsive approach (see e.g. the extensive review by Aiello [1]). We have previously conducted two small studies in this direction, that we will briefly discuss here. Both had a limited sample size and used a Wizard of Oz. In one of them, we set up a conversation such that hearing problems were to be expected and then had the robot use one of two different improvement strategies once certain feedback variables were observed [11]. In the other, we compared conditions in which a robot either did (1) an approach without personal space invasion, (2) an approach with personal space invasion, or (3) a personal space invasion after which it backed up and apologized [8]. The results of both studies suggest that participants appreciate the responsive behavior, perhaps even over directly picking the ‘improved’ action.

6 Discussion

We have given formal definitions of both the responsive and the setting-specific approach. Though in theory capable of finding the optimally appropriate behavior, the setting-specific approach ideally requires the agent to estimate and reason with all relevant setting variables – which may well be infeasible in realistic settings. We showed that the responsive approach can be used to (partly) circumvent these challenges, as it instead looks for behavior that is sufficiently appropriate or improves appropriateness.

Our theoretical discussion of the responsive approach is only a rough starting point for implementations. Since both responsiveness and (online) reinforcement learning need to adapt to feedback, insights from the latter could be used to guide such implementations – though with the responsive approach the adapting is explicitly part of the social dynamic, rather than finite learning. Another challenge will be the social signal processing necessary to detect feedback variables. More so because the expectations one has from an (artificial) agent may influence which feedback variables are used.

If suitable implementations can be created, explicitly considering a responsive approach can offer various opportunities. One such opportunity is to complement a setting-specific with a responsive approach. Another opportunity would be to use responsive-

ness in a more pro-active way, for example by directly asking interactees which actions they would prefer. Further opportunities can be found in the improvement strategy, e.g.; (a) with intelligent reasoning about why the agent got particular feedback, it may be able to respond to it more appropriately, or (b) giving responsive agents different personalities by parametrizing the different factors weighed by the agent when adapting to feedback, such as its own needs and those of the interactee.

Overall, we have introduced an explicit definition of responsiveness, and argued for the potential value of the approach. We hope and expect that this can help to explicitly consider its application in (artificial) social agents, not necessarily as a replacement of the setting-specific approach, but as a potentially valuable addition.

Acknowledgements The work described in this paper has partly been supported by the European Commission under contract number FP7-ICT-611153 (TERESA).

We are grateful for the critical and open-minded comments of Khiat Truong, Dennis Reidsma, Daniel Davison, Bob Schadenberg, Jan Kolkmeier, Michiel Joosse, Roelof de Vries, Jorge Gallego Pérez, Jeroen Linssen and Dirk Heylen.

References

1. Aiello, J.R.: Human spatial behavior. In: Stokols, D., Altman, I. (eds.) *Handbook of environmental psychology*, ch. 12, pp. 389–504. Wiley, New York (1987)
2. Birnbaum, G.E., Mizrahi, M., Hoffman, G., Reis, H.T., Finkel, E.J., Sass, O.: Machines as a source of consolation: Robot responsiveness increases human approach behavior and desire for companionship. In: *10th ACM/IEEE Int. Conf. on Human Robot Interaction*, pp. 165–171. IEEE Press (2016)
3. Bretan, M., Hoffman, G., Weinberg, G.: Emotionally expressive dynamic physical behaviors in robots. *Int. Journal of Human-Computer Studies*, vol. 78, pp. 1–16 (2015)
4. Cappella, J.N.: Mutual influence in expressive behavior: Adult-adult and infant-adult dyadic interaction. *Psychological bulletin*, vol. 89(1), pp. 101–132 (1981)
5. Jung, M.F., Lee, J.J., DePalma, N., Adalgeirsson, S.O., Hinds, P.J., Breazeal, C.: Engaging robots: easing complex human-robot teamwork using backchanneling. In: *Proc. of the 2013 Conf. on Computer supported cooperative work*, pp. 1555–1566. ACM (2013)
6. Mead, R., Matarić, M.J.: Robots have needs too: People adapt their proxemic preferences to improve autonomous robot recognition of human social signals. In: *Proc. of New Frontiers in Human-Robot Interaction*, pp. 100–107 (2015)
7. Patterson, M.L., Mullens, S., Romano, J.: Compensatory reactions to spatial intrusion. *Sociometry*, vol. 34, pp. 114–121 (1971)
8. Snijders, D.: Robot's recovery from invading personal space. Unpublished bachelor's thesis, University of Twente (2015)
9. Takayama, L., Pantofaru, C.: Influences on proxemic behaviors in human-robot interaction. In: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pp. 5495–5502. IEEE (2009)
10. Vroon, J.: Regulated reactive robotics: A formal framework. Unpublished master's thesis, Radboud University Nijmegen (2011)
11. Vroon, J., Kim, J., Koster, R.: Robot response behaviors to accommodate hearing problems. In: *Proc. of New Friends 2015*, pp. 48–49. Windesheim Flevoland University (2015)
12. Walters, M.L., Dautenhahn, K., te Boekhorst, R., Koay, K.L., Syrdal, D.S., Nehaniv, C.L.: An empirical framework for human-robot proxemics. In: *Proc. of New Frontiers in Human-Robot Interaction*, pp. 144–149 (2009)