A World for Information Law

26-28 June 2009
with the support of INSEIT

ΔΗΛΩΣΗ ΕΚΔΟΤΙΚΟΥ ΟΙΚΟΥ

Το περιεχόμενο του παρόντος έργου έχει τύχει επιμελούς και αναλυτικής επιστη-
μονικής επεξεργασίας. Ο εκδοτικός οίκος και οι συντάκτες δεν παρέχουν διά του
παρόντος νομικές συμβουλές ή παρεμφερείς συμβουλευτικές υπηρεσίες, ουδε-
μία δε ευθύνη φέρουν για τυχόν ζημία τρίτου λόγω ενέργειας ή παράλειψης που
βασίστηκε εν όλω ή εν μέρει στο περιεχόμενο του παρόντος έργου.

Proceedings of the Second International Seminar on Information Law*

A World for Information Law

26-28 June 2009
with the support of INSEIT

Ionian University, Corfu, Greece
Department of Archive and Library Science
Department of Informatics

Edited by
Maria Bottis

Classifying Service Robots for Policy

Aimee van Wynsberghe

Introduction

A new generation of robots has entered the scene, one that is designed to interact and cooperate with humans. These robots are known by a number of different names, but are commonly referred to as service robots. Despite the ubiquity of the term, their exact definition or description remains unclear. The term 'robot' was coined by Karl Capek, and referred to mechanical slaves – robots are created to provide services, like slaves, for humans. Thus, all robots are essentially service robots making the term service redundant due to the implication within the word robot. When Joseph Engelberger's book "Robots in Service" (1989) first appeared, he portrayed this class of robots as one that can provide a service to humans outside of industrial applications. Thus, robots which provide a service outside of the factory came to be known as service robots. The presupposition is that this new class of robots is interacting with humans, an attribute these robots do not share with their industrial fathers; however, the current development of industrial robots which can interact with humans prevents us from distinguishing service robots from industrial robots in this way (Bicchi A et al., 2008). Without a clear idea of what this new generation of robots is to be called or how they are to be distinguished from traditional robots in industry, how are we to design policy to ensure their safe and effective implementation; more significantly, how are we to account for the ethical implications of these robots?

Adding to the confusion surrounding service robots is the variety of ways in which such robots are interpreted in various fields and portrayed to the public; as a class of human-controlled robots to embark on tasks in dangerous environments, for example search and rescue robots (Iborra A et al., 2009); as a class of robots to replace humans in mundane tasks, for example the Roomba vacuum (Iborra et al. 2009); as a class of robots that interacts with humans in a human environment, for example robot assistants for the elderly (Kim et al. 2009); as a class of robots that communicate with humans in a social man-

ner, for example MIT's KISMET (Breazeal 2003; Kemp et al. 2008); or, as humanoid robots, for example Honda's ASIMO (Ng-Thow-Hing 2009). Thus, within the classification of service robots there exists a wide range of possibilities for further classifications, or sub-classes, which specify the functions and/or features the robot may have. Only when one is aware of the variety of possible architectures for service robots can one postulate effective policy to guide their implementation.

The development of policy for both the design and introduction of these robots depends on the technical variables of the technology used in their creation. Taking into consideration the meaning of the word 'robot' according to Capek - a slave to humans – we must question how useful the classification of service robots is. This may in part lead to the variety of interpretations of service robots which is also problematic for the creation of policy. Each interpretation assumes a distinct function of the robot which presupposes a certain set of technical details or a structural archetype. What is needed is an understanding of the variety of classes of robots applied in a variety of domains outside of industry. Articulating and understanding the functions, features and applications of service robots is pressing as roboticists and society at large call for international legislation and policy to guide their implementation (Sharkey 2008).

The aim of this paper is to show how the term 'service robots' can be unpacked to reveal important distinctions with respect to robots outside of industrial applications and further to show how these distinctions are relevant for policy development. Firstly, I will highlight the difficulty inherent in the term 'service robots'. Secondly, I will propose a way in which service robots can be broken down into a variety of classes (or sub-classes) of robots with the potential for additional features and a range of applications. Finally, I will defend this classification by illustrating its relevance for policy development.

The problem of service robots

Robots may be one of the most difficult technological innovations to define. This is in part due to the immense technical knowledge required to understand their functioning but also to the role media has played in shaping the image of a robot in the minds of society.

The image given by the media, represented by–Star Wars' C-3PO, Star Treck's Data, Pixar's WALL-E-all represent a class of robots not yet realized by today's technology. These futuristic human-like robots may be part of the future or may never be realized. In today's terms, part of the problem when defining robots is the vast number of qualities/features each one may have. There are, however, common attributes, share by all robots. The first is that they are all man-made and artificial. Second, a robot has physical agency or embodiment which distinguishes them from embedded and/or smart technologies and permits real action in the real world (Haselager 2005). Third, a robot has the ability to sense and perceive information, and to execute a programmed task or action (Bicchi 2008). Robots may perceive a variety of types of information; for search and rescue robots information is in the form of environmental cues, or in the case of social robots, the information may be in the form of verbal or non-verbal cues. The information is then used to execute an action or task.

The category of service robots was introduced by Joseph Engelberger to call attention to the idea that robots could, and should, be applied outside of the factory (Engelberger 1989). In order to create policy guiding the design and/or the implementation of this new generation of robots emerging from the industry we must understand, or unpack, the term service robots. The word robot was conceived by Karel Capek (1890-1938) in his play R.U.R. (Rossum's Universal Robots) where he used it to refer to a race of manufactured humanoid slaves; robots are machines that can do the work of humans. The term robot essentially replaced the terms android and automaton, which had been traditionally used until that time. Looking at the definition of a robot – a machine to do the work of humans – all robots are created to service humans. This makes the prefix 'service' redundant and calls into question the use of the label service robots as a distinct class of robots. When addressing other classes of robots like humanoid robots, we can assume by the title a characteristic or feature of these robots - they will have a human-like appearance in some manner. Additionally, from the class of robots called domestic robots we can assume these robots will fulfill a function in the household, most often cleaning. The title service robots says nothing new about the characteristics or function of the robots that fall into this class.

Perhaps the way to understand the uniqueness of service robots is to distinguish them from industrial robots; the differences between the two may allude to a more concrete description of service robots. According to the International Federation of Robotics (IFR), "a service robot is a robot which operates semi or fully autonomously to perform services useful to the well being of humans and equipment, excluding manufacturing operations". From this, we may conclude that service robots are those employed outside of industrial applications with a (varying) degree of autonomy. Thus, any semi or fully autonomous robot outside industry is a service robot but human-operated robots are not. Industrial robots may be designed with varying degrees of autonomy therefore we may postulate the way in which service robots are distinct from industrial robots is their ability to interact with humans and do so with a degree of autonomy. In the EURON Robotics Roadmap, service robots are distinguished from industrial robots by their capability of executing a task in a human environment while interacting with humans. Distinguishing between service robots and industrial robots in this way is not possible as current industrial robot developments look at ways in which robots can safely interact or cooperate with humans in a direct or indirect way (what is referred to as hands-on or hands-off human-robot interaction). Moreover, robots outside of industry do not have to interact with humans (ex. pharmaceutical robot for sorting medications) and they may also be human-operated. Thus, using this definition our only distinction between service and industrial robots is their application domain.

For Engelberger, service robots were meant to describe a class of robots whose architecture was not predetermined by the class they were in but rather their specific function is determined by their application. This fits the EURON Roadmap's description of service robotics; "service robotics (the study of the field of service robots) is not a basic research topic but rather the science of integrating methods and approaches from the various fields of robotics into real-time capable robot systems, which are customized to specific applications". From this we may conclude that the intention underlying the allocation of the term service robots was not to represent a class per say but to indicate a broad category of robots for use outside of industry. Unfortunately, this point is missed when people speak in terms of service

robots. The assumption is made that these robots make up a class on their own. From Engelberger's description then, we may suggest that service robots are those applied outside of industry which may or may not, be capable of interacting with humans and whose autonomy is not specified. It cannot be denied that the application domain of a robot is of paramount importance for creating policy; however, of equal, and perhaps additional significance, are the functions and features of the robot which are used to determine the safety standards and provide additional insight for policy.

Without an understanding of what the term service robot means and without a clear distinction between service robots and others, namely industrial robots, it follows that service robots may be interpreted in a number of ways. Depending on the literature, service robots are used to describe a variety of robots with diverging functions and features from search and rescue robots, robot vacuums or robot assistants for the elderly. Each of these interpretations refer to robots with very different capabilities and functions and it is these details which are needed to shape policy.

As we have seen, the classification of service robots is problematic for numerous reasons; the term is incoherent/redundant (all robots are designed to provide a service to humans), the term is problematic because there are various interpretations of what a service robot is, and the term says nothing about the functions or features the robot may have. Each of these considerations present significant dilemmas for the development of policy to regulate the implementation of service robots. Contrary to the attempts of other authors I will not use the degree to which a robot interacts with a human as the sole defining criteria for robots outside of industrial applications (EURON Roadmap), nor will I use the application domain alone as the sole means to typify classes of robots (Engelberger 1989; Veruggio 2008; Veruggio 2006). Rather, I propose a taxonomy of robots outside of industrial applications which takes into consideration the technical particulars of robots. The robots will be presented by class (which is determined according to functions, capabilities or appearance of the robot) and discussed according to this classification.

Classifying robots-the sub-classes of service robots

The class of a robot refers to a specific type of robot whereby the function, capability or appearance of the robot shapes the entire architecture of the robot. In other words, the class articulates a grouping of robots with a particular function, capability or appearance that dictate important requirements which determine the design of the robot architecture. Although all robots in a class share the same function, capability or appearance, the classes are not mutually exclusive. This means, a robot from one class is allowed to have features or capabilities of a robot in another class. It is important to note here that within a class of robots, which depicts its function or capability in general, there is a wide range of designs for accomplishing the said function or capability, for example, locomotive or mobile robots. There are numerous ways in which locomotion is achieved through the use of wheels (Campion 2008), legs (Kajita 2008), or snake-like movements (Hirose 2009). Alternatively, the class of social robots is unified in its function to communicate interpersonally with humans. The way in which this communication proceeds may be through visual, auditory or nonverbal communication mechanisms. For practical reasons I will not go into depth on the variety of ways in which a function may be realized I mean to indicate simply (at this point) what the function is.

The current literature discussing robots often presents various applications of robots in terms of classes. For example, robots in healthcare are seen as a class of robots. The number of different types or classes of robots that fall under the domain of healthcare are vast and thus it is paramount to recognize this as an application domain and not a class of robots that all share the same technical fundamentals. As such, the classes I present refer to robots that share a capability, function or appearance and may then be applied in a variety of applications. The classes that I will discuss in this paper are as follows; human-friendly robots, social robots, humanoid robots, human-operated robots, autonomous robots, and mobile robots. A robot from any class can then be applied in one or more of the following domains; domestic, entertainment, education, healthcare, military, dangerous environments, and the service industry (note: industrial applications are intentionally left off this list). The classes discussed in the proceeding section have direct consequences for the

creation of policy involving robots outside of industrial applications in terms of both safety and ethical considerations. These considerations will be discussed afterward.

Human-friendly robots

A human-friendly robot (HFR) is one that is designed with particular software and hardware to ensure it is safe to co-exist and co-operate with humans. Thus, the feature that links all robots within this class (the defining criteria) is safety for human-robot interaction. This class of robots is most likely what many researchers have in mind when using the term service robots. According to the Springer Handbook of Robotics, this class of robots is discussed in terms of best performance where safety is provided throughout task execution (Bicchi 2008). The difficulty with building such 'safe' robots is the tradeoff of safety for speed and accuracy. The dilemma now is to design robots that are safe to interact with humans without having to sacrifice performance criteria. The human-robot interaction may be 'hands-on' or 'hands-off'. The former refers to robots designed intentionally to interact with humans while the latter refers to those robots which will not present a threat to the human if they were to accidentally interact ie. come into contact with a human (Bicchi 2008).

One way to design HFRs is using the concept of intrinsic safety: "a robot will be safe to humans no matter what failure, malfunctioning, or even misuse might happen" (Bicchi 2008, p. 1337). One aspect of intrinsically safe robots is to quantitatively assess the risk of injuries in accidents for comparison with other solutions and for optimizing the robot's design. For this, the severity of a potential impact is linked with the statistical probability of causing a certain level of injury. Other methods for designing intrinsically safe robots take the hardware of the robot into consideration to increase their ability to sense objects in their environment or to add protective layers to manipulators (arms) which may potentially come into contact with humans. Other avenues explored look at introducing mechanical compliance into the design. This means, a motor in one area of the robot (ie. one manipulator) can be decoupled/turned off if an impact occurs (Bensalem 2009; Bicchi, 2008). This design, known as compliant transmission, is thought to diminish performance but this may not be a problem when the robot

is used for an entertainment application. In other applications, speed and accuracy of task execution are more important.

The term "human-friendly robots" is taken from the current European initiatives to test the safety standards of this new generation of robots[111]. For some, human-friendly robots are seen as analogous with social robots because the two share the capability of interacting with humans. Configuring robots for the safety of humans is separate from configuring social robots where the aim is to design the robot for achieving the most intuitive contact possible. I argue that social robots are a class on their own which share a defining criteria of social, interpersonal communication.

Social robots

Closely linked with human-friendly robots are social robots. The function of these robots is to meet the social and/or emotional needs of users, rather than the physical needs, through communication (Breazeal 2003; Breazeal et al., 2008). These robots interact with humans in a social way, meaning they communicate (visually, auditorily or verbally) with humans beyond indicating the initiation or completion of a task. For the purposes of this paper, social interaction is defined in terms of high-level communication between the robot and the human user in a human-like way. Human-like communication refers to the variety of ways in which humans can communicate; robots are responsible for understanding cues from humans but also to communicate as a human would. If a robot does not meet this projected meaning of communication (in a human-like way) then I do not consider them to be communicating.

A social robot must be programmed not only to communicate to the user but also to understand the assortment of ways in which the user may communicate with them. For this reason, social robots are considered among the more sophisticated robots today. Means for, or modalities of, communication range "from whole-body motion, proxemics

---

111. Several European projects (within the 6th and 7th framework) have been initiated to study the mechanics required for safe human-robot interactions; URUS (Ubiquitous networking Robotics in Urban Settings); Robot@CWE Portal; VIACTORS; DEXMART (DEXterous and autonomous dual-arm/hand robotic manipulation with sMART sensory-motor skills.

(ie. interpersonal distance), gestures, facial expressions, gaze behavior, head orientation, linguistic or emotive vocalization, touch-based communication, and an assortment of display technologies" (Breazeal 2008, p. 1350). In addition to the reliance on verbal communication, it is important for a social robot to understand a range of non-verbal cues (or, paralinguistic cues). The robot must be able to perceive this information, interpret it accurately and respond appropriately. Robots like the "child-robot" created in Suita, Japan show how roboticists are trying to train robots to perceive facial cues and to group them into broad categories in the same way a child would (Science 2009).

Communication with a human presupposes that the robot is safe to interact with humans, regardless of any additional features or capabilities the robot may have (like mobility or autonomy). The interpersonal manner in which these robots are meant to engage people in is the crucial distinction between human-friendly robots and social robots; human-friendly robots do not have to communicate with a human at all let alone in the social manner used to define communication in this paper. To illustrate this distinction we may use the da Vinci surgical robot (Intuitive Surgical Inc.); da Vinci is designed to be safe in the presence of humans and under the manipulation of humans for assisting in the surgery but da Vinci does not communicate with the human in a social way. Although the robot must indicate when it is on and the system is ready, there is no social meaning to this communication; the function of this communication is not to meet the emotional or social needs of the user. In contrast, social robots are those designed to engage with humans to fulfill social-emotional goals in diverse applications like education, healthcare, entertainment, communication, and collaboration[112].

As social robots are meant to communicate in a social or human-like way, the embodiment of a social robot is often humanoid or animal-like. There are many examples of social robots which have animal-like features like the baby seal robot, Paro. Paro does not verbally communicate but relies on touch based communication to perceive information with the user but also to communicate to its user (that it wants to be held or petted etc) (Wada 2005). There are mobile social robots

---

112. Further exploration of the meaning of social interaction and communication is needed but goes well beyond the purpose of this paper, see later work.

fitted with a face to enhance social interaction, like the elder-care robot, Pearle, developed at Carnegie Mellon University. For mobile robots, issues of proxemics, as a modality of social communication, are particularly important and culturally dependent (proximity preferences in communication differ between cultures). Other social robots do not have an animal or humanoid appearance, like MIT's Kismet or the Keepon made by the National Institute of Information and Communications Technologies in Japan. Kismet has a mechanical face with anthropomorphic features (large blue eyes) to enhance social communication, while Keepon, the small dancing robot, has a simple face and expresses itself by squashing or stretching its body. There are also social robots with no face or eyes or any anthropomorphic features. These robots resort to language-based communication and proxemics. Therefore, although the robot is assumed to communicate in a human-like way this does not presume the appearance of the robot must also be human-like for social robots.

Humanoid robots (creature/animal-like robots, and android robots)

This class of robots is one in which the form or appearance of the robot, and not a capability, is the defining criteria. A humanoid robot resembles a human in form, meaning it may have arms and/or legs etc like Honda's ASIMO. The robot can then have additional anthropomorphic features like eyes and ears like MIT's KISMET. In contrast to the more mechanical looking humanoid robots are androids; "android robots are designed to have a very human-like appearance with skin, teeth, hair and clothes" (Breazeal 2008, p. 1351). In contrast to both of these, are robots that resemble animals or are creature-like (ex. Sony's AIBO).

For some, humanoid robots refer to those with a human-like appearance or with human-like behavior (Kemp 2008). We have already addressed those robots with human-like behavior (ie. social robots) and I will restrict the classification and discussion of humanoid robots to appearance only. Classification is based on appearance and not capability or function because robots with a shared appearance may have any number of functions or capabilities but the requirements for their appearance determine the majority of their structure (at this time).

Distinguishing humanoid robots – or robots based on appearance – in this way helps to clarify that although the robot may have a human-like form this does not presuppose anything about its function. For example, the PEARL robot for assisting with elderly care has arms but these are used for communication and not for maneuvering objects. A common misunderstanding with humanoid robots is that they also possess the capabilities referred to with HFRs and/or social robots (Kemp 2008; Veruggio 2008; Ng-Thow-Hing 2009). An assumption is being made here that in order for robots to interact with humans, in the social manner described above, they must have a human appearance. Many social robots are humanoid, creature- or animal-like but many are not overtly humanoid nor do they resemble an animal or creature. For this reason, it is important to distinguish between the capabilities of a social or HFR and the appearance of a humanoid robot; one does not presume the other. Therefore, when policy makers are asked to design guidelines for humanoid robots, they may understand that the function and capability of the robot must be explicitly stated.

As mentioned, this representation of humanoid robots, regarding appearance alone, is in contrast to many references made about humanoid robots, which infer capabilities from appearance. Restricting the temptation to attribute capabilities to these robots allows us to address the capabilities separately, to acknowledge the range of capabilities a humanoid robot may have and to maintain a realistic vision of what these robots are capable of (at this time). This classification also allows us to address some of the more deeply rooted philosophical questions regarding; emotional projections humans place on robots of a certain appearance, what happens to the expectations of users when the robot has a humanoid appearance regardless of their capabilities and how will a human respond to a humanoid appearance[113].

---

113. Response to the appearance of a robot is often discussed in terms of Masahiro Mori's Uncanny Valley hypothesis, that there exists a threshold for human comfort with the robots appearance. Below this threshold, the robot does not closely resemble a human and humans are comfortable with the robot. Above this threshold, the robot is commonly thought of as an android and again the human is comfortable with the robot. But right at the threshold, the robot appears more like a zombie than a human which elicits feelings of disgust and unease towards the robot.

Human-operated Robots

Human-operated robots represent one of the earliest classes of robots, one which requires a human to guide the action/movements of the robot; the movements or commands of the human are translated into movements made by the robot. This configuration is often referred to as master-slave - the human operator as master and the robot as slave.

For hands-on human-robot interaction in industrial applications, such robots are often referred to as "cobots". These are "collaborative robots" designed to relieve humans from fatigue or stress and to prevent injuries; "cobots presume a division of control between human and robot, with a robot perhaps supporting a payload and allowing a human to guide it" (Bicchi et al., 2008, pg 1345). In this scenario, the operator is in direct physical contact with the payload. This description may also be used to describe exoskeletons used for rehabilitation purposes; "exoskeletons are also controlled by a human operator, leaving all planning and high-level challenges to the user" (Niemeyer 2008, p. 741; Hayashi 2005). Again, the user is in direct contact with both the robot and the payload.

Within the class of human-operated robots is a subclass known as telerobots whose infrastructure is designed such that a human operator controls the motion/movement/task execution of the robot in the same way as a human-controlled robot with the added condition that the human operator is at a distance (Niemeyer 2008). Again, all planning and cognitive decisions are made by the human user and the robot is used strictly for mechanical completion of a task. The use of 'tele' (derived from the Greek word for distant) presumes a geographical separation between the user and the environment in which the task is being performed. The inaccessibility of the environment may be for any number of reasons; the user cannot or will not physically reach the environment, the environment is dangerous, the environment needs to be scaled. The physical distance between the user and the robot varies depending on the application (ie. for surgical robots the surgeon is often in the same room, for robots in space or underwater the distance is much greater). In most cases there are two sites to speak of; the local site with the human operator and the remote site with the robot. For the information to travel from one side to another, the two sites

must be connected. Traditionally this was done through the use of cables; however, recently computer networks have made it possible to transmit this information from one side to another using a telecommunication system. The use of telecommunication networks implies that more than one user may be able to manipulate the robot, but also the distance between the local and distant site may be increased.

Control of the robot may occur through one of three architectures; direct control, shared control, or supervisory control (Niemeyer 2008). Direct control assumes no autonomy or intelligence on the part of the robot, thus, all the motions of the robot are directly controlled by the user. Shared control refers to a sharing between local and remote sites whereby the human operator decides what to do and how to act while the robot can autonomously refine the command for the environment. For example, in the case of the da Vinci surgical platform, the surgeon performs its movements which the robot autonomously scales down to the appropriate size for the surgical field. Supervisory control is described as analogous with supervising a subordinate staff member whereby the supervisor is responsible for giving orders to the subordinate but in turn receives summary information. This approach is compared with direct control or autonomous robot control by Sheridan, who introduced the concept of supervisory control; "human operators are intermittently programming and continually receiving information from a computer that itself closes an autonomous loop through artificial effectors and sensors" (Niemeyer 2008, p. 746). In other words, "the operator plans activities at a level which can be performed by the robotic system independent of human intervention" (Niemeyer 2008, p. 747). At all times, the human operator may take over control of the task.

Telerobotic systems or human-operated robots are unique in that they provide information to, but also require commands from, the user. These robots are distinguished from autonomous robots in which "a robot executes a motion or other program without further consultation of a user or operator" (Niemeyer p. 746). It may be suggested that autonomous robots evolved from the design of supervisory control robots (Haselager 2005).

Autonomous Robots

This issue of autonomy in robotics is problematic due to the diverse conceptions one may have of the concept of autonomy; "the capacity

for independent (unsupervised) action versus the freedom to choose goals" (Haselager p. 528). While philosophers approach autonomy from the question of why one acts in a certain way, roboticists approach autonomy from the question of how the robot fulfills its task (with or without assistance or supervision); "within robotics, the increase in autonomy of a system is related to the reduction of on-line supervision and intervention of the operator, programmer or designer in relation to the robot's operations in a changing environment" (Haselager 2005, p. 518). Autonomous robots are therefore a class of robots with the capability to fulfill a task without real-time manipulation from a human operator. For Engelberger, "autonomous planning is performed by the machine when sensed data are operated on by application programs with the result that the machine makes navigating (or equivalent) decisions. These decisions do not require human interaction but are, on the robotic side, subject to human supervision and veto" (Engelberger 1989, p. 211-12). This description corresponds with later visions of autonomous robots as a class of robots that "operate under all reasonable conditions without recourse to an outside designer, operator or controller while handling unpredictable events in an environment or niche" (Franklin and Graesser 1996). These two definitions of robot autonomy maintain that the robot is acting according to a pre-programmed set of rules, and the robot is capable of planning their action without referring to a human-operator (or designer or controller) during execution of the task. What the second definition adds is the capability of the autonomous robot to fulfill its task in an environment in which it has not been trained, and/or one that is unpredictable.

More recent definitions of autonomous robots maintain the criteria of the first two definitions presented but add the capability of the robot to fulfill its task within time constraints and with the added component of potential interference by others. For Bensalem et al, autonomous robots must: "operate in highly variable, uncertain, and time-changing environments; meet real-time constraints to work properly; interact with other agents, both humans and other machines" (Bensalem 2009, p. 67). From the description of an autonomous robot today, we see that some of the challenges for their design relate to challenges posed for other classes of robots; meeting real time constraints while at the same time safely interacting with humans is a similar challenge for HFRs. The difference lies

in the robot's capability for autonomous function; predictability of the robot's actions decrease without a human-operator, risk that the robot misinterprets environmental cues and acts improperly increases, or if the robot is unsure how to respond, without the guidance of an operator it may malfunction or shut down. Thus, safety is significant for this class of robots both in terms of interacting with humans but also in terms of reliability - that the robot is capable of accomplishing its task. The third definition also introduces the aspect of human-robot interaction. This, however, is not a defining criterion of an autonomous robot. That is, they may interact with humans but they may not. Autonomous robots, like tele-robots, are intended for situations in which human control is not feasible, not desirable, or perhaps is not the most cost-effective alternative.

Without prescribing what an autonomous robot should refer to, we may suggest that it has the following properties; it can perform its pre-determined task in an unpredictable environment without consulting an outside source for assistance. Both hands-on and hands-off safety criteria apply to these robots depending on the robots function and/ or application.

The class of autonomous robots can be further broken down into: autonomous mobile robots, cluster robots and learning robots. Auton-omous mobile robots refer to a class of robots with the capability of autonomous mobility. This means, the robot does not depend on the human operator to control its mobility. Autonomous mobile robots are considered "tools that human specialists use to perform hazardous tasks in remote environments" (Bensalem 2009; Breazeal 2008). This defini-tion, however, does not address the use of mobile autonomous robots used in applications where the robot is used to relieve humans of the bur-den of mundane, time consuming, boring tasks like the Roomba vacuum cleaner. Autonomous robots which are also designed to be HFRs must pay particular attention to issues of proxemic like how to approach, fol-low or maintain an appropriate distance with a human.

Cluster robots are also a sub-class of autonomous robots. These robots are also known as multirobot systems or networked robots. They are categorized by their ability to work together with other robots in order to accomplish a task. Thus, each robot acts autonomously but

must also autonomously coordinate their actions with the actions of their fellow bots in the system.

Robot learning or learning robots[114] may be used to refer to a property of the robot (Franklin and Graesser, 1996) - it can adapt by changing its behaviour based on its previous experience – or to the way in which the robot is programmed – learning by demonstration, mimicking, or reinforcement (Argall 2009; Billard 2008; Marino 2006; Riebiero 2002). The concept of robot learning invariably increases the degree of autonomy the robot has and increases the success with which the robot will maneuver in a new, unknown environment. With respect to programming robots by learning it is thought that robots learn general rules from their experience in order to meet task assignments in highly variable environments (meaning human environments) (Marino 2006). There are many ways in which roboticists are exploring how to program learning into the robot. The "Child-robot" developed in Suita, Japan, is said to develop social skills by interacting with humans and watching their facial expressions, mimicking a mother-child relationship. The aim of the creators at Osaka University in Japan is to develop this robot to think like a baby, meaning the robot will be able to evaluate facial expressions and cluster them into basic categories like 'happy' or 'sad'. Once referred to as learning-by-demonstration (this approach was targeted towards the use of industrial robots), this term was replaced with imitation learning to reflect the way in which the robot would 'learn' in order to interact with humans in a more natural way (by demonstrating similar skills and processes). These robots are the predecessors of more advanced social robot. On the other hand, learning robots do not always have to be safe for human interaction because they may be applied in military, surveillance, or search and rescue applications.

---

114. Learning robots are problematic for many reasons; how do we transfer the human notion of learning to robots, how can we reliably say when a robot has learned, will the robot be able to act in a way not intended by their designers? Various authors have addressed the issue that these robots may be capable of acting in ways not anticipated in their design and as such there exists a problem of responsibility; who is responsible for these robots if the designers cannot with complete confidence predict the robot's behavior (Marino 2006)? This question is of paramount importance for the creation of policy concerning such a class of robots. Further research on this subclass is required as the technology is still in its beginning stages.

Mobile Robots

The class of robots called mobile robots is distinguished from traditional stationary industrial robots with a fixed platform, or rehabilitation robots that perform a function in the kitchen, on the desktop or by the bed. Mobility of the robot refers to its ability to travel along the x-y planar axis, in other words locomotion. Locomotion is different from a robot which is capable of moving an effector or manipulator (arm or hand). For example, the surgical robot da Vinci does not travel as it operates but its robotic arms must be moveable during the course of the surgery. In contrast, the iRobot Warrier 700 military robot is designed to drag humans to safety. For such a robot, mobility is a defining feature for fulfilling its task.

Locomotion may be accomplished in a variety of ways, the architecture of the robot is then determined by the chosen means for locomotion. Mechanics for mobility vary depending on the institution or company designing the robot and the terrain which the robot is expected to move on. Researchers at the Tokyo Institute of Technology that in a snake-like manner (Hirose 2009). In contrast, researchers at Honda are designing ASIMO to walk like a human using a zero-moment technique (Ng-Thow-Hing 2009). This technique means the robot equally balances all forces so there is no point at which the robot would lose balance and fall. Additionally, this type of motion requires that the robot be on a smooth surface - not an optimal restraint if the robot is to exist in an unstructured environment where these things cannot be accounted for. Other researchers are exploring the use of gravity to propel the 'legs' for moving, a technique referred to as 'passive dynamics' (Kajita 2008). This approach/technique uses little motor power to accomplish walking and is considered a promising, efficient substitute to the zero moment technique used for ASIMO.

Wheels are the most typical means for motion for reasons of simplicity (Campion 2008). A Segway is commonly thought of as a mobile robotic platform which uses wheels for motion. Researchers at Carnegie Mellon are also investigating the use of a ball for locomotion Lauwers 2006). The "ballbot" is a battery operated, omnidirectional robot that balances on a single urethane-coated metal sphere. Because of the

use of the ball it is able to maneuver in tight spaces and has the potential to interact in human environments better than wheeled robots.

Control of the robot's mobility may be human-controlled or autonomous. An example of a human-operated mobile robot is In Touch's RP-7. This robot is aimed at facilitating patient-physician communication when the physician cannot be physically present at the bedside of the patient. The physician, seated at a console in another area of the hospital or in another place entirely, guides the robot through the hallways of the hospital to the patient's bedside. Using a video monitor attached to the mobile autonomous robotic platform, the patient and the physician may communicate directly. In contrast, iRobot's Roomba vacuum cleaner or iRobot's Scooba (pool cleaner) are both mobile robots which operate autonomously; no human manipulation is required to guide the robots locomotion.

In terms of policy, the capability of mobility alone presents distinct safety challenges but when added with the capability of autonomous functions, additional safety considerations must be accounted for; whereas a human-operated mobile robot is less likely to collide with other objects because of the control of the human, an autonomous mobile robot requires redundant (additional) sensors for perceiving their environment. This issue of speed for autonomous mobile robots when traveling and stopping is also significant.

Additional Features, or Properties, of a Robot

In addition to the classes discussed above there are a variety of additional features a robot may possess that are not class defining. Meaning, they are features that do not determine the entire infrastructure of the robot but rather can be added to an existing frame. As technology develops and software programs are created to harmonize the programming of robots, some of the capabilities determining class may evolve into additional properties the robot may posses. Currently, when speaking of features of the robot we may refer to grasping, manipulation, face detection and/or recognition or, voice detection and/or recognition, among others. Grasping is typically associated with hands and refers to the property of a robot to detect and select objects to be positioned and/or oriented (Engelberger 1989). It should be noted that some robots may posses the feature of graspers or

hands but do not use them to grasp objects but rather to communicate. This highlights the importance of understanding the features and functions of the robot separate from any application domain. Manipulation is a feature of a robot which allows it to pick up objects and move them from one place to another (Engelberger 1989). This list is not extensive due to constraints; however, the intention was to indicate an additional array of robot properties to be taken into consideration to further highlight the point that the term service robots leaves too many details unaccounted for.

How this taxonomy aids policy development

In light of the number of classes of robots and the differences between these classes (in terms of functions, capabilities and appearances) it becomes obvious that simply addressing robots that will interact with humans or that will work in a human environment (as opposed to industry) leaves many questions unanswered; will the robot be communicating with the human, what form will this communication take; will the robot be autonomous and thus not require the control of the human; will the robot have a distinct appearance which may change the expectations of the user, what safety considerations must be met? The Roboethics taxonomy classifies robots according to their application domain in order to address ethical considerations associated with their application (Veruggio et al., 2008). While this provides useful insight, it fails to account for the variety of robots (capabilities, functions, properties and appearances) within one application domain. Take robots in healthcare as an example. If one is called upon to create policy for 'robots in healthcare' they will invariably be responsible for creating policy that pertains to robots with a wide range of capabilities and functions each requiring distinct consideration. Robots applied in healthcare range from; stationary human-operated robots (ex. The da Vinci surgical platform); mobile human-operated robots (In Touch's RP-7); stationary autonomous robots (ex. ROBOT-Rx for sorting medications; Jerrard, 2006); and mobile autonomous robots (ex. i-Merc robot that delivers meals; Carreira F, 2006) to name a few. Policy must be tailored to account for the range of robots in healthcare, ipso facto, one must be aware of these differences.

By presenting robots according to this classification, we may; address the considerations pertaining to distinct capabilities or functions the robot may have, we may ask ourselves whether certain classes are appropriate for certain application domains, we may also observe how a combination of capabilities requires different standards than a capability on its own. Additionally, we can tailor the design and introduction of these robots by ethically reflecting on the current state of the art rather than appealing to a futuristic vision of what robots will one day be, or look like.

Safety Considerations

At this point in time, the most substantial concerns for robots outside of industrial applications pertain to the safety and reliability of the robot. Due to the length of time in which robots have been employed in industrial applications, safety standards are already in place[115]. Things become more complicated when speaking of hands-on robots which will be applied in a variety of domains. For hands-on robots, the T-15 committee of the American National Standards Institute (ANSI) is setting safety standards regarding intelligent assist devices (Bicchi 2008). Although these standards cover a wide range of technologies from assistive devices to mobile autonomous robots, these standards are promising in that they may be translated into policy governing domestic applications of robots. For example, one aspect of the standards involves risk assessment replacing fixed rules: "instead of declarations regarding how to accomplish safe operation, risk assessment procedures are advised for assistive devices and physical human-robot interaction robotic technologies, to identify and mitigate risks in proportion to their seriousness and probability". In other words, assessing the chance for injuries take priority over safe operation. This is to ensure that robots are in fact a promising alternative.

Another aspect refers to safety-critical software: under any condition that the robot malfunctions, the entire system will shut down in a safe manner. While on first glance this may seem appropriate,

---

115. Well established nation standards in the US (ANSI-RIA), Canada (CSA), Germany (DIN) etc. These standards are collected and harmonized by the International Organization for Standardization (ISO).

once we address the different classes of robots we may conclude that this standard when applied to autonomous robots working without the supervision or control of a human may do more harm than good. How is one to know if the robot has shut down? If the robot shuts down in the middle of its task (for example, a household chore), who is responsible for addressing the technical, or otherwise, problem which resulted in the robots malfunction?

The standards of the T-15 committee also indicate dynamic limits which restrict the capabilities of robot design such that a human operator must be able to outrun, overpower or turn off the robot. This standard requires detailed knowledge of the intended users of the technology. Robots for elderly or rehabilitative patients will have different dynamic limits from robots in the average household.

Safety issues are of paramount importance when the robot is in direct physical contact with the human user. One application in which this is illustrated is the use of robots for rehabilitation (Kazerooni 2008; Hayashi 2005). These robots come in direct physical contact with patients in a variety of ways. For therapy robots, the robot is in direct contact with the disabled patient and the therapist simultaneously. Roboticists in this area must be sure that the robot is designed in such a way that it cannot cause injury by moving a user's limbs outside their range of motion, with too much strength or with too much speed. In addition to this limits imposed on the robotic apparatus, redundant sensors (additional sensors) are used as back-up such that if one sensor malfunctions another can identify the problem and shut down if necessary. Outside all this, rehabilitation robots must also be designed to be intrinsically safe; "from the systems perspective, when all else fails actively to protect the user, it must be the design itself that makes the robot inherently unable to injure the user" (Bicchi 2008, p. 1244).

Programming safety into the architecture of the robot may also depend on a combination of capabilities and not one exclusively, for example mobility and control. If the robot is mobile but is human-operated, the standards for safety will be somewhat different from a robot which is mobile and autonomous. The autonomous robot must be programmed to stop at a certain speed and distance from the object it approaches. It must also be programmed to shut down immediately, in

a safe manner, in case of malfunction. Mobile human-operated robots rely on the commands on the human-operator for these details.

Ethical Considerations

Above and beyond the safety considerations, policy must also account for the ethical considerations when introducing robots. It is beyond the scope of this paper to address the range of ethical implications pertaining to robots; however, we may focus on one area receiving much attention today where robots are seen as a promising solution; robots for the care of elderly persons. Various European initiatives are currently in place for safeguarding the quality of life of this demographic while exploring the potential use of information and communication technologies[116]. With a lack of healthcare workers and an increase in need of care, robots are seen as a way of compensating for the shortcomings of a given healthcare system, or as a means of proving care for those who would not otherwise have access. Action plans for policy point to the need to maintain independence, autonomy and dignity for this vulnerable demographic when introducing these new technologies. Robots for the elderly draw on all the classes of robots discussed; all must be HFRs; some will be social; some will be autonomous and others human-controlled; some will be mobile and others stationary; and, some will have a humanoid appearance while others resemble a machine. Again we are reminded that a discussion of robots in one application domain invariably begins a discussion of numerous types (or classes) of robots. So, how does one begin?

First, we may look at the differences between autonomous and human-operated robots. Autonomous robots are responsible for completing the task without the guidance or assistance of a human, therefore, the task is entirely delegated to the robot. When speaking of a technological delegation of care, we open a discussion of responsibility. Who is responsible for the robot malfunctioning? Who is responsible to ensure the robot is being used properly? Who is responsible for the actions of the robot? If every-

---

116. "Ageing well in an Information Society" is an action plan made in reaction to the European Union's Riga Ministerial Declaration on e-Inclusion. E-Inclusion was a project to support Information Society policy development. i2010 is the EU policy framework for addressing the positive contribution of information and communication technologies through research and development.

thing goes right (meaning, there are no technical complications with the robot), we must then ask whether we want to delegate the responsibility of care to robots? Is this an appropriate use of the technology? Would we opt for robots if we had enough human care workers? Policy must take these questions into consideration.

Many times it is the intended users of a technology that motivate additional ethical considerations. For this vulnerable demographic, a substantial concern is the risk of stigmatization or discrimination (often referred to as ageism). Another promising robot for the care of elderly persons is one that will bath a person[117]. From an ageist perspective, if we're using robots to wash elderly patients than we should also use the robots for washing young, competent post operative patients in the hospital. This means the type of care should be the same for both groups. If not, then there's an ageist problem. When assessing robots according to the status of the user, the vulnerability of an elderly person (capability and competence) is a specific feature. The danger is that the elderly person is reduced to a physical/bodily being thereby threatening the patient's dignity. Vulnerability increases with dementia and Alzheimer's and therefore robots should not be used because the person cannot say no to it. Thus, we may suggest competence of the user is a necessary condition for implementing robots in a given domain. Safeguarding the patient's dignity can also be incorporated into the design of the robot. For instance, a bathing robot should not have the person on display creating physical and social discomfort. Instead, such robots ought to be designed in such a way that the person's physical integrity and dignity remain intact.

Conclusion

The speed with which robots are developing and the range of robot architectures, properties, functions, capabilities and appearances makes the creation of a taxonomy or classification of robots an arduous one. Regardless, one is needed in order to facilitate the creation of policy for the new generation of robots leaving the industry and entering into human environments. Although these robots have traditionally been, and continue to be, referred to as service robots, the aim of this

---

117. The Avant Santelubain 999, made in Japan, is a robotic bath or human washing machine. The robot will sterilize and clean itself, as well as shampoo one' s body.

paper was to illustrate the many details and variety of robots lost when using this term. Further, this does nothing to assist policy makers in their task. Moreover, this restricts the possibility for interdisciplinary discourse, or societal participation in their design and implementation, as each discipline may have their own interpretation of what a service robot is. Classifying robots according to their application domain is of great value but again misses important details without acknowledging the range of robots applied in any given domain.

The classification presented here is by no means complete; however, this does not undermine the goal to draw attention to the many capabilities a robot may possess. By referencing robots according to class we have seen how each class presents distinct challenges for policy and how when one robot combines capabilities from more than one class the challenges for policy increase. By understanding the capabilities and what they infer for policy (and what a combination of capabilities infers for policy), we may assist policy developers in developing appropriate safety standards. Moreover, we may also encourage ethical reflection pertaining to specific robots, or uses of robots, at this stage in their development, to be incorporated in their design. As an emerging technology, we are afforded the luxury of shaping their design and implementation in a way that safeguards traditions and practices dear to one society or another[118].

---

118. It must be noted that culture influences the acceptance of classes of robots as well as their intended application. For instance, Japan views robots, and technology in general, as the solution to many problems faced within their culture. Alternatively, European cultures are not as inclined to introduce robots in general but also in certain application domains, without exhaustive ethical analysis.

References

Aging well in an Information Society; an i2010 initiative, action plan on information and communication technologies and aging, [Online]. (2007). Available: http://fuhu.dk/filer/DEA/T%E6nketanke/EU-projekt/Ageing_well_in_the_information_society.pdf

Argall B., Chernova S., Veloso M., Browning B. (2009) A survey of robot learning from demonstration. Robotics and Autonomous Systems, 57, 469-483.

Bensalem S., Gallien M., Ingrand F., Kahloul I., Thanh-Hung N. Designing Autonomous Robots; toward a more dependable software architecture. IEEE Robotics and Automation Magazine, 16 (1), 67-77.

Bicchi A., Peshkin M., Colgate J. (2008) Safety for physical human-robot interaction, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp 1335-1348.

Billard A., Calinon S., Dillmann R., Schaal S. (2008) Robot programming by demonstration, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 1371-1394.

Breazeal C. (2003) Emotion and sociable humanoid robots. International Journal of HumanComputer Studies, 59, pp. 119-155.

Breazeal C., Takanishi A., Kobayashi T. (2008) Social robots that interact with people, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 1349-1370.

Campion G., Chung W. (2008) Wheeled Robots, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 391-410.

Carreira F. (2006) i-Merc: a mobile robot to deliver meals inside health services, Proceedings from IEEE Conference on Robotics, Automation and Mechatronics.

Engelberger J. (1989) Robotics in Service, Biddles Ltd, Guildford.

European Robotics Research Network, EURON. [*Online*]. (2008). Available: www.euron.org/activities/benchmarks/motionplan

Franklin S., Graesser A. (1996) Is it an agent, or just a program? A taxonomy for autonomous agents, Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages. Springer, pp. 21-35.

Haselager W. (2005) Robotics, philosophy and the problem of autonomy, Journal of Pragmatics and Cognition, 13 (3), pp. 515-532.

Hayashi T., Kawamoto H., Sankai Y. (2005) Control method of robot suit HAL working as operator's muscle using biological and dynamical information, Proceedings from Intelligent Robots and Systems (IROS), pp. 3063-3068.

Hirose S., Yamada H. Snake-like Robots; machine design of biologically inspired robots, IEEE Robotics and Automation Magazine, 16 (1), pp. 88-98.

Iborra A., Caceres D., Ortiz F., Franco J., Palma P., Alvarez B. Design of Service Robots; experiences using software engineering, IEEE Robotics and Automation Magazine, 16 (1), pp. 24-33.

International Federation of Robotics; classification and definition of service robots [Online], Available at: http://www.ifr.org/modules.php?name=News&file=article&sid=14

Japan child robot mimics infant learning (2009) Science, [Online] Available at: http://www.iol.co.za/index.php?set_id=1&click_id=31&art_id=nw20090411170159511C566377

Jerrard, J. (2006) Robot PharMD; drug dispensing robots drastically decrease medication errors, The Hospitalist.

Kajita S., Espau B. (2008) Legged Robots, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 361-382.

Kazerooni H. (2008) Exoskeletons for human performance augmentation, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 773-881.

Kemp C., Fitzpatrick P., Hirukawa H., Yokoi K., Harada K., Matsumoto Y. (2008) Humanoids, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 1307-1334.

Kim M., Kim S., Park S., Choi M., Kim M., Gomaa H. (2009) Service Robot for the Elderly; software development with the COMET/UML method, IEEE Robotics and Automation Magazine, 16 (1), pp. 34-45.

Lauwers T., Kantor G., Hollis R. (2006) A dynamically stable single-wheeled mobile robot with inverse mouse-ball drive, Proceedings from the IEEE International Conference on Robotics and Automation.

Marino D., Tamburrini G. (2006) Learning robots and human responsibility, International Review of Information Ethics; 6, pp. 47-51.

Niemeyer G., Preusche C., Hirzinger G. (2008) Telerobotics, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 741-758.

Ng-Thow-Hing V., Thorisson K., Sarvadevabhatla S., Wormer J., List T., Cognitive Map Architecture; facilitation of human-robot interaction in humanoid robots, IEEE Robotics and Automation Magazine, 16 (1), pp. 55-66.

Riebiero C. (2002) Reinforcement Learning Agents, Artificial Intelligence Review, 17, pp. 223-250.

Sharkey N. (2008) The Ethical Frontiers of robotics, Science 332, pp. 1800-1801.

Veruggio G., Operto F. (2006) Robotethics: a bottom-up interdisciplinary discourse in the field of applied ethics in robotics, International Review of Information Ethics; 6, pp. 3-8.

Veruggio G., Operto F. (2008) Roboethics: social and ethical implications of robotics, The Springer Handbook of Robotics. (ed. Siciliano B. and Khatib O.), Springer, pp. 1499-1524.

Wada K., Shibata T., Saito T., Sakamoto K., Tanie K. (2005) Psychological and social effects of one year robot assisted activity on elderly people at a health service facility for the aged, Robotics and Automation, pp. 2785-2790.

Biography

Aimee van Wynsberghe has just begun her PhD in Philosophy at the University of Twente, the Netherlands. During her undergraduate degree in Cell Biology at the University of Western Ontario, Canada, she was a research assistant at CSTAR (Canadian Surgical Technologies and Advanced robotics) working on the Telesurgery project which inspired her to continue working with robots. Following her studies in Science, she pursued Applied Ethics and Bioethics in her graduate studies. This gave her the opportunity to reflect on the philosophical issues pertaining to technology in healthcare, with a particular focus on robotics. Her current work focuses on the social implications of human-robot interactions but will specifically address the use of robots in the care of elderly persons.