A survey of socially interactive robots

Terrence Fong [a,b,*], Illah Nourbakhsh [a], Kerstin Dautenhahn [c]

[a] *The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA*
[b] *Institut de production et robotique, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland*
[c] *Department of Computer Science, The University of Hertfordshire, College Lane, Hatfield, Hertfordshire AL10 9AB, UK*

Abstract

This paper reviews "socially interactive robots": robots for which social human–robot interaction is important. We begin by discussing the context for socially interactive robots, emphasizing the relationship to other research fields and the different forms of "social robots". We then present a taxonomy of design methods and system components used to build socially interactive robots. Finally, we describe the impact of these robots on humans and discuss open issues. An expanded version of this paper, which contains a survey and taxonomy of current applications, is available as a technical report [T. Fong, I. Nourbakhsh, K. Dautenhahn, A survey of socially interactive robots: concepts, design and applications, Technical Report No. CMU-RI-TR-02-29, Robotics Institute, Carnegie Mellon University, 2002].
© 2003 Elsevier Science B.V. All rights reserved.

Keywords: Human–robot interaction; Interaction aware robot; Sociable robot; Social robot; Socially interactive robot

1. Introduction

1.1. The history of social robots

From the beginning of biologically inspired robots, researchers have been fascinated by the possibility of interaction between a robot and its environment, and by the possibility of robots interacting with each other. Fig. 1 shows the robotic tortoises built by Walter in the late 1940s [73]. By means of headlamps attached to the robot's front and positive phototaxis, the two robots interacted in a seemingly "social" manner, even though there was no explicit communication or mutual recognition.

As the field of artificial life emerged, researchers began applying principles such as stigmergy (indirect communication between individuals via modifications made to the shared environment) to achieve "collective" or "swarm" robot behavior. Stigmergy was first described by Grassé to explain how social insect societies can collectively produce complex behavior patterns and physical structures, even if each individual appears to work alone [16].

Deneubourg and his collaborators pioneered the first experiments on stigmergy in simulated and physical "ant-like robots" [10,53] in the early 1990s. Since then, numerous researchers have developed robot collectives [88,106] and have used robots as models for studying social insect behavior [87].

Similar principles can be found in multi-robot or distributed robotic systems research [101]. Some of the interaction mechanisms employed are communication [6], interference [68], and aggressive competition [159]. Common to these group-oriented social

* Corresponding author. Present address: The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA.
E-mail addresses: terry@ri.cmu.edu (T. Fong), illah@ri.cmu.edu
(I. Nourbakhsh), k.dautenhahn@herts.ac.uk (K. Dautenhahn).

Fig. 1. Precursors of social robotics: Walter's tortoises, Elmer and Elsie, "dancing" around each other.

robots is maximizing benefit (e.g., task performance) through collective action (Figs. 2–4).

The research described thus far uses principles of self-organization and behavior inspired by social insect societies. Such societies are anonymous, homogeneous groups in which individuals do not matter. This type of "social behavior" has proven to be an attractive model for robotics, particularly because it enables groups of relatively simple robots perform difficult tasks (e.g., soccer playing).

However, many species of mammals (including humans, birds, and other animals) often form individualized societies. Individualized societies differ from anonymous societies because the individual matters.


Fig. 2. U-Bots sorting objects [106].

Fig. 3. Khepera robots foraging for "food" [87].

Fig. 4. Collective box-pushing [88].

Although individuals may live in groups, they form relationships and social networks, they create alliances, and they often adhere to societal norms and conventions [38] (Fig. 5).

In [44], Dautenhahn and Billard proposed the following definition:

> Social robots are embodied agents that are part of a heterogeneous group: a society of robots or humans. They are able to recognize each other and engage in social interactions, they possess histories (perceive and interpret the world in terms of their own experience), and they explicitly communicate with and learn from each other.

Developing such "individual social" robots requires the use of models and techniques different from "group


Fig. 5. Early "individual" social robots: "getting to know each other" (left) [38] and learning by imitation (right) [12,13].

Fig. 6. Fields of major impact. Note that "collective robots" and "social robots" overlap where individuality plays a lesser role.

social" collective robots (Fig. 6). In particular, social learning and imitation, gesture and natural language communication, emotion, and recognition of interaction partners are all important factors. Moreover, most research in this area has focused on the application of "benign" social behavior. Thus, social robots are usually designed as assistants, companions, or pets, in addition to the more traditional role of servants.

1.2. Social robots and social embeddedness: concepts and definitions

Robots in individualized societies exhibit a wide range of social behavior, regardless if the society contains other social robots, humans, or both. In [19], Breazeal defines four classes of social robots in terms of: (1) how well the robot can support the social model that is ascribed to it and (2) the complexity of the interaction scenario that can be supported as follows.

Socially evocative. Robots that rely on the human tendency to anthropomorphize and capitalize on feelings evoked when humans nurture, care, or involved with their "creation".

Social interface. Robots that provide a "natural" interface by employing human-like social cues and communication modalities. Social behavior is only modeled at the interface, which usually results in shallow models of social cognition.

Socially receptive. Robots that are socially passive but that can benefit from interaction (e.g. learning skills by imitation). Deeper models of human social

competencies are required than with social interface robots.

Sociable. Robots that pro-actively engage with humans in order to satisfy internal social aims (drives, emotions, etc.). These robots require deep models of social cognition.

Complementary to this list we can add the following three classes:

Socially situated. Robots that are surrounded by a social environment that they perceive and react to [48]. Socially situated robots must be able to distinguish between other social agents and various objects in the environment.[1]

Socially embedded. Robots that are: (a) situated in a social environment and interact with other agents and humans; (b) structurally coupled with their social environment; and (c) at least partially aware of human interactional structures (e.g., turn-taking) [48].

Socially intelligent. Robots that show aspects of human style social intelligence, based on deep models of human cognition and social competence [38,40].

1.3. Socially interactive robots

For the purposes of this paper, we use the term "socially interactive robots" to describe robots for which social interaction plays a key role. We do this, not to introduce another class of social robot, but rather to distinguish these robots from other robots that involve "conventional" human–robot interaction, such as those used in teleoperation scenarios.

In this paper, we focus on peer-to-peer human–robot interaction. Specifically, we describe robots that exhibit the following "human social" characteristics:

- express and/or perceive emotions;
- communicate with high-level dialogue;
- learn/recognize models of other agents;
- establish/maintain social relationships;
- use natural cues (gaze, gestures, etc.);
- exhibit distinctive personality and character;
- may learn/develop social competencies.

Socially interactive robots can be used for a variety of purposes: as research platforms, as toys, as educational tools, or as therapeutic aids. The common,

---

[1] Other researchers place different emphasis on what *socially situated* implies (e.g., [97]).

underlying assumption is that humans prefer to interact with machines in the same way that they interact with other people. A survey and taxonomy of current applications is given in [60].

Socially interactive robots operate as partners, peers or assistants, which means that they need to exhibit a certain degree of adaptability and flexibility to drive the interaction with a wide range of humans. Socially interactive robots can have different shapes and functions, ranging from robots whose sole purpose and only task is to engage people in social interactions (Kismet, Cog, etc.) to robots that are engineered to adhere to social norms in order to fulfill a range of tasks in human-inhabited environments (Pearl, Sage, etc.) [18,117,127,140].

Some socially interactive robots use deep models of human interaction and pro-actively encourage social interaction. Others show their social competence only in reaction to human behavior, relying on humans to attribute mental states and emotions to the robot [39,45,55,125]. Regardless of function, building a socially interactive robot requires considering the human in the loop: as designer, as observer, and as interaction partner.

1.4. Why socially interactive robots?

Socially interactive robots are important for domains in which robots must exhibit peer-to-peer interaction skills, either because such skills are required for solving specific tasks, or because the primary function of the robot is to interact socially with people. A discussion of application domains, design spaces, and desirable social skills for robots is given in [42,43].

One area where social interaction is desirable is that of "robot as persuasive machine" [58], i.e., the robot is used to change the behavior, feelings or attitudes of humans. This is the case when robots mediate human–human interaction, as in autism therapy [162]. Another area is "robot as avatar" [123], in which the robot functions as a representation of, or representative for, the human. For example, if a robot is used for remote communication, it may need to act socially in order to effectively convey information.

In certain scenarios, it may be desirable for a robot to develop its interaction skills over time. For example, a pet robot that accompanies a child through his childhood may need to improve its skills in order to maintain the child's interest. Learned development of social (and other) skills is a primary concern of epigenetic robotics [44,169].

Some researchers design socially interactive robots simply to study embodied models of social behavior. For this use, the challenge is to build robots that have an intrinsic notion of sociality, that develop social skills and bond with people, and that can show empathy and true understanding. At present, such robots remain a distant goal [39,44], the achievement of which will require contributions from other research areas such as artificial life, developmental psychology and sociology [133].

Although socially interactive robots have already been used with success, much work remains to increase their effectiveness. For example, in order for socially interactive robots to be accepted as "natural" interaction partners, they need more sophisticated social skills, such as the ability to recognize social context and convention.

Additionally, socially interactive robots will eventually need to support a wide range of users: different genders, different cultural and social backgrounds, different ages, etc. In many current applications, social robots engage only in short-term interaction (e.g., a museum tour) and can afford to treat all humans in the same manner. But, as soon as a robot becomes part of a person's life, that robot will need to be able to treat him as a distinct individual [40].

In the following, we closely examine the concepts raised in this introductory section. We begin by describing different design methods. Then, we present a taxonomy of system components, focusing on the design issues unique to socially interactive robots. We conclude by discussing open issues and core challenges.

2. Methodology

2.1. Design approaches

Humans are experts in social interaction. Thus, if technology adheres to human social expectations, people will find the interaction enjoyable, feeling empowered and competent [130]. Many researchers, therefore, explore the design space of anthropomorphic (or zoomorphic) robots, trying to endow their

creations with characteristics of intentional agents. For this reason, more and more robots are being equipped with faces, speech recognition, lip-reading skills, and other features and capacities that make robot–human interaction "human-like" or at least "creature-like" [41,48].

From a design perspective, we can classify how socially interactive robots are built in two primary ways. With the first approach, "biologically inspired", designers try to create robots that internally simulate, or mimic, the social intelligence found in living creatures. With the second approach, "functionally designed", the goal is to construct a robot that appears outwardly to be socially intelligent, even if the internal design does not have a basis in science.

Robots have limited perceptual, cognitive, and behavioral abilities compared to humans. Thus, for the foreseeable future, there will continue to be significant imbalance in social sophistication between human and robot [20]. As with expert systems, however, it is possible that robots may become highly sophisticated in restricted areas of socialization, e.g., infant-caretaker relations.

Finally, differences in design methodology means that the evaluation and success criteria are almost always different for different robots. Thus, it is hard to compare socially interactive robots outside of their target environment and use.

2.1.1. Biologically inspired

With the "biologically inspired" approach, designers try to create robots that internally simulate, or mimic, the social behavior or intelligence found in living creatures. Biologically inspired designs are based on theories drawn from natural and social sciences, including anthropology, cognitive science, developmental psychology, ethology, sociology, structure of interaction, and theory of mind. Generally speaking, these theories are used to guide the design of robot cognitive, behavioral, motivational (drives and emotions), motor and perceptual systems.

Two primary arguments are made for drawing inspiration from biological systems. First, numerous researchers contend that nature is the best model for "life-like" activity. The hypothesis is that in order for a robot to be understandable by humans, it must have a naturalistic embodiment, it must interact with its environment in the same way living creatures do, and

it must perceive the same things that humans find to be salient and relevant [169].

The second rationale for biological inspiration is that it allows us to directly examine, test and refine those scientific theories upon which the design is based [1]. This is particularly true with humanoid robots. Cog, for example, is a general-purpose humanoid platform intended for exploring theories and models of intelligent behavior and learning [140].

Some of the theories commonly used in biologically inspired design are as follows.

Ethology. This refers to the observational study of animals in their natural setting [95]. Ethology can serve as a basis for design because it describes the types of activity (comfort-seeking, play, etc.) a robot needs to exhibit in order to appear life-like [4]. Ethology is also useful for addressing a range of behavioral issues such as concurrency, motivation, and instinct.

Structure of interaction. Analysis of interactional structure (such as instruction, cooperation, etc.) can help focus design of perception and cognition systems by identifying key interaction patterns [162]. Dautenhahn, Ogden and Quick use explicit representations of interactional structure to design "interaction aware" robots [48]. Dialogue models, such as turn-taking in conversation, can also be used in design as in [104].

Theory of mind. Theory of mind refers to those social skills that allow humans to correctly attribute beliefs, goals, perceptions, feelings, and desires to themselves and others [163]. One of the critical precursors to these skills is joint (or shared) attention: the ability to selectively attend to an object of mutual interest [7]. Joint attention can aid design, by providing guidelines for recognizing and producing social behaviors such as gaze direction, pointing gestures, etc. [23,140].

Developmental psychology. Developmental psychology has been cited as an effective mechanism for creating robots that engage in natural social exchanges. As an example, the design of Kismet's "synthetic nervous system", particularly the perceptual and behavioral aspects, is heavily inspired by the social development of human infants [18]. Additionally, theories of child cognitive development, such as Vygotsky's "child in society" [92], can offer a framework for constructing robot architecture and social interaction design [44].

2.1.2. Functionally designed

With the "functionally designed" approach, the objective is to design a robot that outwardly appears to be socially intelligent, even if the internal design does not have a basis in science or nature. This approach assumes that if we want to create the impression of an artificial social agent driven by beliefs and desires, we do not necessarily need to understand how the mind really works. Instead, it is sufficient to describe the mechanisms (sensations, traits, folk-psychology, etc.) by which people in everyday life understand socially intelligent creatures [125].

In contrast to their biologically inspired counterparts, functionally designed robots generally have constrained operational and performance objectives. Consequently, these "engineered" robots may need only to generate certain effects and experiences with the user, rather than having to withstand deep scrutiny for "life-like" capabilities.

Some motivations for "functional design" are:

- The robot may only need to be superficially socially competent. This is particularly true when only short-term interaction or limited quality of interaction is required.
- The robot may have limited embodiment, capability for interaction, or may be constrained by the environment.
- Even limited social expression can help improve the affordances and usability of a robot. In some applications, recorded or scripted speech may be sufficient for human–robot dialogue.
- Artificial designs can provide compelling interaction. Many video games and electronic toys fully engage and occupy attention, even if the artifacts do not have real-world counterparts.

The three techniques most often used in functional design are as follows.

Human–computer interaction (HCI) design. Robots are increasingly being developed using HCI techniques, including cognitive modeling, contextual inquiry, heuristic evaluation, and empirical user testing [115]. Scheeff et al. [142], for example, describe robot development based on heuristic design.

Systems engineering. Systems engineering involves the development of functional requirements to facilitate development and operation [135]. A basic charac-

teristic of system engineering is that it emphasizes the design of critical-path elements. Pineau et al. [127], for example, describe mobile robots that assist the elderly. Because these robots operate in a highly structured domain, their design centers on specific task behaviors (e.g., navigation).

Iterative design. Iterative (or sequential) design, is the process of revising a design through a series of test and redesign cycles [147]. It is typically used to address design failures or to make improvements based on information from evaluation or use. Willeke et al. [164], for example, describe a series of museum robots, each of which was designed based on lessons learned from preceding generations.

2.2. Design issues

All robot systems, whether socially interactive or not, must solve a number of common design problems. These include cognition (planning, decision making), perception (navigation, environment sensing), action (mobility, manipulation), human–robot interaction (user interface, input devices, feedback display) and architecture (control, electromechanical, system). Socially interactive robots, however, must also address those issues imposed by social interaction [18,40].

Human-oriented perception. A socially interactive robot must proficiently perceive and interpret human activity and behavior. This includes detecting and recognizing gestures, monitoring and classifying activity, discerning intent and social cues, and measuring the human's feedback.

Natural human–robot interaction. Humans and robots should communicate as peers who know each other well, such as musicians playing a duet [145]. To achieve this, the robot must manifest believable behavior: it must establish appropriate social expectations, it must regulate social interaction (using dialogue and action), and it must follow social convention and norms.

Readable social cues. A socially interactive robot must send signals to the human in order to: (1) provide feedback of its internal state; (2) allow human to interact in a facile, transparent manner. Channels for emotional expression include facial expression, body and pointer gesturing, and vocalization.

Real-time performance. Socially interactive robots must operate at human interaction rates. Thus, a robot

needs to simultaneously exhibit competent behavior, convey attention and intentionality, and handle social interaction.

In the following sections, we review design issues that are unique to socially interactive robots. Although we do not discuss every aspect of design, we feel that addressing each of the following is critical to building an effective social robot.

2.3. Embodiment

We define embodiment as "that which establishes a basis for structural coupling by creating the potential for mutual perturbation between system and environment" [48]. Thus, embodiment is grounded in the relationship between a system and its environment. The more a robot can perturb an environment, and be perturbed by it, the more it is embodied. This also means that social robots do not necessarily need a *physical* body. For example, conversational agents [33] might be embodied to the same extent as robots with limited actuation.

An important benefit of this "relational definition" is that it provides an opportunity to quantify embodiment. For example, one might measure embodiment in terms of the complexity of the relationship between robot and environment over all possible interactions (i.e., all perturbatory channels).

Some robots are clearly more embodied than others [48]. Consider the difference between Aibo (Sony) and Khepera (K-Team), as shown in Fig. 7. Aibo has approximately 20 actuators (joints across mouth, heads, ears, tails, and legs) and a variety of sensors (touch, sound, vision and proprioceptive). In contrast, Khepera has two actuators (independent wheel control) and an array of infrared proximity sensors. Because Aibo has more perturbatory channels and bandwidth at its disposal than does Khepera, it can be considered to be more strongly embodied than Khepera.

2.3.1. Morphology

The form and structure of a robot is important because it helps establish social expectations. Physical appearance biases interaction. A robot that resembles a dog will be treated differently (at least initially) than one which is anthropomorphic. Moreover, the relative familiarity (or strangeness) of a robot's morphology


Fig. 7. Sony Aibo ERS-110 (top) and K-Team Khepera (bottom).

can have profound effects on its accessibility, desirability, and expressiveness.

The choice of a given form may also constrain the human's ability to interact with the robot. For example, Kismet has a highly expressive face. But because it is designed as a head, Kismet is unable to interact when touch (e.g., manipulation) or displacement (self-movement) is required.

To date, most research in human–robot interaction has not explicitly focused on design, at least not in the traditional sense of industrial design. Although knowledge from other areas of design (including product, interaction and stylized design) can inform robot construction, much research remains to be performed.

2.3.2. Design considerations

A robot's morphology must match its intended function [54]. If a robot is designed to perform tasks for the human, then its form must convey an amount of "product-ness" so that the user will feel comfortable

Fig. 8. Mori's "uncanny valley" (from DiSalvo et al. [54]).

using the robot. Similarly, if peer interaction is important, the robot must project an amount of "humanness" so that the user will feel comfortable in socially engaging the robot.

At the same time, however, a robot's design needs to reflect an amount of "robot-ness". This is needed so that the user does not develop detrimentally false expectations of the robot's capabilities [55].

Finally, if a robot needs to portray a living creature, it is critical that an appropriate degree of familiarity be maintained. Mashiro Mori contends that the progression from a non-realistic to realistic portrayal of a living thing is non-linear. In particular, there is an "uncanny valley" (see Fig. 8) as similarity becomes almost, but not quite perfect. At this point, the subtle imperfections of the recreation become highly disturbing, or even repulsive [131]. Consequently, caricatured representations may be more useful, or effective, than more complex, "realistic" representations.

We classify social robots as being embodied in four broad categories: anthropomorphic, zoomorphic, caricatured, and functional.

2.3.3. Anthropomorphic

Anthropomorphism, from the Greek "anthropos" for man and "morphe" for form/structure, is the tendency to attribute human characteristics to objects with a view to helping rationalize their actions [55]. Anthropomorphic paradigms have widely been used to augment the functional and behavioral characteristics of social robots.

Having a naturalistic embodiment is often cited as necessary for meaningful social interaction [18,82,140]. In part, the argument is that for a robot to interact with humans as humans do (through gaze,

gesture, etc.), then it must be structurally and functionally similar to a human. Moreover, if a robot is to learn from humans (e.g., through imitation), then it should be capable of behaving similarly to humans [11].

The role of anthropomorphism is to function as a mechanism through which social interaction can be facilitated. Thus, the ideal use of anthropomorphism is to present an appropriate balance of illusion (to lead the user to believe that the robot is sophisticated in areas where the user will not encounter its failings) and functionality (to provide capabilities necessary for supporting human-like interaction) [54,79].

2.3.4. Zoomorphic

An increasing number of entertainment, personal, and toy robots have been designed to imitate living creatures. For these robots, a zoomorphic embodiment is important for establishing human–creature relationships (e.g., owner-pet). The most common designs are inspired by household animals, such as dogs (Sony Aibo and RoboScience RoboDog) and cats (Omron), with the objective of creating robotic "companions".

Avoiding the "uncanny valley" may be easier with zoomorphic design because human–creature relationships are simpler than human–human relationships and because our expectations of what constitutes "realistic" animal morphology tends to be lower.

2.3.5. Caricatured

Animators have long shown that a character does not have to appear realistic in order to be believable [156]. Moreover, caricature can be used to create desired interaction biases (e.g., implied abilities) and to focus attention on, or distract attention from, specific robot features.

Scheeff et al. [142] discusses how techniques from traditional animation can be used in social robot design. Schulte et al. [143] describe how a caricatured human face can provide a "focal point" for attention. Similarly, Severinson-Eklund et al. [144] describe the use of a small mechanical character, CERO, as a robot "representative" (see Fig. 9).

2.3.6. Functional

Some researchers argue that a robot's embodiment should first, and foremost, reflect the tasks it must perform. The choice and design of physical features is thus guided purely by operational objectives. This type

Fig. 9. CERO (KTH).

of embodiment appears most often with functionally designed robots, especially service robots.

Health care robots, for example, may be required to assist elderly, or disabled, patients in moving about. Thus, features such as handle bars and cargo space, may need to be part of the design [127].

The design of toy robots also tends to reflect functional requirements. Toys must minimize production cost, be appealing to children, and be capable of facing the wide variety of situations that can experienced during play [107].

2.4. Emotion

Emotions play a significant role in human behavior, communication and interaction. Emotions are complex phenomena and often tightly coupled to social context [5]. Moreover, much of emotion is physiological and depends on embodiment [122,126].

Three primary theories are used to describe emotions. The first approach describes emotions in terms of discrete categories (e.g., happiness). A good review of "basic emotions" is [57]. The second approach characterizes emotions using continuous scales or basis dimensions, such as *arousal* and *valence* [137]. The third approach, componential theory, acknowledges the importance of both categories and dimensions [128,147].

In recent years, emotion has increasingly been used in interface and robot design, primarily because of the recognition that people tend to treat computers as they treat other people [31,33,121,130]. Moreover, many studies have been performed to integrate emotions into products including electronic games, toys, and software agents [8].

2.4.1. Artificial emotions

Artificial emotions are used in social robots for several reasons. The primary purpose, of course, is that emotion helps facilitate believable human–robot interaction [30,119]. Artificial emotion can also provide feedback to the user, such as indicating the robot's internal state, goals and (to an extent) intentions [8,17,83]. Lastly, artificial emotions can act as a control mechanism, driving behavior and reflecting how the robot is affected by, and adapts to, different factors over time [29,108,160].

Numerous architectures have been proposed for artificial emotions [18,29,74,132,160]. Some closely follow emotional theory, particularly in terms of how emotions are defined and generated. Arkin et al. [4] discuss how ethological and componential emotion models are incorporated into Sony's entertainment robots. Cañamero and Fredslund [30] describe an affective activation model that regulates emotions through stimulation levels.

Other architectures are only loosely inspired by emotional theory and tend to be designed in an ad hoc manner. Nourbakhsh et al. [117] detail a fuzzy state machine based system, which was developed through a series of formative evaluation and design cycles. Schulte et al. [143] summarize the design of a simple state machine that produces four basic "moods".

In terms of expression, some robots are only capable of displaying emotion in a limited way, such as individually actuated lips or flashing lights (usually LEDs). Other robots have many active degrees of freedom and can thus provide richer movement and gestures. Kismet, for example, has controllable eyebrows, ears, eyeballs, eyelids, a mouth with two lips and a pan/tilt neck [18].

2.4.2. Emotions as control mechanism

Emotion can be used to determine control precedence between different behavior modes, to

coordinating planning, and to trigger learning and adaptation, particularly when the environment is unknown or difficult to predict. One approach is to use computational models of emotions that mimic animal survival instincts, such as escape from danger, look for food, etc. [18,29,108,160].

Several researchers have investigated the use of emotion in human–robot interaction. Suzuki et al. [153] describe an architecture in which interaction leads to changes in the robot's emotional state and modifications in its actions. Breazeal [18] discusses how emotions influence the operation of Kismet's motivational system and how this affects its interaction with humans. Nourbakhsh et al. [117] discusses how mood changes can trigger different behavior in Sage, a museum tour robot.

2.4.3. Speech

Speech is a highly effective method for communicating emotion. The primary parameters that govern the emotional content of speech are loudness, pitch (level, variation, range), and prosody. Murray and Arnott [111] contend that the vocal effects caused by particular emotions are consistent between speakers, with only minor differences.

The quality of synthesized speech is significantly poorer than synthesized facial expression and body language [9]. In spite of this shortcoming, it has proved possible to generate emotional speech. Cahn [28] describes a system for mapping emotional quality (e.g., sorrow) onto speech synthesizer settings, including articulation, pitch, and voice quality.

To date, emotional speech has been used in few robot systems. Breazeal describes the design of Kismet's vocalization system. Expressive utterances (used to convey the affective state of the robot without grammatical structure) are generated by assembling strings of phonemes with pitch accents [18]. Nourbakhsh et al. [117] describe how emotions influence synthesized speech in a tour guide robot. When the robot is frustrated, for example, voice level and pitch are increased.

2.4.4. Facial expression

The expressive behavior of robotic faces is generally not life-like. This reflects limitations of mechatronic design and control. For example, transitions between expressions tend to be abrupt, occurring suddenly and


Fig. 10. Actuated faces: Sparky (left) and Feelix (right).

rapidly, which rarely occurs in nature. The primary facial components used are mouth (lips), cheeks, eyes, eyebrows and forehead. Most robot faces express emotion in accordance with Ekman and Frieser's FACS system [56,146].

Two of the simplest faces (Fig. 10) appear on Sparky [142] and Feelix [30]. Sparky's face has 4-DOF (eyebrows, eyelids, and lips) which portray a set of discrete, basic emotions. Feelix is a robot built using the LEGO Mindstorms™ robotic construction kit. Feelix's face also has 4-DOF (two eyebrows, two lips), designed to display six facial expressions (anger, sadness, fear, happiness, surprise, neutral) plus a number of blends.

In contrast to Sparky and Feelix, Kismet's face has fifteen actuators, many of which often work together to display specific emotions (see Fig. 11). Kismet's facial expressions are generated using an interpolation-based technique over a three-dimensional, componential "affect space" (arousal, valence, and stance) [18].

Perhaps the most realistic robot faces are those designed at the Science University of Tokyo [81]. These faces (Fig. 12) are explicitly designed to be human-like and incorporate hair, teeth, and a covering silicone skin layer. Numerous control points actuated beneath the "skin" produce a wide range of facial movements and human expression.

Instead of using mechanical actuation, another approach to facial expression is to rely on computer graphics and animation techniques [99]. Vikia, for

Fig. 11. Various emotions displayed by Kismet.

Fig. 12. Saya face robots (Science University of Tokyo).

example, has a 3D rendered face of a woman based on Delsarte's code of facial expressions [26]. Because Vikia's face (see Fig. 13) is graphically rendered, many degrees of freedom are available for generating expressions.

2.4.5. Body language

In addition to facial expressions, non-verbal communication is often conveyed through gestures and body movement [9]. Over 90% of gestures occur during speech and provide redundant information [86,105]. To date, most studies on emotional body movement have been qualitative in nature. Frijda [62],


Fig. 13. Vikia has a computer generated face.

Table 1
Emotional body movements (adapted from Frijda [62])

Emotion	Body movement
Anger	Fierce glance; clenched fists; brisk, short motions
Fear	Bent head, truck and knees; hunched shoulders; forced eye closure or staring
Happiness	Quick, random movements; smiling;
Sadness	Depressed mouth corners; weeping
Surprise	Wide eyes; held breath; open mouth

for example, described body movements for a number of basic emotions (Table 1). Recently, however, some work has begun to focus on implementation issues, such as in [35].

Nakata et al. [113] state that humans have a strong tendency to be cued by motion. In particular, they refer to analysis of dance that shows humans are emotionally affected by body movement. Breazeal and Fitzpatrick [21] contend humans perceive all motor actions to be semantically rich, whether or not they were intended to be. For example, gaze and body direction is generally interpreted as indicating locus of attention.

Mizoguchi et al. [110] discuss the use of gestures and movements, similar to ballet poses, to show emotion through movement. Scheeff et al. [142] describe the design of smooth, natural motions for Sparky. Lim et al. [93] describe how walking motions (foot dragging, body bending, etc.) can be used to convey emotions.

2.5. Dialogue

2.5.1. What is dialogue?

Dialogue is a joint process of communication. It involves sharing of information (data, symbols, context) between two or more parties [90]. Humans employ a variety of para-linguistic social cues (facial displays, gestures, etc.) to regulate the flow of dialogue [32]. Such cues have also proven to be useful for controlling human–robot dialogue [19].

Dialogue, regardless of form, is meaningful only if it is grounded, i.e., when the symbols used by each party describe common concepts. If the symbols differ, information exchange or learning must take place before communication can proceed. Although human–robot communication can occur in many forms, we consider there to be three primary types of dialogue: low-level (pre-linguistic), non-verbal, and natural language.

Low-level. Billard and Dautenhahn [12–14] describe a number of experiments in which an autonomous mobile robot was taught a synthetic proto-language. Language learning results from multiple spatio-temporal associations across the robot's sensor-actuator state space.

Steels has examined the hypothesis that communication is bootstrapped in a social learning process and that meaning is initially context-dependent [150,151]. In his experiments, a robot dog learns simple words describing the presence of objects (ball, red, etc.), its behavior (walk, sit) and its body parts (leg, head).

Non-verbal. There are many non-verbal forms of language, including body positioning, gesturing, and physical action. Since most robots have fairly rudimentary capability to recognize and produce speech, non-verbal dialogue is a useful alternative. Nicolescu and Mataric [116], for example, describe a robot that asks humans for help, communicating its needs and intentions through its actions.

Social conventions, or norms, can also be expressed through non-verbal dialogue. Proxemics, the social use of space, is one such convention [70]. Proxemic norms include knowing how to stand in line, how to pass in hallways, etc. Respecting these spatial conventions may involve consideration of numerous factors (administrative, cultural, etc.) [114].

Natural language. Natural language dialogue is determined by factors ranging from the physical and perceptual capabilities of the participants, to the social and cultural features of the situation. To what extent human–robot interfaces should be based on natural language remains clearly an open issue [144].

Severinson-Eklund et al. [144] discuss how explicit feedback is needed for users to interact with service robots. Their approach is to provide designed natural language. Fong et al. [59,61] describe how high-level dialogue can enable a human to provide assistance to a robot. In their system, dialogue is limited to mobility issues (navigation, obstacle avoidance, etc.) with an emphasis on query-response speech acts.

2.6. Personality

2.6.1. What is personality?

In psychological terms, personality is the set of distinctive qualities that distinguish individuals. Since the late 1980s, the most widely accepted taxonomy of

personality traits has been the "Big Five Inventory" [76]. The "Big Five", which was developed through lexical analysis, described personality in terms of five traits:

- extroversion (sociable, outgoing, confidence);
- agreeableness (friendliness, nice, pleasant);
- conscientiousness (helpful, hard-working);
- neuroticism (emotional stability, adjusted);
- openness (intelligent, imaginative, flexibility).

Common alternatives to the "Big Five" are questionnaire-based scales such as the *Myers–Briggs Type Indicator (MBTI)* [112].

2.6.2. Personality in social robots

There is reason to believe that if a robot had a compelling personality, people would be more willing to interact with it and to establish a relationship with it [18,79]. In particular, personality may provide a useful affordance, giving users a way to model and understand robot behavior [144].

In designing robot personality, there are numerous questions that need to be addressed. Should the robot have a designed or learned personality? Should it mimic a specific human personality, exhibiting specific traits? Is it beneficial to encourage a specific type of interaction?

There are five common personality types used in social robots.

Tool-like. Used for robots that operate as "smart appliances". Because these robots perform service tasks on command, they exhibit traits usually associated with tools (dependability, reliability, etc.).

Pet or creature. These toy and entertainment robots exhibit characteristics that are associated with domesticated animals (cats, dogs, etc.).

Cartoon. These robots exhibit caricatured personalities, such as seen in animation. Exaggerated traits (e.g., shyness) are easy to portray and can be useful for facilitating interaction with non-specialists.

Artificial being. Inspired by literature and film, primarily science fiction, these robots tend to display artificial (e.g., mechanistic) characteristics.

Human-like. Robots are often designed to exhibit human personality traits. The extent to which a robot must have (or appear to have) human personality depends on its use.

Robot personality is conveyed in numerous ways. Emotions are often used to portray stereotype personalities: timid, friendly, etc. [168]. A robot's embodiment (size, shape, color), its motion, and the manner in which it communicates (e.g., natural language) also contribute strongly [144]. Finally, the tasks a robot performs may also influence the way its personality is perceived.

2.7. Human-oriented perception

To interact meaningfully with humans, social robots must be able to perceive the world as humans do, i.e., sensing and interpreting the same phenomena that humans observe. This means that, in addition to the perception required for conventional functions (localization, navigation, obstacle avoidance), social robots must possess perceptual abilities similar to humans.

In particular, social robots need perception that is human-oriented: optimized for interacting with humans and on a human level. They need to be able to track human features (faces, bodies, hands). They also need to be capable of interpreting human speech including affective speech, discrete commands, and natural language. Finally, they often must have the capacity to recognize facial expressions, gestures, and human activity.

Similarity of perception requires more than similarity of sensors. It is also important that humans and robots find the same types of stimuli salient [23]. Moreover, robot perception may need to mimic the way human perception works. For example, the human ocular-motor system is based on foveate vision, uses saccadic eye movements, and exhibits specific visual behaviors (e.g., glancing). Thus, to be readily understood, a robot may need to have similar visual motor control [18,21,25].

2.7.1. Types of perception

Most human-oriented perception is based on passive sensing, typically computer vision and spoken language recognition. Passive sensors, such as CCD cameras, are cheap, require minimal infrastructure, and can be used for a wide range of perception tasks [2,36,66,118].

Active sensors (ladar, ultrasonic sonar, etc.), though perhaps less flexible than their passive counterparts, have also received attention. In particular, active

sensors are often used for detecting and localizing human in dynamic settings.

2.7.2. People tracking

For human–robot interaction, the challenge is to find efficient methods for people tracking in the presence of occlusions, variable illumination, moving cameras, and varying background. A broad survey of human tracking is presented in [66]. Specific robotics applications can be reviewed in [26,114,127,154].

2.7.3. Speech recognition

Speech recognition is generally a two-step process: signal processing (to transform an audio signal into feature vectors) followed by graph search (to match utterances to a vocabulary). Most current systems use Hidden Markov models to stochastically determine the most probable match. An excellent introduction to speech recognition is [129].

Human speech contains three types of information: who the speaker is, what the speaker said, and how the speaker said it [18]. Depending on what information the robot requires, it may need to perform speaker tracking, dialogue management, or emotion analysis. Recent applications of speech in robotics include [18,91,103,120,148].

2.7.4. Gesture recognition

When humans converse, we use gestures to clarify speech and to compactly convey geometric information (location, direction, etc.). Very often, a speaker will use hand movement (speed and range of motion) to indicate urgency and will point to disambiguate spoken directions (e.g., "I parked the car over there").

Although there are many ways to recognize gestures, vision-based recognition has several advantages over other methods. Vision does not require the user to master or wear special hardware. Additionally, vision is passive and can have a large workspace. Two excellent overviews of vision-based gesture recognition are [124,166]. Details of specific systems appear in [85,161,167].

2.7.5. Facial perception

Face detection and recognition. A widely used approach for identifying people is face detection. Two comprehensive surveys are [34,63]. A large number of real-time face detection and tracking systems have been developed in recent years, such as [139,140,158].

Facial expression. Since Darwin [37], facial expressions have been considered to convey emotion. More recently, facial expressions have also been thought to function as social signals of intent. A comprehensive review of facial expression recognition (including a review of ethical and psychological concerns) is [94]. A survey of older techniques is [136].

There are three basic approaches to facial expression recognition [94]. Image motion techniques identify facial muscle actions in image sequences. Anatomical models track facial features, such as the distance between eyes and nose. Principal component analysis (PCA) reduce image-based representations of faces into principal components such as eigenfaces or holons.

Gaze tracking. Gaze is a good indicator of what a person is looking at and paying attention to. A person's gaze direction is determined by two factors: head orientation and eye orientation. Although numerous vision systems track head orientation, few researchers have attempted to track eye gaze using only passive vision. Furthermore, such trackers have not proven to be highly accurate [158]. Gaze tracking research includes [139,152].

2.8. User modeling

In order to interact with people in a human-like manner, socially interactive robots must perceive human social behavior [18]. Detecting and recognizing human action and communication provides a good starting point. More important, however, is being able to interpret and reacting to behavior. A key mechanism for performing this is user modeling.

User modeling can be quantitative, based on the evaluation of parameters or metrics. The stereotype approach, for example, classifies users into different subgroups (stereotypes), based on the measurement of pre-defined features for each subgroup [155]. User modeling may also be qualitative in nature. Interactional structure analysis, story and script based matching, and BDI all identify subjective aspects of behavior.

There are many types of user models: cognitive, attentional, etc. A user model generally contains attributes that describe a user, or group, of users. Models may be static (defined a priori) or dynamic (adapted or learned). Information about users may be acquired

explicitly (through questioning) or implicitly (inferred through observation). The former can be time consuming, and the latter difficult, especially if the user population is diverse [69].

User models are employed for a variety of purposes. First, user models help the robot understand human behavior and dialogue. Second, user models shape and control feedback (e.g., interaction pacing) given to the user. Finally, user models are useful for adapting the robot's behavior to accommodate users with varying skills, experience, and knowledge.

Fong et al. [59] employ a stereotype user model to adapt human–robot dialogue and robot behavior to different users. Pineau et al. discuss the use of a quantitative temporal Bayes net to manage individual-specific interaction between a nurse robot and elderly individuals. Schulte et al. [143] describe a memory-based learner used by a tour robot to improve its ability to interact with different people.

2.9. Socially situated learning

In socially situated learning, an individual interacts with his social environment to acquire new competencies. Humans and some animals (e.g., primates) learn through a variety of techniques including direct tutelage, observational conditioning, goal emulation, and imitation [64]. One prevalent form of influence is local, or stimulus, enhancement in which a teacher actively manipulates the perceived environment to direct the learner's attention to relevant stimuli [96].

2.9.1. Robot social learning

For social robots, learning is used for transferring skills, tasks, and information. Learning is important because the knowledge of the teacher, or model, and robot may be very different. Additionally, because of differences in sensing and perception, the model and robot may have very different views of the world. Thus, learning is often essential for improving communication, facilitating interaction, and sharing knowledge [80].

A number of studies in robot social learning have focused on robot–robot interaction. Some of the earliest work focused on cooperative, or group, behavior [6,100]. A large research community continues to investigate group social learning, often referred to

as "swarm intelligence" and "collective robotics". Other robot–robot work has addressed the use of "leader following" [38,72], inter-personal communication [13,15,149], imitation [14,65], and multi-robot formations [109].

In recent years, there has been significant effort to understand how social learning can occur through human–robot interaction. One approach is to create sequences of known behaviors to match a human model [102]. Another approach is to match observations (e.g., motion sequences) to known behaviors, such as motor primitives [51,52]. Recently, Kaplan et al. [77] have explored the use of animal training techniques for teach an autonomous pet robot to perform complex tasks. The most common social learning method, however, is imitation.

2.9.2. Imitation

Imitation is an important mechanism for learning behaviors socially in primates and other animal species [46]. At present, there is no commonly accepted definition of "imitation" in the animal and human psychology literature. An extensive discussion is given in [71]. Researchers often refer to Thorpe's definition [157], which defines imitation as the "copying of a novel or otherwise improbable act or utterance, or some act for which there is clearly no instinctive tendency".

With robots, imitation relies upon the robot having many perceptual, cognitive, and motor capabilities [24]. Researchers often simplify the environment or situation to make the problem tractable. For example, active markers or constrained perception (e.g., white objects on a black background) may be employed to make tracking of the model amenable.

Breazeal and Scassellati [24] argue that even if a robot has the skills necessary for imitation, there are still several questions that must be addressed:

- *How does the robot know when to imitate*? In order for imitation to be useful, the robot must decide not only when to start/stop imitating, but also when it is appropriate (based on the social context, the availability of a good model, etc.).
- *How does the robot know what to imitate*? Faced with a stream of sensory data, the robot must decide which of the model's actions are relevant to the task, which are part of the instruction process, and which are circumstantial.

- *How does the robot map observed action into behavior*? Once the robot has identified and observed salient features of the model's actions, it must ascertain how to reproduce these actions through its behavior.
- *How does the robot evaluate its behavior, correct errors, and recognize when it has achieved its goal*? In order for the robot to improve its performance, it must be able to measure to what degree its imitation is accurate and to recognize when there are errors.

Imitation has been used as a mechanism for learning simple motor skills from observation, such as block stacking [89] or pendulum balancing [141]. Imitation has also been applied to the learning of sensor–motor associations [3] and for constructing task representations [116].

2.10. Intentionality

Dennett [50] contends that humans use three strategies to understand and predict behavior. The *physical stance* (predictions based on physical characteristics) and *design stance* (predictions based on the design and functionality of artifacts) are sufficient to explain simple devices. With complex systems (e.g., humans), however, we often do not have sufficient information, to perform physical or design analysis. Instead, we tend to adopt an *intentional stance* and assume that the systems' actions result from its beliefs and desires.

In order for a robot to interact socially, therefore, it needs to provide evidence that is intentional (even if it is not intrinsic [138]). For example, a robot could demonstrate goal-directed behaviors, or it could exhibit the attentional capacity. If it does so, then the human will consider the robot to act in a rational manner.

2.10.1. Attention

Scassellati [139] discusses the recognition and production of joint attention behaviors in Cog. Just as humans use a variety of physical social cues to indicate which object is currently under consideration, Cog performs gaze following, imperative pointing, and declarative pointing.

Kopp and Gardenfors [84] also claim that attentional capacity is a fundamental requirement for intentionality. In their model, a robot must be able to identify relevant objects in the scene, direct its sen-

sors towards an object, and maintain its focus on the selected object.

Marom and Hayes [96–98] consider attention to be a collection of mechanisms that determine the significance of stimuli. Their research focuses on the development of pre-learning attentional mechanisms, which help reduce the amount of information that an individual has to deal with.

2.10.2. Expression

Kozima and Yano [82,83] argue that to be intentional, a robot must exhibit goal-directed behavior. To do so, it must possess a sensorimotor system, a repertoire of behaviors (innate reflexes), drives that trigger these behaviors, a value system for evaluating perceptual input, and an adaptation mechanism.

Breazeal and Scassellati [22] describe how Kismet conveys intentionality through motor actions and facial expressions. In particular, by exhibiting proto-social responses (affective, exploratory, protective, and regulatory), the robot provides cues for interpreting its actions as intentional.

Schulte et al. [143] discuss how a caricatured human face and simple emotion expression can convey intention during spontaneous short-term interaction. For example, a tour guide robot might have the intention of making progress while giving a tour. Its facial expression and recorded speech playback can communicate this information.

3. Discussion

3.1. Human perception of social robots

A key difference between conventional and socially interactive robots is that the way in which a human perceives a robot establishes expectations that guide his interaction with it. This perception, especially of the robot's intelligence, autonomy, and capabilities is influenced by numerous factors, both intrinsic and extrinsic.

Clearly, the human's preconceptions, knowledge, and prior exposure to the robot (or similar robots) have a strong influence. Additionally, aspects of the robot's design (embodiment, dialogue, etc.) may play a significant role. Finally, the human's experience over time will undoubtedly shape his judgment, i.e., initial

impressions will change as he gains familiarity with the robot.

In the following, we briefly present studies that have examined how these factors affect human–robot interaction, particularly in the way in which the humans relate to, and work with, social robots.

3.1.1. Attitudes towards robots

Bumby and Dautenhahn conducted a study to identify how people, specifically children, perceive robots and what type of behavior they may exhibit when interacting with robots [27]. They found that children tend to conceive of robots as geometric forms with human features (i.e., a strong anthropomorphic pre-disposition). Moreover, children tend to attribute free will, preferences, emotion, and male gender to the robots, even without external cueing.

In [78], Khan describes a survey to investigate people's attitudes towards intelligent service robots. A review of robots in literature and film, followed by a interview study, were used to design the survey questionnaire. The survey revealed that people's attitudes are strongly influenced by science fiction. Two significant findings were: (1) a robot with machine-like appearance, serious personality, and round shape is preferred; (2) verbal communication using a human-like voice is highly desired.

3.1.2. Field studies

Thus far, few studies have investigated people's willingness to closely interact with social robots. Given that we expect social robots to play increasingly larger roles in daily life, there is a strong need for field studies to examine how people behave when robots are introduced into their activities.

Scheeff et al. [142] conducted two studies to observe how a range of people interact with a creature-like social robot, in both laboratory and public conditions. In these studies, children were observed to be more engaged than adults and had responses that varied with gender and age. Also, a friendly robot personality was reported to have prompted qualitatively better interaction than an angry personality.

In [75], Huttenrauch and Severinson-Eklund describe a long-term usage study of CERO, a service robot that assists motion-impaired people in an office environment (Fig. 9). The study was designed to observe interaction over time, especially after the user had fully integrated the robot into his work routine. A key finding was that robots need to be capable of social interaction, or at least aware of the social context, whenever they operate around people.

In [47], Dautenhahn and Werry describe a quantitative method for evaluating robot–human interactions, which is similar to the way ethologists use observation to evaluate animal behavior. This method has been used to study differences in interaction style when children play with a socially interactive robotic toy versus a non-robotic toy. Complementing this approach, Dautenhahn et al. [49] have also proposed qualitative techniques (based on conversation analysis) that focus on social context.

3.1.3. Effects of emotion

Cañamero and Fredslund [30] performed a study to evaluate how well humans can recognize facial expressions displayed by Feelix (Fig. 10). In this study, they asked test subjects to make subjective judgments of the emotions displayed on Feelix's face and in pictures of humans. The results were very similar to those reported in other studies of facial expression recognition.

Bruce et al. [26] conducted a 2×2 full factorial experiment to explore how emotion expression and indication of attention affect a robot's ability to engage humans. In the study, the robot exhibited different emotions based on its success at engaging and leading a person through a poll-taking task. The results suggest that having an expressive face and indicating attention with movement can help make a robot more compelling to interact with.

3.1.4. Effects of appearance and dialogue

One problem with dialogue is that it can lead to biased perceptions. For example, associations of stereotyped behavior can be created, which may lead users to attribute qualities to the robot that are inaccurate. Users may also form incorrect models, or make poor assumptions, about how the robot actually works. This can lead to serious consequences, the least of which is user error [61].

Kiesler and Goetz conducted a series of studies to understand the influence of a robot's appearance and dialogue [79]. A primary contribution of this work are

measures for characterizing the mental models used by people when they interact with robots. A significant finding was that neither ratings, nor behavioral observations alone, are sufficient to fully describe human responses to robots. In addition, Kiesler and Goetz concluded that dialogue more strongly influences development and change of mental models than differences in appearance.

DiSalvo et al. [54] investigated how the features and size of humanoid robot faces contribute to the perception of humanness. In this study, they analyzed 48 robots and conducted surveys to measure people's perception. Statistical analysis showed that the presence of certain features, the dimensions of the head, and the number of facial features greatly influence the perception of humanness.

3.1.5. Effects of personality

When a robot exhibits personality (whether intended by the designer or not), a number of effects occur. First, personality can serve as an affordance for interaction. A growing number of commercial products targeting the toy and entertainment markets, such as Tiger Electronics Furby (a creature-like robot), Hasbro's My Real Baby (a robot doll), and Sony's Aibo (robot dog) focus on personality as a way to entice and foster effective interaction [18,60].

Personality can also impact task performance, in either a negative or positive sense. For example, Goetz and Kiesler examined the influence of two different robot personalities on user compliance with an exercise routine [67]. In their study, they found some evidence that simply creating a charming personality will not necessarily engender the best cooperation with a robotic assistant.

3.2. Open issues and questions

When we engage in social interaction, there is no guarantee that it will be meaningful or worthwhile. Sometimes, in spite of our best intentions, the interaction fails. Relationships, especially long-term ones, involve a myriad of factors and making them succeed requires concerted effort.

In [165], Woods writes:

It seems paradoxical, but studies of the impact of automation reveal that design of automated systems is really the design of a new human–machine cooperative system. The design of automated systems is really the design of a team and requires provisions for the coordination between machine agents and practitioners.

In other words, humans and robots must be able to coordinate their actions so that they interact *productively* with each other. It is not appropriate (or even necessary) to make the robot as socially competent as possible. Rather, it is more important that the robot be compatible with the human's needs, that it matches application requirements; that it be understandable and believable, and that it provide the interactional support the human expects.

As we have seen, building a social robot involves numerous design issues. Although much progress has already been made to solving these problems, much work remains. This is due, in part, to the broad range of applications for which social robots are being developed. Additionally, however, is the fact that there are many research questions that remain to be answered, including the following.

What are the minimal criteria for a robot to be social? Social behavior includes such a wide range of phenomena that it is not evident which features a robot must have in order to show social awareness or intelligence. Clearly, a robot's design depends on its intended use, the complexity of the social environment and the sophistication of the interaction. But, to what extent does social robot design need to reflect theories of human social intelligence?

How do we evaluate social robots? Many researchers contend that adding social interaction capabilities will improve robot performance, e.g., by increasing usability. Thus far, however, little experimental evidence exists to support this claim. What is needed is a systematic study of how "social features" impact human–robot interaction in the context of different application domains [43]. The problem is that it difficult to determine which metrics are most appropriate for evaluating social "effectiveness". Should we use human performance metrics? Should we apply psychological, sociological or HCI measures? How do we account for cross-cultural differences and individual needs?

What differentiates social robots from robots that exhibit good human–robot interaction? Although

conventional HRI design does not directly address the issues presented in this paper, it does involve techniques that indirectly support social interaction. For example, HCI methods (e.g., contextual inquiry) are often used to ensure that the interaction will match user needs. The question is: are social robots so different from traditional robots that we need different interactional design techniques?

What underlying social issues may influence future technical development? An observation made by Restivo is that "robotics engineers seem to be driven to program out aspects of being human that for one reason or another they do not like or that make them personally uncomfortable" [134]. If this is true, does that mean that social robots will always be "benign" by design? If our goal is for social robots to eventually have a place in human society, should we not investigate what could be the negative consequences of social robots?

Are there ethical issues that we need to be concerned with? For social robots to become more and more sophisticated, they will need increasingly better computational models of individuals, or at least, humans in general. Detailed user modeling, however, may not be acceptable, especially if it involves privacy concerns. A related question is that of user monitoring. If a social robot has a model of an individual, should it be capable of recognizing when a person is acting erratically and taking action?

How do we design for long-term interaction? To date, research in social robot has focused exclusively on short duration interaction, ranging from periods of several minutes (e.g., tour-guiding) to several weeks, such as in [75]. Little is known about interaction over longer periods. To remain engaging and empowering for months, or years, will social robots need to be capable of long-term adaptiveness, associations, and memory? Also, how can we determine whether long-term human–robot relationships may cause ill-effects?

3.3. Summary

As we look ahead, it seems clear that social robots will play an ever larger role in our world, working for and in cooperation with humans. Social robots will assist in health care, rehabilitation, and therapy. Social robots will work in close proximity to humans, serving as tour guides, office assistants, and household staff. Social robots will engage us, entertain us, and enlighten us.

Central to the success of social robots will be close and effective interaction between humans and robots. Thus, although it is important to continue enhancing autonomous capabilities, we must not neglect improving the human–robot relationship. The challenge is not merely to develop techniques that allow social robots to succeed in limited tasks, but also to find ways that social robots can participate in the full richness of human society.

Acknowledgements

References

[1] B. Adams, C. Breazeal, R.A. Brooks, B. Scassellati, Humanoid robots: a new kind of tool, IEEE Intelligent Systems 15 (4) (2000) 25–31.

[2] J. Aggarwal, Q. Cai, Human motion analysis: A review, Computer Vision and Image Understanding 73 (3) (1999) 428–440.

[3] P. Andry, P. Gaussier, S. Moga, J.P. Banquet, Learning and communication via imitation: An autonomous robot perspective, IEEE Transactions on Systems, Man and Cybernetics 31 (5) (2001).

[4] R. Arkin, M. Fujita, T. Takagi, R. Hasekawa, An ethological and emotional basis for human–robot interaction, Robotics and Autonomous Systems 42 (2003) 191–201 (this issue).

[5] C. Armon-Jones, The social functions of emotions, in: R. Harré (Ed.), The Social Construction of Emotions, Basil Blackwell, Oxford, 1985.

[6] T. Balch, R. Arkin, Communication in reactive multiagent robotic systems, Autonomous Robots 1 (1994).

[7] S. Baron-Cohen, Mindblindness: An Essay on Autism and Theory of Mind, MIT Press, Cambridge, MA, 1995.

[8] C. Bartneck, M. Okada, Robotic user interfaces, in: Proceedings of the Human and Computer Conference, 2001.

[9] C. Bartneck, eMuu—an emotional embodied character for the ambient intelligent home, Ph.D. Thesis, Technical University Eindhoven, The Netherlands, 2002.

[10] R. Beckers, et al., From local actions to global tasks: Stigmergy and collective robotics, in: Proceedings of Artificial Life IV, 1996.

[11] A. Billard, Robota: clever toy and educational tool, Robotics and Autonomous Systems 42 (2003) 259–269 (this issue).

[12] A. Billard, K. Dautenhahn, Grounding communication in situated, social robots, in: Proceedings Towards Intelligent Mobile Robots Conference, Report No. UMCS-97-9-1, Department of Computer Science, Manchester University, 1997.

[13] A. Billard, K. Dautenhahn, Grounding communication in autonomous robots: An experimental study, Robotics and Autonomous Systems 24 (1–2) (1998) 71–81.

[14] A. Billard, K. Dautenhahn, Experiments in learning by imitation: Grounding and use of communication in robotic agents, Adaptive Behavior 7 (3–4) (1999).

[15] A. Billard, G. Hayes, Learning to communicate through imitation in autonomous robots, in: Proceedings of the International Conference on Artificial Neural Networks, 1997.

[16] E. Bonabeau, M. Dorigo, G. Theraulaz, Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press, Oxford, 1999.

[17] C. Breazeal, A motivational system for regulating human–robot interaction, in: Proceedings of the National Conference on Artificial Intelligence, Madison, WI, 1998, pp. 54–61.

[18] C. Breazeal, Designing Sociable Robots, MIT Press, Cambridge, MA, 2002.

[19] C. Breazeal, Toward sociable robots, Robotics and Autonomous Systems 42 (2003) 167–175 (this issue).

[20] C. Breazeal, Designing sociable robots: Lessons learned, in: K. Dautenhahn, et al. (Eds.), Socially Intelligent Agents: Creating Relationships with Computers and Robots, Kluwer Academic Publishers, Dordrecht, 2002.

[21] C. Breazeal, P. Fitzpatrick, That certain look: Social amplification of animate vision, in: Proceedings of the AAAI Fall Symposium on Society of Intelligence Agents—The Human in the Loop, 2000.

[22] C. Breazeal, B. Scassellati, How to build robots that make friends and influence people, in: Proceedings of the International Conference on Intelligent Robots and Systems, 1999.

[23] C. Breazeal, B. Scassellati, A context-dependent attention system for a social robot, in: Proceedings of the International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 1999, pp. 1146–1153.

[24] C. Breazeal, B. Scassellati, Challenges in building robots that imitate people, in: K. Dautenhahn, C. Nehaniv (Eds.), Imitation in Animals and Artifacts, MIT Press, Cambridge, MA, 2001.

[25] C. Breazeal, A. Edsinger, P. Fitzpatrick, B. Scassellati, Active vision systems for sociable robots, IEEE Transactions on Systems, Man and Cybernetics 31 (5) (2001).

[26] A. Bruce, I. Nourbakhsh, R. Simmons, The role of expressiveness and attention in human–robot interaction, in: Proceedings of the AAAI Fall Symposium Emotional and Intelligent II: The Tangled Knot of Society of Cognition, 2001.

[27] K. Bumby, K. Dautenhahn, Investigating children's attitudes towards robots: A case study, in: Proceedings of the Cognitive Technology Conference, 1999.

[28] J. Cahn, The generation of affect in synthesized speech, Journal of American Voice I/O Society 8 (1990) 1–19.

[29] L. Cañamero, Modeling motivations and emotions as a basis for intelligent behavior, in: W. Johnson (Ed.), Proceedings of the International Conference on Autonomous Agents.

[30] L. Cañamero, J. Fredslund, I show you how I like you—can you read it in my face? IEEE Transactions on Systems, Man and Cybernetics 31 (5) (2001).

[31] L. Cañamero (Ed.), Emotional and Intelligent II: The Tangled Knot of Social Cognition, Technical Report No. FS-01-02, AAAI Press, 2001.

[32] J. Cassell, Nudge, nudge, wink, wink: Elements of face-to-face conversation for embodied conversational agents, in: J. Cassell, et al. (Eds.), Embodied Conversational Agents, MIT Press, Cambridge, MA, 1999.

[33] J. Cassell, et al. (Eds.), Embodied Conversational Agents, MIT Press, Cambridge, MA, 1999.

[34] R. Chellappa, et al., Human and machine recognition of faces: A survey, Proceedings of the IEEE 83 (5) (1995).

[35] M. Coulson, Expressing emotion through body movement: A component process approach, in: R. Aylett, L. Cañamero (Eds.), Animating Expressive Characters for Social Interactions, SSAISB Press, 2002.

[36] J. Crowley, Vision for man–machine interaction, Robotics and Autonomous Systems 19 (1997) 347–358.

[37] C. Darwin, The Expression of Emotions in Man and Animals, Oxford University Press, Oxford, 1998.

[38] K. Dautenhahn, Getting to know each other—artificial social intelligence for autonomous robots, Robotics and Autonomous Systems 16 (1995) 333–356.

[39] K. Dautenhahn, I could be you—the phenomenological dimension of social understanding, Cybernetics and Systems Journal 28 (5) (1997).

[40] K. Dautenhahn, The art of designing socially intelligent agents—science, fiction, and the human in the loop, Applied Artificial Intelligence Journal 12 (7–8) (1998) 573–617.

[41] K. Dautenhahn, Socially intelligent agents and the primate social brain—towards a science of social minds, in: Proceedings of the AAAI Fall Symposium on Society of Intelligence Agents, 2000.

[42] K. Dautenhahn, Roles and functions of robots in human society—implications from research in autism therapy, Robotica, to appear.

[43] K. Dautenhahn, Design spaces and niche spaces of believable social robots, in: Proceedings of the International Workshop on Robots and Human Interactive Communication, 2002.

[44] K. Dautenhahn, A. Billard, Bringing up robots or—the psychology of socially intelligent robots: From theory to implementation, in: Proceedings of the Autonomous Agents, 1999.

[45] K. Dautenhahn, C. Nehaniv, Living with socially intelligent agents: A cognitive technology view, in: K. Dautenhahn (Ed.), Human Cognition and Social Agent Technology, Benjamin, New York, 2000.

[46] K. Dautenhahn, C. Nehaniv (Eds.), Imitation in Animals and Artifacts, MIT Press, Cambridge, MA, 2001.

[47] K. Dautenhahn, I. Werry, A quantitative technique for analysing robot–human interactions, in: Proceedings of the International Conference on Intelligent Robots and Systems, 2002.

[48] K. Dautenhahn, B. Ogden, T. Quick, From embodied to socially embedded agents—implications for interaction-aware robots, Cognitive Systems Research 3 (3) (2002) (Special Issue on Situated and Embodied Cognition).

[49] K. Dautenhahn, I. Werry, J. Rae, P. Dickerson, Robotic play-mates: Analysing interactive competencies of children with autism playing with a mobile robot, in: K. Dautenhahn, et al. (Eds.), Socially Intelligent Agents: Creating Relationships with Computers and Robots, Kluwer Academic Publishers, Dordrecht, 2002.

[50] D. Dennett, The Intentional Stance, MIT Press, Cambridge, MA, 1987.

[51] J. Demiris, G. Hayes, Imitative learning mechanisms in robots and humans, in: Proceedings of the European Workshop on Learning Robotics, 1996.

[52] J. Demiris, G. Hayes, Active and passive routes to imitation, in: Proceedings of the AISB Symposium on Imitation in Animals and Artifacts, 1999.

[53] J.-L. Deneubourg, et al., The dynamic of collective sorting robot-like ants and ant-like robots, in: Proceedings of the International Conference on Simulation of Adaptive Behavior, 2000.

[54] C. DiSalvo, et al., All robots are not equal: The design and perception of humanoid robot heads, in: Proceedings of the Conference on Designing Interactive Systems, 2002.

[55] B. Duffy, Anthropomorphism and the social robot, Robotics and Autonomous Systems 42 (2003) 177–190 (this issue).

[56] P. Ekman, W. Friesen, Measuring facial movement with the facial action coding system, in: Emotion in the Human Face, Cambridge University Press, Cambridge, 1982.

[57] P. Ekman, Basic emotions, in: T. Dalgleish, M. Power (Eds.), Handbook of Cognition and Emotion, Wiley, New York, 1999.

[58] B. Fogg, Introduction: Persuasive technologies, Communications of the ACM 42 (5) (1999).

[59] T. Fong, C. Thorpe, C. Baur, Collaboration, dialogue, and human–robot interaction, in: Proceedings of the International Symposium on Robotics Research, 2001.

[60] T. Fong, I. Nourbakhsh, K. Dautenhahn, A survey of socially interactive robots: concepts, design, and applications, Technical Report No. CMU-RI-TR-02-29, Robotics Institute, Carnegie Mellon University, 2002.

[61] T. Fong, C. Thorpe, C. Baur, Robot, asker of questions, Robotics and Autonomous Systems 42 (2003) 235–243 (this issue).

[62] N. Frijda, Recognition of emotion, Advances in Experimental Social Psychology 4 (1969).

[63] T. Fromherz, P. Stucki, M. Bichsel, A survey of face recognition, MML Technical Report No. 97.01, Department of Computer Science, University of Zurich, 1997.

[64] B. Galef, Imitation in animals: History, definition, and interpretation of data from the psychological laboratory, in:

Social Learning: Psychological and Biological Perspectives, Erlbaum, London, 1988.

[65] P. Gaussier, et al., From perception–action loops to imitation processes: A bottom-up approach of learning by imitation, Applied Artificial Intelligence Journal 12 (7–8) (1998).

[66] D. Gavrilla, The visual analysis of human movement: A survey, Computer Vision and Image Understanding 73 (1) (1999).

[67] J. Goetz, S. Kiesler, Cooperation with a robotic assistant, in: Proceedings of CHI, 2002.

[68] D. Goldberg, M. Mataric, Interference as a tool for designing and evaluating multi-robot controllers, in: Proceedings AAAI-97, Providence, RI, 1997, pp. 637–642.

[69] D. Goren-Bar. Designing model-based intelligent dialogue systems, in: M. Rossi, K. Siau (Eds.), Information Modeling in the New Millennium, Idea Group, 2001.

[70] E. Hall, The hidden dimension: Man's use of space in public and private, The Bodley Head Ltd., 1966.

[71] C. Heyes, B. Galef, Social Learning in Animals: The Roots of Culture, Academic Press, 1996.

[72] G. Hayes, J. Demiris, A robot controller using learning by imitation, in: Proceedings of the International Symposium on Intelligent Robotic Systems, 1994.

[73] O. Holland, Grey Walter: The pioneer of real artificial life, in: C. Langton, K. Shimohara (Eds.), Proceedings of the International Workshop on Artificial Life, MIT Press, Cambridge, MA.

[74] E. Hudlicka, Increasing SIA architecture realism by modeling and adapting to affect and personality, in: K. Dautenhahn, et al. (Eds.), Socially Intelligent Agents: Creating Relationships with Computers and Robots, Kluwer Academic Publishers, Dordrecht, 2002.

[75] H. Hüttenrauch, K. Severinson-Eklund, Fetch-and-carry with CERO: Observations from a long-term user study, in: Proceedings of the International Workshop on Robots and Human Communication, 2002.

[76] O. John, The 'Big Five' factor taxonomy: Dimensions of personality in the natural language and in questionnaires, in: L. Pervin (Ed.), Handbook of Personality: Theory and Research, Guilford, 1990.

[77] F. Kaplan, et al., Taming robots with clicker training: A solution for teaching complex behaviors, in: Proceedings of the European Workshop on Learning Robots, 2001.

[78] Z. Khan, Attitudes towards intelligent service robots, Technical Report No. TRITA-NA-P9821, NADA, KTH, Stockholm, Sweden, 1998.

[79] S. Kiesler, J. Goetz, Mental models and cooperation with robotic assistants, in: Proceedings of CHI, 2002.

[80] V. Klingspor, J. Demiris, M. Kaiser, Human–robot-communication and machine learning, Applied Artificial Intelligence Journal 11 (1997).

[81] H. Kobayashi, F. Hara, A. Tange, A basic study on dynamic control of facial expressions for face robot, in: Proceedings of the International Workshop on Robots and Human Communication, 1994.

[82] H. Kozima, H. Yano, In search of otogenetic prerequisites for embodied social intelligence, in: Proceedings of the

Workshop on Emergence and Development on Embodied Cognition; International Conference on Cognitive Science, 2001.

[83] H. Kozima, H. Yano, A robot that learns to communicate with human caregivers, in: Proceedings of the International Workshop on Epigenetic Robotics, 2001.

[84] L. Kopp, P. Gardenfors, Attention as a Minimal Criterion of Intentionality in Robotics, vol. 89, Lund University of Cognitive Studies, 2001.

[85] D. Kortenkamp, E. Huber, P. Bonasso, Recognizing and interpreting gestures on a mobile robot, in: Proceedings of the AAAI-96, Portland, OR, 1996, pp. 915–921.

[86] R. Krauss, P. Morrel-Samuels, C. Colasante, Do conversational hand gestures communicate? Journal of Personality and Social Psychology 61 (1991).

[87] M. Krieger, J.-B. Billeter, L. Keller, Ant-like task allocation and recruitment in cooperative robots, Nature 406 (6799) (2000).

[88] C. Kube, E. Bonabeau, Cooperative transport by ants and robots, Robotics and Autonomous Systems 30 (2000) 85–101.

[89] Y. Kuniyoshi, et al., Learning by watching: Extracting reusable task knowledge from visual observation of human performance, IEEE Transactions of the Robotics and Automation 10 (6) (1994).

[90] M. Lansdale, T. Ormerod, Understanding Interfaces, Academic Press, New York, 1994.

[91] S. Lauria, G. Bugmann, T. Kyriacou, E. Klein, Mobile robot programming using natural language, Robotics and Autonomous Systems 38 (2002) 171–181.

[92] V. Lee, P. Gupta, Children's Cognitive and Language Development, Blackwell Scientific Publications, Oxford, 1995.

[93] H. Lim, A. Ishii, A. Takanishi, Basic emotional walking using a biped humanoid robot, in: Proceedings of the IEEE SMC, 1999.

[94] C. Lisetti, D. Schiano, Automatic facial expression interpretation: Where human–computer interaction, artificial intelligence, and cognitive science intersect, Pragmatics and Cognition 8 (1) (2000).

[95] K. Lorenz, The Foundations of Ethology, Springer, Berlin, 1981.

[96] Y. Marom, G. Hayes, Preliminary approaches to attention for social learning, Informatics Research Report No. EDI-INF-RR-0084, University of Edinburgh, 1999.

[97] Y. Marom, G. Hayes, Attention and social situatedness for skill acquisition, Informatics Research Report No. EDI-INF-RR-0069, University of Edinburgh, 2001.

[98] Y. Marom, G. Hayes, Interacting with a robot to enhance its perceptual attention, Informatics Research Report No. EDI-INF-RR-0085, University of Edinburgh, 2001.

[99] D. Massaro, Perceiving Talking Faces: From Speech Perception to Behavioural Principles, MIT Press, Cambridge, MA, 1998.

[100] M. Mataric, Learning to behave socially, in: Proceedings of the International Conference on Simulation of Adaptive Behavior, 1994.

[101] M. Mataric, Issues and approaches in design of collective autonomous agents, Robotics and Autonomous Systems 16 (1995) 321–331.

[102] M. Mataric, et al., Behavior-based primitives for articulated control, in: Proceedings of the International Conference on Simulation of Adaptive Behavior, 1998.

[103] T. Matsui, H. Asoh, J. Fry, et al., Integrated natural spoken dialogue system of Jijo-2 mobile robot for office services, in: Proceedings of the AAAI-99, Orlando, FL, 1999, pp. 621–627.

[104] Y. Matsusaka, T. Kobayashi, Human interface of humanoid robot realizing group communication in real space, in: Proceedings of the International Symposium on Humanoid Robotics, 1999.

[105] D. McNeill, Hand and Mind: What Gestures Reveal About Thought, University of Chicago Press, Chicago, IL, 1992.

[106] C. Melhuish, O. Holland, S. Hoddell, Collective sorting and segregation in robots with minimal sensing, in: Proceedings of the International Conference on Simulation of Adaptive Behavior, 1998.

[107] F. Michaud, S. Caron, Roball—An autonomous toy-rolling robot, in: Proceedings of the Workshop on Interactive Robot Entertainment, 2000.

[108] F. Michaud, et al., Artificial emotion and social robotics, in: Proceedings of the International Symposium on Distributed Autonomous Robotic Systems, 2000.

[109] F. Michaud, et al., Dynamic robot formations using directional visual perception, in: Proceedings of the International Conference on Intelligent Robots and Systems, 2002.

[110] H. Mizoguchi, et al., Realization of expressive mobile robot, in: Proceedings of the International Conference on Robotics and Automation, 1997.

[111] I. Murray, J. Arnott, Towards the simulation of emotion in synthetic speech: A review of the literature on human vocal emotion, Journal of Acoustic Society of America 93 (2) (1993).

[112] I. Myers, Introduction to Type, Consulting Psychologists Press, Palo Alto, CA, 1998.

[113] T. Nakata, et al., Expression of emotion and intention by robot body movement, in: Proceedings of the 5th International Conference on Autonomous Systems, 1998.

[114] Y. Nakauchi, R. Simmons, A social robot that stands in line, in: Proceedings of the International Conference on Intelligent Robots and Systems, 2000.

[115] W. Newman, M. Lamming, Interactive System Design, Addison-Wesley, Reading, MA, 1995.

[116] M. Nicolescu, M. Mataric, Learning and interacting in human–robot domains, IEEE Transactions on Systems, Man and Cybernetics 31 (5) (2001).

[117] I. Nourbakhsh, An affective mobile robot educator with a full-time job, Artificial Intelligence 114 (1–2) (1999) 95–124.

[118] A. Rowe, C. Rosenberg, I. Nourbakhsh, CMUcam: a low-overhead vision system, in: Proceedings of the International Conference on Intelligent Robots and Systems (IROS 2002), Lausanne, Switzerland, 2002.

[119] T. Ogata, S. Sugano, Emotional communication robot: WAMOEBA-2R emotion model and evaluation experiments,

in: Proceedings of the International Conference on Humanoid Robots, 2000.

[120] H. Okuno, et al., Human–robot interaction through real-time auditory and visual multiple-talker tracking, in: Proceedings of the International Conference on Intelligent Robots and Systems, 2001.

[121] A. Paiva (Ed.), Affective interactions: Towards a new generation of computer interfaces, in: Lecture Notes in Computer Science/Lecture Notes in Artificial Intelligence, vol. 1914, Springer, Berlin, 2000.

[122] J. Panksepp, Affective Neuroscience, Oxford University Press, Oxford, 1998.

[123] E. Paulos, J. Canny, Designing personal tele-embodiment, Autonomous Robots 11 (1) (2001).

[124] V. Pavlovic, R. Sharma, T. Huang, Visual interpretation of hand gestures for human–computer interaction: A review, IEEE Transactions on Pattern Analysis and Machine Intelligence 19 (7) (1997).

[125] P. Persson, et al., Understanding socially intelligent agents—A multilayered phenomenon, IEEE Transactions on SMC 31 (5) (2001).

[126] R. Pfeifer, On the role of embodiment in the emergence of cognition and emotion, in: Proceedings of the Toyota Conference on Affective Minds, 1999.

[127] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, S. Thrun, Towards robotic assistants in nursing homes: Challenges and results, Robotics and Autonomous Systems 42 (2003) 271–281 (this issue).

[128] R. Plutchik, Emotions: A general psychoevolutionary theory, in: K. Scherer, P. Ekman (Eds.), Approaches to Emotion, Erlbaum, London, 1984.

[129] L. Rabiner, B. Jaung, Fundamentals of Speech Recognition, Englewood Cliffs, Prentice-Hall, NJ, 1993.

[130] B. Reeves, C. Nass, The Media Equation, CSLI Publications, Stanford, 1996.

[131] J. Reichard, Robotics: Fact, Fiction, and Prediction, Viking Press, 1978.

[132] W. Reilly, Believable social and emotional agents, Ph.D. Thesis, Computer Science, Carnegie Mellon University, 1996.

[133] S. Restivo, Bringing up and booting up: Social theory and the emergence of socially intelligent robots, in: Proceedings of the IEEE Conference on SMC, 2001.

[134] S. Restivo, Romancing the robots: Social robots and society, in: Proceedings of the Robots as Partners: An Exploration of Social Robots Workshop, International Conference on Intelligent Robots and Systems (IROS 2002), Lausanne, Switzerland, 2002.

[135] A. Sage, Systems Engineering, Wiley, New York, 1992.

[136] A. Samal, P. Iyengar, Automatic recognition and analysis of human faces and facial expressions: A survey, Pattern Recognition 25 (1992).

[137] H. Schlossberg, Three dimensions of emotion, Psychology Review 61 (1954).

[138] J. Searle, Minds, Brains and Science, Harvard University Press, Cambridge, MA, 1984.

[139] B. Scassellati, Investigating models of social development using a humanoid robot, in: B. Webb, T. Consi (Eds.), Biorobotics, MIT Press, Cambridge, MA, 2000.

[140] B. Scassellati, Foundations for a theory of mind for a humanoid robot, Ph.D. Thesis, Department of Electronics Engineering and Computer Science, MIT Press, Cambridge, MA, 2001.

[141] S. Schall, Robot learning from demonstration, in: Proceedings of the International Conference on Machine Learning, 1997.

[142] M. Scheeff, et al., Experiences with Sparky: A social robot, in: Proceedings of the Workshop on Interactive Robot Entertainment, 2000.

[143] J. Schulte, et al., Spontaneous, short-term interaction with mobile robots in public places, in: Proceedings of the International Conference on Robotics and Automation, 1999.

[144] K. Severinson-Eklund, A. Green, H. Hüttenrauch, Social and collaborative aspects of interaction with a service robot, Robotics and Autonomous Systems 42 (2003) 223–234 (this issue).

[145] T. Sheridan, Eight ultimate challenges of human–robot communication, in: Proceedings of the International Workshop on Robots and Human Communication, 1997.

[146] C. Smith, H. Scott, A componential approach to the meaning of facial expressions, in: J. Russell, J. Fernandez-Dols (Eds.), The Psychology of Facial Expression, Cambridge University Press, Cambridge, 1997.

[147] R. Smith, S. Eppinger, A predictive model of sequential iteration in engineering design, Management Science 43 (8) (1997).

[148] D. Spiliotopoulos, et al., Human–robot interaction based on spoken natural language dialogue, in: Proceedings of the European Workshop on Service and Humanoid Robots, 2001.

[149] L. Steels, Emergent adaptive lexicons, in: Proceedings of the International Conference on SAB, 1996.

[150] L. Steels, F. Kaplan, AIBO's first words. The social learning of language and meaning, in: H. Gouzoules (Ed.), Evolution of Communications, vol. 4, No. 1, Benjamin, New York, 2001.

[151] L. Steels, Language games for autonomous robots, IEEE Intelligent Systems 16 (5) (2001).

[152] R. Stiefelhagen, J. Yang, A. Waibel, Tracking focus of attention for human–robot communication, in: Proceedings of the International Conference on Humanoid Robots, 2001.

[153] K. Suzuki, et al., Intelligent agent system for human–robot interaction through artificial emotion, in: Proceedings of the IEEE SMC, 1998.

[154] R. Tanawongsuwan, et al., Robust tracking of people by a mobile robotic agent, Technical Report No. GIT-GVU-99-19, Georgia Institute of Technology, 1999.

[155] L. Terveen, An overview of human–computer collaboration, Knowledge-Based Systems 8 (2–3) (1994).

[156] F. Thomas, O. Johnston, Disney Animation: The Illusion of Life, Abbeville Press, 1981.

[157] W. Thorpe, Learning and Instinct in Animals, Methuen, London, 1963.

[158] K. Toyama, Look, Ma—No hands! Hands-free cursor control with real-time 3D face tracking, in: Proceedings of the Workshop on Perceptual User Interfaces, 1998.

[159] R. Vaughan, K. Stoey, G. Sukhatme, M. Mataric, Go ahead, make my day: Robot conflict resolution by aggressive competition, in: Proceedings of the International Conference on SAB, 2000.

[160] J. Velasquez, A computational framework for emotion-based control, in: Proceedings of the Workshop on Grounding Emotions in Adaptive Systems; International Conference on SAB, 1998.

[161] S. Waldherr, R. Romero, S. Thrun, A gesture-based interface for human–robot interaction, Autonomous Robots 9 (2000).

[162] I. Werry, et al., Can social interaction skills be taught by a social agent? The role of a robotic mediator in autism therapy, in: Proceedings of the International Conference on Cognitive Technology, 2001.

[163] A. Whiten, Natural Theories of Mind, Basil Blackwell, Oxford, 1991.

[164] T. Willeke, et al., The history of the mobot museum robot series: An evolutionary study, in: Proceedings of FLAIRS, 2001.

[165] D. Woods, Decomposing automation: Apparent simplicity, real complexity, in: R. Parasuraman, M. Mouloua (Eds.), Automation and Human Performance: Theory and Applications, Erlbaum, London, 1996.

[166] Y. Wu, T. Huang, Vision-based gesture recognition: A review, in: Gesture-Based Communications in HCI, Lecture Notes in Computer Science, vol. 1739, Springer, Berlin, 1999.

[167] G. Xu, et al., Toward robot guidance by hand gestures using monocular vision, in: Proceedings of the Hong Kong Symposium on Robotics Control, 1999.

[168] S. Yoon, et al., Motivation driven learning for interactive synthetic characters, in: Proceedings of the International Conference on Autonomous Agents, 2000.

[169] J. Zlatev, The Epigenesis of Meaning in Human Beings and Possibly in Robots, Lund University Cognitive Studies, vol. 79, Lund University, 1999.

Terrence Fong is a joint postdoctoral fellow at Carnegie Mellon University (CMU) and the Swiss Federal Institute of Technology/Lausanne (EPFL). He received his Ph.D. (2001) in Robotics from CMU. From 1990 to 1994, he worked at the NASA Ames Research Center, where he was co-investigator for virtual environment telerobotic field experiments. His research interests include human–robot interaction, PDA and web-based interfaces, and field mobile robots.

Illah Nourbakhsh is an Assistant Professor of Robotics at Carnegie Mellon University (CMU) and is co-founder of the Toy Robots Initiative at The Robotics Institute. He received his Ph.D. (1996) degree in computer science from Stanford. He is a founder and chief scientist of Blue Pumpkin Software, Inc. and Mobot, Inc. His current research projects include robot learning, believable robot personality, visual navigation and robot locomotion.

Kerstin Dautenhahn is a Reader in artificial intelligence in the Computer Science Department at University of Hertfordshire, where she also serves as coordinator of the Adaptive Systems Research Group. She received her doctoral degree in natural sciences from the University of Bielefeld. Her research lies in the areas of socially intelligent agents and HCI, including virtual and robotic agents. She has served as guest editor for numerous special journal issues in AI, cybernetics, artificial life, and recently co-edited the book "Socially Intelligent Agents—Creating Relationships with Computers and Robots".