



Hortensius, R. and Cross, E. S. (2018) From automata to animate beings: the scope and limits of attributing socialness to artificial agents. *Annals of the New York Academy of Sciences*, 1426(1), pp. 93-110.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

This is the peer reviewed version of the following article Hortensius, R. and Cross, E. S. (2018) From automata to animate beings: the scope and limits of attributing socialness to artificial agents. *Annals of the New York Academy of Sciences*, 1426(1), pp. 93-110. <http://dx.doi.org/10.1111/nyas.13727>

This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#).

<http://eprints.gla.ac.uk/160888/>

Deposited on: 11 May 2018

From automata to animate beings:

The scope and limits of attributing socialness to artificial agents

Ruud Hortensius* and Emily S. Cross*

Wales Institute for Cognitive Neuroscience, School of Psychology, Bangor University and Institute of Neuroscience and Psychology, School of Psychology, University of Glasgow, Scotland

* both authors contributed equally to this work

Correspondence:

Wales Institute for Cognitive Neuroscience, School of Psychology, Bangor University,

Adeilad Brigantia, Ffordd Penrallt, Bangor, Gwynedd LL57 2AS, Wales, UK.

and

Institute of Neuroscience and Psychology, School of Psychology, University of Glasgow, 58

Hillhead Street, Glasgow, G12 8QB, Scotland, UK

R. Hortensius: ruud.hortensius@glasgow.ac.uk / +44 (0) 1248 38 3013

E.S. Cross: emily.cross@glasgow.ac.uk / +44 (0) 1248 38 3274

Running head: Socialness attribution and artificial agents

Accepted version of Hortensius, R. and Cross, E. S. (2018) From automata to animate beings: the scope and limits of attributing socialness to artificial agents. *Annals of the New York Academy of Sciences*, (<http://dx.doi.org/10.1111/nyas.13727>).

Abstract

Understanding the mechanisms and consequences of attributing socialness to artificial agents has important implications for how we can use technology to lead more productive and fulfilling lives. Here, we integrate recent findings on the factors that shape behavioural and brain mechanisms that support social interactions between humans and artificial agents. We review how visual features of an agent, as well as knowledge factors within the human observer, shape attributions across dimensions of socialness. We explore how anthropomorphism and dehumanization further influence how we perceive and interact with artificial agents. Based on these findings, we argue that the cognitive reconstruction within the human observer is likely to be far more crucial in shaping our interactions with artificial agents than previously thought, while the artificial agent's visual features are possibly of lesser importance. We integrate these findings to provide an integrative theoretical account based on the "like me" hypothesis, and discuss the key role played by the Theory-of-Mind network, especially the temporal parietal junction, in the shift from mechanistic to social attributions. We conclude by highlighting outstanding questions on the impact of long-term interactions with artificial agents on the behavioural and brain mechanisms of attributing socialness to these agents.

Keywords: socialness attribution, animacy, anthropomorphism, artificial agents, human-robot interaction, social cognition

Introduction

Humans readily attribute socialness to artificial agents¹. We erupt in anger at the computer that ‘knowingly’ crashes in the middle of an important task, sympathize with a robot character in a film, or attribute a personality to an artificial personal assistant. The ease with which we ascribe agency and socialness to artificial entities has been exploited by writers, artists, and filmmakers for nearly a century, resulting in a rich fiction exploring the relationship between man and sentient machine. How do we make the transition from seeing a robot as a simple automaton to a sentient social being? This and related questions have puzzled scholars for centuries, from the discussion of the uniqueness of human social nature by Aristotle to the study of early automatons by Leonardo da Vinci and most recently the detailed empirical investigations by roboticists, psychologists, and cognitive neuroscientists¹⁻³.

These days, robots not only work alongside people on factory floors, but can increasingly be found in health care, education, and service industry settings as well. Similarly, with the increasing use of virtual and augmented reality, interactions with virtual humans are also likely to play a key role in the social fabric of society in the near future. Amidst this fourth industrial revolution, as social robots and other artificial agents become increasingly sophisticated and resolutely move from fiction to reality, important questions regarding the flexibility and adaptability of human social cognition when interacting with these entities require urgent attention.

Why might we perceive a robot as merely an automaton in some situations, while in other situations we see the same robot as an engaging social partner? Is this process of attributing

¹ We use the term artificial agents to refer to robots (including those that are machine-like, pet-like or human-like), virtual agents (including avatars of oneself or other virtual humans or characters), and artificial personal assistants (such as Siri, Cortana, Alexa, etc.).

socialness to artificial agents similar to that for attributing social characteristics to biological agents? To what extent is the same neural machinery we use to navigate our social world, refined over millennia of interacting with other people, co-opted when we reason about and interact with artificial agents? In this review, we aim to formulate answers to these complex and increasingly relevant social cognition questions by using an integrative approach that synthesises the latest findings in psychology, social robotics, virtual reality and neuroscience. We first review work on the role of artificial agents’ physical appearance in eliciting responses at the brain and behavioural level. While this work documents how variations in an agents’ visual features can shape human social engagement to a certain extent, we discuss noteworthy new findings incorporating insights from social cognition and social neuroscience on the impact of a person’s prior knowledge, beliefs, or expectations about an artificial agent. In addition, we consider the processes of anthropomorphism and dehumanization on social engagement at brain and behavioural levels. We discuss a model that takes into account both perceptual and cognitive factors of human interactions with artificial agents and shows the functional convergence of cognitive factors driving socialness attribution to artificial agents within a specialized neural network. We conclude by discussing current challenges and future directions.

The attribution of socialness

A social interaction between two agents involves a complex cascade of expressions and reactions to social and emotional signals. During these interactions, we are not passive observers, but instead we actively construct the social nature of the other agent. We try to understand and explain the behaviour and internal states of the other agent in terms of emotions, intentions and beliefs. As we attribute emotions or intentions to other agents, artificial or not, we infer and attribute *socialness* to these agents. Socialness can be defined as

the presence of intentional goal-directed recursive interactions with other beings. This process of attribution is closely related to the perception of other agents' minds⁴, or the representation of their mental state, also referred to as mentalizing or Theory-of-Mind⁵. If we attribute socialness to an agent, we adopt what Dennett^{6,7} calls the intentional stance. We view and treat the agent as rational, with beliefs, desires and behavioural consistency. This is in contrast to the design stance, in which we view and treat an agent based on knowledge of its function or design (for example, predicting the response of a robot based on the knowledge of its software or actuators), or physical stance, in which we view and treat an agent based on knowledge derived from physics and chemistry (for example, predicting the trajectory of a ball or a self-driving car based on its mass and velocity).

The attribution of socialness to artificial agents should not be viewed as a binary decision or an all-or-nothing process. An agent's socialness or the presence of intentional goal-directed recursive interactions with other beings is not dichotomous, but is instead a continuum constructed across multiple dimensions^{4,8}. Researchers have described at least two dimensions that summarize different parts of socialness. People use experience, or the ability to sense and feel, and agency, the ability to plan and act, to distinguish the minds of agents⁹. For instance, data suggest that people rate an adult human being as high on both experience and agency, a dog or a chimp as high on experience and low on agency, and a robot as low on experience and medium on agency. Other researchers suggest a different, but related, distinction between warmth (similar to experience) and competence (similar to agency)¹⁰. Yet other researchers discuss a distinction based on human nature (aspects related to emotion that allow for a human-inanimate distinction), and human uniqueness (aspects related to morality that allow for a human-animal distinction)^{11,12}. Regardless of which distinction is used, the important point is that socialness can be distinguished on multiple

dimensions. In order to be seen as a social being, an agent does not need to qualify as social across all dimensions (for example having both agency and experience). While other researchers have discussed humanness or animacy, derived from the Latin word *animat*, meaning 'instilled with life', here we mainly use socialness to highlight this notion of a continuum and multidimensionality, and to capture the full dimensions of these attributions. Socialness comprises capacities like actions, emotions, and intentions, and some scholars have argued that at its essence, a social agent is an agent that is capable of influencing the behaviour of another agent¹³.

Crucially, the attribution of socialness is an ongoing, dynamic process between the perceived agent and the observing agent⁴, and is composed of many cues¹⁴. Some of these cues are derived from features of the artificial agent, such as its form and motion, and are referred to as bottom-up or stimulus cues to socialness. Above and beyond these cues, recent studies, as reviewed here, show that the prior knowledge of the observer, based upon beliefs, expectations, and experience, is key in the attribution of socialness. These observer-oriented factors are collectively referred to as top-down or knowledge cues to socialness. The distinction between these two types of cues provides the necessary framework to distil and disentangle the factors that influence the attribution of socialness to artificial agents.

At a basic level, brain regions associated with the person perception network (PPN), the action observation network (AON) and the Theory-of-Mind (ToM) network have been shown to selectively respond to animate agents^{15,16}. The PPN includes regions in the occipital and temporal cortex, such as the fusiform face area (FFA) and body area (FBA), occipital face area (OFA), extrastriate body area (EBA) and posterior superior temporal sulcus (STS)¹⁷⁻¹⁹. Of course, these regions do not represent discrete animate and inanimate categories, but

instead are responsive to a wide variety of stimuli^{20,21}. Activity within these regions appears to index an observed agent's features (e.g., facial features, body posture, motion)^{17,18,22}, the interactive nature of the situation²³, context^{24,25} and perceived animacy and socialness^{26,27}. The AON comprises frontoparietal regions spanning the posterior inferior frontal gyrus and inferior parietal lobule, which are engaged in a similar manner for executed and observed actions²⁸⁻³⁰. Activity within this network is also modulated by animacy and socialness. The very first work on this system demonstrated a distinction in parietal and premotor neural firing when non-human primates observed animate (a grasp being performed by a human hand) compared to inanimate actions (a grasp being performed by pliers)^{31,32}. Social processing performed by the PPN and AON is further informed by the ToM network, which comprises cortical regions spanning the medial prefrontal cortex (MPFC), temporoparietal junction (TPJ), precuneus and temporal pole³³⁻³⁵. This network is crucial for the inferring the mental state of other agents, including even inanimate agents³⁶.

As we show in the next sections, a mechanistic understanding of the attribution of socialness to artificial agents can be advanced through use of a social and cognitive neuroscientific lens. Some researchers have even suggested that measures of brain activity can serve as a 'neural Turing test'^{37,38}, a way of assessing the ability of an artificial agent to be indistinguishable from a human being³⁹. Whether or not this is (yet) feasible, findings from social and cognitive neuroscience can nonetheless illuminate the factors underpinning the attribution of socialness to artificial agents.

Impact of Artificial Agent's Visual Features

Form

Form follows function, not just in the world of architecture, but also in the design of artificial agents. The first cue towards socialness is the form and shape of an observed agent. The perception of the face and body of a human agent provide access to a rich set of cues to socialness that facilitate subsequent behaviour. Besides identity, the human face and body communicate emotions and intentions¹⁷⁻¹⁹. The perception of this information is partly influenced by stimulus cues, such as the shape and gender of an agent. Importantly, the perception of animacy is at the core of face processing⁴⁰. By directly evaluating the perception of artificial compared to human agents, important first insights of the impact of artificial agents' visual features have come to light.

Several studies have looked at the pattern of activation in the PPN when observing emotions expressed by artificial agents. So far, both electroencephalography and neuroimaging studies suggest that activity within this network is not necessarily decreased by the appearance of the artificial agent⁴¹⁻⁴⁷. Two influential studies provided the first insight into the effect of stimulus cues of socialness^{45,46}. In both studies, participants observed a wide variety of emotional facial expressions (e.g., happiness, disgust, anger) made by a humanoid robot or a human. Regardless of the instruction to either passively observe or actively rate the expressions, similar findings emerged in the neuroimaging data. Specifically, both studies reported no attenuation in activity within the PPN when observing robotic facial expressions. In terms of the response profile of individual regions, activity within the superior temporal gyrus did not discriminate between humanoid robot or human facial expressions, while activity in the fusiform face area was increased for the robotic face compared to the human face^{45,46}. However, the study by Gobbini and colleagues⁴⁶ reported the first evidence of

decreased activity in the ToM network, specifically the right MPFC and the right TPJ, when observing artificial agents. This finding has recently been corroborated by Wang and Quadflieg⁴⁷ in a study on the perception of human-robot interactions. In this study, participants observed a human interacting with either another human or a humanoid robot, and were instructed to indicate if one agent was helping the other agent. This instruction was given to focus participants' attention on the relational aspect of the interaction between the two agents. Similar to previous findings, no robust differences were observed between the perception of human-robot interactions compared with the perception of human-human interactions within the PPN. Only three out of ten regions in this network, the right fusiform face area and bilateral posterior superior temporal sulcus, showed greater activation for perception of human-human interaction compared with human-robot interaction. However, this study also showed sensitivity to the socialness of the agent within the ToM network, with more activation for human-robot interaction in the vMPFC and precuneus, but less activation in the left TPJ during the observation of these interactions compared to human-robot interaction.

These findings of overlap between artificial and human agents at the level of the PPN are complimented by studies probing this network using schematic faces or bodies^{18,22}. With minimal cues present, people readily see faces in face-like objects or even random patterns, with similar activation patterns observed within dedicated brain areas implicated in person perception^{48,49}. Differences in how an agent's physical appearance impacts person perception compared to the cognitive processes of Theory-of-Mind are further borne out by behavioural findings. For instance, emotions expressed by artificial agents, especially in the case of negative emotions, are sometimes difficult to recognize by human observers⁵⁰. Martini, Gonzalez and Wiese⁵¹ directly investigated the role of a humanlike appearance of an

artificial agent on the attribution of different states to the artificial agent, including emotions, goals and agency. Findings from two experiments suggest that the attribution of these states is a two-step process. While observers seldom attribute socialness to artificial agents below a threshold of human-likeness, this linearly increases as more human features become part of the artificial agent. Studies examining these effects on social interactions with artificial agents show decreased human cooperation during direct economic interactions with a small humanoid robot compared to a person⁵², and that people are prone to punish artificial agents more than people^{53,54}. In sum, these findings suggest a differential impact of artificial agents' form on processing in the PPN and ToM network. Whereas the PPN might not rely as heavily on them being human-like, the ToM network and related behaviours might. It should be noted that while the study of perceiving artificial agents is an emerging topic, the few studies published have mostly been limited to facial expressions, with the exception one study on perceiving whole-body interactions between humans and humanoid robots⁴⁷. This reflects a similar face-centric bias observed in studies on the perception of human social signals⁵⁵. As such, it will be valuable for future studies to use a larger variety of social signals to probe the effect of artificial agent's form on engagement of the PPN and ToM network, in order to build a more complete picture of social perception.”

Movement

Numerous studies document how the human brain reliably extracts a wealth of socially-relevant information from simple motion cues. Since the seminal work by Johansson on point-light displays⁵⁶, it has been shown that videos featuring a handful of points following a biological motion profile, containing no further information on the form of the agent, can be used to distil not only the direction and type of actions performed by an actor, but also the actor's emotions, gender, and identity⁵⁷. Some researchers have argued that biological motion might serve as a 'life-detector'⁵⁸⁻⁶⁰, which helps us to detect conspecifics and other

animals. In an important early study, Pelphrey and colleagues⁶¹ showed that the STS is selective to biological motion cues but not the form of the agent. While biological motion is clearly an important social cue, with some researchers arguing for a biological tuning of the motion node of the PPN⁶², as well as the human motor system and the AON⁶³⁻⁶⁵, open questions remain concerning whether biological motion is necessary for engagement of each of these networks.

First insights into such questions can be found via studies on the attribution of socialness to simple animated shapes. Since the influential work by Heider and Simmel⁶⁶, multiple studies found that the observation of animations of simple shapes or social animations featuring non-biological, but self-propelled, motion not only trigger the attribution of goals and intentions to these shapes^{67,68}, but also robustly activate the posterior part of the STS (pSTS)^{36,69-71}, a core node of the PPN. Activity in this region increases when movement parameters suggest an interaction between animated shapes and decreases when movement parameters suggest less interactive and more random motion, implying a ‘perception of animacy’ response gradient⁷². A recent study provides further insights into the role of pSTS that extend beyond motion *per se*⁷³. In two experiments, the authors presented participants with short clips of point-light displays with two agents interacting or completing individual actions, as well as animations of simple shapes engaging in helping or hindering social interactions. Results showed that pSTS maximally respond to social interactions between point-light figures as well as simple shapes. Further analysis revealed that activity in this region does not depend on shape, goal or animacy of the agent, *per se*. Crucially, the pSTS appears to be specifically sensitive to decoding the *nature* of these interactions, whether the interaction was helpful or hindering. Together these findings suggest that a role for the pSTS that is more flexible, moving beyond mere selectivity for biological motion⁶². While these findings complement

previously discussed findings of the role of the PPN, it is important to note that perception of social animations also engages brain activity beyond this network. For example, the goal-directed movement of simple shapes triggers activation in the anterior intraparietal sulcus, part of the AON⁷⁴, similar to human goal-directed movements⁷⁵, and social animations can be used to functionally localize the ToM network⁷⁶ and robustly activate the TPJ⁷⁷.

These findings are corroborated by studies examining automatic imitation during human-robot interactions. Automatic imitation studies seek to quantify the reflexive imitation of observed behaviour, in this case, the interference of an observed robot’s movements on the human perceiver’s executed movements. Some evidence suggests that a biological motion profile of an observed agent impacts an observer’s ongoing or subsequent movement more than a non-biological motion profile^{64,78,79}, and that automatic imitation is greater for robotic movements with quasi-biological motion^{80,81}. However, other studies call into question the necessity of biological motion in an observed action in order for an observer’s motor performance to show interference effects. While automatic imitation of robotic actions is smaller compared to human actions in absolute value, it is not completely absent, and several studies document automatic imitation of movements made by real and virtual full-body humanoid robots^{80,82-84}, regardless of the presence of biological motion. A key factor driving automatic imitation appears to be the presence of human-like joint configuration, and not human-like motion *per se*⁸⁵. Interference effects of observed movements on executed movements are reported for both humanoid robot and mechanical robot arms, as long as the latter had human-like joint configurations. No interference effects are observed if the mechanical robot arm had a non-human joint configuration, despite having quasi-biological motion. Interference effects are also found for apparent motion movements made by robotic hands⁸⁶⁻⁸⁹. Thus, little direct evidence exists for biological tuning at the behavioural level.

In contrast to an early study ⁶⁵, the majority of studies report that activity in the AON is not reliably decreased when observing actions performed by artificial agents compared to humans ^{37,90-93}. Indeed, motion parameters and an agent's appearance appear to not impact the AON in isolation, but rather in combination, and based on context ^{90,91}. In an innovative study, Saygin and colleagues ⁹⁰, compared the observation of simple actions performed by a human, an android, or a robot. Importantly, the android and robot shared the identical motion profile, but differed in appearance. This was achieved by removing or replacing all of the external 'human-like' features of the android so the appearance looked far more mechanical. Results showed most AON engagement for android compared to human or robot movements. Thus, activity in the AON appears to be mediated by an interaction between form (human-like) and motion (machine-like). The authors interpreted their finding in terms of a predictive coding model ⁹⁴. In this model ⁹⁵, observing unfamiliar actions can lead to increased activity in the AON due to greater prediction error. Further evidence for the importance of familiarity above and beyond the effect of motion parameters and appearance comes from a study by Cross and colleagues ⁹¹. Across two experiments, they showed that AON activity was reliably greater when participants watched unfamiliar robotic dancing movements compared to natural dancing movements, regardless of whether these movements were performed by a person or a robot. These studies provide critical evidence that it is not simply stimulus cues, like an agent's form and motion, that drive engagement of brain regions involved in social perception, but also an observer's previous experience, familiarity, and expectations about how an artificial or human agent moves. In sum, while biological motion is an important cue to socialness, a number of lines of evidence suggest it is not necessary. Instead, the core networks implicated in social perception can be flexibly engaged when observing an artificial agent in action, depending on a number of other mostly stimulus-independent factors.

Presence

The first visual encounter with an artificial agent provides a human observer with a first, albeit partial, indication of the agent's socialness. This understanding is only partial at first, since socialness attribution is a dynamic, emergent property of the active social interaction between two or more agents ⁹⁶. One important feature that supports social interactions is the physical co-location of both or all agents in the same environment. However, to maximise experimental control and efficiency, most studies so far merely explore how people *perceive* other (artificial) agents, which is a far leap from active social interaction during which each agent's ongoing behaviour has the potential to influence and be influenced by the behaviour of the other agent. Investigations that use agent observation as the main measure of interest focus on offline social cognition, while investigations employing reciprocal social interaction can delve more deeply into online social cognition, thereby tapping into distinct psychological and neural processes ⁹⁷. Physical embodiment and agent presence are crucial features for studying social interaction between human and artificial agents ^{98,99}. Besides sharing a virtual environment in virtual reality ¹⁰⁰, physical embodiment is another way to ensure artificial and human agents share the same space ¹⁰¹. A physically embodied artificial agent is a real, physical agent that is physically present in the same room as the human agent and allows for physical and face-to-face interaction between the two agents. A physically embodied agent can also be physically present in another room but presented on a screen, thereby reducing the presence of the artificial agent but maintaining the potential for face-to-face interaction. Lastly, a virtually embodied artificial agent is a virtual construction of an artificial agent presented on a screen, thus having neither presence in the real world nor physical embodiment. Initial evidence documents the impact of embodiment and presence on the attribution of socialness.

Engaging in mutual gaze, compared to averted gaze, with a physically embodied robot increases engagement and can drive perceived human-likeness ¹⁰². Confirming previous observations ¹⁰³, humans perceive a physically embodied and collocated robot more positively and persuasively than a visual representation of the same robot ¹⁰⁴. In addition, people also recognize physically present robots' emotions more accurately ¹⁰⁵, and even report higher levels of empathy for robots with whom they share the same space ^{106,107}. A recent study showed that the presence of an android robot directly influences perceived humanness and spontaneous mimicry by participants ¹⁰⁸. First, participants rated the android higher on human-likeness when it was collocated with them, compared to being presented on a computer screen. Second, while spontaneous mimicry was robustly observed across participants for a collocated android, only participants who rated the visually presented android higher on human-likeness showed spontaneous mimicry for this agent.

While evidence is so far limited to behavioural studies, indirect evidence of the impact of an artificial agent's physical embodiment and presence on social engagement at the brain level can be distilled from studies on gaze interaction ^{98,109}. Displaying the gaze behaviour of a human agent via a virtual avatar can result in increased feelings of presence of the human agent and increased positive evaluations of this agent ¹¹⁰. Schilbach and colleagues ¹¹¹ used an interactive gaze task to tease apart the effect of self-initiated and other-initiated joint attention at behavioural and brain levels. They found that other-initiated joint attention increased activation in the MPFC, a core region of the ToM network, and self-initiated joint attention increased activation in the ventral striatum, a region associated with reward processing. Interestingly, activity in these regions is decreased when participants believe that the gaze behaviour of the avatar has a computer origin ¹¹², suggesting less involvement of

social neurocognitive processing when an interactant has artificial origins. These findings thus contribute to our understanding of how the physical embodiment and presence of an agent might potentially shape attributions of socialness, in addition to the agent's form and motion characteristics.

In sum, a number of studies have attempted to address the extent to which stimulus cues, including the physical features, movement parameters and presence of an artificial agent, influence behavioural and brain measures of socialness. To date, we argue that there is not enough evidence to suggest a reliable, clear impact of any of these features on socialness attributions. Instead, the work reviewed demonstrates that both the PPN and the AON are flexibly engaged when perceiving a diverse array of artificial agents, from schematic faces to social animations of shapes to mechanical and humanoid robots. Regarding the AON, several studies already suggest that it is not the appearance and motion of an artificial agent, but instead expectations or familiarity that might be more important in driving engagement of this network. In contrast, activity in the ToM network appears to be sensitive to the physical presence of an artificial agent, and whether or not there is a social narrative that can be ascribed to groups of animated shapes, but is not particularly sensitive to the presence or absence of biological motion. In contrast, several researchers have concluded that visual features of the artificial agent influence the attribution of socialness at the brain and behavioural level to some extent ^{63,113}. Naturally, it seems likely that the prospect of social behaviours being present, and thus the potential for a human interaction partner to attribute socialness, is higher in a robot designed to look like a human than a robot with a more machine-like appearance. However, the attribution of socialness will not necessarily follow from a human-like visual appearance of the artificial agent only. As we describe below, beliefs, expectations and the prior experience of the individual also all contribute to

attributions of socialness above and beyond the artificial agent's appearance. For instance, the previous experience of the human agent with a sophisticated mechanical robot or the unrealistic high expectation of the individual for the humanoid robot can counteract any effect of visual appearance ¹¹⁴.

Impact of Knowledge Cues the Human Observer

Belief and expectations

Turning our focus from the artificial agent to the human observer or interaction partner, in this section we explore how the knowledge, thoughts, beliefs, and expectations that a person brings to an interaction with an artificial agent shapes the extent to which the agent is perceived in a social or non-social manner. In the classic cognitive psychology literature, such factors fall under the category of top-down processes. Top-down cognitive processes are endogenous to the perceiving/acting individual, driven by contexts, knowledge or goals. Such processes can help facilitate perception, in that our past experience or knowledge can help us to formulate predictions about what is going to happen next^{115,116}. Internal knowledge representations guide visual processing ¹¹⁷, and importantly, the attribution of mental states influences the perception of social cues ¹¹⁸. Below, we examine how research insights from the behavioural and brain sciences advance our understanding of the human side of human-robot interaction, and the importance of knowledge cues in attributing socialness to artificial agents and fostering social connections with these agents.

Some of the earliest work on the impact of knowledge cues on social perception comes from Stanley, Gowen and Miall ¹¹⁹. In this study, the authors used an elegant paradigm that required participants to follow the trajectory of a bouncing dot with their arm. The dot followed either a biologically plausible or biologically impossible (i.e., mechanical) velocity

profile, and participants were instructed that the movement they were watching was either pre-recorded human movement or computer-generated movement. The authors found that participants' belief about the human origins of the moving dot stimulus had a stronger impact on their actions than whether or not the velocity profile was biological or mechanical in nature. This led Stanley and colleagues to conclude that even when a cue is a simple bouncing dot, our beliefs about the human origins of a dot's movements can shape behaviour more than the whether or not the dot moves in a biologically plausible manner. A related study by Liepelt and Brass ¹²⁰ also examined the extent to which an observer's actions are influenced by their beliefs about an observed action's humanness, but this time, in a further step toward ecological validity, the stimuli featured hands performing finger lifting movements. Critically for our purposes, before this task, half the participants received instructions that the hand they were about to see was a human hand wearing a glove, while the other half were shown that the same gloved hand was in fact a wooden hand artist's model wearing a glove. The authors reported stronger automatic imitation in participants who were told they were watching a human hand wearing a glove, suggesting here again that our beliefs about the human origins of an action strongly shape the extent to which we behaviourally respond to them in a social manner ¹²⁰.

Converging evidence comes from a number of other studies employing various belief manipulations and paradigms ¹²¹⁻¹²⁸. For example, Wykowska, Wiese and colleagues showed across a series of gaze cueing experiments that an observer's belief also shapes gaze following ^{126,127}. When people were told they were observing an intentional agent (a human or a human-controlled robot), gaze following was stronger than when they were told the agent had no intention (a robot or a human-like mannequin). While these studies contrast the belief that an agent is human versus a robot or machine, one study comparing a 'human-like'

robot versus a ‘machine-like’ robot found similar effects during a joint action task ¹²². Crucially, the effect of belief manipulations is related to true imitation ¹²⁵, rather than attentional effects in combination with general stimulus-response compatibility effects that can confound automatic imitation studies ¹²⁹. It is of note, however, that earlier work contrasting knowledge and stimulus cues found that only stimulus cues appeared to impact automatic imitation of hand movements ⁸⁸.

On balance, behavioural findings build a largely (though not completely) consistent case for a strong influence of knowledge cues on shaping perceptions of socialness. In the past several years, a number of brain imaging studies have further advanced our understanding of the influence of knowledge cues to socialness by revealing marked differences in neural processing based on participants’ expectations about the human or artificial origins of a perceived agent. The first study along these lines was again performed by Stanley, Gowen and Miall ¹³⁰. This time, the authors asked participants undergoing fMRI scanning to watch a number of point light animations of simple actions (such as walking, kicking a ball, lifting a box, etc), which could either be presented as originally recorded, or with different levels of noise introduced so that the dots appeared to be moving in an increasingly random manner. Participants were notified before each video whether it featured human or computer-generated movement, and their task was to decide whether this label was accurate or inaccurate. The authors found that participants were more likely to agree that a stimulus looked like a person moving if they were told a video had human origins, while watching the identical video paired with a computer-generated label led to different behavioural responses, as well as different patterns of neural activation. Specifically, the ventral paracingulate cortex was most strongly engaged when watching those videos believed to have human origins, while the dorsal paracingulate cortex was most active when participants viewed ambiguous

stimuli (such as scrambled dots with human instructions, or human action-like dots with computer generated instructions).

A subsequent functional neuroimaging study used a similar automatic imitation paradigm to Liepelt and Brass’s ¹²⁰ paradigm, described above, with a few important differences ¹³¹. First, this study employed a within-subjects design to determine the extent to which differences in stimulus and knowledge cues to human animacy can be observed in the same group of participants. Second, the authors asked participants to watch a 20-minute bespoke documentary wherein two kinds of cutting-edge film making techniques were described: human motion capture, and computer keyframe animation. They were then shown two different hand stimuli – one that featured an avatar of realistic-looking human hand, and the other featured two blocky robot-like fingers. Participants performed the same imitation for both kinds of “hand” stimuli, but half of the trials were preceded with instructions that the following stimuli were made from human motion capture, and the other half with the instructions that the following videos were made with computer keyframe animation. In reality, all stimuli followed the identical motion profile, regardless of instructions. Somewhat in contrast to what Liepelt and Brass ¹²⁰ reported with their between-subjects behavioural study, Klapper and colleagues found that *any* cue to humanness led to greater motor priming compared to when no cues to humanness were present (in other words, more interference of executed movements was observed when participants observed a human hand avatar and/or either kind of hand paired with the instruction that the video had human motion capture origins). In contrast, the brain imaging data demonstrated that the right TPJ, a brain region that is often associated with cognitive processes that involve self-other distinctions, was most strongly engaged when participants were performing the imitation task when both knowledge and stimulus cues to human animacy were present ¹³¹. The authors suggest that this finding

underscores the critical role the right TPJ plays in mediating interactions with other human (but not artificial) agents. These findings are corroborated by a number of other previous and subsequent findings¹³²⁻¹³⁶. For instance, a recent study using an established gaze cueing paradigm^{126,127,137} reported increased activation in bilateral TPJ for gaze following when participants thought the gaze had human origins compared to a pre-programmed, computer-determined origin¹³⁵.

Another recent functional neuroimaging study helped to further illuminate the differential roles played by stimulus and knowledge cues in social perception¹⁴. In this study, Cross and colleagues made use of the same documentary used by Klapper and colleagues¹³¹, but this time paired the documentary with an elaborate cover story that the researchers were working for the German Film Commission and were tasked with evaluating participants' perceptions of the smoothness and likeability of different simple action movies, based on their human versus computer origins. Again, the motion profiles were the same for each video and half the videos featured a human avatar performing a number of simple goal-directed actions (such as tidying a table, stacking blocks, or hammering a nail), or an abstract-looking robotic avatar performing the same actions. Participants reported liking and finding most smooth those videos that were paired with human origin instructions (regardless of whether the actions were performed by a human or robot avatar). The brain data demonstrated that visual features of an agent appear to primarily influence ventral temporal brain regions, such as the fusiform gyrus, when observing a robot compared to a human avatar. This increase in activity in the fusiform gyrus for robotic compared to human agents is in accordance with previous studies on the perception of facial expressions of emotions in robotic agents^{45,46}. Crucially, brain regions associated with Theory-of-Mind processes, including the precuneus, are more

strongly engaged when videos are paired with human compared to computer-generated instructions.

Taken together, the studies examined in this section begin to build a compelling case that the knowledge cues that a participant has when perceiving or interacting with another agent, whether human or artificial in appearance, have a strong impact on behaviour and underlying brain circuits.

Anthropomorphism and Dehumanization

Corroborating evidence of the central role played by knowledge cues to socialness comes from research on anthropomorphism. Anthropomorphism can be defined as our propensity to attribute human-like qualities, characteristics and behaviours, to non-human agents, entities, or objects, by the observer. The tendency to anthropomorphise animals, robots and computers is a stable characteristic that is robust over time¹³⁸, and already present in young children¹³⁵. Research on anthropomorphism suggests that it is not the visual features of the agent, but instead knowledge factors and the social motivation of the observer that are central in attributing socialness to non-human agents¹³⁹.

Several studies have investigated the process of anthropomorphising and corresponding brain networks. A recent study reported that an individual's disposition to anthropomorphize is related to grey matter volume of the left TPJ¹⁴⁰. Specifically, they found that a greater tendency to anthropomorphize non-human animals is positively correlated with TPJ volume. However, the authors did not observe any correlation between brain structure and a disposition to anthropomorphize non-animal agents or objects. While this already provides evidence on the role of the ToM network in anthropomorphism, further details come from

studies measuring brain activity during the active process of anthropomorphising^{141,142}. An early study by Chaminade, Hodgins and Kawato¹⁴¹ investigated how the visual appearance of an animated character influenced the perception of motion. The agents ranged from an animated simple or complex point-light display, to a stick figure, robot, alien, clown or a human-like jogger. Participants' task was to indicate if the agents' movement was biological or computer-generated. Importantly, all agents moved with the same motion parameters, which came in one condition directly from motion capture data of a human actor's movements (biological motion), while in the other condition they were computer-generated movements (non-biological motion). Crucially, brain activity did not serve as a function of characters' appearance, but was instead modulated by participants' anthropomorphic bias (in this study, a participant's tendency to report the observed motion as biological). The authors reported a positive correlation between this response bias and activity in the ToM network, specifically within left TPJ and bilateral precuneus. Thus, a tendency to perceive the motion of the less human-like agents as biological in nature tracked with increased engagement of the left TPJ. Interestingly, however, this response bias was negatively correlated with activity in regions of the AON¹⁴¹.

Besides the features of an artificial agent and an observer's belief, anthropomorphism is also influenced by the motivation to understand and predict the social environment¹⁴³. Waytz and colleagues¹⁴² showed that this so-called effectance motivation is directly related to a tendency to anthropomorphise artificial agents as well as objects. In a series of five experiments, they showed that when uncertainty and unpredictability, two key determinants of effectance motivation, increase, so does anthropomorphism. For example, when participants interacted with an unpredictable robot, they rated this robot higher on anthropomorphic aspects, such as having its own mind, intentions, free will, consciousness,

desires, beliefs, and the ability to experience emotions, compared to participants who interacted with a predictable robot. Similarly, when asked to predict the behaviour of an unfamiliar robot, participants' anthropomorphic judgements also increased, compared to when participants were not asked to predict its behaviour. The authors showed that the tendency to anthropomorphise the actions/behaviours of unpredictable objects is related to increased activity within brain regions associated with the ToM network. Participants first read about gadgets that were either predictable or unpredictable. Next, they answered the question to what extent the gadget had a mind of its own. Behavioural results showed that again participants were more likely to attribute mind-like qualities to unpredictable gadgets compared to predictable gadgets. Importantly, activity in the ventral part of MPFC was increased for unpredictable compared to predictable gadgets. The involvement of the ToM network in these anthropomorphism judgements was further confirmed in a connectivity analysis, showing functional connectivity with the precuneus as well as the anterior cingulate cortex. Furthermore, activity in the ventral part of the MPFC directly covaried with anthropomorphism judgements across participants. In sum, a greater tendency to anthropomorphise is related to structural and functional changes within the ToM network. An area ripe for future exploration concerns how predictability and appearance of the artificial agent interact with the belief and tendency to anthropomorphise of the human observer, and the role of the ToM network, especially the MPFC in these processes¹⁴⁴. The engagement of brain regions associated with the ToM network is also robustly implicated in studies on dehumanizing behaviour, a similar but opponent process to the attribution of socialness to artificial agents.

Dehumanization is the process by which a human agent or group of individuals is seen to possess fewer human-like qualities compared to another person or group of individuals.

People not only view themselves as more human than others¹⁴⁵, but can view other humans or groups as closer to animals or automata¹⁴⁶. An early study by Harris and Fiske¹⁴⁷ found that activity in the MPFC was decreased when participants viewed pictures of members of an extreme outgroup that was rated low on both warmth (similar to experience or human nature) and competence (similar to agency or human uniqueness), such as homeless individuals or drug addicts. Activity in the MPFC, as well as other regions of the ToM network, have been documented to show similar effects for other aspects of dehumanization, such as sexual objectification of women¹⁴⁸, or viewing people as products¹⁴⁹. A recent study showed that activity in the ToM network, namely the MPFC and right TPJ, is decreased when participants are told that the observed person is close to a machine¹⁵⁰. Findings from research on dehumanization will inform us further on the factors that might (negatively) influence the attribution of socialness to artificial agents.

Eyssel and colleagues showed across a series of experiments¹⁵¹⁻¹⁵³ that similar biases are present in human-robot interaction as in human-human interaction. That is, people favour a robot that belongs to their in-group and ascribe more human-like characteristics to these robots. Similar biases were observed for people interacting with a computer¹⁵⁴⁻¹⁵⁶ or a virtual human^{157,158}. Thus, an increase in the human-likeness of the artificial agent might increase the potential for biases or stereotypes to enter the socialness equation. Another influence on socialness attribution comes from the perspective put forward by Ferrari, Paladino and Jetten¹⁵⁹. These authors suggest that artificial agents' threat to distinctiveness also has the potential to influence socialness attribution and social behaviour towards artificial agents. The authors argue that robots with human-like appearances threaten to blur the boundaries between humans, machines and other artificial agents. This perception of threat to human uniqueness can therefore reduce the attribution of socialness to artificial agents. This process could be

similar to dehumanization of human agents. Thus, the attribution of socialness to artificial agents is a highly contextual process that depends on the observer, the artificial agent, and the environment¹⁶⁰.

In sum, these studies confirm the vital role played by the ToM network, in particular the TPJ and vMPFC (but see¹⁶¹) in the attribution of socialness and extend our understanding of factors that influence these attributions. It is not so much the features of the artificial agent that drive the attribution of socialness to artificial agents, but the belief and expectations of the human observer as well as her tendency to attribute human-like qualities to human and artificial agents as seen in the process of anthropomorphism and dehumanization.

Integrative perspective

When attempting to link together the broad range of findings covered in this review, it is instructive to ask whether a common theoretical thread links much of this work examining whether, how, and when people engage with artificial agents in a social manner. In light of this, we argue that a particular theoretical position borrowed from developmental psychology is useful for framing past findings and future questions about humans' social future with artificial agents. This theoretical position, termed the 'like me' hypothesis, states that understanding the basic similarity between self and other forms the foundation of social cognition, and that humans have evolved to seek out self-other equivalence in others^{162,163}. This account further proposes that actions performed by oneself and another are represented in common cognitive codes¹⁶², and early neurophysiological work establishing the existence of mirror neurons in the non-human primate brain^{31,32} has helped inspire much of the foundational work in this domain.

However, attributing socialness to another agent involves far more than linking perceived and executed actions between oneself and an observed other, as we have seen throughout this review. The representation of other agents' minds is a three-stage developmental process¹⁶⁰, which involves not just linking observed and executed acts (via imitation), but also the first-person everyday experience of action-intention coupling, and finally the ultimate step of understanding other agents by projection of one's own state onto another agent. This correspondence between self and other at multiple levels is vital to the attribution of socialness to artificial agents. As we have seen, this process is flexible in nature, and the 'like me' hypothesis provides a useful point of departure for contextualising brain and behavioural findings about the impact of form, motion, knowledge, experience, anthropomorphising and dehumanisation on the attribution of socialness to artificial agents. For instance, some of the studies reviewed above adhere neatly to the 'like me' hypothesis by demonstrating evidence of behavioural overlap^{63,64,78,79} and increased engagement of the AON when participants observe familiar actions or interact with agents similar to themselves^{65,130,131}, while others do not. For example, other studies have demonstrated similar or greater AON engagement when participants observe very much *unlike-me* robotic actions compared to more familiar human actions^{37,90-93}, with similar findings for behaviour engagement^{80,82-84,86-89}. Likewise, almost no correspondence with the 'like me' hypothesis is found at the level of the PPN.

However, neuroimaging findings consistently support the 'like-me' nature of ToM network engagement, especially within the MPFC and the TPJ, when interacting with socially similar others. Stimulus cues^{46,47,112}, knowledge cues¹³¹⁻¹³⁶, anthropomorphism^{141,142}, and dehumanization¹⁴⁷⁻¹⁵⁰ processes all impact activity in these regions (**Figure 1**). Further evidence for MPFC-mediated self-other equivalence comes from studies on self/other identification with a virtual avatar¹⁶⁴ or trust during interactions with an artificial agent¹⁶⁵.

Similarly, a wealth of evidence from human-human interaction supports the notion of a crucial role of the TPJ in inferring the mental states of others³³⁻³⁵, in differentiating the self from others during joint attention and perspective taking tasks^{77,166}, as well when making judgements of in-group versus out-group members¹⁶⁶.



Figure 1. Functional convergence of cognitive factors driving socialness attribution to artificial agents within the Theory-of-Mind network. Studies on knowledge factors, anthropomorphism and dehumanization robustly report engagement of the temporoparietal junction, precuneus and dorsal and ventral medial prefrontal cortex, regions within the Theory-of-Mind network. The dots illustrate the clusters of activation found in the studies and do not reflect exact coordinates. The automated term-based meta-analytic brain activation map for the Theory-of-Mind network was created and downloaded from the 'Neurosynth' database (<http://neurosynth.org>, December 1 2017)¹⁷⁷. The maps are based on 124 studies using the term 'mentalizing' and are corrected at FDR < .01.

A recent perspective on the TPJ suggests that it codes information about social context^{166,167}. This perspective dovetails with a mechanistic model on downstream effects of mental state attribution¹¹⁸, whereby the attribution of a mind in an observed agent influences perception of social cues associated with this agent. A study by Carter and colleagues¹⁶⁸ provides evidence for this social bias model during interactions with artificial agents. In this study, participants played a game of poker against human or computer opponents. The authors used activity from 55 bilateral brain regions to predict the decisions of the participant throughout

the game. While activity in regions associated with the ToM network were predictive for these decisions, this was independent of social context. Only activity within the TPJ predicted future decisions while taking into account the social context. That is, only when the participant deemed the human opponent superior to the computer opponent, and therefore more socially relevant, was activity in the TPJ informative of subsequent decisions. Combined, these findings suggest that the highly context-dependent implicit or explicit decision of the observer that the agent is ‘like me’, partly coded in the TPJ, is key in fostering the attribution of socialness to artificial agents.

Conclusions and Future Directions

The studies reviewed here support the notion that knowledge cues consistently impact engagement of behavioural and brain mechanisms supporting the attribution of socialness to artificial agents. The data also suggest that knowledge cues play an extremely (if not more) important role in socialness attributions than stimulus cues. Knowledge cues, as well as the process of anthropomorphism and dehumanization, influence ToM network engagement and determine the attribution of socialness. While the research reviewed here begins to provide answers as to why we can sometimes perceive a robot as an automaton, and at other times as a social agent, several questions remain.

One important question concerns how socialness attributions unfold over time. Most studies so far have examined passive observation of artificial agents or one-off interactions. A challenge for future work is to investigate real, repeated, and ongoing interactions with artificial agents in order to map functional changes in socialness attribution at behavioural

and brain levels across time^{98,99}. Several examples already exist that show how incorporating the temporal dimension of human-robot *relationships* can enrich our understanding of the mechanisms of socialness attribution¹⁶⁸. For instance, repeated exposure to robotic actions was shown to increase automatic imitation of these actions to levels comparable to human actions⁸⁹. Moreover, several recent studies on trust during interactions with artificial agents or machines provide compelling first evidence on how expectations and changes in the involvement of the ToM network shape these interactions over time^{165,169,170}. An exciting opportunity for future research is to study in more detail the dynamic, temporal dimensions of these processes.

Work in this domain already builds upon a solid foundation of findings from psychology and neuroscience on social cognition during human-human interaction. Integrating findings and approaches from additional related fields stands to advance understanding of not only the factors that drive the attribution of socialness to artificial agents, but also the temporal dynamics. To this end, future work could benefit from considering work on human attachment and relationship formation¹⁷¹, as well as the emerging field of human-animal interactions¹⁷²⁻¹⁷⁴. Studies on these latter interactions not only provide converging evidence on behavioural and brain mechanisms of socialness attribution¹⁷³, but also can help us to understand how long-term interaction with non-human agents shapes these attributions over time by studying pet owners versus non-pet owners^{172,174}. Thus, thinking openly and creatively about how work from distinct but complementary disciplines might inform our understanding about humans’ evolving relationship with socially-savvy technology.

Given the importance of knowledge factors during interactions with artificial agents, as we highlight here, it will also be enormously important for future research to acknowledge and

investigate inter-individual and group differences, as well as developmental changes in the attributions of socialness. Dispositional levels of anthropomorphism and dehumanization work in concert with knowledge and stimulus cues. Similarly, not only are developmental considerations important to further our understanding of the behavioural and brain mechanisms supporting the attribution of socialness, they are crucial given that both children and the elderly are target groups for the deployment of social robots, and age-dependent effects have been reported with regard to socialness attribution^{175,176}. Lastly, group and cultural differences that determine the scope of attribution of socialness are likely at play. With the potential for far-reaching consequences, for instance on perceived human uniqueness, a nuanced approach that takes into account dispositional, situational and cultural factors is warranted to truly capture the mechanisms and dynamics of socialness attribution across space and time.

Acknowledgements

The authors thank Kami Koldewyn, Anna Henschel, and Katie A. Riddoch, as well as the anonymous reviewers for valuable comments on an earlier version of the manuscript and gratefully acknowledge funding from the European Research Council to E. S. C. (H2020 ERC-2015-StG-67720-SOCIAL ROBOTS).

References

- Broadbent E. 2017. Interactions With Robots: The Truths We Reveal About Ourselves. *Annu Rev Psychol* **68**: 627–652.
- Chaminade T. & G. Cheng. 2009. Social cognitive neuroscience and humanoid robotics. *J. Physiol. Paris* **103**: 286–295.
- Dautenhahn K. 2007. Socially intelligent robots: dimensions of human-robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences* **362**: 679–704.
- Waytz A., K. Gray, N. Epley, et al. 2010. Causes and consequences of mind perception. *Trends Cogn. Sci. (Regul. Ed.)* **14**: 383–388.
- Frith U. & C.D. Frith. 2009. The social brain: allowing humans to boldly go where no other species has been. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**: 165–176.
- Dennett D.C. 1989. *“The Intentional Stance.”* Cambridge, Massachusetts: MIT Press.
- Dennett D.C. 1971. Intentional systems. *The Journal of Philosophy* **68**: 87–106.
- Haslam N. & S. Loughnan. 2014. Dehumanization and infrahumanization. *Annu Rev Psychol* **65**: 399–423.
- Gray H.M., K. Gray & D.M. Wegner. 2007. Dimensions of mind perception. *Science* **315**: 619.
- Fiske S.T., A.J.C. Cuddy, P. Glick, et al. 2002. A model of (often mixed) stereotype content: competence and warmth respectively follow from perceived status and competition. *J Pers Soc Psychol* **82**: 878–902.
- Haslam N. 2006. Dehumanization: an integrative review. *Pers Soc Psychol Rev* **10**: 252–264.
- Kagan J. 2004. The uniquely human in human nature. *Daedalus* **133**: 77–88.
- Mason P. & H. Shan. 2017. A valence-free definition of sociality as any violation of inter-individual independence. *Proceedings of the Royal Society B: Biological Sciences* **284**: 20170948–6.
- Cross E.S., R. Ramsey, R. Liepelt, et al. 2016. The shaping of social perception by stimulus and knowledge cues to human animacy. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.* **371**: 20150075.
- Adolphs R. 2009. The social brain: neural basis of social knowledge. *Annu Rev Psychol* **60**: 693–716.
- Frith C.D. & U. Frith. 2012. Mechanisms of social cognition. *Annu Rev Psychol* **63**: 287–313.
- de Gelder B. 2006. Towards the neurobiology of emotional body language. *Nat. Rev. Neurosci.* **7**: 242–249.
- Peelen M.V. & P.E. Downing. 2007. The neural basis of visual body perception. *Nat. Rev. Neurosci.* **8**: 636–648.
- Adolphs R. 2017. Emotion Perception from Face, Voice, and Touch: Comparisons and Convergence. *Trends Cogn. Sci. (Regul. Ed.)* **21**: 216–228.
- Huth A.G., S. Nishimoto, A.T. Vu, et al. 2012. A continuous semantic space describes the representation of thousands of object and action categories across the human brain. *Neuron* **76**: 1210–1224.
- Kriegeskorte N., M. Mur, D.A. Ruff, et al. 2008. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* **60**: 1126–1141.
- Grill-Spector K., K.S. Weiner, K. Kay, et al. 2017. The Functional Neuroanatomy of Human Face Perception. *Annu Rev Vis Sci* **3**: 167–196.
- Quadflieg S. & K. Koldewyn. 2017. The neuroscience of people watching: how the human brain makes sense of other people's encounters. *Ann. N. Y. Acad. Sci.* **1396**: 166–182.
- Van den Stock J., M. Vandenbulcke, C.B.A. Sinke, et al. 2014. Affective scenes influence fear perception of individual body expressions. *Hum Brain Mapp* **35**: 492–502.
- Van den Stock J., M. Vandenbulcke, C.B.A. Sinke, et al. 2014. How affective information from faces and scenes interacts in the brain. *Soc Cogn Affect Neurosci* **9**: 1481–1488.
- Shultz S. & G. McCarthy. 2014. Perceived animacy influences the processing of human-like surface features in the fusiform gyrus. *Neuropsychologia* **60**: 115–120.
- Looser C.E., J.S. Guntupalli & T. Wheatley. 2013. Multivoxel patterns in face-sensitive temporal regions reveal an encoding schema based on detecting life in a face. *Soc Cogn Affect Neurosci* **8**: 799–805.
- Caspers S., K. Zilles, A.R. Laird, et al. 2010. ALE meta-analysis of action observation and imitation in the human brain. *Neuroimage* **50**: 1148–1167.
- Keysers C. & V. Gazzola. 2009. Expanding the mirror: vicarious activity for actions, emotions, and sensations. *Curr. Opin. Neurobiol.* **19**: 666–671.
- Molenberghs P., R. Cunnington & J.B. Mattingley. 2012. Brain regions with mirror properties: a meta-analysis of 125 human fMRI studies. *Neurosci Biobehav Rev* **36**: 341–349.
- Gallese V., L. Fadiga, L. Fogassi, et al. 1996. Action recognition in the premotor cortex. *Brain* **119 (Pt 2)**: 593–609.
- Rizzolatti G., L. Fadiga, V. Gallese, et al. 1996. Premotor cortex and the recognition of motor actions. *Brain Res Cogn Brain Res* **3**: 131–141.
- Frith U. & C.D. Frith. 2003. Development and neurophysiology of mentalizing. *Philosophical Transactions of the Royal Society B: Biological Sciences* **358**: 459–473.
- Molenberghs P., H. Johnson, J.D. Henry, et al. 2016. Understanding the minds of others: A neuroimaging meta-analysis. *Neurosci Biobehav Rev* **65**: 276–291.

35. Schurz M., J. Radua, M. Aichhorn, *et al.* 2014. Fractionating theory of mind: a meta-analysis of functional brain imaging studies. *Neurosci Biobehav Rev* **42**: 9–34.
36. Castelli F., F. Happé, U. Frith, *et al.* 2000. Movement and mind: a functional imaging study of perception and interpretation of complex intentional movement patterns. *Neuroimage* **12**: 314–325.
37. Oberman L.M., J.P. McCleery, V.S. Ramachandran, *et al.* 2007. EEG evidence for mirror neuron activity during the observation of human and robot actions: Toward an analysis of the human qualities of interactive robots. *Neurocomputing* **70**: 2194–2203.
38. Saygin A.P., I. Cicckli & V. Akman. 2000. Turing Test: 50 Years Later. *Minds and Machines* **10**: 463–518.
39. Turing A.M. 1950. Computing Machinery and Intelligence. *Mind* **59**: 433–460.
40. Koldewyn K., P. Hanus & B. Balas. 2014. Visual adaptation of the perception of “life”: animacy is a basic perceptual dimension of faces. *Psychon Bull Rev* **21**: 969–975.
41. Dubal S., A. Foucher, R. Jouvett, *et al.* 2011. Human brain spots emotion in non humanoid robots. *Soc Cogn Affect Neurosci* **6**: 90–97.
42. Chammam M., A. Foucher, J. Nadel, *et al.* 2010. Reading sadness beyond human faces. *Brain Research* **1348**: 95–104.
43. Craig R., R. Vaidyanathan, C. James, *et al.* 2010. Assessment of human response to robot facial expressions through visual evoked potentials. In 647–652.
44. Moser E., B. Derntl, S. Robinson, *et al.* 2007. Amygdala activation at 3T in response to human and avatar facial expressions of emotions. *Journal of Neuroscience Methods* **161**: 126–133.
45. Chaminade T., M. Zecca, S.-J. Blakemore, *et al.* 2010. Brain response to a humanoid robot in areas implicated in the perception of human emotional gestures. *PLoS ONE* **5**: e11577.
46. Gobbin M.I., C. Gentili, E. Ricciardi, *et al.* 2011. Distinct neural systems involved in agency and animacy detection. *J Cogn Neurosci* **23**: 1911–1920.
47. Wang Y. & S. Quadflieg. 2015. In our own image? Emotional and neural processing differences when observing human-human vs human-robot interactions. *Soc Cogn Affect Neurosci* **10**: 1515–1524.
48. Hadjikhani N., K. Kveraga, P. Naik, *et al.* 2009. Early (M170) activation of face-specific cortex by face-like objects. *NeuroReport* **20**: 403–407.
49. Liu J., J. Li, L. Feng, *et al.* 2014. Seeing Jesus in toast: neural and behavioral correlates of face pareidolia. *CORTEX* **53**: 60–77.
50. Hortensius R., F. Hekele & E.S. Cross. 2017. The perception of emotion in artificial agents.
51. Martini M.C., C.A. Gonzalez & E. Wiese. 2016. Seeing Minds in Others - Can Agents with Robotic Appearance Have Human-Like Preferences? *PLoS ONE* **11**: e0146310.
52. Sandoval E.B. 2016. Reciprocity in Human-Robot Interaction: A Quantitative Approach Through the Prisoner’s Dilemma and the Ultimatum Game. *Int J of Soc Robotics* **8**: 303–317.
53. Bartneck C. & J. Hu. 2008. Exploring the abuse of robots. *Interaction Studies* **9**: 415–433.
54. Slater M., A. Antley, A. Davison, *et al.* 2006. “A virtual reprise of the Stanley Milgram obedience experiments.” Rustichini, A., Ed.: Public Library of Science.
55. de Gelder B. & R. Hortensius. 2014. The Many Faces of the Emotional Body. In *New Frontiers in Social Neuroscience* Decety J. & Christen Y., Eds. 153–164. Berlin: Springer, Cham.
56. Johansson G. 1973. Visual perception of biological motion and a model for its analysis. **14**: 201–211.
57. Puce A., A. Rossi & F.J. Parada. 2015. Biological Motion. *Brain Mapping An Encyclopedic Reference* **3**: 125–130.
58. Simion F., E. Di Giorgio, I. Leo, *et al.* 2011. The processing of social stimuli in early infancy: from faces to biological motion perception. *Prog. Brain Res.* **189**: 173–193.
59. Troje N.F. & C. Westhoff. 2006. The inversion effect in biological motion perception: evidence for a “life detector”? *Curr. Biol.* **16**: 821–824.
60. Johnson M.H. 2006. Biological motion: a perceptual life detector? *Curr. Biol.* **16**: R376–7.
61. Pelphrey K.A., T.V. Mitchell, M.J. McKeown, *et al.* 2003. Brain activity evoked by the perception of human walking: controlling for meaningful coherent motion. *J. Neurosci.* **23**: 6819–6825.
62. Blake R. & M. Shiffrar. 2007. Perception of human motion. *Annu Rev Psychol* **58**: 47–73.
63. Press C. 2011. Action observation and robotic agents: learning and anthropomorphism. *Neurosci Biobehav Rev* **35**: 1410–1418.
64. Kilner J.M., Y. Paulignan & S.-J. Blakemore. 2003. An interference effect of observed biological movement on action. *Curr. Biol.* **13**: 522–525.
65. Tai Y.F., C. Scherfler, D.J. Brooks, *et al.* 2004. The Human Premotor Cortex Is “Mirror” Only for Biological Actions. *Curr. Biol.* **14**: 117–120.
66. Heider F. & M. Simmel. 1944. An experimental study of apparent behavior. *The American Journal of Psychology* **57**: 243–259.
67. Johnson S.C. 2003. Detecting agents. *Philosophical Transactions of the Royal Society B: Biological Sciences* **358**: 549–559.
68. Opfer J.E. 2002. Identifying living and sentient kinds from dynamic information: the case of goal-directed versus aimless autonomous movement in conceptual change. *COGNITION* **86**: 97–122.
69. Wheatley T., S.C. Milleville & A. Martin. 2007. Understanding animate agents: distinct roles for the social network and mirror system. *Psychological Science* **18**: 469–474.
70. Schultz J., H. Imamizu, M. Kawato, *et al.* 2004. Activation of the human superior temporal gyrus during observation of goal attribution by intentional objects. *J Cogn Neurosci* **16**: 1695–1705.
71. Blakemore S.-J., P. Boyer, M. Pachot-Clouard, *et al.* 2003. The detection of contingency and animacy from simple animations in the human brain. *Cerebral Cortex* **13**: 837–844.
72. Schultz J., K.J. Friston, J. O’Doherty, *et al.* 2005. Activation in posterior superior temporal sulcus parallels parameter inducing the percept of animacy. *Neuron* **45**: 625–635.
73. Isik L., K. Koldewyn, D. Beeler, *et al.* 2017. Perceiving social interactions in the posterior superior temporal sulcus. *Proc. Natl. Acad. Sci. U.S.A.* **114**: E9145–E9152.
74. Ramsey R. & A.F. de C. Hamilton. 2010. Triangles have goals too: understanding action representation in left aIPS. *Neuropsychologia* **48**: 2773–2776.
75. Hamilton A.F. de C. & S.T. Grafton. 2006. Goal representation in human anterior intraparietal sulcus. *J. Neurosci.* **26**: 1133–1137.
76. Jacoby N., E. Bruneau, J. Koster-Hale, *et al.* 2016. Localizing Pain Matrix and Theory of Mind networks with both verbal and non-verbal stimuli. *Neuroimage* **126**: 39–48.
77. Schurz M., M.G. Tholen, J. Perner, *et al.* 2017. Specifying the brain anatomy underlying temporo-parietal junction activations for theory of mind: A review using probabilistic atlases from different imaging modalities. *Hum Brain Mapp* **38**: 4788–4805.
78. Kilner J.M., A.F. de C. Hamilton & S.-J. Blakemore. 2007. Interference effect of observed human movement on action is due to velocity profile of biological motion. *Soc Neurosci* **2**: 158–166.
79. Bisio A., A. Sciutti, F. Nori, *et al.* 2014. Motor Contagion during Human-Human and Human-Robot Interaction. *PLoS ONE* **9**: e106172–10.
80. Cook J.L., D. Swapp, X. Pan, *et al.* 2014. Atypical interference effect of action observation in autism spectrum conditions. *Psychol Med* **44**: 731–740.
81. Chaminade T., D.W. Franklin, E. Oztop, *et al.* 2005. Motor interference between Humans and Humanoid Robots: Effect of Biological and Artificial Motion. In 96–101. IEEE.
82. Chaminade T. 2005. Motor interference between humans and humanoid robots: Effect of biological and artificial motion. In 96–101. IEEE.
83. Oztop E., D.W. Franklin & T. Chaminade. 2005. Human–humanoid interaction: is a humanoid robot perceived as a human? *International Journal of Humanoid Robotics* **2**: 537–559.
84. Hofree G., B.A. Urgen, P. Winkielman, *et al.* 2015. Observation and imitation of actions performed by humans, androids, and robots: an EMG study. *Front Hum Neurosci* **9**: 364.
85. Kupferberg A., M. Huber, B. Helfer, *et al.* 2012. Moving Just Like You: Motor Interference Depends on Similar Motility of Agent and Observer. *PLoS ONE* **7**: e39637–8.
86. Press C., G. Bird, R. Flach, *et al.* 2005. Robotic movement elicits automatic imitation. *Brain Res Cogn Brain Res* **25**: 632–640.
87. Bird G., J. Leighton, C. Press, *et al.* 2007. Intact automatic imitation of human and robot actions in autism spectrum disorders. *Proceedings of the Royal Society B: Biological Sciences* **274**: 3027–3031.
88. Press C., H. Gillmeister & C. Heyes. 2006. Bottom-up, not top-down, modulation of imitation by human and robotic models. *Eur. J. Neurosci.* **24**: 2415–2419.
89. Press C., H. Gillmeister & C. Heyes. 2007. Sensorimotor experience enhances automatic imitation of robotic action. *Proceedings of the Royal Society B: Biological Sciences* **274**: 2509–2514.
90. Saygin A.P., T. Chaminade, H. Ishiguro, *et al.* 2012. The thing that should not be: predictive coding and the uncanny valley in perceiving human and humanoid robot actions. *Soc Cogn Affect Neurosci* **7**: 413–422.
91. Cross E.S., R. Liepelt, A.F. de C. Hamilton, *et al.* 2012. Robotic movement preferentially engages the action observation network. *Hum Brain Mapp* **33**: 2238–2254.

92. Gazzola V., G. Rizzolatti, B. Wicker, *et al.* 2007. The anthropomorphic brain: The mirror neuron system responds to human and robotic actions. *Neuroimage* **35**: 1674–1684.
93. Urgen B.A., M. Plank, H. Ishiguro, *et al.* 2013. EEG theta and Mu oscillations during perception of human and robot actions. *Front Neurobot* **7**: 19.
94. Saygin A.P. & W. Stadler. 2012. The role of appearance and motion in action prediction. *Psychol Res* **76**: 388–394.
95. Kilner J.M., K.J. Friston & C.D. Frith. 2007. Predictive coding: an account of the mirror neuron system. *Cogn Process* **8**: 159–166.
96. Gangopadhyay N. & L. Schilbach. 2012. Seeing minds: A neurophilosophical investigation of the role of perception-action coupling in social perception. *Soc Neurosci* **7**: 410–423.
97. Schilbach L. 2014. On the relationship of online and offline social cognition. *Front Hum Neurosci* **8**: 278.
98. Schilbach L., B. Timmermans, V. Reddy, *et al.* 2013. Toward a second-person neuroscience. *Behavioral and Brain Sciences* **36**: 393–414.
99. Wykowska A., T. Chaminade & G. Cheng. 2016. Embodied artificial agents for understanding human social cognition. *Philos. Trans. R. Soc. Lond., B, Biol. Sci.* **371**.
100. Sanchez-Vives M.V. & M. Slater. 2005. From presence to consciousness through virtual reality. *Nat. Rev. Neurosci.* **6**: 332–339.
101. Jung Y. & K.M. Lee. 2004. Effects of Physical Embodiment on Social Presence of Social Robots. In 80–87. Valencia.
102. Kompatsiari K., V. Tikhonoff, F. Ciardo, *et al.* 2017. The Importance of Mutual Gaze in Human-Robot Interaction. In *Intelligent Virtual Agents* Brinkman W.-P., Broekens J., & Heylen D., Eds. 443–452. Cham: Springer International Publishing.
103. Kiesler S. 2008. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition* **26**: 169–181.
104. Li J. 2015. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *International Journal of Human-Computer Studies* **77**: 23–37.
105. Lazzeri N., D. Mazzei, A. Greco, *et al.* 2015. Can a Humanoid Face be Expressive? A Psychophysiological Investigation. *Frontiers in Bioengineering and Biotechnology* **3**: 21.
106. Seo S.H. 2015. Poor Thing! Would You Feel Sorry for a Simulated Robot?: A comparison of empathy toward a physical and a simulated robot. In 125–132.
107. Kwak S.S., Y. Kim, E. Kim, *et al.* 2013. What makes people empathize with an emotional robot?: The impact of agency and physical embodiment on human empathy for a robot. In 180–185. IEEE.
108. Hofree G., P. Ruvolo, M.S. Bartlett, *et al.* 2014. Bridging the mechanical and the human mind: spontaneous mimicry of a physically present android. *PLoS ONE* **9**: e99934.
109. Schilbach L. 2015. Eye to eye, face to face and brain to brain: novel approaches to study the behavioral dynamics and neural mechanisms of social interactions. *Current Opinion in Behavioral Sciences* **3**: 130–135.
110. Cross E.S., D.J.M. Kraemer, A.F. de C. Hamilton, *et al.* 2009. Sensitivity of the Action Observation Network to Physical and Observational Learning. *Cerebral Cortex* **19**: 315–326.
111. Schilbach L., M. Wilms, S.B. Eickhoff, *et al.* 2010. Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *J Cogn Neurosci* **22**: 2702–2715.
112. Pfeiffer U.J., L. Schilbach, B. Timmermans, *et al.* 2014. Why we interact: on the functional role of the striatum in the subjective experience of social interaction. *Neuroimage* **101**: 124–137.
113. Cracco E., L. Bardi, C. Desmet, *et al.* 2018. Automatic imitation: A meta-analysis. *Psychol Bull.*
114. Dautenhahn K. 1998. The art of designing socially intelligent agents: Science, fiction, and the human in the loop. *Applied Artificial Intelligence* **12**: 573–617.
115. Gregory R.L. 1968. Perceptual illusions and brain models. *Proc. R. Soc. Lond., B, Biol. Sci.* **171**: 279–296.
116. Summerfield C., T. Egner, M. Greene, *et al.* 2006. Predictive codes for forthcoming perception in the frontal cortex. *Science* **314**: 1311–1314.
117. Smith M.L., F. Gosselin & P.G. Schyns. 2012. Measuring internal representations from behavioral and brain data. *Curr. Biol.* **22**: 191–196.
118. Teufel C., P.C. Fletcher & G. Davis. 2010. Seeing other minds: attributed mental states influence perception. *Trends Cogn. Sci. (Regul. Ed.)* **14**: 376–382.
119. Stanley J., E. Gowen & R.C. Miall. 2007. Effects of agency on movement interference during observation of a moving dot stimulus. *J Exp Psychol Hum Percept Perform* **33**: 915–926.
120. Liepelt R. & M. Brass. 2010. Top-Down Modulation of Motor Priming by Belief About Animacy. *Experimental Psychology* **57**: 221–227.
121. Tsai C.-C., W.-J. Kuo, D.L. Hung, *et al.* 2008. Action Co-representation is Tuned to Other Humans. *J Cogn Neurosci* **20**: 2015–2024.
122. Stenzel A., E. Chinellato, M.A.T. Bou, *et al.* 2012. When humanoid robots become human-like interaction partners: corepresentation of robotic actions. *J Exp Psychol Hum Percept Perform* **38**: 1073–1077.
123. Stenzel A., T. Dolk, L.S. Colzato, *et al.* 2014. The joint Simon effect depends on perceived agency, but not intentionality, of the alternative action. *Front Hum Neurosci* **8**: 371–10.
124. Longo M.R. & B.I. Bertenthal. 2009. Attention modulates the specificity of automatic imitation to human actors. *Exp Brain Res* **192**: 739–744.
125. Gowen E., E. Bolton & E. Poliakoff. 2016. Believe it or not: Moving non-biological stimuli believed to have human origin can be represented as human movement. *COGNITION* **146**: 431–438.
126. Wiese E., A. Wykowska, J. Zwickel, *et al.* 2012. I See What You Mean: How Attentional Selection Is Shaped by Ascribing Intentions to Others. *PLoS ONE* **7**: e45391.
127. Wykowska A., E. Wiese, A. Prosser, *et al.* 2014. Beliefs about the minds of others influence how we process sensory information. *PLoS ONE* **9**: e94339.
128. Shen Q. 2011. The impact of participants' beliefs on motor interference and motor coordination in human-humanoid interactions. In 6–16.
129. Cho Y.S. & R.W. Proctor. 2003. Stimulus and response representations underlying orthogonal stimulus-response compatibility effects. *Psychon Bull Rev* **10**: 45–73.
130. Stanley J., E. Gowen & R.C. Miall. 2010. How instructions modify perception: an fMRI study investigating brain areas involved in attributing human agency. *Neuroimage* **52**: 389–400.
131. Klapper A., R. Ramsey, D.H.J. Wigboldus, *et al.* 2014. The control of automatic imitation based on bottom-up and top-down cues to animacy: insights from brain and behavior. *J Cogn Neurosci* **26**: 2503–2513.
132. Gallagher H.L., A.I. Jack, A. Roepstorff, *et al.* 2002. Imaging the intentional stance in a competitive game. *Neuroimage* **16**: 814–821.
133. Krach S., F. Hegel, B. Wrede, *et al.* 2008. Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PLoS ONE* **3**: e2597.
134. Chaminade T., D. Rosset, D. Da Fonseca, *et al.* 2012. How do we think machines think? An fMRI study of alleged competition with an artificial intelligence. *Front Hum Neurosci* **6**: 103.
135. Özdem C., E. Wiese, A. Wykowska, *et al.* 2016. Believing androids - fMRI activation in the right temporo-parietal junction is modulated by ascribing intentions to non-human agents. *Soc Neurosci* **7**: 1–12.
136. McCabe K., D. Houser, L. Ryan, *et al.* 2001. A functional imaging study of cooperation in two-person reciprocal exchange. *Proc. Natl. Acad. Sci. U.S.A.* **98**: 11832–11835.
137. Caruana N., P. de Lissa & G. McArthur. 2017. Beliefs about human agency influence the neural processing of gaze during joint attention. *Soc Neurosci* **12**: 194–206.
138. Waytz A., J.T. Cacioppo & N. Epley. 2010. Who Sees Human? The Stability and Importance of Individual Differences in Anthropomorphism. *Perspect Psychol Sci* **5**: 219–232.
139. Waytz A., N. Epley & J.T. Cacioppo. 2010. Social Cognition Unbound: Insights Into Anthropomorphism and Dehumanization. *Curr Dir Psychol Sci* **19**: 58–62.
140. Cullen H., R. Kanai, B. Bahrami, *et al.* 2014. Individual differences in anthropomorphic attributions and human brain structure. *Soc Cogn Affect Neurosci* **9**: 1276–1280.
141. Chaminade T., J. Hodgins & M. Kawato. 2007. Anthropomorphism influences perception of computer-animated characters' actions. *Soc Cogn Affect Neurosci* **2**: 206–216.
142. Waytz A., C.K. Morewedge, N. Epley, *et al.* 2010. Making sense by making sentient: effectance motivation increases anthropomorphism. *J Pers Soc Psychol* **99**: 410–435.
143. Epley N., A. Waytz & J.T. Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychol Rev* **114**: 864–886.
144. Gowen E. & E. Poliakoff. 2012. How does visuomotor priming differ for biological and non-biological stimuli? A review of the evidence. *Psychol Res* **76**: 407–420.
145. Haslam N., P. Bain, L. Douge, *et al.* 2005. More human than you: attributing humanness to self and others. *J Pers Soc Psychol* **89**: 937–950.
146. Loughnan S. & N. Haslam. 2007. Animals and androids: implicit associations between social categories and nonhumans. *Psychological Science* **18**: 116–121.
147. Harris L.T. & S.T. Fiske. 2006. Dehumanizing the lowest of the low: neuroimaging responses to extreme out-groups. *Psychological Science* **17**: 847–853.

148. Cikara M., J.L. Eberhardt & S.T. Fiske. 2011. From agents to objects: sexist attitudes and neural responses to sexualized targets. *J Cogn Neurosci* **23**: 540–551.
149. Harris L.T., V.K. Lee, B.H. Capestany, *et al.* 2014. Assigning economic value to people results in dehumanization brain response. *Journal of Neuroscience, Psychology, and Economics* **7**: 151–163.
150. Jack A.I., A.J. Dawson & M.E. Norr. 2013. Seeing human: distinct and overlapping neural signatures associated with two forms of dehumanization. *Neuroimage* **79**: 313–328.
151. Eyssele F.A. & D. Kuchenbrandt. 2011. My robot is more human than yours: Effects of group membership on anthropomorphic judgments of social robots. *Proceedings of the 24th IEEE/RSJ, International Conference on Intelligent Robots and Systems (IROS 2011)*.
152. Eyssele F. & D. Kuchenbrandt. 2012. Social categorization of social robots: anthropomorphism as a function of robot group membership. *Br J Soc Psychol* **51**: 724–731.
153. Kuchenbrandt D., F. Eyssele, S. Bobinger, *et al.* 2011. Minimal group-maximal effect? evaluation and anthropomorphization of the humanoid robot NAO. *Social robotics*.
154. Nass C., Y. Moon & N. Green. 1997. Are Machines Gender Neutral? Gender-Stereotypic Responses to Computers With Voices. *Journal of Applied Social Psychology* **27**: 864–876.
155. Nass C. 1996. Can computers be teammates? *International Journal of Human-Computer Studies* **45**: 669–678.
156. Nass C. & Y. Moon. 2000. Machines and Mindlessness: Social Responses to Computers. *Journal of Social Issues* **56**: 81–103.
157. Sacheli L.M., A. Christensen, M.A. Giese, *et al.* 2015. Prejudiced interactions: implicit racial bias reduces predictive simulation during joint action with an out-group avatar. *Sci Rep* **5**: 8507.
158. Sacheli L.M., M. Candidi, E.F. Pavone, *et al.* 2012. And yet they act together: interpersonal perception modulates visuo-motor interference and mutual adjustments during a joint-grasping task. *PLoS ONE* **7**: e50223.
159. Ferrari F., M.P. Paladino & J. Jetten. 2016. Blurring Human-Machine Distinctions: Anthropomorphic Appearance in Social Robots as a Threat to Human Distinctiveness. *Int J of Soc Robotics* **8**: 287–302.
160. Waytz A. & M.I. Norton. 2014. Botsourcing and outsourcing: Robot, British, Chinese, and German workers are for thinking—not feeling—jobs. *Emotion* **14**: 434–444.
161. Kühn S., T.R. Brick, B.C.N. Müller, *et al.* 2014. Is This Car Looking at You? How Anthropomorphism Predicts Fusiform Face Area Activation when Seeing Cars. *PLoS ONE* **9**: e113885–14.
162. Meltzoff A.N. 2007. “Like me”: a foundation for social cognition. *Developmental Science* **10**: 126–134.
163. Meltzoff A. & W. Prinz. 2003. “The imitative mind: development, evolution, and brain bases.” Cambridge, UK: Cambridge University Press.
164. Ganesh S., H.T. van Schie, F.P. de Lange, *et al.* 2012. How the human brain goes virtual: distinct cortical regions of the person-processing network are involved in self-identification with virtual agents. *Cereb. Cortex* **22**: 1577–1585.
165. Riedl R., P.N.C. Mohr, P.H. Kenning, *et al.* 2014. Trusting humans and avatars: A brain imaging study based on evolution theory. *Journal of Management Information Systems* **30**: 83–114.
166. Schuwerk T., M. Schurz, F. Müller, *et al.* 2017. The rTPJ’s overarching cognitive function in networks for attention and theory of mind. *Soc Cogn Affect Neurosci* **12**: 157–168.
167. Carter R.M. & S.A. Huettel. 2013. A nexus model of the temporal–parietal junction. *Trends Cogn. Sci. (Regul. Ed.)* **17**: 328–336.
168. Tanaka F., A. Cicourel & J.R. Movellan. 2007. Socialization between toddlers and robots at an early childhood education center. *Proc. Natl. Acad. Sci. U.S.A.* **104**: 17954–17958.
169. Goodyear K., R. Parasuraman, S. Chernyak, *et al.* 2016. Advice Taking from Humans and Machines: An fMRI and Effective Connectivity Study. *Front Hum Neurosci* **10**: 181–15.
170. Goodyear K., R. Parasuraman, S. Chernyak, *et al.* 2017. An fMRI and effective connectivity study investigating miss errors during advice utilization from human and machine agents. *Soc Neurosci* **12**: 570–581.
171. Vrtička P. & P. Vuilleumier. 2012. Neuroscience of human social interactions and adult attachment style. *Front Hum Neurosci* **6**: 212.
172. Hayama S., L. Chang, K. Gumus, *et al.* 2016. Neural correlates for perception of companion animal photographs. *Neuropsychologia* **85**: 278–286.
173. Spunt R.P., E. Ellsworth & R. Adolphs. 2017. The neural basis of understanding the expression of the emotions in man and animals. *Soc Cogn Affect Neurosci* **12**: 95–105.
174. Stoeckel L.E., L.S. Palley, R.L. Gollub, *et al.* 2014. Patterns of Brain Activation when Mothers View Their Own Child and Dog: An fMRI Study. *PLoS ONE* **9**: e107205–12.
175. McLoughlin N. & H. Over. 2017. Young Children Are More Likely to Spontaneously Attribute Mental States to Members of Their Own Group. *Psychological Science* **28**: 1503–1509.
176. McLoughlin N., S.P. Tipper & H. Over. 2017. Young children perceive less humanness in outgroup faces. *Developmental Science* **12**: e12539.
177. Yarkoni T., R.A. Poldrack, T.E. Nichols, *et al.* 2011. Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* **8**: 665–670.