



Received: 5 November 2022

Accepted: 2 May 2023

Published: 15 June 2023

Affective Responses to Embodied Intelligence. The Test-Cases of Spot, Kaspar, and Zeno

Alessandra Fussi (University of Pisa) – alessandra.fussi@unipi.it

Abstract

The design, planning, and public presentation of service robots is gradually eroding the traditional distinction between social robots and service robots. This paper concentrates on one service robot (Spot) and two humanoid robots employed in the therapy of children diagnosed with ASD (Kaspar and Zeno). The paper argues that 1) Spot, which looks like a dog but is clearly identifiable as a machine, is less likely than Kaspar and Zeno to provoke the so-called uncanny valley effect. 2) The symbolic connotations associated with the robot's appearance are positive in the case of the two humanoid robots, while Spot elicits opposite responses depending on the context of use (e.g. trust in an industrial facility/fear and outrage when employed by the police in an underfunded neighbourhood). 3) Kaspar and Zeno play a positive role in a therapeutic setting, as affective mediators between neurotypical adults and children diagnosed with ASD, but the claim that they help children to recognise emotions is overstated. 4) Fears connected with weaponised robots built by other companies in imitation of Spot invite reflections on the mutual shaping between technology, society and politics.

Keywords: Social robots, service robots, ASD, emotion recognition, uncanny valley, situated affectivity

1. Introduction

It is well known that humans tend to project sociality when confronted with natural or mechanical entities that *appear* to be endowed with autonomous movement, either for the purpose of understanding their movement or in order to interact with them. For this to happen, a robot does not need actual social capabilities. It is enough that it seems to have them. Brazeal (2003) distinguishes three types of social robots, from those that are less sophisticated to those that are more advanced. To the first class belong machines that are “socially evocative.” They have no real social capabilities, but they evoke the impression in users that they may have them. To the second class belong robots that have a “social interface.” These machines know how to communicate by adopting relatively fluid language, matched with appropriate facial expressions and gestures. However, they are not receptive. Hence, they cannot properly interact with users. To the more sophisticated class belong machines that are “socially receptive.” They can learn from exchanges with users, and are able to refine their communicative and interactive capabilities accordingly.

In this paper I argue that the traditional distinction between social robots and service robots is gradually being eroded in the design, planning, and public presentation of service robots. I do not concentrate on the actual interactive capabilities of robots (i.e., their strictly social skills). Rather, I discuss the affective responses elicited by two kinds of robots. My first example is a service robot named Spot, produced by Boston Dynamics, a company at the forefront of robot manufacturing, based in Waltham, Massachusetts. As I argue in sections 4 and 5, Spot is socially evocative in different important (and problematic) respects. My second example, which I discuss in section 3, concerns two robots with a social interface: Kaspar and Zeno. Kaspar is a child-sized humanoid robot developed by the University of Hertfordshire's Adaptive Systems Research Group; it was first produced in 2009, but has undergone several transformations since then. Zeno is a 43 cm tall humanoid robot developed by Hanson Robotics in 2007. They are both used for therapeutic purposes and research with children diagnosed with autism spectrum disorder (ASD). They show emotions, utter sounds and words, and move some body parts, but their features are simplified.

I will consider three factors that play a key role in the affective responses robots may elicit in the public. The first factor, which I will discuss in section 2, is the “uncanny valley effect,” i.e., the cognitive and affective dissonance that may be provoked by machines that are very close to reaching lifelike appearance, but fail to replicate it with respect to behavior, fine motor skills, or perceptual cues. I will maintain that service robots like Spot are much less likely to elicit the uncanny valley effect than machines expressly built for social interaction, such as Kaspar and Zeno.

The second factor is the capacity to evoke affective responses thanks to the symbolic connotations that may be associated with the robot's appearance in different situations. By “robot's appearance” I mean the way it looks (for example: Zeno and Kaspar look like children), the kind of creature it evokes without necessarily being shaped like it (Spot evokes the appearance of a dog, but looks like a mechanical entity), and the way it behaves (does its movement appear autonomous? Does it inspire trust, or is it likely to elicit fear and anxiety? Can it be perceived as sociable?). As I will argue, Spot's association with a dog can give rise to opposing reactions. It can stimulate friendly feelings in a context such as an industrial facility, in which it takes the role of a docile and helpful companion, while it may elicit anxiety and fear in situations in which one could associate a dog's behavior with hostility and aggression.

The third factor is also contextual and situational. The production of social robots has heavily relied on research that aims to identify the emotions as episodic events that find expressions in certain facial configurations. This is the kind of research on which robots as therapeutic tools for individuals with ASD are based. However, philosophers of emotions and psychologists have long stressed that we cannot really attune to the emotion someone is expressing without taking into account his or her situation (Barrett et al. 2011; Griffiths and Scarantino 2009; Hess and Hareli 2016). Contextual information is essential to get the expression right (is someone about to cry because she is happy or because she is sad?), to understand why a certain emotion occurs (does she show anger because she wants to obtain compensation after being insulted, or because she missed her train?), and to grasp whether something in the environment is used as emotional scaffolding (is she listening to rhythmic music in order to get through her workout? Is she dimming the lights in her apartment to calm down after a rough day? Is she feeding her nostalgia by looking at certain childhood pictures? Is she feeling “on the wrong planet” as an autistic person?).¹

¹ There is an extensive literature on situated affectivity. See Stephan, Walter, and Wilutzky 2014 and Colombetti and Roberts 2014 for the application of the conceptual framework of situated cognition to situated affectivity; on emotional offloading see Colombetti and Krueger 2015; for music as a tool for emotional offloading see Krueger 2014. Scaffoldings for nostalgia are discussed in Massantini 2020. Krause-Jensen and Rodogno (2023) discuss the affective dynamics involved in the feeling of “being wrong” in people on the autism spectrum. Slaby (2016) argues that the user-resource model often employed in the analysis of situated affectivity does not sufficiently illuminate cases in which a normative domain constrains affective responses through attunement and habituation. He proposes “mind invasion” as a better conceptual tool, especially with respect to certain patterns in corporate workplaces.

As I argue, the thesis that therapeutic robots such as Kaspar and Zeno help autistic children recognise emotions is not supported by present research, besides being problematic at the theoretical level. The problem of understanding others' emotions cannot be reduced to guessing a robot's expressions, especially considering the cognitive style of individuals diagnosed with ASD (see Happé and Frith 2006). By analysing the results of some empirical studies, I infer that the benefits provided by these robots concern, at the perceptual level, the possibility to avoid overstimulation effects. With regard to the affective sphere broadly understood, benefits seem mainly to derive from the fact that these machines can be used as mediators between different cognitive styles (i.e., between that of the child with ASD and that of the neurotypical adult in a therapeutic setting). In this role, robots can alleviate social anxiety and facilitate shared attention, bodily awareness, or the exchange of roles in a game. I suggest that robots might also be able to bring benefits to the therapeutic setting by regulating the anxiety involved at both sides of the interaction (in this sense, they can be seen as affective scaffoldings). Precisely because their role as mediators would be lost if such machines were used outside of a therapeutic setting, for example as "companions" to leave alone with children diagnosed with ASD, in my concluding remarks I suggest that the potential disadvantages in this different situation should be carefully evaluated.

In my view, situational aspects are also key with respect to service robots. A problem that seems to be under the radar of public representatives and engineers is that affective reactions to embodied intelligence are not best studied in the abstract, or in a lab. One needs to consider the specific social environments in which service robots might be employed. I argue this point by contrasting the affective reactions elicited by Spot in different instances. In section 4 I analyse two videos posted at different times on the manufacturers' website. I argue that the first video intentionally blurs the distinction between Spot's actual capacities as a service robot and the social capacities it appears to have. The second video, by contrast, presents it as a powerful tool for gathering information, transporting materials, and moving into environments that would be dangerous or inaccessible to humans, but it focuses mainly on conveying a different message: Spot does not have a social interface and is incapable of autonomous movement. In section 5 I explain why, in my view, the second video can be read as a response to the events that led the New York Police Department to discontinue its contract with Boston Dynamics, following the reactions of fear and anger elicited by Spot when it was used by the police in the Bronx. In the remainder of the section and in the conclusion I highlight a further problem: machines like Spot can be weaponised. In this case, the context of interest is wider than the manufacturer of this specific machine in this particular time and place. My concluding point is that the question of whether service machines like Spot will bring more benefits than dangers must necessarily involve considerations about future scenarios in a globalised world. Hence the need to involve community members, users, philosophers of emotions, psychologists, and sociologists, so that problematic issues may become apparent before it is too late.

2. Anthropomorphic Projections and the Uncanny Valley Effect

As Hegel et al. (2009, 171) noted, certain conditions favour anthropomorphic projection: "Humans attribute human-like qualities to nonhuman agents depending on several parameters like appearance and social context – and consequently expect human-like actions."

There are limits, however, to anthropomorphic projections. Mori (1970) drew scholarly attention to a psychological phenomenon that hinders them. It became known, after the title of Mori's article, as the "uncanny valley" effect. One might expect a proportional relationship between technical refinement in the

production of robots that increasingly resemble human beings and the empathy and sense of trustworthiness they may generate in response. The relationship would be comparable to the one we experience when, pressing our car's accelerator, we expect it to increase its speed. However, Mori argued, the production of humanoid robots is better understood by analogy with mountain climbing: it is not true that the further you walk, the higher you climb, because, after reaching a peak, you are bound to descend into a valley. In the case of robots and prosthetics, the risk one runs after reaching the peak of extreme similarity to human shapes is to fall down into the uncanny valley. Minor differences indeed elicit feelings of repulsion. The consequence is that humanoid robots may end up hindering anthropomorphic projections precisely because they are too similar to the humans they are meant to resemble.

Mori identified two areas in which progressive similarity can reach critical peaks, with the transformation of a feeling of affinity into an eerie sensation:

- 1) Physical (static) resemblance. A wooden prosthetic hand is less similar to a human limb than a technologically advanced prosthetic hand in which nails and veins are carefully emphasised. A person shaking a hand of the latter type can be disturbed by the soft, boneless texture, or by the lack of warmth. All of a sudden the sense of affinity initially elicited by the way the hand looks is lost, and one finds it uncanny. This affective reaction is unlikely to happen in reaction to a well-crafted wooden prosthesis.
- 2) Functional mobility. A robotic arm unloading packages in a warehouse may be far less upsetting than a near-perfect humanoid robot that laughs while moving only certain facial muscles, speaks in a slowed-down manner, or does not move its eyes as one would expect.

The *uncanny valley* effect has been of great concern to those involved in digital animation. For example, the problem was found in the 2004 animated film *Polar Express*, directed by Robert Zemeckis. The movement and appearance of the characters is extremely realistic (much more so than in a typical Disney animated film of the same period). The accuracy and realism are achieved through the technique of "performance capture," in which an actor (in this case Tom Hanks) is filmed by a large number of computerised cameras as he acts in the scene. His movements are then transferred to the 3D models of the characters. The result in *Polar Express* is mostly excellent, but at times it causes a creepy effect. As is apparent in some of the frames, the protagonists are lifelike throughout, except for the detail in their eyes, which appear wide open, motionless, like those of the dead. In some cases, reviewers have complained about the lack of adequate facial muscle movement when characters speak. The realism of the figure is not matched by the micro-movements that characterise a living face. The motionless gaze, the reflection of light absent from the pupils, the eyelids not blinking at regular intervals, the rigid cheeks, are indeed disturbing.

Mori's basic suggestion is that if designers want their creations to elicit a feeling of affinity, trust, or even tenderness, as in the case of certain dolls or animated characters, they should avoid high degrees of human likeness, because in that case minimal elements of dissonance are sufficient to elicit the eerie effect. The original paradigm of the uncanny, according to Mori, is that of the corpse. In every way similar to the living body, it arouses deeply uncomfortable feelings once one notices its pallor or feels its icy temperature. It is no accident that the zombie and its variations are ubiquitous elements in horror films. And it is precisely the zombie effect to which some critical commentators trace certain images from *Polar Express*.

Damiano and Dumouchel (2018, 2), who echo Mori (1970), agree with him that a robot may look uncanny when a great resemblance to human traits is matched by uncoordinated or oddly mechanical movement. They point out, however, that the reverse is also true, namely, that if an object that does not physically resemble a human being behaves autonomously and in coordination with human movement, it provokes a social response, and people tend to interact with it as if it were human.

This asymmetry in anthropomorphic projection leads Damiano and Dumouchel (2018, 3) to contrast the affective response elicited by traditional dolls and that generated by robots. Traditional dolls are *ascribed* human traits thanks to a perceived physical resemblance. Robots, on the other hand, may look nothing like humans in physical appearance, but their movements may provoke the *inference* that they have human traits.

I find the authors' distinction between a projection based on static similarity and a projection based on motion similarity very interesting. It should be noted, however, that the second type of projection is not always traceable to inferential reasoning. Urquiza-Hass and Kortschal (2015), summarising the research of cognitive scientists and neuroscientists, observe that the cognitive system involved in identifying others as agents (and not just as external objects) is mostly not reflexive in nature. Rather, what is mostly at play is a (pre-reflexive, we might say) mechanism that involves areas of the brain deputed to social cognition. Citing the studies of Gallese and Rizzolatti, the authors recall that, being deputed to the imitation of movement, the mirror neuron system is also responsible for the synchronisation of movement and emotions in groups. Moreover, it serves to attribute intentionality not only within the human species, but also to animals, up to and including robots, especially when the observed actions belong to the repertoire of actions that agents can relate back to themselves. The recognisability of actions does not depend on the observed species, but simply on the similarity of their movement to one's own movement:

The generality of this process indicates that the MN [mirror neuron] system may be less dependent on species-specific shape features than on general motor properties of subjects, animated such as in animals, or evoking the impression of animacy, such as in robots. Results of a functional magnetic resonance imaging (fMRI) study showed that when people observed motor actions of humans (talking, reading and biting), monkeys (lip smacking and biting) or dogs (barking and biting), the difference in activation of the motor and visual areas depended not on the species but on the actions shown (Buccino et al., 2004). Actions that are part of the observer's motor repertoire (talking, reading and biting) are processed by their motor system, including MN, while actions that are not in the observer's repertoire (lip smacking and barking) are processed based only on their visual properties. (Urquiza-Hass and Kortschal 2015, 169).

The thesis that actions such as eating, drinking, talking, and running are recognised by a system that does not reflect on the specific qualities of the object (e.g., similarity of shape), but automatically mirrors its movement, prompts the hypothesis that anthropomorphic identification is best achieved when reflexive modes of interpretation are not activated. Simply put, the less one reflects, the fewer obstacles there are to anthropomorphic projection. We can infer that anthropomorphic projection is favoured when, for a variety of reasons, conscious reflection is hindered, but at the same time there is also a sense of urgency to ascribe intentionality to certain movements.

Epley, Waytz, and Cacioppo (2007) argue that the urgency toward anthropomorphic mirroring is facilitated by three conditions: the first is cognitive, the second and third motivational. On the cognitive side, the authors

point out that one can normally rely on a wide range of knowledge and beliefs about humans in general, while little is known about nonhuman agents. It comes naturally to interpret problematic phenomena one encounters in different species, or even in mechanical objects, as if they belong to the human domain. The second and third conditions are *effectance* and *sociality*. *Effectance* consists in the need to act effectively in response to one's environment, and this also involves interpreting present stimuli and making correct predictions about the future. The attribution of human characteristics to nonhuman agents reduces the sense of uncertainty about the future behaviours of objects with which one interacts, and consequently increases one's confidence in the ability to respond appropriately.

The third condition, *sociality*, consists of the need to establish and maintain social relationships. It is not necessary to dwell at length on this type of need, which came dramatically to the fore during the COVID-19 pandemic. Since interaction and physical closeness are basic psychological needs, isolation breeds suffering, and a lack of social interaction provides a strong stimulus to anthropomorphism. The relationship of the main character in the film *Cast Away* with his personified friend offers an emblematic example of such a mechanism. Kelly (Tom Hanks), shipwrecked on a desert island, finds a Wilson Sporting Goods volleyball along with other debris from the crashed plane. At some point he draws a rudimentary face with his bloody handprints on the volleyball and starts treating it as his only companion. Talking to Wilson, expressing to it and sometimes at its expenses, the frustration of the castaway, saves Kelly from madness by allowing him to retain a connection to the world he has momentarily lost. Most importantly, anthropomorphic projection allows Kelly to cling to hope. He also persists in his many disappointing and frustrating attempts to go back to a life shared with others because his constant interaction with Wilson allows him to keep his social skills and desires alive (on the active power of hope and the difference between hoping and daydreaming, see Bovens 1999; Walker 2006; Meirav 2009; Miceli and Castelfranchi 2010).

Epley, Waytz, and Cacioppo (2007, 866) comment on the third motivation for anthropomorphism thus:

In the absence of social connection to other humans. . . . People create human agents out of nonhumans through anthropomorphism to satisfy their motivation for social connection. This predicts that anthropomorphism will increase when people feel a lack of social connection to other humans and decrease when people feel a strong sense of social connection.

While it is true that anthropomorphic projection is most powerful when there is a suspension of the reflective process, when the burden of cognitive uncertainty is strong but one feels the urgency of certain answers, and when there is a deep need for companionship, it is evident that the manipulation of emotions by means of sensory stimulation has a key role to play in robotics.

Those who find themselves in the conditions described by Epley, Waytz, and Cacioppo (2007) are more ready than others to see, in certain movements similar to their own (such as walking, raising an object, following a moving object with their gaze, etc.), purposeful and affectively charged actions. Arguably, this is why the social robot industry often targets as potential users those who have not yet developed forms of complex rationality (children), those who suffer from isolation and affective deficiencies (the elderly), those who are cognitively impaired (dementia patients and some residents of nursing homes), those who feel isolated, and those who have problems in social interaction (for example, children diagnosed with ASD).

3. Learning Emotions from Kaspar and Zeno?

Given these premises, we might assume that the way social robots look (their static properties) has a less significant impact on users than the way they move (their dynamic properties). In general, however, producers of social robots tend to stimulate users' anthropomorphic projections by building machines that look humanlike, zoomorphic, or have a mixture of animal, human, and fantastic features.

Kaspar and Zeno are interesting from the point of view of their affective impact. They supposedly help children on the autistic spectrum to recognise emotions. Furthermore, they are used to enhance capacities for interactive and collaborative games, and to help modulate force used in interaction with others.

What does it mean that they help children on the autistic spectrum to recognise emotions? Behind this claim lies the theory originally developed by Ekman (see Ekman 1992; Ekman and Davidson 1994), according to which certain basic emotions are expressed by all of us without appreciable cultural variation thanks to certain micro-movements of the face. Basic emotions are schematised on Russell's scale along two axes: the horizontal axis of valence (pleasure/pain) and the vertical axis of activation (major or minor) (see Russell 1980; Russell and Barrett 1999). On this model (or *circumplex*) we find on the point of least arousal fatigue (or drowsiness), and on the point of greatest arousal surprise. At the negative valence extreme we see sadness; at the positive extreme, contentment. Along the quadrant we find some intermediate conditions: categorised with positive valence/low arousal are emotional states such as feeling calm, serene, relaxed; with strong arousal/positive valence joy, elation, excitement; with strong arousal/negative valence disgust, anger, fear; with low arousal/negative valence sadness and feeling depressed.

Kasper and Zeno essentially exhibit the emotions we find on the circumplex. One problem, however, stands out: their faces lack the complexity of micro-movements that characterises a living face, so that they are likely to elicit the "uncanny valley" effect. Videos posted on the web show therapy sessions in which children are asked by educators to recognise the emotions expressed by the robot's face. In a video recorded with Zeno, an educator asks a child to guess whether the robot is happy, sad, angry, or scared by comparing the machine's expressions with corresponding *emoticons* drawn on *flashcards* (ITV News 2017). The robot's facial movements are rudimentary and repetitive: the sad expression is always the same, as is the cheerful expression, the angry expression, etc. The premise is that the child, who might be disturbed by the frustrating complexity of the human face, can learn, in a safe environment and through repeated experience, to distinguish the robot's facial movements by giving them the exact name found on the flashcard. However, the limited range of facial movements may not only produce a perturbing effect (at least in neurotypical adults). It may make emotion recognition difficult. In one study dedicated to Zeno, Schadenberg et al. (2018) affirm:

The emotions that can be modelled with Zeno's expressive face can be difficult to recognize, even by typically developing individuals. This can be partly attributed to the limited degrees of freedom of Zeno's expressive face, resulting in emotional facial expressions that may not be legible, but more importantly because facial expressions are inherently ambiguous when they are not embedded in a situational context. (Schadenberg et al., 2018).

The experiment described in the article consisted of adding to Zeno's facial expressions some vocal outbursts (such as "Wow!"; "No!"; "Ouch!"; etc.). As a result, the children's capacity to identify the emotions correctly increased significantly. One challenge that people with ASD face is integrating different sources of sensory

information (Kaliouby et al. 2006). In order to be helpful, the therapeutic setting needs to avoid on the one hand stimulation overload, and on the other hand extreme simplification. Sartorato et al. (2017) refer to several scientific articles documenting altered sensory perception in individuals with ASD (including overstimulation aversions). Problems seem to arise, more specifically, at the level of integration of multiple sensory modalities, while the processing of “simple” stimuli can be heightened (for example, the perception of pure tones). From a neurological point of view, multisensory stimuli are processed by allowing a “temporal binding window” (Sartorato et al. 2017, 3) so that stimuli that are not perfectly synchronous can still be ascribed to the same source. Apparently, in children with ASD the temporal binding window is larger compared to neurotypical children. This may explain the difficulty in accurately filtering stimuli from the environment, which might lead to sensory overload and anxiety, especially in crowded and noisy places (social gatherings, restaurants, theatres, etc.).

Affective cues in everyday life complement each other (e.g., vocal outbursts may accompany facial movements and changes in posture). Therapeutic machines cannot help in identifying emotions if they exhibit just one sensory cue (e.g., frowning), but they would be confusing if they perfectly matched the complexity of an emotional exchange in real life (e.g., frowning, shaking head, joking, saying something that is contradicted by a facial cue, using metaphors, moving eyes and hands, etc.).

The idea that a robot may teach emotion recognition raises two problems. The first concerns the cognitive style typical of children on the autism spectrum; the second concerns more generally the understanding of emotions in social relationships.

Regarding the first point, research has shown that a cognitive style described as “weak coherence account” is often encountered on the autism spectrum (Happé and Frith 2006). A person who can very accurately distinguish the parts of a whole may have difficulties when trying to identify the whole. For example, a situation may not be judged to be the same as a previous one if even the smallest detail changed, just as an object is not identified as that same object if a part of it is missing.² Different objects of the same type are only categorised as such with difficulty, because the perception of difference overrides that of similarity. For this particular cognitive style, each different example of a certain type (table) must be explicitly identified as such (this is a table, and also this, and this, and this, etc.), so as to generate a kind of large repository of images all referable to the same concept (I refer here to visual images, but the same can apply to sounds, tastes, etc.). Similarly, a typical situation can become a source of great stress if it changes with respect to one of its aspects (if you often go out to dine at a specific place, choosing a new restaurant may not be seen at all as fulfilling the promise “tonight we will dine at the restaurant”).

Now, if a significant problem is recognising a whole when even minute parts have changed, it can be assumed that a child with this cognitive style will not be helped to recognise an emotion already seen when it reappears in the complexity of a living face, because in the movements of the new face many differences from those of the robot will be perceptible. The difficulties in this regard are numerous: a person’s face changes very quickly when expressing emotions. Moreover, it certainly does not express only so-called basic emotions: it will be rare to find a person who expresses, in the same way as the robot, only contentment, only sadness, etc. Finally,

2 With respect to the integrity of objects, one may notice that it is not uncommon even among neurotypical individuals to be bothered when a favourite object undergoes minor accidents. For example, if one’s favourite mug, still perfectly capable of holding liquids, falls to the ground and is a little bit chipped, some may react as if they could no longer recognise it as the object they were fond of using for tea.

people can express different emotions with the same expressions. One may certainly learn that when the robot raises its eyebrows in a certain way it expresses surprise, but someone who raises his or her eyebrows in daily life may communicate different things, just as someone who laughs may express joy, contempt, irony, surprise, nervousness, frustration, serenity.

Then there is a further problem related to the theory of emotions presupposed in the construction and therapeutic use of these robots. The majority of our affective responses do not happen in a social vacuum. While it is true that anger means different things depending on context, eliminating context also implies eliminating a fundamental element of understanding. Someone who is angry feels offended, but understanding someone means understanding *in what sense* they may feel offended. It involves sensing whether the other person feels belittled because their needs are not taken seriously, whether they think that values with which they identify are treated with contempt, whether their gender identity is not recognised, whether they find a certain behaviour unacceptable because they feel that their authority is threatened, etc. One certainly also makes contact with others by noticing the micro-movements of their faces, but one does not understand what they feel unless one encounters their world. Without a sensitivity exercised to understand the other's values, the imagination cannot explore the reasons that might be involved when someone displays expressions attributable to anger. To say, therefore, that the emotions manifested by the robot help in recognising similar emotions when they occur in reality seems to me to overstate the usefulness of the exercise.

I am not denying that interaction with the robot may bring benefits. Indeed, research shows that Kaspar can help with imitation, body awareness, and appropriate physical interaction (Costa et al. 2015), joint attention, and turn taking in a game (Huijnen 2016). Such benefits, however, could arguably derive from aspects other than the cognitive aid offered by recognising similar expressions in different faces. For example, it is possible to imagine that the anxiety of social interaction might be mitigated by the repetitiveness and reliability of the robot. If among the things that generate frustration and a sense of failure in the social world stands out its mutability and unpredictability, knowing that there is at least one playmate who manifests recognisable emotions and who does not get angry if these are not immediately understood or mirrored can provide peace of mind, and perhaps even help develop confidence and self-esteem. With stress removed, it may become less difficult to concentrate on certain tasks. Understood in this way, the robot functions as a transitional object (cf. Winnicott 1971): it does not make intrusive demands – for example, it does not demand to be looked in the eyes – and it offers stability, reliability, solidity, predictability. Furthermore, the proven benefits often presuppose a triadic social interaction. Costa et al. (2015, 269–70) describe the room arrangement with child and therapist sitting at the desk, both looking at Kaspar, which is placed on the desk at the apex of an ideal triangle:

The arrangement of the actors involved in the session (robot, child and experimenter) had taken into consideration a cooperative position. In this arrangement of the room, two people work together on the same task, which provides an opportunity for eye contact and mirroring. The experimenter is able to move without the child feeling as if his territory has been invaded. Most importantly, this arrangement in a triangle allows the experimenter to encourage the child to engage in the interaction, without threatening his space and forcing eye contact.

Turn taking, joint attention, etc. are facilitated by the robot in a therapeutic setting, but humans set up, initiate, and maintain the interaction. It would be interesting to find out if the mediation provided by robots such as Kaspar and Zeno helps not just the children but also the therapists to regulate the anxiety involved in the interaction.

4. Dancing Robots, Busy Robots

In early 2021, a New Year's greeting video produced by Boston Dynamics appeared on several websites. The video, which was covered by major news outlets, depicts robots dancing to the sound of a 1962 song by the Contours called "Do You Love Me?" (2020).

The song's narrative is about a man rejected by his lover, apparently because he was clumsy and not fun to be around. He comes back, having become a skilled dancer, and provocatively asks, "Do you love me now that I can dance?"

Meanwhile, in the video, a lone humanoid robot begins to move a few timid steps to the rhythm of the music, and then gradually picks up the pace as robots of different kinds and shapes enter the scene one after another. In the end they all dance together. Their movements follow the music smoothly in a well-coordinated choreography. Arguably, the rejection evoked by the first words of the song is used in the video to hint at the sense of unease and fear inspired by the most primitive, awkward, and dangerous robots, which were the causes of even fatal workplace accidents. Bekey (2012, 20) cites a number of accidents caused by robotic arms used in industry, starting in 1979. The general public may have justified concerns regarding dangers posed by robots moving freely in an environment shared with humans. In response, the video seeks to inspire confidence by presenting robots that exhibit agility and coordination in a relatively small space.

The scene constitutes a perfect match to declarations one finds on the Boston Dynamics website. When asked if the company's robots use artificial intelligence, the answer takes as an example Spot, a steel quadruped robot, quite heavy (it weighs about 60 pounds), that is equipped with five cameras: two in front (the eyes), two on the sides, and one in the back. It resembles a dog and is a big hit for its mobility and ability to collect data in extreme physical situations:

Out-of-the-box, Spot has an inherent sense of balance and perception that enables it to walk steadily on a wide variety of terrains. This form of AI that we call "athletic intelligence" allows Spot to walk, climb stairs, avoid obstacles, traverse difficult terrain, and autonomously follow preset routes with little or no input from users. (2023)

The expression "athletic intelligence" is philosophically significant, since it calls into play the claim that human thinking is not confined to intracranial activity (with the body conceived merely as an external tool that may or may not be employed). Rather, cognitive activity is situated, i.e., dependent on or even co-constituted by corporeality and the environment with which it interacts.³

Not only does the video reflect the image of athletic intelligence that is also made explicit elsewhere on the site, but it illuminates other significant aspects. Atlas, the humanoid robot, can dance far better than most of us, but, more importantly, it can do so with other robots (including Spot) by synchronising with their movements. As the title of the song suggests, robots not only ask for our love, but show that they deserve it, since they can build, together, a pluralistic, cheerful, open society capable of valuing different forms of situated corporeality and affectivity.

³ There is still no agreement on the different options regarding situated intelligence (co-constitution, co-dependence, and embeddedness). For an overview of the debate concerning embodied intelligence see Shapiro and Spaulding 2021. For a discussion of situated conceptions of intelligence in robot construction, see Duffy et al. 1999.

Two important points emerge from the video: on the one hand, artificial intelligence has reached unparalleled levels of development by working on movement, balance, and responsiveness to the environment. On the other, companies such as Boston Dynamics are developing service robots that are meant to appear sociable. In this regard, it is interesting to note that in May 2022 the company released a new video entitled “No Time to Dance” (2022). The primary purpose was to advertise new capabilities of Spot, which this time was pictured working in an industrial plant. Its cameras, as seen in the video, have been upgraded, and now use colour images instead of black and white as before. Spot is equipped with sensors and gyroscopes that allow it to take readings from gauges and send all the information it gathers back to the computer. It can be guided with a tablet or programmed to perform a mission on its own. Significant upgrades include a new tablet, more sensitive and rugged than the previous one, and better batteries, with faster charging times. Payloads may add additional sensors and cameras, as well as the possibility of 5G connectivity.

The video is interesting beyond the information it gives about Spot’s new capabilities, since it refers back to the previous narrative and slightly modifies it. Spot has so much work to do that it cannot waste its time dancing – or so we learn from two employees who guide it remotely with their tablets, while a technician, in a different room, analyses the data collected by the robot. The two employees are amused by a man they observe from above, a worker who (like many of us, one may assume) was mesmerised by the previous Boston Dynamic video and follows the robot around the industrial facilities trying to persuade it to dance with him. Spot is plainly incapable of interacting with the man, who refuses to accept that what he has before him is neither a real dog nor a social robot: he pats it on the back, does a few dance steps hoping it will imitate him, even offers it some bolts to eat (what food can a robot dog like if not something made of steel?). Eventually Spot starts dancing with him, but this time we are made aware of what is really going on. Without the worker realising it, the robot is remotely guided by the two employees, who exchange amused and patronising remarks about the worker, while he is left under the illusion of dancing with his new friend. Viewers who may have formed opinions similar to those ascribed to the naive worker are informed by this new video that Spot’s autonomy is in fact very limited: it can be programmed to repeatedly follow a certain path or be guided by people using a tablet. It is not a machine meant for social interactions. Furthermore, it is plainly incapable of doing anything on its own initiative.

Before asking ourselves why the company might have chosen to send this new message to the public, let us remind ourselves of the difference between the capacity to socially interact with human beings and the capacity to engage with other robots. A “society of robots” was, after all, evoked in the “Do You Love Me?” video.⁴

The idea of a society of robots is discussed by Bicchi and Tamburrini (2015, 239–40), who interpreted the phenomenon in light of a sociology of subculture(s). They proposed two examples: the first concerns a group of robots employed in museums as night and day guards, the second is about a group of robot butlers working in a shopping mall. The robots must be able to communicate with each other while at the same time avoiding collisions with objects, moving carts, and the crowds of visitors. It is sufficient for such robots to behave as social groups, without being socially responsive towards human visitors. They do not communicate other than within their group; they have their own language and their own codes of behaviour.

4 On the difference between *societal robotics* and *social robotics*, Duffy et al. (1999) maintain that “the former represents the integration of robotic entities into the human environment or society, while the latter deals specifically with the social empowerment of robots permitting opportunistic goal solution with fellow agents”; Fong et al. (2003) describe the evolution from societies of robots inspired by insect societies (self-organised, homogeneous groups formed by anonymous individuals) to societies formed by heterogeneous groups (robots and humans) in which individuals can recognise each other, communicate, socially interact, and learn from each other. This model requires the construction of “individual social” robots (see Duffy 2004). For further reflections on social robots, see Hegel et al. 2009; Naneva et al. 2020.

From a certain point of view, the image of a subculture may cause apprehension: a subculture may seem a short step away from becoming a counterculture that does not share the dominant values, opposes them, and may seek to subvert them. On the other hand, precisely the unease that this image evokes reflects the current affective ambivalence toward embodied artificial intelligence. If it is becoming more and more usual to see individual robots used in industry or in the service of individuals (think of the millions of Roombas vacuuming up dust and moving around in our homes), the idea of a robot society may be frightening. One can imagine, then, that a first step toward generating acceptance in the public would be for manufacturers to envision groups of robots that refrain from explicit interactions with humans and make the communicative relationships they have with each other as inconspicuous as possible.

Such robots must have an appealing appearance, fluid movements, the ability to collect and transport data and objects, and, above all, the ability to avoid humans. They need to be able to calculate their trajectories in order to change direction whenever there is a risk of getting too close or even colliding with humans or cutting them off. Significantly, one of the applications Spot works with causes it to abort its mission as soon as it gets (too) close to a person. Still at the beginning of the path to their integration, robots must move in the background of human life. They need to look like dumb, docile, harmless servants.

5. Friendly Dogs, Scary Dogs

The interactive capacities of social robots seem to be on the whole still underdeveloped (Henschel et al. 2021), while in some cases service robots have very strong potential for eliciting anthropomorphic projections. Spot is a service robot endowed with characteristics that, as we have seen so far, tend to elicit positive reactions. It does not produce the uncanny valley effect because it does not aim to be mistaken for a living dog: it is clunky, rather noisy, it is made of steel, and does not hide its mechanical structure. For all these reasons it does not give rise to dissonant effects at the level of perception.

However, we noted that the “No Time to Dance” video issued by Boston Dynamics addresses precisely some unrealistic expectations Spot may provoke. This machine, we learn from the video, is incapable of autonomous movement. In contrast with a real dog, it cannot act without human control. It can either be moved from a distance by a tablet, or it can be programmed to follow a certain path and to perform a definite set of actions. Furthermore, Spot is clearly not a machine designed to interact with humans: it is not a social robot.

Why did Boston Dynamics feel the need to lower the public’s expectations concerning this incredibly efficient machine?

The reason may lie in the polemics that followed the employment of Spot by a few police departments in the United States. In particular, when the New York police department bought a few exemplars (renamed Digidog) and started using them in areas of the city that usually fall under the radar of public funding, the backlash produced by their employment ultimately led to their dismissal. Public outrage began when, in February 2021, the NYC police employed Digidog during an operation in the Bronx, as reported by the New York Times (Cramer and Hauser 2021). Apparently, two men armed with guns were holding two other men hostage in an apartment. When policemen arrived at the scene, they employed Digidog to explore the building. The inhabitants reacted with outrage and fear when they saw the mechanical dog let loose in their neighbourhood.

Political reactions immediately followed. Representative Ocasio-Cortez, for example, called into question the factoring around \$70,000) were given precedence over meeting much more urgent needs in the Bronx. On Twitter, she said:

Please ask yourself: When was the last time you saw next-generation, world-class technology for education, health care, housing, etc. consistently prioritized for underserved communities like this? (2021)

In a New York Times article published on April 14, 2021, journalist Mihir Zaveri reported on a new incident involving the use of Digidog in a public housing building in Manhattan. Apparently, the robotic dog did not play any active role in the apprehension of a man who, armed with a gun, was keeping a woman and a baby hostage in one of the apartments. However, people who saw the robot in the lobby reacted with fear and outrage at the use of public money for this kind of device (Zaveri 2021).

The political debate that ensued focused on the lack of transparency concerning the police department's surveillance tools, the likely breach of privacy, the possibility that tools of this sort could reinforce police bias against underprivileged groups, and the likelihood that such machines could be weaponised and developed to initiate autonomous action.

One of the most frequent associations evoked by commentators was with the robotic dogs appearing in a widely popular 2017 episode of Black Mirror, "Metalhead." The episode portrayed a dystopian society in which robotic dogs lacking any capacity for empathy were chasing and killing humans who were trying to recover a teddy bear for a child who desperately missed it. The resemblance between Digidog and the robot dogs appearing in "Metalhead" is unmistakable. Indeed, the series creator, Charlie Brooker, affirmed in an interview (Hibberd 2017) that he took inspiration from a Boston Dynamics video (posted on YouTube in 2015) portraying a robot dog being kicked.⁵ While the video was meant to showcase Spot's capacity to maintain balance in extreme situations, several viewers posted comments showing empathic distress at the machine being abused by humans. "Metalhead" takes the opposite affective perspective: it shows how threatening a machine like that can be.

In response to the polemics following the employment of Digidog by the New York Police Department, representatives of Boston Dynamics clarified that their devices are mostly used in industrial facilities and in commercial places where they can replace humans in dangerous settings. In February 2021, however, an art collective mounted a paintball gun on a Boston Dynamics dog, allowed participants to take control of it remotely, and let it shoot up a gallery (the event was significantly named "Spot's Rampage"). In response, Boston Dynamics circulated a statement on Twitter condemning the portrayal of its products in ways that promote "violence harm and intimidation" and reminding the public that all buyers of the company's products must comply with Terms of Agreement that prevent their robots from being used "to harm or intimidate people or animals." Violations of the Terms of Sale would void the warranty and make it impossible to update, repair, or replace the products which had been misused.⁶

5 This video (accessible at <https://www.youtube.com/watch?v=aR5Z6AoMh6U>), titled "Watch robot dog 'Spot' run, walk...and get kicked" (runtime 1:11), was posted by On Demand News on 11 February, 2015.

6 Boston Dynamics (@BostonDynamics), Twitter, February 20, 2021, <https://twitter.com/BostonDynamics/status/1362921918781943816>.

In light of events that led the New York Police Department to terminate its contract with Boston Dynamics in April 2021, we may understand why the company's most recent videos aim to lower people's expectations about Spot, underlie its usefulness mostly in industrial facilities, and stress that it lacks the capacity for autonomous movement.

The company's terms of use follow ethical principles posted on the website (n.d.). One of the principles states the following:

We will not authorize nor partner with those who wish to use our robots as weapons or autonomous targeting systems. If our products are being used for harm, we will take appropriate measures to mitigate that misuse.

Since science does not happen in a vacuum, however, other companies started building quadruped robots with Spot's capabilities. Some did indeed weaponise their products. A company called Ghost Robotics (based in Philadelphia) presented a robot dog with a gun mounted on it at the Association of the United States Army's 2021 annual conference in Washington DC. Similar weaponised dogs were produced in China and in Russia. In an open letter to the robotics industry, Boston Dynamics pledges neither to produce nor to support weaponised mobile robots, and issues a call "on every organization, developer, researcher, and user in the robotics community to make similar pledges not to build, authorize, support, or enable the attachment of weaponry to such robots." One of the reasons adduced in the letter is that weaponised robots hinder public trust in robotic machines and obscure the fact that they can bring great benefits to society if used properly. I have no reason to question the sincerity of Boston Dynamics' pledge. However, I doubt that it will seriously influence the process of weaponisation of robot dogs.

6. Conclusions

Anthropomorphic projections towards embodied intelligence can be sympathetic, but they can also be very hostile. Industries like Boston Dynamics focus on the technological capabilities of their machines and issue products that are both carefully designed and skilfully promoted. However, something clearly went wrong in the case of Spot.

We saw three problems: the first arose from the associations that, via "Metalhead," made Spot look scary. The second was due to the lack of a fruitful exchange between the police, the company, and political and social representatives concerning public policies. Dogs can be good to friends and nasty to enemies, but of course they can be trained to identify as enemies anybody one may choose as a target of aggression. This much was known since Plato's *Republic*, in which the warrior class was associated with dogs, both in its positive qualities and in its more dangerous side (education of the warrior class was meant to limit the danger of an armed class that could attack its own citizens instead of defending them; see Fussi 2022). The third problem concerns the possibility that machines like Spot may be weaponised.

While the uncanny valley effect plays a role at the level of individual perception, the second and third problems are clearly situational and contextual. A service dog like Spot can elicit sympathy in an industrial setting, fear when it is enrolled by the police, outrage when it is used as a model for new instruments of destruction and social control.

The first conclusion we can draw is that industries developing embodied intelligence should consider carefully the symbolic connotations that may be associated with the designs chosen for their robots. The second is that the introduction of embodied intelligence in human environments should involve not just scientists, legal experts, and ethicists during the production process, but also social scientists, political representatives, as well as psychologists and philosophers of emotions.

The third conclusion is that it is necessary to think long and hard before arriving at the production process, about the possibility of future uses and misuses of technologies. It cannot be beyond imagination that a robot dog, guided from a distance, endowed with cameras and refined sensors, capable of carrying weight and moving in all kinds of terrain and temperatures, may become an instrument of death. Endowed with software for facial recognition and powerful cameras, a machine like that can be used by organisations both interested in surveillance and indifferent to privacy laws or civil rights. The context in which scientists and companies make decisions cannot be confined to their labs, to their good intentions, or to their own wish to instil public confidence in the machines they produce here and now.

At first sight, the examples we discussed in this paper may appear very different. Spot is a service robot, while Kaspar and Zeno are therapeutic robots. Spot does not have a social interface, while Kaspar and Zeno do. However, Spot is less likely to elicit the uncanny valley effect than the two humanoid robots that are used with children diagnosed with ASD. From one point of view, Spot is also much more likely to invite social projections (as we have seen both by analysing the videos provided by Boston Dynamics and by following the events that led to its dismissal from the New York Police Department).

An overly simplified interpretation of what it means to understand emotions can lead us to overstate the role played by therapeutic robots like Kaspar and Zeno. Their main function may be found, not in enhancing the capacity to recognise occurring emotions, but in facilitating affective interactions with others in a triadic social setting. As I suggested, in such settings it is likely that not only the children involved, but also their therapists, may regulate the anxiety involved. If this intuition proved accurate, we could say that therapeutic robots are affective scaffoldings and function as mediators between different cognitive and affective styles.

Would it be as advantageous for a child who has difficulties with social interaction to be able to spend a lot of time alone with Kaspar? Could this not reinforce isolation, repetitiveness, aversion to the complexity of social involvement with other human beings? This is another problem that needs urgently addressing, given the announcement recently posted by the University of Hertfordshire:

Following positive results from field trials in schools and homes – and ahead of an upcoming NHS trial – the Kaspar research team is looking to redesign and reengineer its Kaspar prototype into an advanced, robust and commercially viable robot that can be made available to any child around the world that needs it. (n.d.)

Ultimately, I agree with Šabanović (2010) when she points out that it is too often the case that scientists consider how societies can adjust to technological improvements instead of considering the mutual shaping between technology and society. Producers of robots should involve citizens and their political representatives in the early stages of design, rather than expect them to passively adapt to their products once they are on the market. Cultural and political connotations should be taken into account, and the different contexts in which robots will be used ought to play a significant role in a bottom-up approach to robot production. Šabanović

(2010) mainly addresses the issue of social robots. However, as I hope becomes clear from my examples, the affective responses elicited by sophisticated service robots like Spot suggest that this sort of strategy ought to be extended to embodied intelligence in general, especially when public funding and public policies are at stake.

Emotions such as fear, hope, and anxiety play a role in the relatively small social settings we examined when we considered Kaspar and Zeno, as well as in the differentiated public settings we evoked in the case of Spot (industrial facilities, police forces, city neighbourhoods, military conventions, art galleries). Such emotions concern the future. Inspired by them, we urgently need rational, imaginative, collaborative thinking on future ecological settings inhabited by humans and robots.

References

- Barrett, L. F., B. Mesquita, and M. Gendron. 2011. "Context in Emotion Perception." *Current Directions in Psychological Science* 20 (5): 286–90.
- Bekey, G. A. 2012. "Current Trends in Robotics: Technology and Ethics." In *Robot Ethics: The Ethical and Social Implications of Robotics*, edited by P. Lin, K. Abney, and G. A. Bekey, 17–34. MIT Press.
- Bicchi, A. and G. Tamburrini. 2015. "Social Robotics and Societies of Robots." *The Information Society: An International Journal* 31 (3): 237–43.
- Boston Dynamics. n.d. "About." Accessed May 20, 2023, <https://www.bostondynamics.com/about#Q3>.
- . n.d. "Boston Dynamics Ethical Principles." Accessed May 20, 2023. <https://www.bostondynamics.com/ethics>.
- . 2020. "Do You Love Me?" 29 December, 2020. Video, 2:54. <https://www.youtube.com/watch?v=fn3KWM1kuAw>.
- . 2022. "No Time to Dance." May 3, 2022. Video, 2:57. <https://youtu.be/6U-bI3On1Ww>.
- Bovens, L. 1999. "The Value of Hope." *Philosophy and Phenomenological Research* 59, 667–81.
- Brazeal, C. 2003. "Emotion and Sociable Humanoid Robots." *International Journal of Human Computer Interaction* 59: 119–55.
- Buccino, G., F. Lui, N. Canessa, I. Patteri, G. Lagravinese, F. Benuzzi, C. A. Porro, and G. Rizzolatti. 2004. "Neural Circuits Involved in the Recognition of Actions Performed by Non-Cospecifics: An FMRI Study." *Journal of Cognitive Neuroscience* 16 (1): 114–126.
- Colombetti, G. and J. Krueger. 2015. "Scaffoldings of the Affective Mind." *Philosophical Psychology* 28 (8): 1157–76.
- Colombetti, G. and T. Roberts. 2014. "Extending the Extended Mind: The Case for Extended Affectivity." *Philosophical Studies* 172: 1243–63.
- Costa, S., H. Lehmann, K. Dautenhahn, B. Robins, and F. Soares. 2015. "Using a Humanoid Robot to Elicit Body Awareness and Appropriate Physical Interaction in Children with Autism." *International Journal of Social Robotics* 7: 265–78.
- Cramer, M. and C. Hauser. 2021. "Digidog, a Robotic Dog Used by the Police, Stirs Privacy Concerns." *New York Times*, February 27, 2021. <https://www.nytimes.com/2021/02/27/nyregion/nypd-robot-dog.html>.
- Damiano, L. and P. Dumouchel. 2018. "Anthropomorphism in Human-Robot Co-Evolution." *Frontiers in Psychology* 9: 468.
- Duffy, B. 2004. "Robots Social Embodiment in Autonomous Mobile Robotics." *International Journal of Advanced Robotic Systems* 1 (3).
- Duffy, B. R., C. Rooney, G. M. P. O'Hare, and R. O'Donoghue. 1999. "What is a Social Robot?" Talk presented at the 10th Irish Conference on Artificial Intelligence & Cognitive Science, University College Cork, Ireland, 1–3 September, 1999. Text accessed May, 2, 2023. <http://hdl.handle.net/10197/4412>.
- Ekman, P. 1992. "An Argument for Basic Emotions." *Cognition and Emotion* 6: 169–200.
- Ekman, P. and R. J. Davidson, eds. 1994. *The Nature of Emotion: Fundamental Questions*. Oxford University Press.
- Epley, N., A. Waytz, and J. T. Cacioppo. 2007. "On Seeing Human: A Three-Factor Theory of Anthropomorphism." *Psychological Review* 114 (4): 864–86.
- Fong, T., I. Nourbakhsh, and K. Dautenhahn. 2003. "A Survey of Socially Interactive Robots: Concepts, Design, and Applications." *Robotics and Autonomous Systems* 42: 143–66.
- Fussi, A. 2022. "Anger, Hatred, Prejudice. An Aristotelian Perspective." In *Empathy, Intersubjectivity, and the Social World: The Continued Relevance of Phenomenology. Essays in Honour of Dermot Moran*, edited by A. Bortolan and E. Magri, 389–412. De Gruyter.
- Griffiths, P. E. and A. Scarantino. 2009. "Emotions in the wild. The situated perspective on emotion." In *The Cambridge Handbook of Situated Cognition*, edited by P. Robbins and A. Murat, 437–453. Cambridge University Press.
- Happé, F. and U. Frith. 2006. "The Weak Coherence Account: Detail-Focused Cognitive Style in Autism Spectrum Disorders." *Journal of Autism and Developmental Disorders* 36: 5–25.
- Hegel, F., C. Muhl, B. Wrede, M. Hielscher-Fastabend, and G. Sagerer. 2009. "Understanding Social Robots." In 2009 Second International Conferences on Advances in Computer-Human Interactions, edited by S. Dascalu and I. Pupyrev, 169–74.

- Henschel, A., G. Laban, and E. S. Cross. 2021. "What Makes a Robot Social? A Review of Social Robots from Science Fiction to a Home or Hospital Near You." *Current Robotics Report* 2: 9–19.
- Hess, U. and S. Harel. 2016. "The Impact of Context on the Perception of Emotions." In *The Expression of Emotion: Philosophical, Psychological and Legal Perspectives*, edited by C. Abell and J. Smith, 199–218. Cambridge University Press.
- Hibberd, J. 2017. "Black Mirror Creator Explains that 'Metalhead' Robot Nightmare." *Entertainment Weekly*, December 29, 2017. <https://ew.com/tv/2017/12/29/black-mirror-metalhead-interview/>.
- Huijnen, C. A. G. J., M. A. S. Lexis, and L. P. de Witte. 2016. "Matching Robot KASPAR to Autism Spectrum Disorder (ASD) Therapy and Educational Goals." *International Journal of Social Robotics* 8: 445–55.
- ITV News. 2017. "Zeno: The Robot Teaching Autistic Children how to Communicate." 16 February, 2017. News report, 2:23. <https://www.youtube.com/watch?v=8-q9tsZE0Js>.
- Kaliouby, R., R. Picard, and S. Baron-Cohen. 2006. "Affective Computing and Autism." *Annals of the New York Academy of Sciences* 1093: 228–48.
- Krause-Jensen, K. and R. Rodogno. 2023. "Shame on Wrong Planet: A Study of Shame among People on the Autism Spectrum." In *The Moral Psychology of Shame*, edited by A. Fussi and R. Rodogno, 221–242. Rowman & Littlefield.
- Krueger, J. 2014. "Affordances and the Musically Extended Mind." *Frontiers in Theoretical and Philosophical Psychology*, 4 (1003): 1–13.
- Massantini, L. 2020. "Affective Scaffolds of Nostalgia." *Philosophical Inquiries* 8 (2): 47–70.
- Meirav, A. 2009. "The Nature of Hope." *Ratio* 22 (2), 216–33.
- Miceli, M. and C. Castelfranchi. 2010. "Hope: The Power of Wish and Possibility." *Theory and Psychology* 20 (2): 251–76.
- Mori, M. 1970. "The Uncanny Valley." *Energy* 7: 33–5.
- Naneva, S., M. Sarda Gou, and T. L. Webb. 2020. "A Systematic Review of Attitudes, Anxiety, Acceptance, and Trust towards Social Robots." *International Journal of Social Robotics* 12: 1179–201.
- Ocasio-Cortez, A. (@AOC). 2021. "Please ask yourself: when was the last time you saw next-generation, world class technology for education, healthcare, housing, etc consistently prioritized for underserved communities like this?" Twitter, February 25, 2021. <https://twitter.com/AOC/status/1365023067769098245>.
- Russell, J. A. 1980. "A Circumplex Model of Affect." *Journal of Personality and Social Psychology* 39 (6): 1161–78.
- Russell, J. A. and L. F. Barrett. 1999. "Core Affect, Prototypical Emotional Episodes and Other Things Called Emotions: Dissecting the Elephant." *Journal of Personality and Social Psychology* 76 (5): 805–19.
- Šabanović, S. 2010. "Robots in Society, Society in Robots." *International Journal of Social Robotics* 2: 439–50.
- Sartorato, F., L. Przybyłowski, and D. K. Sarko. 2017. "Improving Therapeutic Outcomes in Autism Spectrum Disorders: Enhancing Social Communication and Sensory Processing Through the Use of Interactive Robots." *Journal of Psychiatric Research* 90: 1–11.
- Schadenberg, B. R., D. K. J. Heylen, and V. Evers. 2018. "Affect Bursts to Constrain the Meaning of the Facial Expressions of the Humanoid Robot Zeno." In *Joint Proceedings of the Workshop on Social Interaction and Multimodal Expression for Socially Intelligent Robots and the Workshop on the Barriers of Social Robotics take-up by Society*, edited by C. Tsiourti, S. Hanke, and L. Santos, 30–9. CEUR.
- Slaby, J. 2016. "Mind Invasion: Situated Affectivity and the Corporate Life Hack." *Frontiers in Psychology* 7.
- Stephan, A., S. Walter, and W. Wilutzky. 2014. "Emotions beyond Brain and Body." *Philosophical Psychology* 27 (1): 65–81.
- University of Herfordshire. n.d. "A Kaspar in every School." University of Hertfordshire. Accessed April 10, 2022. <https://www.herts.ac.uk/kaspar/a-kaspar-in-every-school>.
- Urquiza-Hass, E. G. and K. Kortschal. 2015. "The Mind behind Anthropomorphic Thinking." *Animal Behaviour* 109: 167–76.
- Walker, M. 2006. *Moral Repair. Reconstructing Moral Relations after Wrongdoing*. Cambridge University Press.
- Winnicott, D. W. 1971. *Playing and Reality*. Routledge.
- Zaveri, M. 2021. "N.Y.P.D.'s Robot Dog Returns to Work, Touching Off a Backlash." *New York Times*, April 14, 2021. <https://www.nytimes.com/2021/04/14/nyregion/robot-dog-nypd.html>.