Science and Engineering Ethics (2020) 26:2169–2188
https://doi.org/10.1007/s11948-019-00172-6

ORIGINAL RESEARCH/SCHOLARSHIP

Can a Robot Be a Good Colleague?

Sven Nyholm[1] · **Jilles Smids**[2]

Received: 19 January 2019 / Accepted: 5 December 2019 / Published online: 23 December 2019

Abstract

This paper discusses the robotization of the workplace, and particularly the question of whether robots can be good colleagues. This might appear to be a strange question at first glance, but it is worth asking for two reasons. Firstly, some people already treat robots they work alongside as if the robots are valuable colleagues. It is worth reflecting on whether such people (e.g. soldiers giving "fallen" military robots military funerals and medals of honor) are making a mistake. Secondly, having good colleagues is widely regarded as a key aspect of what can make work meaningful. In discussing whether robots can be good colleagues, the paper compares that question to the more widely discussed questions of whether robots can be our friends or romantic partners. The paper argues that the ideal of being a good colleague has many different parts, and that on a behavioral level, robots can live up to many of the criteria typically associated with being a good colleague. Moreover, the paper also argues that in comparison with the more demanding ideals of being a good friend or a good romantic partner, it is comparatively easier for a robot to live up to the ideal of being a good colleague. The reason for this is that the "inner lives" of our friends and lovers are more important to us than the inner lives of our colleagues.

Keywords Robots · Colleagues · Meaningful work · Human–robot interaction · Friendship and love

Introduction

When Boomer "died" in the battlefield in Iraq, the US soldiers in his team gave him an improvised military funeral.[1] They also gave him two medals of honor: a Purple Heart and a Bronze Star. These soldiers regarded Boomer as a highly valued

[1] Different kinds of more or less advanced bomb disposal robots have been used for the last 40 years. For a brief history and account of what they do, see Allison (2016). See also Garreau (2007).

✉ Sven Nyholm
s.r.nyholm@uu.nl

[1] Utrecht University, Utrecht, The Netherlands

[2] Eindhoven University of Technology, Eindhoven, The Netherlands

team-member and good colleague. Unlike the rest of the team, however, Boomer was not a human being. Boomer was a robot, whose job was to seek out and disarm bombs. Prior to his destruction, Boomer had saved many lives. Not only was he a life-saver; Boomer's coworkers also thought he had "develop[ed] a personality" of his own (Garber 2013; Carpenter 2016). So, it is perhaps no wonder that Boomer was given these honors when he was destroyed or that he was regarded as a good colleague worthy of an honorable funeral.[2]

This is a real-life story, not a philosophical thought experiment. But this real-life story does raise philosophical questions. For example, can a robot be a good colleague? In other words, could a robot live up to the ideal of being a good colleague in the way that a fellow human is able to? This is the question we will be discussing in what follows. As more and more robots are entering the workplace across different domains, this is a question we need to take seriously. Whether we are concerned with military robots like Boomer, care robots in a hospital, logistics robots in a warehouse setting, or any other robot in the workplace, the future of work satisfaction and meaningful human work partly depend on whether robots can be good colleagues (Smids et al. 2019). One of the most important factors that determine work satisfaction and whether people find their work meaningful is whether they are members of a team with good colleagues (Lysova et al. 2018; Madden and Bailey 2016; Martela and Riekki 2018; Ward and King 2017). Accordingly, as more and more robots are introduced into the workplace, this prompts the question of whether robots can be good colleagues that can make work meaningful.

Our aim here is not to give a categorical "yes" or "no" answer that applies to all cases in the same way. Our aim is rather to consider how to philosophically approach this question of whether a robot can be a good colleague and to take initial steps towards answering it. In discussing this topic, we will compare the question of whether a robot can be a good colleague to the more widely discussed questions of (i) whether a robot can be a friend and (ii) whether a robot can be a romantic companion. We come down in favor of the following thesis: whether or not a robot can be as a good of a colleague as a human can be, it is comparatively speaking "easier" for a robot to be a good colleague than it is for a robot to be a friend or a loving romantic partner. The capacities a robot needs in order to live up to the criteria it appears plausible to associate with being a good colleague are easier to realize in a robot than are the criteria commonly associated with being a true friend or a loving romantic partner. Moreover, some advanced robots already have capabilities, and are likely to soon start acquiring more capabilities, that make them fit to be regarded as good colleagues by many of the key standards we typically apply to good human colleagues. For example, robots can be designed to communicate with human beings and reliably help to bring about good work-related outcomes. Those are just two quick examples of capabilities relevant to being viewed as a good colleague.

We start below by first saying a little more about what we mean by our main question and why this is an important question (Sect. 2). We then briefly look at how

[2] We say "he" here, we should also note, because Boomer's human collaborators viewed Boomer in a gendered way, whereby Boomer was thought of as a "he".

philosophers of technology have recently discussed the questions of whether robots can be our friends and whether there can be mutual love between humans and robots (Sect. 3). This will provide us with the general approach we use to investigate our question of whether robots can be good colleagues. Our next step will be to set out criteria for being a good colleague, which seem to us to apply to most lines of work, and which we think will be commonly recognized as plausible criteria. We start by asking whether a robot can live up to these criteria on a performative or behavioral level and argue that this appears possible (Sect. 4). Returning thereafter to the comparison with friendship and love, we end by discussing whether what goes on "on the inside" matters as much in the case of good colleagues as it does in the cases of good friends or romantic partners. We suggest that it matters less in the case of colleagues. This is one of the reasons why it is easier for robots to live up to the ideal of being a good colleague than it is for them to live to the ideals of friendship and love (Sect. 5). Thus, in addition to providing us with our general method, the existing discussions about whether robots can be our friends or lovers also enable us to make a comparison between robotic colleagues on the one hand, and robotic friends and romantic partners, on the other hand.

Our Main Question and Why to Take it Seriously

Notably, many of us spend a lot of time with our colleagues at work. Some people even spend more time with their work colleagues than with their friends and loved ones. Accordingly, having good colleagues is a crucial aspect of having good and valuable relationships with those around us. Not surprisingly, having good colleagues and being part of a team one values are commonly cited as factors making an important difference to whether one's work is satisfying and meaningful (Lysova et al. 2018).[3]

Many commentators discussing the increased robotization of the workplace are interested in whether robots pose a threat to meaningful work. In particular, this often takes the form of a worry that robots will replace workers with the result of widespread so-called technological unemployment (Ford 2015). Work is often thought to be a meaningful aspect of life. So if robots "take over" our jobs, they can thereby pose a threat to one key aspect of a meaningful life. According to one line of argument, this might mean that many people will need to find other sources of meaning in life than work (Danaher 2017).

We are here interested in a scenario that we find to be more realistic in the short- or mid-term perspective: namely, not that robots "take over" all work, but rather that humans will increasingly work alongside robots. The robots that enter the workplace may replace human workers, or may be added to the workforce, but the result will be the same: part of what was previously workers' collaboration with other fellow

[3] More generally, the fabric of society is crucially dependent on good work communities. For example, work-related illnesses put a strain on society, in addition to being burdensome for the working people themselves.

humans will instead be with robotic co-workers (cf. Decker et al. 2017; Savela et al. 2018). For example, this is already happening in logistics warehouses (Smids et al. 2019). What we are concerned with here is whether working alongside robots will make work less meaningful by removing or diminishing the valuable component of work that consists in having good colleagues around us. One interesting way that this "threat" to meaningful work can potentially be neutralized is if robots—at least some robots—could themselves be good and valuable colleagues in the workplace. That idea motivates our main question, namely, whether robots could ever live up to the ideal of being a good colleague.

By a "robot" we here understand an embodied automated system, with some degree of functional autonomy and some degree of artificial intelligence (Royakkers and van Est 2015; Nyholm 2020; cf. Gunkel 2018). The paradigmatic robot in the workplace we are primarily concerned with would be a so-called "social robot" (Darling 2016). This refers to a robot that interacts, to some extent, with the humans around it and is able to adjust its functioning and respond in intelligent ways to the humans that the robot interacts with.[4] Could social robots be equipped with capacities that could enable them to live up to the ideal of being a good colleague?

We are most concerned with social robots that interact to a significant extent with the humans in the workplace in order to help getting the work done, so that it can potentially make sense to think of the robot as a member of the team. Robots that interact with humans in the workplace in very limited ways fall outside of the scope of what we are focusing on here. Such robots are closer to mere tools, which are used by human workers, and are unlikely to give rise to the experience of there being any team-work or cooperation together with the robot. Examples of the types of robots we are interested in include care robots, autonomous vehicles (e.g. an autonomous long haul truck), robots in a logistics setting, police robots, and military robots (Royakkers and Van Est 2015). These are all robots in the workplace that may have extensive interaction with the humans in the workplace. Notably, some of the most interesting on-going empirical research and development of such robots is focused on issues such as mutual understanding of behavioral intentions, coordination, joint action, and the like, all of which make most sense as part of a phenomenology of working together with robots that goes beyond the experience of using robots merely as tools (cf. Iqbal and Riek 2017.). Accordingly, in the relevant types of empirical research such robots are more and more often being referred to as "co-workers", "team-members", and "colleagues" (see e.g. Gombolay et al. 2015; Ljungblad et al. 2012; Sauppé and Mutlu 2015; You and Robert Jr 2018).

[4] Darling also understands a social robot as a robot specifically designed to interact with human beings on a "social level", such that it can potentially be a "companion" to the human beings it interacts with (Darling 2016, 215–216). The idea of a robot specifically designed to be some sort of companion is a little stronger than what we have in mind when we are asking whether a robot can be a good colleague—at least on an understanding of "companion" that suggests some sort of friendship. But we follow Darling in understanding a social robot as being one that can interact and communicate with human beings on a "social level", to some extent. Those are the kinds of robots, we think, that stand the best chance of being perceived as colleagues by humans who might work with them.

Of course, it is hard to predict what further types of future robots will fall under our paradigmatic specification of a social robot in the workplace. For example, perhaps robots will eventually play a larger role in educational settings, such as in the university classroom (Belpaeme et al. 2018). Our prediction is that robots in the workplace will become increasingly interactive in their capacities and be used for more and more purposes. In our assessment, the question about whether a robot can be a good colleague is already a question with real-world relevance. There are already robots in many workplaces. However, this question will become even more relevant and increasingly pressing as we move into the future.

It is worth noting here as well that some philosophers argue that there exists a general duty of justice for society as a whole to make sure that people have access to meaningful work (e.g. Gheaus and Herzog 2016; Roessler 2012; Schwartz 1982). We are not taking any definite stance on that issue here. Instead, we simply wish to note that if the writers who argue for that position are right, this gives us even more reason to investigate whether and how robots can be good colleagues, given that having good colleagues is typically seen as a key component of meaningful work.

The Philosophy of Friendship and Mutual Love Between Humans and Robots

As far as we know, there are no philosophy publications about whether robots can be good colleagues yet. In organizational psychology and human–robot interaction research, there are publications about whether robots can function and be perceived as members of teams (e.g. Groom and Nass 2007) or as co-workers (Sauppe and Mutlu 2015). As far as we know, however, none of those empirical papers address head-on the normative question of whether robots can satisfy the ideal of being good colleagues. We are here specifically interested in the philosophical question of whether robots could realize the ideal of being a good colleague.

Existing philosophical literature about robotics and work covers other topics, such as whether robots will take over people's jobs or whether robots will open up responsibility gaps (e.g. Danaher 2017; Sparrow 2007; Nyholm 2018).[5] However, there are publications about whether robots can be our friends (or "companions") (e.g. Elder 2017; Coeckelberg 2010; Danaher 2019; Nyholm 2020). There are also publications about whether there can be love between humans and robots (e.g. Levy 2008; Nyholm and Frank 2017; Hauskeller 2017). To have a model to base our current argument on, we will take inspiration from those other discussions. We will also later compare the different cases, to see whether it appears easier for robots to be our colleagues than it is for them to be our friends or romantic partners.

[5] The idea of a responsibility gap refers to a situation in which some morally significant outcome has been brought about for which it appears appropriate to find somebody to hold responsible, but where it is unclear whether there is any particular person or persons who could justifiably be held responsible. For example, if a robot with a significant form of functional autonomy harms a human being, it seems right that somebody should be held responsible for this. But it will not always be clear who exactly it is appropriate to hold responsible. See, e.g., Sparrow (2007) and Nyholm (2018).

Friendship: Most general philosophical writings about friendship refer back to Aristotle's influential treatment of the subject in his *Nicomachean Ethics* (Aristotle 1999). Discussions about whether robots can be our friends do this as well. In particular, it is common to refer to Aristotle's division among three types of friendships:

> Utility friendships: the parties are useful to each other.
>
> Pleasure friendships: the parties get pleasure from their interactions.
> Virtue friendships: the relationship is based on shared values, common interests, and the perceived virtue(s) in one another, and the parties develop and go through life together.

It is commonly thought that the third type of friendship is the most valuable one, and also the most interesting one from a philosophical perspective. When this type of friendship is discussed in relation to robots, it is tempting to rather quickly conclude that robots cannot be our friends in this sense. After all, the third type of friendship requires that a friend can have values and virtues, and thus be a moral agent. And robots—at least current robots—seemingly lack the required capacities (Himma 2009). That is, robots lack the agency required to be a moral agent and to be able to have virtues (Elder 2017; Nyholm 2020, chapter seven).

John Danaher notes that this idea of Aristotelian virtue friendship is also typically interpreted as requiring the presence of a rich inner life (Danaher 2018, 2019). One needs to be able, for example, to have the sort of inner mental life that makes it possible to either be honest or dishonest in our dealings with those around us. That is to say, our outward behavior as human beings is typically thought to either correspond with our actual inner feelings, thoughts, and intentions, or to not do so. And a true friend is—among other thing—somebody whose inner life matches their displayed behavior towards us. Somebody might pretend to be our friend, and we may or may not be aware of that fact about them (Nyholm 2020, chapter five). The true friend has the feelings, thoughts, and intentions that their behavior indicates that they have. They are honest in their dealings with us. As Danaher notes, if we think that robots lack an inner life of this kind, this provides another reason why a robot cannot be a friend (Danaher 2018, 2019).

Of course, when it comes to both the type of agency required and the inner life just discussed, we should not be too quick to conclude that a robot could never come to have such capacities. Suppose we take a functionalist understanding of agential and mental capacities, on which we understand such capacities in terms of the performance of various different functions (Levin 2018). There is then no principled reason why robots in the future could not eventually come to acquire the capacities needed to be moral agents or to have the appropriate form of inner life (Frank and Nyholm 2017). This is not necessarily to suggest that future robots will become "super-intelligent" in the senses that authors like Nick Bostrom or John Harris are concerned about in their work (Bostrom 2014; Harris 2019). Rather, it is to say that future robots might achieve at least, functionally speaking,

minimal levels of some of the key capacities often associated with being a moral agent or with having an inner life. So perhaps future robots could be our friends in this Aristotelian virtue-based sense (Danaher 2019). Danaher himself—interestingly—expresses some skepticism about the focus on a need for a rich inner life. He argues that, in general, perhaps we should adopt a more performative and behaviorist understanding of friendship. On such a view, if an entity—such as a robot—can perform the outward actions of a friend and its behavior conforms to that of a friend, this may be enough for it to make sense for us to regard the entity as a potential friend (Danaher 2018, 2019).

We will not engage further with that interesting suggestion here, other than to note that this would be rather revisionary. Our intention here is to stay fairly close to common sense, whenever possible. We grant that new technologies—whether it is robots, artificial intelligence, social media, or whatever—sometimes change our ideals and values (Kudina and Verbeek 2019). But our focus is on widely shared current ideals, neither revisionist nor potential future alternative ideals.

Love: There are some humans who want to have love relationships with robots (Beck 2013). And there are some companies that try to develop robots that people will fall in love with.[6] But as one of us has argued elsewhere, there is reason to be skeptical about the possibility of mutual love between humans and robots, at least with respect to any robots we will be able to create in the near-term future (Nyholm and Frank 2017). The reasons are similar to those offered above for why Aristotelian virtue friendships with robots are hard to realize.

Consider some of the most commonly shared ideas about what constitutes a valuable romantic relationship involving mutual love. One first idea is that of being a "good match" (Plato 1997). When we seek a romantic partner, we hope to find somebody who is a good match for us. And we also hope to be a good fit for them. This would require of a robot that it had values, preferences, likes and dislikes. It would require that whereas some humans could be a good match for a robot with those values etc., other humans would not be. The ideal of being a good match (a) involves an idea of mutuality and (b) also involves the idea that not everyone is a good match for everyone else. A robot cannot simply be custom-made for you. You also need to be a good fit for it. This is a challenge.

Consider next another commonly discussed idea: namely, that human lovers value each other in their distinctive particularity. Love requires, according to this ideal, that lovers value their unique relationship and each other as the individual persons they are (Kolodny 2003). So a robot would need to be able to value a human person. Specifically, it would need to be able to value that person for who they are, in his or her distinctive particularity. Can robots value anything? If we understand valuing as a complex phenomenon involving a robust concern, certain patterns of thought, emotional vulnerabilities, and motivations—and if we think of it as involving seeing certain considerations as reasons for action—this can be seen as involving agential and mental capacities robots cannot (yet!) possess (Nyholm and Frank 2017).

[6] The website truecompanion.com, for example, claims to sell a sex robot called "Roxxxy", which can become a "true companion".

Consider lastly the idea of commitment, which is another widely valued aspect of the ideal of a loving relationship. A lover is not obsessed with, nor a slave to their beloved. Rather, they freely commit to their beloved. They have the capacity to choose to do otherwise (Kühler 2014). Of course, the lover may not desire to choose otherwise. But we typically think that they nevertheless possess that capacity. This means that we attribute free will to our human lovers. This holds for our friends as well. With both our lovers and friends, we appreciate the commitments they make to us, partly because we think of them as being able to choose otherwise. We view their commitment to us as a choice in our favor. Accordingly, if a robot is to realize this aspect of love and friendship, the robot would need to be equipped with a capacity for choice, or some sort of free will. This is another capacity one can imagine that future robots will perhaps eventually come to have or approximate. However, we typically do not think that any current robots have a capacity for choice we would think of as a form of free will. The idea of commitment is therefore another obstacle in the way of valuable human–robot relationships.

What these arguments related to friendship and love have in common is this: first, certain key aspects of what we value in valuing human friendship and love are identified. Second, certain underlying capacities are identified that robots would need to have in order to be able to realize the given valuable aspects of these relationships. It is then argued that robots—at least current robots!—lack these abilities. Hence it is concluded that robots are unable to be our friends or romantic partners in the senses we value the most. This gives us a template for how to approach the question of whether robots could be good colleagues. It also gives us a point of comparison.

Criteria for Being a Good Colleague (and Whether a Robot Can Live Up to Them)

Unlike the cases of friendship and love, there is not a rich and long tradition within philosophy discussing what it is to be a good colleague that we can draw on. This is not a topic that philosophers have typically discussed in general terms, like they have discussed in general terms what it is to be a good friend or loving partner. Again, we are not interested in the empirical question of whether humans sometimes treat some robots like co-workers or team-members, but rather the normative question of whether robots could live up to whatever criteria it is plausible to associate with the ideal of being a good colleague. And that is not a topic in relation to which there are philosophical resources to draw on comparable to those available in the cases of friendship and love. Accordingly, in coming up with criteria to work with, we need a different method than what one can use in an article about friendship or love with robots.

What we do here is to suggest what appears to us to be a number of key aspects of what makes somebody a good human colleague that we expect will resonate with

common sense. Our hope is that readers will agree with these criteria.[7] At any rate, we will assume for the sake of our argument that these are some key features of a good colleague. We do not seek to compile an exhaustive list of criteria. Instead, we leave it open that there may also be other criteria that we are overlooking here.[8]

Our suggestions for what makes somebody a good colleague are:

(a) Working well together to achieve desired outcomes and goals specifically related to the work in question (e.g. caring for patients properly, creating good products and useful services, etc.)
(b) Being able to engage in pleasant, informal conversations, to help to keep work pleasant
(c) Not harassing or bullying colleagues, but instead treating each other respectfully
(d) Providing mutual help and support as needed
(e) Learning and developing together
(f) "Consistency", meaning that the colleagues are not constantly replaced with new colleagues, but that they work together over extended periods of time
(g) Being reliable and trustworthy[9]
(h) Being sensitive to how one's colleagues are doing (e.g. whether they are tired or energized, well or unwell, happy or sad, and so on), and adjusting one's work-related interaction accordingly
(i) Sharing work-related values and being motivated to honor these values in the joint work
(j) Potentially also doing some socializing (in a collegial, respectful, and friendly way), which may or may not lead to a "deeper" friendship or closer relationship. Such socializing needs to be sensitive to the wishes and interests of all involved parties; it also needs to take power-balances into account in a responsible way

[7] Having come up with a draft list of criteria, we ran our initial list by three work psychologists at our university, to see whether these conditions fit with what is usually understood as good collegial relationships in workplace psychology. Subsequently, we presented a revised list at a philosophy conference, asking for feedback from the audience attending our presentation (which consisted of around 40–50 people). The audience at that particular conference—a large Dutch philosophy conference—found our list intuitively plausible, and did not suggest any further criteria.

[8] One last general remark: we intend this list to have wide application, across various different types of work. But we recognize that depending on what type of work is in question, different criteria may have different importance or priority in terms of what makes for a good colleague within the particular line of work in question. A more specialized discussion—e.g. of what makes somebody a good colleague in an intensive care unit or in a large restaurant kitchen—would make it appropriate to try to rank or assign weights to these different criteria. A more general discussion, such as our present discussion, appears best conducted without an attempt to assign specific weights or rankings to these different criteria.

[9] We understand being reliable and being trustworthy as two distinct, but to some extent related, criteria for being a good colleague. Being trustworthy is, for example, a more demanding criteria than being reliable is. For more on the issue of robots and trust, see footnote 16.

Two initial comments about these criteria and how they relate to the criteria for being friends or romantic partners[10]: firstly, we have tried to formulate criteria pertaining specifically to the relationships between colleagues. We do not wish to suggest that good colleagues could not also be friends or romantic partners. If the colleagues in question are also friends or romantic partners, our intention is for the criteria above to identify aspects that make the friends or romantic partners in question into good colleagues as well. Of course, some people work together with their friends or romantic partners. But what we are particularly interested in here are aspects of work interaction that make people fall under the general umbrella of being a good colleague (as opposed to what makes somebody a good friend or a good romantic partner).[11]

Secondly, unlike in the cases of friendship and love, we think of the interaction with colleagues as being more context-sensitive or domain-specific. Friends and romantic partners interact together across different contexts and within different domains of life. Colleagues, in contrast, primarily interact—and perform their duties as good colleagues—within circumscribed domains and within specific contexts: namely, the sphere of work. If somebody wants to interact a lot with you outside of work, this had better be somebody you would be willing to also be friends with (or potentially romantic partners with). Otherwise, it is inappropriate for that person, we take it, to want to interact with you a lot outside of the work context.

Now, what about robots and their ability to live up to the above-listed criteria for being a good colleague? In considering this question, we will primarily look at current robot capacities to assess whether some robots already possess capabilities that could enable them to function as good colleagues. However, we will also attempt to give a reasonable estimate of potential future robotic capacities, in order to assess whether current limitations to being a good colleague are merely technical

[10] Another comment one might make about these suggested criteria is that they bring up the question of whether robot colleagues should be given some form of moral and/or legal status. We will very briefly comment on the issue of whether robot colleagues should be treated with some degree of moral concern in our concluding remarks below. But since our main focus in this paper is on whether robots can live up to the ideal of being good colleagues, we will save a more thorough discussion of the moral and legal status—or potential lack thereof—of robot colleagues for another occasion. For related discussion, see Bryson (2018) and Gunkel (2018).

[11] When we say that we are especially interested in criteria associated with being a good colleague rather than criteria for being a good friend, we are referring specifically to what we called "virtue friendships" above. There may be significant overlap between what is involved in being a good colleague and what is involved in being a good utility friend: good colleagues are useful to each other, for example, just like good utility friends are useful to each other. At the same time, though, there are also differences between being good colleagues and being good utility friends. Collegial relationships, for example, are had within the context of workplaces, where the colleagues have contracts specifying what their work tasks are, which might include specifications concerning ways in which they need to work together with their colleagues. Utility friendships, as we understand them, are typically not be governed by any explicit contracts.

limitations, or rather in-principle limitations, which cannot be overcome even with the most intricate technology.[12]

The first thing that strikes one when one glances over this list of dispositions and behaviors is that it appears possible for a robot to either perform or approximate many of them. We will dig deeper into this in the next section. But let us here first do a superficial evaluation of a robot's potential to live up to these ideals. Our focus here is on what might be called the behavioral level (cf. Danaher 2019).

Robots can, for example, clearly perform well together with humans to achieve some of the outcomes and goals associated with different types of work. In that sense, they can be efficient colleagues. A good example is robotics used in hospital settings, which ranges from telemedicine robots, to transport robots, AI diagnosing systems, companion robots, and lifting robots (Kontzer 2016; Ljungblad et al. 2012; Su et al. 2014). Such hospital robots can be very helpful in achieving the desired outcomes and goals.

But what about something like engaging in pleasant conversation, a less goal-oriented aspect of being a good colleague?[13] Interestingly, some of the most well-known developers of so-called sex robots are trying to enable those robots to be able to talk with their users so that those robots can become companions, and not only be used for purely sexual purposes (Devlin 2018). Of course, some of the conversational possibilities that sex robot creators like Matt McMullen with his robot "Harmony" or Sergio Santos with his "Samantha" are exploring are a type of sexual banter not appropriate in a workplace setting. However, these inventors are also working specifically to enable their robots to talk about a wide variety of topics in a way that users can enjoy. The general message here is that developers of robot companions are working on enabling their robots to engage in pleasant conversation. If successful, their achievements could be exported into the robots we work together with, thus enabling our robotic colleagues to engage in pleasant conversation.[14]

Presumably, we can control robots in the workplace to make them abstain from any kinds of bullying or harassing behaviors. If these robots function like Microsoft's chatbot Tay—which had to be turned off because Tay became too outrageous[15]—robots in the workplace might become bullies or harassers (Gunkel 2018).

[12] We are inspired here by a distinction that John Danaher draws between technical and metaphysical obstacles to the prospects for robots to be able to be our friends. For Danaher's discussion of the distinction between technical and metaphysical possibilities as those relate to human–robot friendship, see Danaher (2019, p. 11).

[13] Some workers interviewed in an empirical study by Sauppé and Mutlu (2015) explicitly asked for the collaborative manufacturing robots they worked with to be equipped with the capability for small talk. In that way, their interaction with their robotic co-workers would be more like working with human colleagues.

[14] Human–robot conversation has been extensively studied for several decades. A recent review concludes that "we seem to be still far from our goal of fluid and natural verbal and non-verbal communication between humans and robots" (Mavridis 2015, p. 31). Nevertheless, according to the review, considerable progress is being made. And although there are some tough challenges, there appear to be no in-principle, or metaphysical, obstacles toward fluent human–robot conversation.

[15] Tay was trained on the inputs from human users. Some of the inputs from human users were racist or otherwise highly inappropriate in nature. The result was that Tay started generating morally inappropriate sentences, based on the human inputs in the training data. See Gunkel (2018).

However, it is quite feasible that robots in the workplace can be made to not engage in any behaviors commonly regarded as forms of bullying or harassment. So their human colleagues should be able to feel safe around these robots. The robots can be programmed to interact with people in ways that most people experience as respectful. Human–machine interaction researchers are studying how care robots should interact with people in order for people to feel comfortable and non-threatened. They can provide developers of other robots with insights about how robots should interact with humans in the workplace so that the people in question will feel comfortable around the robots (Torta et al. 2012).

What about mutual help and support? This appears to be another thing we can easily imagine in human–robot interaction. Robots can help and support humans in all sorts of ways. And, of course, humans can also do the same for the robots. If the robots are equipped with the right kind of machine learning capabilities, they can also interact with the humans in the workplace in ways that enable the robots and the humans to learn and develop together. The robots do then not need to be replaced all the time, but can become "consistent" long-term collaborators in the workplace. This can all also help them to become reliable and trustworthy, like we want good colleagues to be. Research shows that humans can in fact become willing to trust robots, at least to some extent. Here, robot performance appears to be the most important factor (Hancock et al. 2011; Robinette et al. 2017).[16]

Consider next the idea of being sensitive to how one's colleagues are doing, and adjusting one's interaction with them accordingly. A good human colleague, for example, will be less demanding around a colleague who appears tired or unwell than around a colleague who appears energized and ready for a challenge. Could a robot be similarly sensitive to how people are doing and adjust their interaction with their human colleagues accordingly? Researchers working on affective computing and emotion-recognition are working on enabling robots to become responsive to people's emotional states (Calvo et al. 2014; Cavallo et al. 2018). Drawing on such research, developers of robots for the workplace could also work to make their robots responsive to how people are doing and disposed to interact with the people in a way that takes this into account.

[16] Notably, Groom and Nass (2007) challenge the conceptualization of robots as full-fledged team-members, arguing that the humans do not sufficiently trust robots. They argue that robots do not have humanlike mental models and consequently cannot share in the team's mental model that enables the team to work well together. As a result, robots cannot engage in the relevant trust-building interactions in ways that can make human team members come to trust the robots. Especially in safety-critical situations, human will feel unable to rely on robots, and Groom and Nass seem to view this as an in-principle limitation of robots. In response, we would like to make a few brief remarks. In the first place, a good robotic colleague is not necessarily a full-fledged team member. For example, we can imagine that the manufacturing robots mentioned above, studied by Sauppé and Mutlu (2015), only interact with their direct operator and are a good colleague merely to them. Secondly, it could turn out to be difficult to design robots that will be sufficiently trusted in contexts where the life of human workers is at considerable risk. In that case, the application context of robotic colleagues would be somewhat restricted, but this would not settle this paper's question, since there would be many other contexts where robots potentially could be good colleagues. However, the most sensible approach seems to be to suspend judgment and see how trust in robots will develop in future work practices. Lots of research is being done on human-robot trust, for example on ways in which robots could repair human trust (Robinette et al. 2015) or even help teams to moderate interpersonal conflict (Jung et al. 2015). See also Coeckelbergh (2012) and Alaieri and Vellino (2016).

What about sharing work-related values and trying to respect these values in joint work? And what about potentially socializing to some extent also outside of work? Both of these aspects of being a good colleague can seem a little trickier than some of the other aspects considered above. We will here set aside the idea of possibly socializing outside of the work setting. We can certainly imagine that some people will become so attached to the robots they work with that they want to interact with these robots also outside of work. There are many real-world cases that one can point to of people getting very attached to robots (Carpenter 2016). But we think that this is a complex enough issue in the human case—just what is a good, appropriate, balanced amount of socializing with one's colleagues outside of work?—that we are unwilling to say anything general about it in relation to robots in the workplace.

However, when it comes to sharing work-related values and trying to work together in a way that honors those values, this does seem like something that we can imagine a robot being able to live up to, at least on a behavioral or dispositional level. Robots can be programmed to work toward certain goals and to follow rules and restrictions (Nyholm 2018). And those are both important aspects of what is involved in having and trying to live up to work-related values. So at least at first glance, it appears imaginable that this is another aspect of being a good colleague that robots could live up to or approximate.

At first glance, then, it appears that a robot could live up to—or at least approximate—a lot of the aspects we associate with being a good colleague. Should this lead us to conclude that robots can be good colleagues? Let us now dig a little deeper. We will consider what capacities help human beings to realize the above-listed types of criteria for being a good colleague. We will ask two questions: firstly, could a robot exercise similar or equivalent capacities to realize these aspects of being a good colleague? Secondly, are those underlying capacities in a good human colleague essential to what we value in valuing them as good colleagues? These questions will take us back to the comparison between robots as friends or romantic partners, on the one hand, and robots as good colleagues, on the other. It will take us back to this comparison since in both cases, the question arises of what underlying capacities—in particular mental capacities—are needed in order to qualify as a good friend, romantic partner, or colleague.

Underlying Capacities and the Comparison with Friendship and Love

When it comes to friendship and love, we understand the behavior of our friends and romantic partners as springing from various aspects of our friends' and romantic partners' inner lives. We attribute their actions and behaviors to their inner thoughts, emotions, feelings, attitudes, decisions, and the free will we view them as having. Not only that; we also attribute a value to the relation between our friends' and romantic partners' behavior and their underlying inner life (Pettit 2015; Nyholm 2020). This is not the only part of a friendship or loving human relationship that people value. But it is one of the things commonly valued as part of a good friendship or romantic relationship. When it comes to both friends and lovers, then, what goes on "on the inside" matters: it helps to distinguish true friends and lovers from

false friends and lovers, and it is part of what people value (Nyholm and Frank 2017). How is it in the case of good colleagues?

Firstly, in the case of human colleagues, it is clear that we understand our colleagues as displaying the kinds of behaviors and dispositions associated with being a good colleague because they have the inner life and capacities we associate with ordinary human agency. That is, we spontaneously interpret our fellow human beings as decision-making agents, and we constantly engage in so-called "mind-reading" (also known as "theory of mind") in our interaction with those around us (Heyes 2018; Marraffa 2019). In other words, we attribute mental states to those around us, and we interpret their interactions with us in terms of the mental states that we attribute to them. Most of the time, this mind-reading is spontaneous and effortless. Sometimes, we find those around us "hard to read", and we may need to think and reason about what might be going on in their minds. In either case, understanding other human beings' actions is nearly always—if not always—done in terms of mind-reading or the attribution of an inner life to them (Nyholm 2020, chapter six). So of course this also goes for the actions and behaviors of our colleagues, including those whom we regard as good colleagues.

This does not yet mean that we necessarily value what goes on "on the inside" of our colleagues in the same ways that we value what goes on "on the inside" of our friends and romantic partners. That we inevitably understand people's actions in terms of mental states or an inner life does not by itself mean that in all types of interactions with other people, we always attribute the same importance to people's inner life or mental states. In different kinds of cases or situations, people's inner states matter in different ways, to different degrees.

To see this, think about the difference between civil law and criminal law and the importance, or lack thereof, of the mindset of people who are being sued or prosecuted (Nyholm 2020, chapter two). In civil law, if somebody is being sued and expected to be held liable for some damage they have caused, the mindset of the person is not what is most relevant. What matters is instead whether it can be settled that the person performed the act in question that caused the damage. In criminal law, in contrast, whether something counts as (say) murder or manslaughter or whatever it might be will depend on whether the accused is judged to have had a criminal mindset or not. Was the act premeditated? Was the person sane at the time of the incident? Was the person sober and in full control, etc.? In this case, the internal mindset matters a lot.

What that helps to illustrate is that sometimes the mindset we associate with somebody's actions or behaviors matters a lot, whereas sometimes it is judged to be less important in the evaluation of their actions or behaviors. So this raises the question of how it is in the case of the ideal of being a good colleague. Is it more like the cases of friendship or love, or more like the case of evaluations made within the context of civil law? This will bear importantly on whether and to what degree a robot could realize the ideal of being a good colleague in a way approximating how human beings can realize this ideal.

Let us first note that if a robot is able to do a lot of the things we discussed in the foregoing section, this will depend on the presence within the robot of what might be termed a robotic form of "inner life". By this we do not mean that the robot will

have conscious subjective experiences. Nor do we mean that the robot will have a will that functions in the same way that a human being's free will can be understood as functioning. We simply mean that the robot needs to be equipped with internal hardware and software that can take in information and perform calculations and computations that help to select the robot's responses to the situation it is in. The robot will form a model of its surroundings and operate on that basis, meaning that there is a form of "representation" of the world around it that helps to regulate the robot's behaviors. In short, the robot's outward behavior will depend on its internal states, processing, and what goes on "inside" the robot. It is just that the robot's inner life is rather different from that of a human (Nyholm 2020, chapter six). At the very least in a metaphorical sense, then, we can attribute internal states to the robot without making a terrible mistake, so long as we do not think that these internal states are like a human being's inner life.[17]

The question arises, therefore, of whether the "inner life" of a robot is different enough from that of a human being that this undermines the robot's potential to live up to the ideal of being a good colleague. This depends on whether what we value in valuing people as good colleagues is intimately linked with our associating their outward behaviors with a distinctively human inner life. Does the good colleague's inner life matter as much as it does in the cases of friendship and love? And does it matter a lot that it is the particularly human kind of inner life that our colleagues have?

These questions are hard to answer. We are inclined to think that there is a difference between the cases of friends and romantic partners, on the one hand, and the case of good colleagues, on the other. The difference is that in the former case, the specifics of the inner lives of friends and lovers appear to matter more than the specifics of the inner lives of apparently good colleagues matter. The closer the relationship is with somebody, the more we typically care about what they actually think and feel. And since colleagues—even good colleagues—are typically less "close" to us than our friends and loved ones are, it matters less what exactly is going on in their inner lives as they are behaving in the ways we associate with their being good colleagues.[18]

To test this hypothesis, part of what we can do is to consult the list of criteria for being a good colleague that we sketched above and ask whether the items on the list appear to require the presence of an inner life. Certainly the last item on the list, which is about the potential for possibly becoming friends outside of work, can seem to clearly require an inner life if we understand friends as having a valued inner life. But this last item on the list is different from the other items, which are all more closely related to the things directly expected of a good colleague. And those

[17] For related discussion, see Richard Bright's interview with the philosopher Keith Frankish (Frankish 2018) "AI and Consciousness", *Interalia Magazine,* Issue 39, February 2018, available here: https://www.interaliamag.org/interviews/keith-frankish/ (Accessed on August 21, 2019).

[18] We want to emphasize that our claim here is not that people typically have no concern about the inner lives of their colleagues. Our claim is, instead, a comparative claim according to which the inner lives of those who are considered to be good colleagues typically matter less—perhaps even much less—to us than the inner lives of our friends or romantic partners matter to us.

other things primarily appear to require behaviors and reliable performances, rather than the presence of an active inner life. This makes it plausible to accept the suggestion that, in general, the inner lives of colleagues are less important than those of friends or romantic partners.

What about the issue of whether it matters a great deal to us that it is a distinctively *human* inner life that prompts the outward behavior that we associate with a good colleague? This too is a very tough question. This is a tough question, we take it, because common sense does not have a widely shared opinion on this issue. But we are here again inclined to suggest that this matters less in the case of the ideal of a good colleague than it does in the cases of a friend or a romantic partner. That is what we suggest. But we wish to make it clear that we base this suggestion primarily on what philosophers call "intuition", rather than any more elaborate argument we are able to articulate. We are putting this forward as a suggestion that can be criticized and as an assumption we base our argument on.

Let us now bring the robots back in and see what is implied by the just-made suggestions and assumptions. As we just noted, we take it that, firstly, the inner life of a good colleague matters less than the inner life of a close friend or romantic partner. We are also taking it that in the case of a good colleague, it matters less whether the inner life of the good colleague is of a distinctively human kind than it matters whether the inner life of a close friend or romantic partner is distinctively human. Moreover, we suggested above that robots are likely be able to perform many of the behaviors and have many of the dispositions it is compelling to associate with the ideal of being a good colleague. If these claims are accepted—or at least assumed for the sake of the argument—where does this lead to? What conclusion can we draw? The tentative conclusion we can draw from such an argument is that at the very least, realizing the ideal of being a good colleague is easier for a robot than it is for a robot to realize the ideal of being a friend or a good romantic partner.

Concluding Remarks

We have argued that robots can live up to—and will soon be even better at living up to—many of the key criteria it is plausible to associate with the ideal of being a good colleague. We have also just argued that there is a general reason why it is easier for a robot to be a good colleague than it is a for a robot to be a friend or a romantic partner. That reason is that in the cases of friends and romantic partners, we typically place greater importance on the inner lives of our friends and romantic partners than we do in the case of our colleagues. The ideal of being a good colleague is more focused on reliable behaviors and performances than it is on the thinking and motives behind those behaviors and performances.

It might also be added here that friendship and love require more flexibility than the ideal of being a good colleague does in terms of what different types of situations one needs to be able to behave in the right ways in. As noted above, friends and romantic partners interact with each other in many different types of situations and in different domains of life (cf. Danaher 2019). Colleagues typically confine their interaction to the context of work and the types of situations associated with the

work in question. This is another thing that makes it easier for robots to be good colleagues than to be good friends or lovers. The ideal of being a good colleague requires less domain-general flexibility of a robot than the ideals of friendship and love do.

Now, of course it is not going to be the case that all robots are going to be able, or equally able, to live up to the ideal of being a good colleague. Robots will differ with respect to how many aspects of this complex ideal that they can live up to. Robots will also differ from humans in the ways that they live up to the criteria associated with being a good colleague. So even if the criteria are the same on a general level, a good robotic colleague is going to be different than a good human colleague. The robot might also do better than humans with respect to some criteria associated with being a good colleague, while doing worse with respect to other criteria. For example, a robot might more reliably help to produce good results (at least in certain lines of work) than human colleagues, but be a worse conversation partner than a human.

We have focused in this article on the question of whether a robot can be a good colleague. The last thing we will note is that this question should be distinguished from the question of whether we should create robots that would live up to the ideal of being good colleagues or at least many aspects of this ideal. It might be argued that even if it is possible to create robots that can live up to many of the criteria associated with being a good colleagues, we should not do so. The reason for this might be that it is better to use our resources to create jobs for more humans. Another reason might be that robots with highly advanced capacities would be confusing in terms of whether such robots should be treated with moral consideration or be given some sort of rights. It might be thought that, as Joanna Bryson (2018) argues, it is best to avoid creating too advanced robots so that we avoid any troubling ambiguity with respect to the status of the robots. Is that a convincing argument? Should we avoid trying to create robots that can be good colleagues for us? Or is there perhaps a moral imperative to try to create robots that can be good colleagues, if people are going to be working alongside robots anyway?

Those are also important questions that also need to be discussed. We hope to do so at greater length elsewhere. Here, however, are some very brief remarks about this. It is worth noting that if we start conceptualizing robots as "colleagues", this might create some pressure to assign at least some minimal moral status to these robotic colleagues.[19] The reason for this is that we typically understand the relationship of being colleagues as involving having various duties to one another. For that reason, treating what are thought of as robotic colleagues with at least some small amount of moral consideration can be a way of reaffirming to our human colleagues that we think that colleagues have moral duties to treat each other well. In other words, treating what are conceptualized as robotic "colleagues" with some degree of

[19] David Gunkel argues that anytime that we apply human labels to robots (including labels like "slave" or "servant") this creates pressure to ask whether any rights—even minimal rights—associated with those labels in the human case would also need to be extended to the robots. See Gunkel's critical discussion of Bryson (2010) in Gunkel (2018).

ethical concern can be a way of showing respect for our human colleagues, whether or not we think that the robots have any independent moral status on their own (cf. Nyholm 2020, chapter eight).

The question of whether robots at all can be good colleagues in the first place needs to be discussed more thoroughly than we have been able to do here. Our aim here has primarily been to raise that question, to take some initial steps towards answering it, and to briefly compare this question with those of whether robots can be our friends or romantic partners. Another highly relevant question here is whether there is something problematically deceptive about trying to create a robot that is supposed to function like a good colleague. That question also deserves further discussion. There is more work to do here.[20]

References

Alaieri, F., & Vellino, A. (2016). Ethical decision making in robots: Autonomy, trust and responsibility. In Agah, et al. (Eds.), *Social robotics* (pp. 159–168). Berlin: Springer.

Allison, P. R. (2016). *What does a bomb disposal robot actually do*? BBC Future. https://www.bbc.com/future/article/20160714-what-does-a-bomb-disposal-robot-actually-do.

Aristotle. (1999). *Nicomachean ethics*. Indianapolis: Hackett.

Beck, J. (2013). Married to a doll: Why one man advocates synthetic love, *The Atlantic,*https://www.theatlantic.com/health/archive/2013/09/married-to-a-doll-why-one-man-advocates-synthetic-love/279361/

Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018). Social robots for education: A review. *Science Robotics, 3*(21), eaat5954.

Bryson, J. J. (2010). Robots should be slaves. In Y. Wilks (Ed.), *Close engagements with artificial companions* (pp. 63–74). London: John Benjamins.

Bryson, J. J. (2018). Patiency is not a virtue: The design of intelligent systems and systems of ethics. *Ethics and Information Technology, 20*(1), 15–26.

Calvo, R. A., D'Mello, S., Gratch, J., & Kappas, A. (2014). *Oxford handbook of affective computing*. Oxford: Oxford University Press.

Carpenter, J. (2016). *Culture and human–robot interactions in militarized spaces*. London: Routledge.

Cavallo, F., Semeraro, F., Fiorini, L., Magyar, G., Sinčák, P., & Dario, P. (2018). Emotion modelling for social robotics applications: A review. *Journal of Bionic Engineering, 15*(2), 185–203.

Coeckelberg, M. (2010). Artificial companions: Empathy and vulnerability mirroring in human–robot relations. *Studies in Ethics, Law, and Technology, 4*(3), 1–17.

Coeckelbergh, M. (2012). Can we trust robots? *Ethics & Information Technology, 14*(1), 53–60.

Danaher, J. (2017). Will life be worth living in a world without work? *Science and Engineering Ethics, 23*(1), 41–64.

Danaher, J. (2018). Embracing the robot. *Aeon* The MIT Press, https://aeon.co/essays/programmed-to-love-is-a-human-robot-relationship-wrong.

Danaher, J. (2019). The philosophical case for robot friendship. *Journal of Posthuman Studies, 3*(1), 5–24.

Darling, K. (2016). Extending legal protection to social robots: the effects of anthropomorphism, empathy, and violent behavior towards robotic objects. In M. Froomki, R. Calo, & I. Kerr (Eds.), *Robot law* (pp. 213–232). Cheltenham: Edward Elgar.

[20] Many thanks to Hannah Berkers, Pascale Le Blanc, Sonja Rispens, Jason Borenstein, the anonymous reviewers for this journal, and an audience at the sixth Annual OZSW Philosophy Conference in 2018 at Twente University for valuable feedback on this material. This work is part of the research program "Working with or Against the Machine? Optimizing Human–Robot Collaboration in Logistics Warehouses" with Project Number 10024747, which is (partly) financed by the Dutch Research Council (NWO).

Decker, M., Fischer, M., & Ott, I. (2017). Service robotics and human labor: A first technology assessment of substitution and cooperation. *Robotics and Autonomous Systems, 87*, 348–354.
Elder, A. (2017). *Friendships, robots, and social media*. London: Routledge.
Ford, M. (2015). *Rise of the robots: Technology and the threat of a jobless future*. New York: Basic Books.
Frank, L., & Nyholm, S. (2017). Robot sex and consent: Is consent to sex between a robot and a human conceivable, possible, and desirable? *Artificial Intelligence and Law, 25*(3), 305–323.
Frankish, K. (2018). *AI and consciousness* (R. Bright, Interviewer). Retrieved from https://www.interaliamag.org/interviews/keith-frankish/
Garber, M. (2013, September 20). Funerals for fallen robots. Retrieved 7 December 2018, from https://www.theatlantic.com/technology/archive/2013/09/funerals-for-fallen-robots/279861/
Garreau, J. (2007). Bots on the ground. *Washington Post*. Retrieved from https://www.washingtonpost.com/wp-dyn/content/article/2007/05/05/AR2007050501009.html.
Gheaus, A., & Herzog, L. (2016). The goods of work (other than money!). *Journal of Social Philosophy, 47*(1), 70–89.
Gombolay, M. C., Gutierrez, R. A., Clarke, S. G., Sturla, G. F., & Shah, J. A. (2015). Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams. *Autonomous Robots, 39*(3), 293–312.
Groom, V., & Nass, C. (2007). Can robots be teammates? Benchmarks in human–robot teams. *Interaction Studies, 8*(3), 483–500.
Gunkel, D. (2018). *Robot rights*. Cambridge, MA: The MIT Press.
Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human–robot interaction. *Human Factors, 53*(5), 517–527.
Harris, J. (2019). Reading the minds of those who never lived. Enhanced beings: The social and ethical challenges posed by super intelligent AI and reasonably intelligent humans. *Cambridge Quarterly of Healthcare Ethics, 8*(4), 585–591.
Hauskeller, M. (2017). Automatic sweethearts for transhumanists. In J. Danaher & N. McArthur (Eds.), *Robot sex: Social and ethical implications* (pp. 203–218). Cambridge, MA: The MIT Press.
Heyes, C. (2018). *Cognitive gadgets*. Cambridge, MA: Harvard University Press.
Himma, K. E. (2009). Artificial agency, consciousness, and the criteria for moral agency: What properties must an artificial agent have to be a moral agent? *Ethics and Information Technology, 11*(1), 19–29.
Iqbal, T., & Riek, L. D. (2017). Human–robot teaming: Approaches from joint action and dynamical systems. In A. Goswami & P. Vadakkepat (Eds.), *Humanoid robotics: A reference* (pp. 1–20). Berlin: Springer.
Jung, M. F., Martelaro, N., & Hinds, P. J. (2015). Using robots to moderate team conflict: The case of repairing violations. *Proceedings of the tenth annual ACM/IEEE international conference on human–robot interaction* (pp. 229–236). ACM, https://doi.org/10.1145/2696454.2696460.
Kolodny, N. (2003). Love as valuing a relationship. *Philosophical Review, 112*(2), 135–189.
Kontzer, T. (2016). Deep learning cuts error rate for breast cancer diagnosis | NVIDIA Blog. Retrieved 27 October 2018, from https://blogs.nvidia.com/blog/2016/09/19/deep-learning-breast-cancer-diagnosis/
Kudina, O., & Verbeek, P. P. (2019). Ethics from within: Google glass, the collingridge dilemma, and the mediated value of privacy. *Science, Technology, & Human Values, 44*(2), 291–314.
Kühler, M. (2014). Loving persons: activity and passivity in romantic love. In C. Maurer (Ed.), *Love and its objects* (pp. 41–55). London: Palgrave MacMillan.
Levin, J. (2018). Functionalism. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), https://plato.stanford.edu/entries/functionalism/
Levy, D. (2008). *Love and sex with robots*. London: Harper.
Ljungblad, S., Kotrbova, J., Jacobsson, M., Cramer, H., & Niechwiadowicz, K. (2012). Hospital robot at work: Something alien or an intelligent colleague? In *Proceedings of the ACM 2012 conference on computer supported cooperative work* (pp. 177–186). New York: ACM.
Lysova, E. I., Allan, B. A., Dik, B. J., Duffy, R. D., & Steger, M. F. (2018). Fostering meaningful work in organizations: A multi-level review and integration. *Journal of Vocational Behavior.*. https://doi.org/10.1016/j.jvb.2018.07.004.
Madden, C., & Bailey, A. (2016). What makes work meaningful—Or meaningless. *MIT Sloan Management Review, 57*(4), 52–61.

Marraffa, M. (2019). Theory of mind. *The Internet Encyclopedia of Philosophy* ISSN 2161-0002, https://www.iep.utm.edu/theomind/

Martela, F., & Riekki, T. J. J. (2018). Autonomy, competence, relatedness, and beneficence: A multicultural comparison of the four pathways to meaningful work. *Frontiers in Psychology, 9*, 1157. https://doi.org/10.3389/fpsyg.2018.01157.

Mavridis, N. (2015). A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems, 63*, 22–35. https://doi.org/10.1016/j.robot.2014.09.031.

Nyholm, S. (2018). Attributing agency to automated systems: reflections on human–robot collaborations and responsibility-loci. *Science and Engineering Ethics, 24*(4), 1201–1219.

Nyholm, S. (2020). *Humans and robots: Ethics, agency, and anthropomorphism.* London: Rowman & Littlefield International.

Nyholm, S., & Frank, L. (2017). From sex robots to love robots: Is mutual love with a robot possible? In J. Danaher & N. McArthur (Eds.), *Robot sex: Social and ethical implications* (pp. 219–244). Cambridge, MA: The MIT Press.

Pettit, P. (2015). *The robust demands of the good*. Oxford: Oxford University Press.

Plato. (1997). Symposium. In J. Cooper (Ed.), *Complete works* (pp. 457–505). Indianapolis: Hackett.

Robinette, P., Howard, A. M., & Wagner, A. R. (2015). Timing is key for robot trust repair. In A. Tapus, E. André, J.-C. Martin, F. Ferland, & M. Ammi (Eds.), *Social robotics* (pp. 574–583). New York: Springer.

Robinette, P., Howard, A. M., & Wagner, A. R. (2017). Effect of robot performance on human–robot trust in time-critical situations. *IEEE Transactions on Human-Machine Systems, 47*(4), 425–436.

Roessler, B. (2012). Meaningful work: Arguments from autonomy. *Journal of Political Philosophy, 20*(1), 71–93.

Royakkers, L., & Van Est, R. (2015). *Just ordinary robots: Automation from love to war*. London: CRC Press.

Sauppé, A., & Mutlu, B. (2015). The social impact of a robot co-worker in industrial settings. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (pp. 3613–3622). New York: ACM.

Savela, N., Turja, T., & Oksanen, A. (2018). Social acceptance of robots in different occupational fields: A systematic literature review. *International Journal of Social Robotics, 10*(4), 493–502.

Schwartz, A. (1982). Meaningful work. *Ethics, 92*(4), 634–646.

Smids, J., Nyholm, S., & Berkers, H. (2019). Robots in the workplace: A threat—Or opportunity for—Meaningful work? *Philosophy & Technology*. https://doi.org/10.1007/s13347-019-00377-4.

Sparrow, R. (2007). Killer robots. *Journal of Applied Philosophy, 24*(1), 62–77.

Su, N. M., Liu, L. S., & Lazar, A. (2014). Mundanely miraculous: The robot in healthcare. In *Proceedings of the 8th nordic conference on human-computer interaction: fun, fast, foundational* (pp. 391–400). New York: ACM.

Torta, E., Oberzaucher, J., Werner, F., Cuijpers, R. H., & Juola, J. F. (2012). Attitudes towards socially assistive robots in intelligent homes: Results from laboratory studies and field trials. *Journal of Human-Robot Interaction, 1*(2), 76–99.

Ward, S. J., & King, L. A. (2017). Work and the good life: How work contributes to meaning in life. *Research in Organizational Behavior, 37*, 59–82.

You, S., & Robert Jr., L. P. (2018). Human–robot similarity and willingness to work with a robotic coworker. In *Proceedings of the 2018 ACM/IEEE international conference on human–robot interaction* (pp. 251–260). New York: ACM.