AI & SOCIETY (2025) 40:1403–1415
https://doi.org/10.1007/s00146-024-01993-1

MAIN PAPER

The meaningfulness gap in AI ethics: a guide on how to think through a complex challenge

Markus Rüther[1,2]

Received: 21 March 2024 / Accepted: 3 June 2024 / Published online: 12 July 2024

Abstract
Technological outsourcing is increasingly prevalent, with AI systems taking over many tasks once performed by humans. This shift has led to various discussions within AI ethics. A question that was largely ignored until recently, but is now increasingly being discussed, concerns the meaningfulness of such a lifestyle. The literature largely features skeptical views, raising several challenges. Many of these challenges can be grouped under what I identify as the "meaningfulness gap". Although this gap is widely acknowledged, there is a notable absence of systematic exploration in the literature. This paper aims to fill this void by offering a detailed, step-by-step guide for systematically exploring the different instances of the meaningfulness gap and aids in navigating their complexities. More specifically, it proposes differentiating the gaps according to their realms and objects, normative nature, scope, and severity. To make these areas manageable, the paper takes several taxonomies and distinctions on board. Finally, the guide is summarized, and some skeptical replies are anticipated and countered by clarificatory remarks.

Keywords AI ethics · Meaningfulness · Meaning in life · Meaningfulness gap · Technological outsourcing

1 Introduction

It is not an overstatement to assert that artificial intelligence (AI) technologies have become a fundamental part of many individuals' daily lives, significantly enhancing or entirely substituting a wide range of tasks. This dependency on AI covers a vast array of decisions and activities, from routine daily chores to complex decisions that affect both professional and personal spheres. AI systems are now capable of assisting with diet management, executing daily grocery shopping, autonomously driving vehicles, monitoring health conditions, and scheduling medical appointments when necessary. In addition, the influence of AI extends into the professional domain: an increasing number of workers find their roles supported or even replaced by AI technologies, signifying a major shift in the nature of work and how it is performed. Furthermore, in the social and leisure contexts, AI-driven recommendations are becoming increasingly influential in shaping human interactions and leisure activities, for instance, by influencing partner selection and personal engagements through digital platforms and smartphone apps.

This scenario describes the current state of outsourcing to AI; and this is not merely a passing trend expected to fade away. On the contrary, it is reasonable to predict that such developments will not only continue but also intensify, having an even greater impact on our lives. This prospect forces us to examine the essence of living in an AI-driven world. Should we admire a life so intertwined with AI? Does it align with our deepest values? Or, in short, is such a life meaningful?

In the field of AI ethics, the issue of life's meaningfulness or meaning in life, as it usually has been called, has recently

✉ Markus Rüther
m.ruether@fz-juelich.de

[1] Permanent Researcher, INM-7, Group: Ethics of Artificial Intelligence, Research Center Jülich GmbH, Jülich, Germany

[2] Associate Professor ("Privatdozent") of Philosophy, Department of Philosophy, University of Bonn, Bonn, Germany

gained attention.[1] In this paper, I intend to contribute to this burgeoning discourse. More specifically, my focus will be on a certain challenge that is almost omnipresent in the discourse, but that has not been sufficiently addressed yet. I will call this challenge: the meaningfulness gap. This gap, very crudely put, opens up because technological outsourcing takes place: If certain tasks are not available to us due to the fact that AI systems have taken them over, and if those tasks have been identified as meaning-conferring, we suffer an axiological loss; we experience a gap between having the possibility of performing meaningful tasks and the reality that they are not available to us anymore.

Although the AI literature acknowledges the presence of the meaningfulness gap, there has been insufficient systematic exploration of its specifics. Not all meaningfulness gaps are the same. They differ in terms of the realms where they occur, their normative nature, their scope, and their severity. Or to put it in slightly different words: many in the field agree about the diagnosis that there is a challenge but did not make the effort to spell out the details. But "details" matter in this regard if one wants a justifiable assessment of the overall importance of a meaningfulness gap, avoiding both exaggerating and downplaying it. But how can we determine the details of a certain meaningfulness gap? The following paper is my answer to this question. It proposes a step-by-step guide which can be used to think through cases in which a meaningfulness gap might occur. Given that this enterprise is a fairly new endeavor the guide is not meant as the final word on this matter. It is not conclusive and fixed in its used parameters and some distinctions I make along the way might need further justification or even adjustments.[2] But even given these small restraints, I am inclined to think that it is still a promising starting point for future research to think one's way through the complexities of a certain meaningfulness gap one might be confronted with.

To develop the guide, the paper proceeds in four steps. The first step involves defining the key terms and concepts. The second step presents an overview of the meaningfulness gap, acknowledges its presence in the existing literature, and differentiates it from other gaps in AI ethics. The third step explores the complexities of various meaningfulness gaps, employing different parameters to determine their relevance and importance. In a fourth and concluding step, the paper will briefly summarize the developed stages of the guide, addresses challenges for my proposal, and explore avenues for future research.

[1] For more details on the meaning in life discourse in general and in AI ethics specifically see Sects. 2 and 3.

[2] I will address both challenges in the last section of the paper in more detail, see Sect. 5.

2 Some clarifications: AI systems, technological outsourcing, and meaningfulness

In this paper, I aim to explore a particular challenge, the meaningfulness gap, as it is created by AI systems through the outsourcing of tasks that we traditionally performed ourselves. To better understand the scope and limitations of this endeavor, it is necessary to clarify what is meant by the basic phenomena, particularly the terms "technological outsourcing" and "meaningfulness."

Concerning the former, I broadly define the phenomenon of technological outsourcing as a process of replacement, where AI systems partially or fully take over tasks that humans previously executed using their natural intelligence (see for similar definitions, Coeckelbergh, 2020; Müller, 2020; Gordon & Nyholm, 2021). The realms of outsourcing are diverse, and it is not feasible to address all of them. Therefore, I will primarily focus on those that have already created an impact in the literature, which, in turn, aids the purpose of this paper in better understanding the complexity of various meaningfulness gaps. Some of them have been already flagged in the introduction, namely the realms of everyday life projects, healthcare, the workplace, and relationships.

In addition, I intend to limit the specific object of outsourcing: I only focus on outsourcing processes directly related to the tasks of the individual whose life we are evaluating, such as their cognitive processes, actions, or outcomes. I do not delve into the outsourcing of external factors that might be indirectly relevant for those tasks. An example of this would be the outsourcing of personal relationships (e.g., family, friends, romantic partners or colleagues) to AI systems which do not replace certain tasks of an individual directly, but impacting them indirectly by replacing their enabling conditions. Focusing on these and other indirect outsourcing processes would open a new discourse with its own literature, potentially distracting from the main purpose of this paper.[3] Therefore, I omit this discourse but acknowledge that the developed framework for the meaningfulness gap could be applied in future work.

Regarding the dimension of meaningfulness, ethical explorations have become increasingly common. In normative ethics, influential authors like Susan Wolf (2010) and Thaddeus Metz (2013) have made important contributions, significantly shaping the field now known as the discourse on meaning in life.[4] In fact, ethical theorizing on

[3] For a good starting point on the issue, see the overview of the discourse on personal relationships in AI Ethics in Nyholm, 2023, ch. 9.

[4] Numerous articles and handbooks provide guidance through the field of meaning in life. The most comprehensive and widely recognized among these is likely the Stanford Encyclopedia of Philosophy article by Metz, 2023, which offers an extensive overview. In addition, the handbook edited by Landau, 2022 is dedicated to this topic, featuring contributions from many prominent experts in the field.

meaningfulness did not stop there but has even expanded further. In more recent developments, ethical considerations of meaningfulness are also present in various debates within applied ethics, such as medical ethics (Metz, 2021), animal ethics (Purves & Delon, 2018; Monsó et al., 2018), and climate ethics (Kauppinen, 2014; Campbell & Nyholm, 2015; Scheffler, 2018). This trend has also spilled over into AI ethics, including discussions on meaningful self-improvement, work, and relationships (see for an overview, Nyholm & Rüther, 2023).

In this paper, I will build on the literature by adopting at least two conditions of meaningfulness commonly used as a starting point. The first is a semantic claim, viewing meaningfulness as inherently good and applicable to varying degrees in a person's life or its aspects. It is something "that would be intrinsically good to have in one's life…" (Metz, 2001a, p. 140) and which can be more or less a part of it. This does not imply that we cannot understand meaningfulness differently, for instance, as a subject of empirical inquiry. However, these and other understandings, although might be justifiable and worthy to use as a starting point, are spared out this paper.[5] Here, I am focusing on meaningfulness in ethical contexts, where we can interpret it as a gradable intrinsic value that is an important factor in people's lives.

The second is an axiological claim, suggesting that meaningfulness is distinct from standard axiological dimensions, like self-interest or morality, and therefore is at least partly autonomous. Following this claim, meaningfulness becomes, as Susan Wolf expresses in a subtitle of a paper, "a third dimension of the good life" (Wolf, 2016). That meaningfulness is a dimension in its own right, though, does not entail that the other axiological dimensions are irrelevant to it.[6] It might turn out that, depending on the theory, the standard dimensions do have an impact, being either necessary, sufficient, or contributive. But, given its partly autonomous character, there always must remain a residue that cannot be captured by them.

Note that this minimal definition, including the semantic and axiological claims, does not encompass thoughts on the content of meaningfulness. Indeed, the nature of meaningfulness is a hotly debated topic in the discourse on meaning in life, reflecting broader and classic philosophical debates (e.g., subjectivism vs. objectivism, consequentialism vs. deontology, and monism vs. pluralism) (cf. Metz, 2023). However, in this paper, I will not commit to any of these views[7]. This is because presuming what constitutes an appropriate theory of meaningfulness would narrow the analysis of the meaningfulness gap. Therefore, I aim to be as neutral as possible, recognizing that complete neutrality is unattainable, to fully capture the challenge's complexity and also develop thoughts that may appeal to those with varying views on meaningfulness.[8]

3 The meaningfulness gap - the core structure

As elaborated in the previous section, the discourse on meaningfulness is already a part of AI ethics, specifically the debate on technological outsourcing. It is noteworthy, however, that the contributions are not symmetrical in terms of their evaluation of whether AI technologies pose a threat or present an opportunity for meaningfulness. Remarkably, only a very few are explicitly exploring the positive effects. Systematically, they can be categorized based on their views on the specific types of benefits that AI technologies might bring us. A first group of contributors focus on technological outsourcing as meaning enablers. For example, some highlight the negative effects of certain tasks and infer that if AI technologies take them over, it would finally enable us to connect to a meaningful existence.[9] A second group does not focus on outsourcing solely as a meaning enabler but also as a meaning booster. Here, one might consider the ways in which outsourcing, for instance, by AI extender systems or others, might expand or support what we are able to do or achieve using new AI technologies (Vold, 2015; Hernández-Orallo & Vold 2019; Smids et al., 2020). A third group of contributors do not only emphasize the enabling or boosting effects of technological outsourcing but also stress that it could lead to new ways of leading a meaningful life. This could be the case, for example, if the project of "outsourcing" is pursued to the extent that traditional meaning-conferring tasks are put in the background and long forgotten, or new sources of meaning become more apparent.[10]

[5] For an overview of the methodological claims of an empirical inquiry into meaningfulness, which also provides a summary of the existing literature, see Schnell, 2020.

[6] To gain an understanding of the various ways in which different axiological dimensions can impact meaningfulness, see the nuanced picture developed and summarized in Metz, 2013, Sect. 12.2. My own thoughts on how well-being, morality, and other dimensions are connected to meaning in life are summarized in Kipke & Rüther, 2019 and Rüther, 2023.

[7] For my own views on that matter refer to Rüther, 2023.

[8] See also the final section of the paper, where I further expand upon the presuppositions I embrace and the type of theoretical neutrality I strive to achieve.

[9] An example is the so-called anti-work-critique which claims basically that work is intrinsically bad or at least has high opportunity costs. There is an extensive literature on the positions and arguments on this movement. The locus classicus is Black, 1986. For more recent thoughts, see Crary, 2014 and Srnicek & Williams, 2015.

[10] For instance, in the context of workplace automation, Knell & Rüther (2023) argue that outsourcing traditional labor could lead to

Again, this rather lengthy portrayal of the positive effects should not divert from the status quo of the debate. The focus on the positive effects of technological outsourcing remains an exception. This is, in my opinion, a regrettable fact. I am unsure if further exploration will ultimately lead to an overall optimistic attitude or even to "cautious optimism" (Danaher et al., 2018), but I assume that this project will at least help to achieve a more multifaceted and balanced view on the topic.

That being said, in this paper, I want to align with the vast majority of the discourse and focus on the threats to meaningfulness posed by technological outsourcing. This is, of course, a very worthwhile project too, as I agree with most in the field that there are in fact good reasons to consider the negative effects. As noted in the introduction, we are surrounded by an increasing number of technologies that seem to take over tasks from us, making it important to clarify how such technologies can pose serious threats to the values we associate with living meaningful lives. Given the perspective on threats, it may come as a surprise that their nature and the variety of how they are presented is very limited. In fact, many of the considerations found in the literature boil down to the same line of thought. In its most basic version, it runs as follows:

(1) Certain of our tasks have been identified as meaning-conferring.
(2) Given the process of outsourcing, those meaningful tasks are not available to us in the same way as before.
(3) Thus, the inability to perform meaningful tasks leads to a loss of meaning.

In the following, I will label this challenge as the "meaningfulness gap". This terminology nicely captures the central idea: technological outsourcing creates a gap between our aspirations to perform meaningful tasks and the reality that we are unable to do so because AI systems have taken over fully or in part. Yet, as far as I am aware, the term "meaningfulness gap" is not widely used to describe this challenge in the literature.[11] The basic idea is, as previously mentioned, widely recognized. This will be explored more extensively in the following sections when I examine the different aspects that can be taken into account to describe and ultimately evaluate a certain gap. At this point, however, I already want to cite just a few instances to convey its prevalence and how the threat is formulated in the literature.[12]

Some contributors, for example, point to a meaningfulness gap by claiming that technological outsourcing makes certain types of feelings, which are necessary for meaningfulness, unlikely or even impossible. In this regard, Nicholas Carr (2015) draws on research by Csikszentmihalyi and colleagues, and seems to think that outsourcing at the workplace leads to fewer opportunities for "flow experiences". These types of experiences are necessary parts of leading a meaningful life. On another note, others focus less on feelings and more on skills and virtues. Shannon Vallor (2015, 2016) is well-known for adopting a virtue-theoretic approach, emphasizing the importance of general virtues such as 'situation assessment' and 'critical reflection,' as well as 'communicative virtues' like patience, honesty, empathy, and tolerance. She argues that these skills are at risk of erosion since AI systems have taken over tasks relevant for building them. In addition, another angle suggests that neither feelings nor virtues are at stake, but a detachment from meaningful outcomes. John Danaher (2017), for example, stresses that in an outsourced world, individuals would no longer be responsible for achieving meaningful outcomes, most notably being responsible for them. Likewise, some are not concerned about feelings, skills, or outcomes, but about a loss in our connection to "something bigger". Jilles Smids and colleagues (2020, p. 511) focus on automation at the workplace, recognizing that the technological outsourcing might have the consequence of isolating individuals, making it very challenging for them to experience themselves as part of overarching projects that aim to achieve a worthwhile goal.

Considering these challenges, it seems obvious that they refer to different aspects that we lose when we outsource certain tasks we usually do ourselves. Carr focuses on certain feelings; Vallor on skills and virtues; Danaher on certain outcomes; and Smids and colleagues on being part of a worthwhile project. Despite these differences, they have one crucial thing in common: they rely on the same core structure in their argument, namely that technological outsourcing prevents us from performing certain tasks essential for aspects of meaningfulness. In other words, they describe different versions of the meaningfulness gap that opens up when technological outsourcing occurs.

At this point, some clarifications are needed to avoid misunderstandings. First, it is relevant to recognize that the described meaningfulness gap differs from a meaningfulness decline. A meaningfulness gap is a negative occurrence that happens when we cannot perform tasks that would confer

Footnote 10 (continued)
a greater appreciation of contemplative sources of meaningfulness, specifically those available only through more passive and receptive modes.

[11] The only place that I am aware of is the overview paper by Nyholm & Rüther, 2023 in which the terminology is used in passing to describe a logical option (p. 20).

[12] I briefly mention instances of the gap here. A more detailed description is provided in Sect. 4.2, where I also contextualize them alongside other formulations of the gap.

meaningfulness, resulting in a loss of meaning by not being able to act in a certain way. In contrast, a meaningfulness decline occurs when we act in a way that subtracts meaning. In the literature, it is disputed whether such actions can be plausibly assumed, as it seems to suggest the existence of negative meaning, or as some call it, anti-meaning.[13] Notwithstanding this discussion, the meaningfulness gap does not entail such a conception of negative or anti-meaning. Rather, it relies on the intuitively plausible thought that not being able to perform tasks that create meaning results in a loss of meaning.

Second, it is also crucial to highlight that the meaningfulness gap is not alone in AI Ethics literature in claiming a certain "gapiness". For instance, the so-called responsibility gap is likely one of the most discussed themes in AI ethics. In addition, mostly in connection with this, some discuss the so-called retribution gap (Danaher, 2016; Kraaijeveld, 2020). Some additionally anticipate what they call an accountability gap, which opens up with further progress in AI technologies (cf. for an overview and analysis Lechterman, 2022). Also, more recently, in the context of automation, some have concerns about the rightful ascription of achievements to humans and call this challenge an achievement gap (Danaher & Nyholm, 2021). Why summarize the above-described challenge of meaningfulness under the umbrella of all these gaps? I think this can be justified by stressing the parallels between them. Taking the most prominent one, the responsibility gap, as an example. This gap can be defined as follows: if AI systems take over tasks that require human intelligence and for which human beings had previously been responsible, then this might open up a gap in responsibility, since tasks for which intelligent and morally responsible humans were previously accountable are now taken over by artificially intelligent but not morally responsible AI technologies (Nyholm, 2023, ch. 5). Now, the parallel to the meaningfulness gap becomes obvious since it can be described in an analogous way: if AI systems take over tasks that were meaningful for human beings, then unless those human beings can engage in other meaningful activities instead, a gap in meaning may have arisen as a result. Thus, the responsibility gap and the meaningfulness gap are different in their diagnoses of what goes wrong, namely a loss of responsibility or a loss of meaningfulness. But they are similar in their basic structure of how this loss is recognized, namely by our inability to perform certain tasks due to technological outsourcing. Hence, in both cases, a gap opens up, which justifies addressing them under the same conceptual framework.

[13] For a discussion on anti-meaning, see Landau, 2011 and Campbell & Nyholm, 2015. For an application of these and related thoughts on anti-meaning in AI ethics, specifically on the debate on sustainable AI, see Rüther, 2024.

Third, some readers also might have the impression that certain gaps mentioned previously have a close relationship with the meaningfulness gap, maybe even to the extent that it seems doubtful to recognize any differences.[14] This impression might get some support from the literature. For instance, the mentioned achievement gap is *explicitly* described as a loss of meaningfulness by its authors (Danaher & Nyholm, 2021, p. 227). Thus, is the meaningfulness gap simply another term for the achievement gap and therefore superfluous? As I will discuss in one of the following sections, I believe the connection is indeed close, yet the two are not interchangeable.[15] More specifically, I argue that the achievement gap represents a manifestation of the meaningfulness gap, occurring specifically within the context of workplace automation and characterized by a certain normative loss, namely achievements, that constitutes a gap in meaningfulness. In other words, the meaningfulness gap offers a broader description of the challenge, while the achievement gap represents a specific instance of it.

4 How to think about a meaningfulness gap: patterns of analysis

In the last section, I explored the core structure of the meaningfulness gap. In the following section, I will delve into the details of the challenge. Not all meaningfulness gaps are created equal. They differ in their realms and objects, normative nature, scope, and severity. To recognize this, I will describe a four-step process that can help us navigate through the complexities of the different instances of the gap.

4.1 Step 1: realms and objects

Clearly, one of the initial steps is to precisely identify the realms and objects of the specific meaningfulness gap in question. Where does this gap occur, and which specific tasks have been assumed by an AI system?

Starting with the specific realm, the range from which one has to choose is equal to the one that has been given if

[14] Apart from delineating the differences, there may also be also some other relationships to explore. Is it not plausible to suggest that the responsibility or accountability gap implies the meaningfulness gap or is, at least, a necessary condition for it? At this juncture, I do not aim to make any definitive claims, as it falls outside the scope of this paper. However, that exploring this relationship could constitute a promising research project is evident from the beginning. Such an endeavor would bridge ethical discussions on meaningfulness with those on responsibility and accountability in the philosophy of action, thereby fostering a deeper understanding of the gap and potentially its significance.

[15] Refer to Sect. 4.2 on "normative nature", where I will discuss the achievement gap in greater detail and contextualize it within the existing literature.

one explores the range of technological outsourcing. This is because logically the gap can open up at every instance where outsourcing can take place. We experience outsourcing in the realms of our everyday life endeavors, at the workplace, in healthcare, in our relationships, and organizing our leisure time, among many others. Thus, in principle, in all these realms, a meaningfulness gap can occur. Apart from the logical space where a gap may open up, one may also differentiate between more usual and less usual ones, given the discourse on outsourcing and meaningfulness. At this point, I cannot browse through the literature extensively but can only summarize my own impressions.[16] Keeping in mind that the term "meaningfulness gap" is not commonly used in the discourses, the core idea of the meaningfulness gap seems to be heavily discussed in the context of work and automation; it seems heavily discussed in the context of personal relationships and social AI technologies; it is frequently discussed in the realm of healthcare and the debate on the replacement of skills and abilities, and it is only occasionally discussed for our everyday life endeavors and leisure time.

Turning to the object of outsourcing, in the previous sections, I used the term "task" to identify what the replacement is about. However, when confronted with a specific meaningfulness gap, this is obviously not enough to clarify its content. One must also differentiate between the kinds of tasks that are affected by the outsourcing process. Here, one easily ventures into philosophically difficult territory since the interpretation depends on preconceptions one takes on board in the philosophy of science and other areas of theoretical philosophy. For matters of brevity, I will not delve into detail with this but only highlight this challenge. In this regard, I want to suggest that it is helpful to break down the term "task" into at least three categories, namely cognitive processes, action types, and outcomes.

Cognitive processes can be interpreted very broadly, encompassing almost everything humans can cognitively perform. Most discussed in the literature are processes such as reflection, evaluation, or deciding. One might also focus on subcategories of each of them to make the meaningfulness gap more nuanced. For example, as I have stressed on another occasion, understanding what makes something a "decision" is a complex matter, and we can differentiate different aspects of a decision which can be relevant for the outsourcing process (cf. Rüther, forthcoming).

What about action types and outcomes? These can also be interpreted very broadly, but some have been discussed more than others. Here, one may suspect that the presupposed theory of meaningfulness determines the action type or outcome that one finds relevant. I will elaborate more on that in the next subsection, but many in the field focus on actions that belong to the so-called realm of the true, good, and beautiful to differentiate the relevant from the non-relevant action types and outcomes.[17]

4.2 Step 2: normative nature

After exploring the realms and objects of a gap, the next clarificatory step consists in describing the normative loss we suffer when a certain task is taken over by an AI system. What is the aspect that is missing, creating the loss of meaning? Here, the answer depends on the normative theory of meaning in life, namely, on what one thinks meaningfulness consists of. We got a glimpse of the different options while I described various instances of the meaningfulness gap in Sect. 3. Now, it is time to be more systematic to get an overview of the available options. In this regard, I rely on distinctions often used in the discourse, which differentiate between subjective and objective theories of meaningfulness (see, for this standard distinction Metz, 2023, sec. 3). Let us explore these two options in defining the normative loss one by one.

a. The basic idea of a subjectivist theory of meaningfulness is the assumption that a person's life depends wholly on her concerns—or, more precisely, that a component of a person's life contributes meaning to it just in case she cares about the element in some suitable way. The literature on meaning in life knows a vast amount of options to spell out this general idea what "suitable caring" amounts to. Here, I am only concerned with the debate on technological outsourcing and the loss that is created by it. In this context, we can find the subjectivists basic idea, for instance, in claims that technological outsourcing hinders certain feelings or emotions which are connected to as Thomas Nagel famously puts it, a certain 'what-it's-like-ness' of experience (Nagel, 1974). The diagnosed loss of "flow experiences" that was mentioned in the meaningfulness gap of Carr previously is an example of such strategies.[18] Another often explored experience that is supposed to be lost is the

[16] For an overview of the literature on meaning in life in AI ethics, from which one might infer the following, see Nyholm & Rüther, 2023.

[17] Of course, when referring to the true, the good, and the beautiful, one might want to make further sub-distinctions in formulating relevant action types and outcomes. This might be necessary because some aspects can be more prone to gap formation while others are more resistant to it. For example, consider Danaher's claim that certain realms, such as the true and the good, are more likely to be severed from humanity compared to the beautiful or the ludic (cf. Danaher, 2017). If this is true, it suggests a need to focus more on the true and the good when addressing a meaningfulness gap. This position could be viable, and I do not intend to argue against it in this paper. Here, again, it is more important to note that this focus is not self-explanatory. It is influenced by (a) the underlying theory of meaningfulness and (b) an assessment of the extent of a meaningfulness gap. This will become clearer in the following sections of the paper.

[18] See Sect. 3.

feeling of purposiveness understood as having the feeling that one's life has an aim or a deep emotional connection with valuable projects. For instance, longitudinal studies by the Copenhagen-based Happiness Research Institute (HRI) have identified a correlation between time spent online and increased life dissatisfaction (Wiking, 2020). These findings have often been replicated, providing even a more nuanced view of the experiences at stake.[19] The expressed concern seems to be that some systems might inhibit the feeling of purposefulness or even exacerbate feelings of confusion, aimlessness, or emptiness.

However, feelings and emotions are not the only subjective factors that can create a loss of meaningfulness. Another one is the lack of personal autonomy. This can manifest itself in different ways, also strongly depending on the specific theory of autonomy. Here, I just recapitulate some examples of the discourse which might be instructive to get an idea about the options: Frischmann and Selinger (2018) discuss how such algorithms in social media and advertising promote reactive behaviors like 'See it, click it' responses, inhibiting critical thinking. Similarly, Susser et al. (2019a, 2019b) analyze 'online manipulation', highlighting how these systems can compromise our ability to make decisions based on personal convictions. In addition, Yeung (2017) introduces the concept of 'hypernudging', where decision-making algorithms, in tandem with big-data analysis, create 'highly personalized choice environments' (2017, p. 122) that exploit cognitive biases to influence user decisions. The threats to autonomy presented by these systems are significant and form a critical area of concern in the field.[20]

Apart from feelings and personal autonomy, some also mention subjective factors that are widely concerned with the lack of abilities and skills. In the last section, I described already Shannon Vallor's virtue ethical approach that highlights the process of deskilling due to outsourcing in the several context that leads to a lack of a certain set of virtues. The same line of reasoning has been also put forward by others also being more specific on the context, e.g., on the deskilling effects of social media and online communication (Dennis & Harrison, 2020). Recently, the field also has incorporated virtue ethics perspectives from Non-Western philosophies, such as Confucianism and Buddhism. Hershock (2021) applies Buddhist principles to contemporary technology discussions, noting how mental skills like attention and presence are negatively influenced by decision-making algorithms in various life aspects. Wong (2020) explores Confucian concepts, particularly the role of rituals ('Li') in technology discourse and how it is mitigated by processes of outsourcing.

b. The basic idea of an objectivist theory of meaningfulness consists in the claim that meaningfulness—*pace* subjectivism—cannot be found in self-focused concerns like personal feelings, autonomous decisions, skills, or capacities. Moreover, it must be located somewhere else, namely beyond the subject. In line with this idea, Robert Nozick (1981, p. 595) eloquently states, "the particular things or causes people find make their life feel meaningful all take them beyond their own narrow limits and connect them up with something else." But what is this something else? Very common is the claim that self-transcendent actions are the right way to answer this question. Some even narrow the variety of the relevant actions, including the pursuit of scientific truth, helping others, or creating beauty, leading to the concept of the axiological triad of "the true, the good, and the beautiful," shaping meaningful lives (cf. Metz, 2011 and Rüther, 2023). Others are less optimistic in narrowing them down and strive for broader definitions (for a broad enumeration see Wielenberg, 2005, p. 34). The basic idea of objectivists, taking into account self-transcendent actions can also be found in the debate on technological outsourcing. One example of this is the previously mentioned rationale by Danaher who claimed that certain outcomes are mitigated by this process. The general idea is further developed in Danaher and Nyholm (2021) by specifying the outcomes as achievements. They define achievement as "outcomes that are brought about by coordinated human activity in such a way that those outcomes can be linked to the efforts of individual human agents" (ibid., 230). In exploring the phenomenon of technological outsourcing at the workplace, they argue for a so-called "achievement gap" which opens up when workers are partly or fully replaced by AI systems. Some have questioned this diagnosis or the severity of the issue (Tigard, 2021; Parmer, 2023; Scripter, 2024). Others have followed them adding to their consequentialist approach also that the lack of achievements can be explained from a deontologist's perspective and is, thus, even more encompassing than they thought (cf. Knell & Rüther, 2023). A further candidate for defining the lack of meaningfulness is the lack of embeddedness of our actions and endeavors in a larger context. This concept varies in interpretation. Some discuss a narrative theory of meaningfulness, suggesting life's meaning partly depends on our overall personal narrative (Kauppinen, 2012, 2015; Weber-Guskar, 2023). Others explore metaphysical interpretations of "being part of something bigger", touching on cosmic harmonization, the meaning of life, and God (Craig, 2013; Baier, 2000; Cottingham, 2003). Another less metaphysical perspective, like that of Bramble (2015), views a meaningful life as contributing to significant projects that improve the world. Given this understanding, many observe

[19] For a comprehensive overview, refer to the short literature report in Dennis, 2022, which also includes an evaluation of the studies conducted to date.

[20] A notable and recent overview is provided by de Bonicalzi et al., 2023, who distinguish between various models of autonomy and correlate them with different decision-making systems.

that technological outsourcing algorithms might diminish this dimension of meaningfulness, primarily by potentially reducing our roles as contributors to larger projects. Smids et al. (2020), as already mentioned previously, are pointing into this direction in the context of workplace automation. They illustrate this with radiologists, whose interpretative roles are increasingly assumed by machine learning systems, possibly diminishing their contribution to diagnostic projects (ibid., p. 511.). This rationale can also apply to many other contexts. In a follow-up paper, Smids et al. (2023) describe the situation in a warehouse where the jobs of order packers and order pickers are increasingly being divided into subtasks, because some subtasks are taken over by technology. As a result, order packers no longer finalized the process in the form of a package ready to be sent to the customer, and are less connected to the overarching project of the whole enterprise. Or to add a further example in another area, consider educators who integrate AI tools for personalized learning may find their roles shifting from content delivery to more of a facilitative position, potentially altering their role and contribution to students' learning journeys (Guilherme, 2019; Schiff, 2021). Similarly, writers and artists who use AI to generate content might experience a change in how their creativity is perceived, both by themselves and their audiences, impacting their roles and relevance in the creation process (Démuth, 2020).

With these clarifications in place, one should be in a better position to determine the normative loss anticipated by a certain meaningfulness gap. Does it consist of factors concerning the subject—its feelings, autonomy, or certain skills and virtues—or is it about the loss of self-transcendent actions, such as the subject's achievements or its embeddedness in larger projects? Delving into this question edges us closer to discerning the specific nature of the meaningfulness gap under consideration.

It is noteworthy, however, that these factors contributing to normative loss are not conclusive. There are likely many more factors in the discourse and beyond which were not covered and are worth exploring as well. The aim of this section was not to catalog every possible scenario but to lay a foundation for navigating the complexities of meaningfulness gaps, a task for which the previously discussed considerations should suffice.

Furthermore, it is crucial to note that the diverse factors contributing to meaningfulness need not be viewed in isolation. They often intersect or blend together. Take, for example, the theory of hybridism proposed by Wolf (2010), which argues that meaningfulness is the product of a confluence of elements. More precisely, she suggests that meaningfulness emerges 'when subjective attraction meets objective attractiveness' (Wolf, 2010, p. 26), implying that a specific subjective state (for her, this involves love and engagement) must converge with an object of inherent value (for her, a project of objective worth).[21] This perspective is particularly relevant to the discourse on outsourcing and meaningfulness, indicating that a gap in meaningfulness might manifest across multiple dimensions, depending on the mixed theory one adopts.

4.3 Step 3: scope

At the earlier stages, we gained more clarity on the realms and objects of the meaningfulness gap, namely where it occurs and which tasks have been taken over. Also, we gained a better understanding of the kind of loss in meaningfulness we are suffering if that happens. Nonetheless, this is not enough for a detailed analysis of the challenge we are dealing with. We also need to estimate the scope of a meaningfulness gap. In the following, I want to suggest that one has to consider at least three aspects.

The *first aspect* that is relevant for the scope of a meaningfulness gap is how "big" the gap is in terms of its size. This can sound initially like an irritating question since one might have thought that diagnosing a gap is an all or nothing thing. Either there is a gap or there is none. Unfortunately, matters are not so simple since the gap is based on the process of technological outsourcing which itself is not solely an all or nothing process. More specifically, if we outsource a certain task, we can do that in a way it is taken over fully by an AI system, but we can also outsource in a way that the task is not done by the AI system alone. This situation is often described as collaborative work between humans and AI systems.[22] The general idea here is that certain tasks are done together to reach completion. "Doing something together" can mean different things. One option could be that collaborative work might consist of sharing a certain task. For instance, if a healthcare app wants to recommend us a certain behavior or lifestyle, it obviously needs some data on which the recommendation is based. In this case, the task of data gathering can be shared between the AI system built into the app, for instance, tracking our footsteps, and humans who complete the data if necessary (about weather conditions, elevations, etc.) or add some new information (for example, about daily nutrition) to improve or even complete the process. Another option is that the task is not being shared but overseen by humans. Here, the contribution of humans in the process is obviously reduced and does not contain the performance of the tasks, but its supervision. Typically, this option is discussed in the literature in connection with the idea that humans are able to give a veto if

[21] I am not intending laying down the complex hybridist's theory of Wolf and its reception at this point. For a start, refer to Evers & Smeden, 2016.

[22] Cf. for a sketch of the different versions of collaborations, Danaher & Nyholm, 2021, p. 229.

the task is not performed in the right way by the AI systems. For instance, the health app might want access to information from us which we do not want to share, e.g., our GPS data or other personally relevant data, or it wants to send us notifications which we do not want to receive. Those two options, task sharing and supervision, are solely examples of collaborative work between humans and AI systems. There is a variety of different ways to allow human contributions. But what do these possibilities imply for estimating the size of the meaningfulness gap? I suggest that it implies that we have to consider not only the possibility of a fully extended gap but also the possibility that smaller gaps open up. This is due to the fact that the outsourcing processes can also be limited to certain aspects of a task so that collaborative work is possible. In those cases, the relevant task is not fully replaced and thus reduces the size of the meaningfulness gap that opens up as a result of outsourcing. Or to put the point slightly differently: the more an outsourcing process allows human contribution in the task, the more it reduces the size of the gap. If we, therefore, want to evaluate the extent of a gap, we should consider the extent to which humans are still involved.

A *second aspect* that is relevant for understanding the scope of a meaningfulness gap does not have to do with the extent of outsourcing but to its duration. It might seem obvious, but the duration significantly impacts the scope of a gap, whether it will be short term and closed in the near future or long term and possibly extend for a long period. In this context, I suggest that we can differentiate between two distinct endpoints on a continuum when evaluating the meaningfulness gap. On one end of the spectrum, we identify gaps that endure over a long period. The term "long" requires interpretation, but the general idea about the duration is quite clear. The meaningfulness gap under consideration is not ephemeral but a prolonged concern, substantially outlasting other events. This does not imply that a gap with a long duration must be a single event that opens and closes at some point. Furthermore, the gap could open and close periodically. The crucial point here is that the overall duration is considered to be significantly long.

On the other end of the spectrum, we find meaningfulness gaps that are short-term events. Again, I do not wish to specify a fixed period, but it is assumed from the outset that a gap with a short duration is inherently limited. This limitation could pertain to a single event or to recurrent episodes that arise and are resolved periodically. In the case of periodic occurrences, the durations should be aggregated, cumulatively representing an event with a relatively short lifespan.

At this point, it merits attention that frequently a meaningfulness gap often falls between these extremes: it is neither fleeting nor eternal. Even in such cases, distinctions can be made between those leaning toward one end of the spectrum; it might also matter whether the gap is a singular event or occurs in a cyclical manner. While I am not making definitive claims, it seems evident that positioning a gap within this spectrum facilitates a deeper understanding of its scope.

The *third aspect* that will help to explore the scope of a meaningfulness gap is not directed at its size or duration, but at its range. With "range," I mean the extent to which the gap applies to the life of people. For how many people does it affect? As to my knowledge, there is no taxonomy that gives some guidance to these kinds of questions.[23] That the range of a meaningfulness gap does make a difference in estimating its scope, though, seems to be obvious. Thus, I want to propose at least a threefold division that might help to discern the range of a meaningfulness gap.

At first, one might think of a gap with a very high-range or almost universal range. These are those that apply to everyone or almost everyone. Typically, such a meaningfulness gap is concerned about the outsourcing of universal meaningful tasks, namely those that play a meaning-conferring role in almost everyone's life. Let us, for instance, take the example of tasks at the workplace. Of course, not all of them are meaning-conferring. But it seems fair to say that some of them are, even varying from person to person, since tasks at work seem to be a relevant aspect of most of our lives (see for the "goods of work" the considerations in Ghaeus & Herzog, 2016). Now, if we assume that automation at the workplace takes place and all or almost all tasks will be outsourced, we have a meaningfulness gap of almost universal range. People are not able to perform those tasks at work anymore since AI systems have taken over them.

Second, one might differentiate universal gaps from those with a middle range that apply to some but not to all people. Those gaps are typically created by outsourcing processes of tasks that some people perform on their own but are now at risk of being outsourced. Imagine the scenario of workplace automation again. In this regard, technological outsourcing might not affect work per se, but only several instances of it, for example, those tasks performed by a certain group of people. Widely discussed is, for instance, the thesis that jobs with certain characteristics are more likely to be outsourced. Again, I am not committing myself to this thesis. I am only stating that these kinds of thoughts could lead to a middle-range meaningfulness gap because it affects some, but not all people.

Third, there is a distinction to be made between broad-range gaps and those affecting solely a narrow segment of

[23] An exception lies in the thoughts presented by Nyholm & Rüther (forthcoming), which attempt to summarize different scenarios regarding the distribution of AI's burdens and benefits. Although the paper focuses on work, it may serve as a starting point for future research aimed at exploring 'range' as an important factor in determining the scope of a meaningfulness gap.

the population. These gaps typically involve the outsourcing of tasks undertaken by only a handful of individuals. In the context of the workplace, this might concern niche groups engaged in highly specialized tasks. The outsourcing in these instances leads to a meaningfulness gap of limited scope, relevant primarily to those directly affected by such changes.

Here, it might seem that the range of a gap correlates with a certain severity of a gap. However, I want to emphasize that this is not the case. There is no direct connection between "This gap concerns everyone" to "This gap is deeply severe", or from "This gap only concerns a few people" to "This gap is marginal". There might be a gap that applies to everyone, but which is less severe than a gap that applies to only a very few people. This does not mean either, that the range does not play a role in discerning the severity at all. I only want to point to the fact that range and severity are not normative substitutes. How we are dealing with the severity more precisely is explored in the now following section.

4.4 Step 4: severity

The first three steps serve mostly for a clarificatory purpose to understand the version of a certain meaningfulness gap that we are confronted with. The last and final step deals with an approximate evaluation of its severity. Given the specific shape of the gap: how bad is it? Are we dealing with a marginalizable or a severe challenge? To give a sufficient answer to these questions, one needs to lay out the facets of the favored ethical theory. I am not inclined to do that as my basic aim of this section is solely to flag certain starting points for estimating the severity. Ideally, those are supposed to be designed in a way that many contributors with various backgrounds can subscribe to them. Given this modest goal, I suggest that navigating toward the following three sub-steps could be beneficial, as it may place us in a more favorable position to evaluate the severity of a meaningfulness gap.

First, I assume that it is promising to approach an internal weighing of the gap in comparison with other meaningfulness gaps. The idea behind this would not be to get an all-things-considered judgment about its severity, but a first impression of the magnitude of the problem at hand. This is due to the fact that sometimes problems can appear bigger or smaller than they really are if they are considered in isolation. Here, it might turn out that a certain meaningfulness gap, first considered as a challenge, turns out to be marginalizable in comparison to other gaps or—vice versa—that a gap that seems to be small is more severe than initially thought if seen in the context of other gaps.

To get a better idea of how a rough and ready comparison might work, let us go very swiftly through the different steps that I developed in the previous sections with an example: consider the case of a certain recommender system in the area of everyday life activities such as selecting movies and TV series, which makes us recommendations about watching certain things based on our past experiences and explicit inputs of the viewer. In doing that, we outsource our reflection about those contents to that system (step 1). The normative losses of meaning can be determined differently. Given a first clue, one might consider the loss of autonomy or certain skills as good candidates (step 2). One might also be inclined to think that the scope is limited in size (the decision is not outsourced; input is possible), has high duration (unlikely that the platform changes its policy), and is limited in range (it affects solely a minority of people) (step 3). Now compare this AI recommender system with an AI system in the area of the workplace which is designed to automate most of our tasks and those of many others in different areas too by taking over almost all of our decisions and also actions (step 1). Here, the normative loss can be determined variably plus piled up from lacking certain emotional states to losing our autonomy in decision-making, no longer being able to perform meaningful actions (step 2). Also, the overall scope is high in size (it is an instance of full outsourcing), also with a long duration (long-term), and wide range (concerns many people in wage work) (step 3).

Which one is the more severe one? I assume that many would agree that the second system is more severe in terms of the overall loss that we suffer if we implement it. Of course, to make a more qualified judgment about this, we need more information about the two systems. However, again, this is not the aim at this stage. In the first step of the process, it is enough to make a rough and ready comparison to get a first clue about the severity of a meaningfulness gap. The following steps will bring that first clue to a more sophisticated level.

Second, after making the comparison, I suggest that we can proceed with an internal weighing of the gap. This means to weigh the loss that a gap presumably causes against other aspects in the dimension of meaningfulness. The purpose of this would be to get a better understanding of the all-things-considered impact of the gap. Here, several constellations might apply. A very bleak constellation might be that the gap in meaningfulness creates a situation which is unbearable in terms of meaningfulness. These scenarios can be described as dystopian, characterized by significant meaningfulness gaps—situations where losses are so substantial that they lead to a meaningless existence. In essence, even when considering potential benefits (should they exist), life in such scenarios could be compared to the archetypal story often cited in discussions: that of Sisyphus, a figure doomed to the futile task of rolling a stone up a hill, only for

it to roll back down each time. Another little less depressing scenario might be that our life would be, if confronted with the meaningfulness gap, be on balance less meaningful, but not meaningless. In that case, the loss of meaningfulness might be really an essential loss, but it might be on balance not enough to condemn us to a Sisyphean existence. However, it might be not all dome and gloom. There may be positive constellations too which indicate that a meaningfulness gap would create a loss, but this can be outweighed by positive meaning enhancers that might be enforced if the gap is present. Also, to complete the picture, we might also get to an utopist scenario, in which the meaningfulness gap is deeply welcomed because the normative gain that is following it as a consequence is so huge. We suffer a possibly small loss and get a huge reward in meaning back.

Third, in determining the overall severity of a gap, we should also weigh it against aspects beyond meaningfulness. Normative ethics often discusses variables like well-being, morality, or cosmic meaning, among potentially others, depending on the ethical framework. Although it is impractical to cover all conceivable scenarios, I aim to highlight a few intuitive ones, concentrating on well-being while omitting others for specificity. It is assumed, however, that these considerations could extend to additional variables.

Reflecting on earlier scenarios, we might encounter situations where a meaningfulness gap not only represents a loss in itself but also in well-being. In such cases, the erosion of meaning accompanies, as an unfortunate bonus, a reduction in well-being. Conversely, there are instances where a gap, despite being an overall loss in terms of meaningfulness, coincides with an increase in well-being. This scenario suggests a trade-off, requiring a balance between the detriments of reduced meaning and the benefits of enhanced well-being. On a more positive note, some gaps might boost both overall meaningfulness and well-being simultaneously. In those cases, meaningfulness and well-being go hand in hand. Alternatively, a scenario could emerge where the gap benefits overall meaningfulness yet detracts from well-being. Here, the meaningfulness gap is welcomed in terms of overall-net meaningfulness, but implies some burdens in terms of one's own welfare.

Recall again, that I do not intend to make any judgements about the plausibility of a certain constellation, nor am I making a final call about how the weighing needs to be done. In this section, I am only stating the possible options and constellations. This should be enough for having some guideline to explore the severity of a certain meaningfulness gap. This must, of course, be done on a case-by-case basis and from the background of one's own ethical theory. In this section, I tried to be neutral toward this, so that the thoughts are necessarily vague and broad, but hopefully in consequence accessible and relatable to many.

5 Summary and further remarks

In the last sections, I tried to give some guidance on how we can think our way through the complexities of the different meaningful gaps. To achieve this, I proposed an elaborate four-step process which entails several main topics and distinctions. To make this step-by-step guide more applicable and transparent, we can summarize the bare bones in the following table:

Step	Details	Options
Step 1: realms and objects	a.) Where does a gap occur? b.) Which tasks have been taken over?	ad a.) Everyday life, workplace, healthcare, relationships, leisure time ad b.) Cognitive processes (e.g., reflection, evaluation, decision), action types, outcomes
Step 2: normative nature	What is the normative loss we are suffering if confronted with a meaningfulness gap?	subjective aspects (e.g., feelings, autonomy, skills, virtues) vs. objective elements (e.g., transcendent actions, being part of bigger projects)
Step 3: scope	What is the scope of a certain meaningfulness gap?	1. Size (full replacement vs. collaborative endeavors) 2. Duration (long-term vs. short-term) 3. Range (universal range, broad range, low range)
Step 4: severity	How severe is a certain meaningfulness gap?	1. Rough and ready in comparison to other meaningfulness gaps 2. Internal weighing against other aspects of meaning 3. External weighing against non-meaning aspects, e.g., well-being

To avoid misunderstandings, I want to add a few clarifications on the status of this step-by-step guide.

First, it is important to understand that while this guide provides a structured approach to analyzing meaningfulness gaps, it is not intended to be exhaustive. This holds true for all aspects of the guide, namely its steps, guiding questions, and the given taxonomies on how to answer them. In fact, some might be inclined to modify or add further aspects. In principle, I see this as a worthwhile endeavor. The primary aim of the guide is to stimulate further discussion and exploration, acknowledging that the field is evolving and our

understanding of these gaps may change as new insights are gained. Hence, if it turns out that the proposed guide can be improved, I do not see any general obstacles.

Second, while I have tried to remain as neutral as possible regarding the presuppositions I have taken on board, it is important to note that the paper does not approach the issue from a standpoint of complete neutrality. It makes, for example, some claims about which steps, questions, and taxonomies are relevant and which are not. Within these pages, I cannot provide extensive justifications for all of them. However, the careful reader would have noted that I aim at least for transparency, mentioning these presuppositions directly and providing references for those interested in further exploration. For instance, consider my discussions on "realms and objects", where I relied on some previous writings to justify the usage of certain distinctions; or consider the taxonomies in the section "normative nature " which I derived from the ongoing discourse in the field of meaning in life. In those and other cases, while not giving an in depth justification, I tried at least to indicate on which theoretical ground I am standing and therefore open the door for further discussion.

Finally, I wish to close this paper by spending a few words on the promising nature of the proposed guide. As shown, the concept of a meaningfulness gap is a recognized reality in the discourse. Some contributors are explicitly sharing and developing the core idea, naturally without labeling it as a meaningfulness gap. However, given its presence in the discourse, it is surprising that there is next to no systematic reflection on its specifics. Meaningfulness gaps differ from each other, for instance in terms of their realms and objects, their normative nature, their scope, and their severity. To assess the importance and relevance of a certain meaningfulness gap, we must reflect on these dimensions. This paper serves as a starting point for anyone who wants to contribute to this project, proposing a method to navigate this complex challenge now and in the future.

Acknowledgements I would like to thank my colleagues in Bonn und Juelich for helpful comments at earlier stages of the manuscript. Also, I would like to thank Roland Kipke for making the paper more readable and pressed me to be detailed and nuanced my presentation of the different arguments and positions.

Authors' contributions Not applicable.

Funding Open Access funding enabled and organized by Projekt DEAL. Not applicable.

Availability of data and material Data sharing is not applicable to this article as no datasets were generated or analyzed during the current study.

Declarations

Competing interests The authors declare that they have no competing interests.

Ethics approval and consent to participate Not applicable.

Consent for publication Not applicable.

References

Baier K (2000) The meaning of life. In: Klemke ED (ed) The meaning of life, 2nd edn. Oxford University Press, Oxford, pp 101–132
Black B (1986) The abolition of work and other essays. Loompanics Unlimited, Port Townshend
Bonicalzi, Sofia, De Caro, Mario & Giovanola, Benedetta (2023) Artificialintelligence and autonomy: on the ethical dimension of recommender systems. Topoi 42(3):819–832.
Bramble, Ben (2015) Consequentialism about meaning in life. Utilitas 27(4):445–459
Campbell SM, Nyholm S (2015) Anti-meaning and why it matters. J Am Philos Assoc 1(4):694–711
Carr N (2015) The glass cage: where automation is taking us. Bodley Head, London
Coeckelbergh M (2020) AI ethics. MIT Press, Cambridge
Cottingham J (2003) On the meaning of life. Taylor & Francis: London.
Craig WL (2013) The absurdity of life without god. In: Seachris JW (ed) Exploring the meaning of life: an anthology and guide. Wiley-Blackwell, pp 153–172
Crary J (2014) 24 / 7: late capitalism and the end of sleep. Verso, London
Danaher J (2016) The threat of algocracy: reality, resistance and accommodation. Philos Technol 29(3):245–268
Danaher J (2017) Will life be worth living in a world without work? Technological unemployment and the meaning of life. Sci Eng Ethics 23(1):41–64
Danaher J (2019) Automation and utopia. Harvard University Press, Cambridge
Danaher J, Nyholm S (2021) Automation, work and the achievement gap. AI Ethics 1(3):227–237
Danaher J, Nyholm S, Earp BD (2018) The quantified relationship. Am J Bioeth 18(2):3–19
Démuth A (2020) Art and artificial intelligence—challenges and dangers. Espes 9(1):26–35
Dennis MJ, Harrison T (2020) Unique ethical challenges for the 21st century: online technology and virtue education. J Moral Educ 2:1–16
Dennis MJ (2021) Digitalwell-being under pandemic conditions: catalysing a theory of online flourishing. Ethics and InformationTechnology 23(3):435–445
Evers D, van Smeden GE (2016) Meaning in life: in defense of the hybrid view. Southern J Philos 54(3):355–371
Frischmann BM, Selinger E (2018) Re-engineering humanity. Cambridge University Press, Cambridge
Gheaus A, Herzog L (2016) The goods of work (other than money!). J Soc Philos 47(1):70–89

Gordon J-S, Nyholm S (2021) Ethics of artificial intelligence. Internet encyclopedia of philosophy. Retrieved from https://iep.utm.edu/ethics-of-artificial-intelligence/
Guilherme A (2019) AI and education: the importance of teacher and student relations. AI Soc 34(1):47–54
Hammerton M (2022) Well-being and meaning in life. Can J Philos 52(5):573–587
Hernández-Orallo J, Vold K (2019) AI extenders: the ethical and societal implications of humans cognitively extended by AI. In AIES '19: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, 507–513
Hershock PD (2021) Buddhism and intelligent technology: toward a more humane future. Bloomsbury Academic, New York
Kauppinen A (2012) Meaningfulness and time. Philos Phenomenol Res 84(2):345–377
Kauppinen A (2014) Flourishing and finitude. J Ethics Soc Philos 2:1–6
Kauppinen A (2015) Meaningfulness. In: Fletcher G (ed) Routledge handbook of the philosophy of well-being. Routledge
Kipke, Roland, Rüther, Markus (2019) Meaning and morality. Social Theory and Practice 45(2):225–247
Knell S, Rüther M (2023) Artificial intelligence, superefficiency and the end of work: a humanistic perspective on meaning in life. AI Ethics 4, 363–373. https://doi.org/10.1007/s43681-023-00273-w
Kraaijeveld SR (2020) Debunking (the) retribution (gap). Sci Eng Ethics 26(3):1315–1328
Landau I (2011) Immorality and the meaning of life. J Value Inquiry 45(3):309–317
Landau I (ed) (2022) The oxford handbook of meaning in life. Oxford University Press, Oxford
Lechterman TM (2022) The concept of accountability in AI ethics and governance. In: Bullock J, Chen YC, Himmelreich J, Hudson V, Korinek M, Young M, Zhang B (eds) Oxford handbook of AI governance. Oxford University Press, Oxford
Metz T (2001) The concept of a meaningful life. Am Philos Q 38(2):137–153
Metz T (2011) The good, the true and the beautiful: towards a unified account of great meaning in life. Relig Stud 47(4):389–409
Metz T (2013) Meaning in life. An analytic study, Oxford
Metz T (2021) Meaning and medicine: an underexplored bioethical value. Ethik in der Medizin 33:439–453. https://doi.org/10.1007/s00481-021-00662-x
Metz T (2023) The meaning of life. In: Zalta EN, Nodelman U (eds) The stanford encyclopedia of philosophy (fall 2023 edition). Retrieved from https://plato.stanford.edu/archives/fall2023/entries/life-meaning/
Monsó S, Benz-Schwarzburg J (n.d.). Must animal lives be meaningless? Unpublished manuscript.
Müller V (2020) Ethics of artificial intelligence and robotics. Stanford encyclopedia of philosophy. Retrieved from https://plato.stanford.edu/entries/ethics-ai/
Nagel T (1974) What is it like to be a bat? Philos Rev 83(4):435–450
Nozick R (1981) Philosophical explanations. Harvard University Press
Nyholm S, Rüther M (2023) Meaning in Life in AI Ethics—Some Trends and Perspectives. Philos. Technol. 36: 20.https://doi.org/10.1007/s13347-023-00620-z
Nyholm S (2023) This is technology ethics: an introduction. Wiley-Blackwell, New York
Nyholm S, Danaher J, Earp BD (2022) The technological future of love. In: McKeever N, Grahlé A, Saunders J (eds) Philosophy of love in the past, present, and future. Routledge, pp 224–239
Nyholm, Sven, Rüther, Markus (forthcoming) Artificial intelligence, equality, and the future of work. In: Murphy, Coleen; Hansson, Sven Ove (eds) Technology and equality. shifting roles in a complex relationship. Rowman and Littlefield International
Parmer WJ (2023) Meaningful work and achievement in increasingly automated workplaces. J Ethics 2:1–25
Purves D, Delon N (2018) Meaning in the lives of humans and other animals. Philos Stud 175(2):317–338
Rüther, Markus (2023) Sinn im Leben. Eine ethische Theorie. Suhrkamp, Berlin
Rüther, Markus (2024) Why care about sustainable AI? Some Thoughts From The Debate on Meaning in Life. Philosophy and Technology 37(1):1–19
Rüther; Markus (forthcoming): Decision-making algorithms and meaning in life: A framework for future research. In B. Heinrichs, U. Ettinger, & C. Murawski (Eds.), Decision Making: Determinants, Mechanisms and Consequences (Studies in Neuroscience, Psychology and Behavioral Economics). Springer
Scheffler S (2018) Why worry about future generations? Oxford University Press, Oxford
Schiff D (2021) Out of the laboratory and into the classroom: the future of artificial intelligence in education. AI Soc 36(1):331–348
Schnell T (2020) The psychology of meaning in life. Routledge
Scripter L (2024) The achievement gap thesis reconsidered: artificial intelligence, automation, and meaningful work. AI Soc
Smids J, Nyholm S, Berkers H (2020) Robots in the workplace: a threat to—or opportunity for—meaningful work? Philos Technol 33(3):503–522
Smids J, Berkers H, Le Blanc P et al (2023) Employers have a duty of beneficence to design for meaningful work: a general argument and logistics warehouses as a case study. J Ethics. https://doi.org/10.1007/s10892-023-09442-9
Srnicek N, Williams A (2015) Inventing the future: postcapitalism and a world without work. Verso, London
Susser D, Roessler B, Nissenbaum H (2019) Online manipulation: hidden influences in a digital world. Georgetown Law Technol Rev 4(1):1–45
Susser D, Roessler B, Nissenbaum H (2019a) Online manipulation: hiddeninfluences in a digital world. Georgetown Law Technol Rev 4(1):1–45
Susser D, Roessler B, Nissenbaum H (2019b) Technology, autonomy, and manipulation. Internet Policy Review. 8(2):1–22
Tigard D (2021) Workplace automation without achievement gaps: a reply to danaher and nyholm. AI Ethics 1(4):611–617
Vallor S (2015) Moral deskilling and upskilling in a new machine age: reflections on the ambiguous future of character. Philos Technol 28(1):107–124
Vallor S (2016) Technology and the virtues: a philosophical guide to a future worth wanting. Oxford University Press, Oxford
Vold K (2015) The parity argument for extended consciousness. J Conscious Stud 22(18):16–33
Weber-Guskar E (2023) Temporal textures. time, meaning, and the good life. Eur J Philos 2:1–14
Wiking M (2020) A happiness report in the age of Corona. Happiness Research Institute. Retrieved May 20, 2020, from www.happinessresearchinstitute.com/news4
Wielenberg E (2005) Value and Virtue in a goodless universe: Cambridge: CambrigeUniversiyt Press.
Wolf SR (2010) Meaning in life and why it matters. Princeton University Press
Wolf S (2016) Meaningfulness: a third dimension of the good life. Found Sci 21(2):253–269
Wong P-H (2020) Why confucianism matters in ethics of technology. In: Vallor S (ed) Oxford handbook of philosophy of technology. Oxford University Press, Oxford
Yeung K (2017) 'Hypernudge': big data as a mode of regulation by design. Inf Commun Soc 20(1):118–136