



AI deskilling is a structural problem

Avigail Ferdman¹

Received: 24 June 2025 / Accepted: 6 October 2025 / Published online: 5 November 2025
© The Author(s) 2025

Abstract

Many artificial intelligence tools replace or stand to replace human activity, via automated decision-making, recommender systems and content generation. The more artificial intelligence (AI) replaces valuable human activity, the more it risks deskilling humans of their human capacities. This paper argues for applying a *structural* perspective to this phenomenon. It introduces the concept of ‘capacity-hostile environments’ to identify instances where AI mediation impedes human capacity cultivation. The analysis moves beyond individual responsibility that agents have to cultivate their human capacities, demonstrating how AI’s influence creates systemic conditions that could inhibit the development and exercise of human capacities by undermining the process of capacity cultivation. Drawing on the philosophy of skill as well as social epistemology, this paper argues that capacity cultivation (skilling) includes acquiring agential control over the capacities, inculcated through a long, gradual process of habituation. Habituation, in turn, depends on learning from others: the ‘know how’ of the skill, as well as a shared understanding of the value of the skill. AI mediation risks undermining the quality of the conditions for capacity habituation, leading to capacity impoverishment. By exploring the role of AI in mediating human activity, the paper highlights the need to evaluate AI applications based on their conduciveness or hostility to capacity cultivation. Ultimately, it calls for a critical reflection on the values inherent in AI socio-technical systems and emphasizes the societal obligation to foster capacity-conducive environments in the age of AI.

Keywords Skill · AI · Flourishing · Affordances · LLMs · Perfectionism

1 Introduction

In Plato’s *Phaedrus*, Socrates cautions that writing will deskill the capacity for memory in people’s minds: trusting an external entity will discourage the use of their own memory within them. A similar warning about technology-induced deskilling has come into sharp relief recently with the rise of Large Language Models (LLMs). Philosophers, educators, artists, and journalists are increasingly raising concerns about the potential decline in our epistemic and social capacities (Vallor 2015; Wong 2019; Fricker 2021; Chiang 2023; Kudina and de Boer 2024; Rosen 2024; Tacca and Gilbert 2024; Marin and Steinert 2025).¹ There is growing evidence of the erosion of critical thinking (Bonicalzi et al. 2023; Gerlich 2025; Lee et al. 2025), decision-making and analytical reasoning (Zhai et al. 2024) as a result of overreliance on artificial intelligence (AI). In the professions,

researchers have found that AI dependence leads to the erosion of activity awareness, competence maintenance, and output assessment (Rinta-Kahila et al. 2023), while in healthcare, AI dependence has been found to diminish diagnostic reasoning, declining clinical judgment, reduced retention of tacit knowledge, declining ethical sensitivity and weakened moral judgment (NHS England 2023; Natali et al. 2025; Budzyń et al. 2025).

A key reason why deskilling is harmful is that it diminishes us as human beings, undermining our abilities in the ‘arts of personhood’ (Vallor 2015): the epistemic, social, moral and creative capacities required to achieve competency in our exercise of practical reason, our self-worth, as well as mutual respect between persons, honesty and justice, developed through relationships of dependency (Hacker-Wright 2009). Yet given AI’s rapid integration into our socio-technical environment, it’s unrealistic to expect persons to be ‘virtuous superheroes’ (Mitcham 2024) and

✉ Avigail Ferdman
avigailf@technion.ac.il

¹ Technion – Israel Institute of Technology, Haifa, Israel

¹ The worry isn’t consensual. Some think that LLMs might be beneficial. <https://thewalrus.ca/i-used-to-teach-students-now-i-catch-chatgpt-cheats>; <https://dailynous.com/2025/03/07/a-philosophers-reflections-on-teaching-in-a-world-with-ai/>.

will themselves to resist deskilling. In other words, framing deskilling as a matter of individual responsibility—expecting people to simply will themselves to cultivate ‘the arts of personhood’—overlooks crucial structural aspects of the problem. To a large extent, we develop and exercise our epistemic, moral, social, and creative capacities in response to our social and material environment (Ferdman 2019). A deficient environment can lead to deskilling of our core human capacities, and to the subsequent devaluation of their worth as they continue to weaken. Rather than focusing only on how human agents ought to engage with AI tools, we should also ask what kind of socio-technical environments would be best for the cultivation of our personhood.

Before addressing this challenge, I want to clarify how I use the concept of skill. “Deskilling” is a useful analogy for the deterioration in the arts of personhood. It’s important, however, to distinguish between skill and capacity. Generally, a skill is a kind of disposition to form knowledge appropriate to a particular domain of human activity (Stanley and Williamson 2017), as well as having the know-how for executing the knowledge (Pavese 2022). Yet ‘skill’ is sometimes used to denote expertise in domains, such as hunting/gathering, basket-weaving, writing, coding, navigating, and composing music. On this use, deskilling is not always negative: skill obsolescence has always been part of human history, where technological developments free up time and resources for the cultivation of other skills (Coeckelbergh 2019; Aylsworth and Castro 2024). The deskilling of our core human capacities—to know, create, form meaningful relationships, and use our willpower—is a different matter. In this sense, capacities are more analogous to virtues (Sherman 1991; Stichter 2020). The deskilling of our core human capacities is bad because inadequate development of these capacities can lead to “capacity impoverishment”, which, in turn, diminishes us as humans and undermines our flourishing. Therefore, from a normative perspective, whereas deskilling of activities like hunting, basket-weaving or coding may be innocuous, capacity-deskilling should always be a cause for worry. As such, we have reason to ensure that socio-technical environments we inhabit are conducive, rather than hostile, to the cultivation of our human capacities.

I begin by introducing the core capacities to set the stage for why their deskilling would be bad for humans (Sect. 2). Next, I draw on philosophy of skill to construct a framework for understanding the human capacities as skills (Sect. 3). The structural lens is then added to the framework, to distinguish between environments that are conducive to capacity cultivation and environments that are hostile to it (Sect. 4). In Sect. 5, I apply this framework to artificial personal assistants (APAs)—one of the fastest growing uses of large language models (LLMs). Section 6 concludes with reflections

on the moral implications of AI systems for capacity-skilling and human flourishing.

2 Core human capacities: Developmental Perfectionism

‘Developmental Perfectionism’ is a neo-Aristotelian approach to human flourishing (Kauppinen 2025), where flourishing is understood as a complex, unified, and well-rounded life (Hurka 1993). Humans exercise their *capacities* in *activities* that yield *outcomes* that have intrinsic value like knowledge, beauty, morality, friendship (Kraut 2013; Bradford 2016, 2021; Fletcher 2016). The competent exercise of the human capacities is the manifestation of the achievement of the intrinsically valuable goods, making the capacities intrinsically valuable too (Bradford 2016). For example, the value of being on the summit of Mt. Everest is gained from the competent exercise of a combination of capacities involved in climbing it, and not from simply being there (say, if one arrives there by helicopter, Hurka 1993; Bradford 2015). Developmental Perfectionism neatly captures why deskilling is bad: capacity-deskilling is bad because our human capacities are constitutive of our flourishing. Impoverished capacities (i.e., capacity-deskilling), therefore, lead to impoverished lives (Ferdman 2024).

Perfectionist philosophers argue for a set of core human capacities such that when humans competently develop and exercise them, they lead a flourishing life²:

The epistemic capacities include theoretical rationality: thinking, considering reasons to believe, forming beliefs, contemplating (Bradford 2025); and practical rationality (Bradford 2015, 2021), manifesting in *phronesis* (common sense or wisdom): doing the right thing in the right way for the right reason.

The social capacities include the moral capacity as well as the social and caring capacities for developing meaningful relationships of friendship and love (Hurka 1993).

The capacity for creativity involves making unfamiliar combinations of familiar ideas, transforming a conceptual space (Boden 2004), without relying merely on luck, accidents, or mechanical procedures (Gaut, 2010).

The capacity to will (volitional capacity): the ability to exert effort, overcome difficulty, persevere and grow, in order to achieve an intrinsically valuable goal (Bradford 2015). Importantly, developing and exercising any of the human capacities is difficult, which therefore requires the competent exercise of the capacity to will. I will therefore treat the capacity to will as a meta-capacity of sorts.

² The physical capacities are part of this set, but I do not discuss them in this paper.

3 What does it mean to be skilled? A capacity-skill framework

The concept of skill is a useful analogy for the cultivation and excellent exercise of the human capacities, as it explains both what constitutes skill, as well as the process by which persons become skilled.³

Talents, knacks, or abilities are not purposeful or goal-directed (e.g., breathing or digesting), rather they are things that persons have. Skills, on the other hand, manifest in action, that is, in what agents *do*, and in the acquisition of mastery and control over an activity (e.g., playing basketball, the violin or chess, Shepherd 2021). Skills are also distinct from competences: competences are acquired, but leave no room for any improvement, talent, or flair, and are not gradable: one can either do it or one cannot (Kremer 2020). Skills are also learnable, in contrast to instincts. Skilled agents' capacity for novel behavior is precisely what distinguishes a skilled agent from a highly reliable automaton (Pavese 2023, p. 592).

Rather than providing an exegesis of the rich literature on the philosophy of skill, I pick out two properties of skill I consider relevant: agential control and habituation,⁴ as these properties will be important later for understanding the structural dimensions of capacity-deskilling.

3.1 Agential control

To be skilled in something is to have agential control over it (Shepherd 2021), the ability to adjust the execution of the action as the performance unfolds. As agency develops, the balance between automaticity and control shifts through a process of tuning and attuning (Sherman 1991, p. 175).. Thus, to cultivate a skill is to develop flexible links between thought and action, such that continue to evolve even after expertise has been achieved (Vigani 2021).

Skills are therefore distinct from habits in an important respect: skills require control over *initiating* the task, whereas mere habits do not: they are elicited in response to environmental triggers (Pavese 2024). A habitual activity is an activity that is triggered by something in the environment, where the agent performs the activity in response to the trigger. A skilled activity is one where the agent has full control over initiating the activity. The implication is that an agent who performs out of habit, in response to triggers in

the environment, is less skilled compared to the agent who initiated the activity, who has full agential control over it.

3.2 Habituation

The process of becoming skilled is a gradual and slow process of *habituation*. Skill generally cannot be learnt through testimony, and requires practice. As we move through stages of skill development, from novice through to expert, we rely on concrete experience rather than abstract principles (Dreyfus and Dreyfus 1986). Habituation involves inculcation through continued attunement to the demands of individual cases; refinement of the action through successive trials and learning from mistakes (Ryle 2009, p. 234). This includes encountering and overcoming disappointment and failure.⁵ Through habituation, we gain experience, which requires considerable time and living (Sherman 1991).

Habituation has two important components: intersubjectivity and embodiment. I take them in turn.

3.2.1 Intersubjectivity

Skill habituation has a 'shared understanding' dimension. First, becoming skilled typically requires a mentor to guide the novice into becoming an expert that engages in practice of the right kind. A good mentor is typically necessary in order to cultivate the proper *motivational structure* for becoming skilled (Sherman 1991). In the relationship between mentor and novice, the mentor instills not only the mechanical aspects of the skill, but, as importantly, the *value* of the skill, such that in the process of habituation, the skilled agent eventually comes to view the value of the skill as their own. In the relationship, both mentor and mentee have shared intentions and they come to have "shared valuing" of the skill. This is an intersubjective mentor–mentee relationship.

The process of habituation is intersubjective also on a societal scale. Social practices performed by groups or whole societies can provide a mentoring environment: the family, colleagues, citizens and so on. Indeed, as Kate Manne argues, social practices are often what generates moral reasons for action, as they have a constitutive aim or telos, which is ultimately the aim of helping people fare well

³ For example, Nancy Sherman (1991) uses the metaphor to examine how persons develop moral virtue.

⁴ Cf Schuster and Lazar (2024) who posit three constituents of skill: having success conditions; developed through practice guided by those success conditions; shaped by the tools used in its exercise. My account differs from theirs in that it puts much more emphasis on the habituation process.

⁵ Dreyfus and Dreyfus (1986) characterize the process of becoming skilled as a process comprised of 5-stages: novice, advanced beginner, competent, proficient and expert. Rule-following is more dominant in the early stages, while gradually the agent exercises more sophistication and sensitivity to varying contexts, so that the later stages are characterized by rapid, fluid, involved engagement. On this account, while becoming skilled in the capacities requires successive attempts at an activity, the process itself is one of growth rather than simple repetition. This point will become relevant for my discussion automation and repetition later on.

rather than badly (Manne 2013, p. 69). Alasdair MacIntyre noted that learning is carried out through the practices of the common life: becoming the kind of agent who is responsive to whatever it is that skillful action requires in order to achieve some common good, “has to be learned from teachers and exemplars who know how to communicate this kind of responsiveness to others through the habits of a *common life*” (MacIntyre 2017 emphasis added). As such, social practices play an important role in the formation of shared valuing regarding both the outputs (the product of the skilled activity) and the skill itself.

That skill has an intersubjective dimension is important for the prospects of capacity development and exercise: to cultivate a capacity requires learning from others, and a shared understanding that the capacity is in itself valuable.

3.2.2 Embodiment

Successful habituation must include an embodied component. According to the Enactive Approach in cognition, embodied ‘participatory sense-making’ (De Jaegher and Di Paolo 2007) is constitutive of social cognition (Di Paolo and Thompson 2014), for example in paying moral attention to the other, because the process of habituation of moral attention requires perceiving the other as an embodied being, who is similar to ourselves, through their voice and their gaze. One comes to apprehend other perceived bodies as being similar loci of embodied subjectivity, creating ‘embodied understanding’ (Gallagher and Zahavi 2013). The social interaction itself and its emergent dynamics constitute an essential part of the cognitive process that is unfolding between the participants: “meaning is generated and transformed in the interplay between the unfolding interaction process and the individuals engaged in it” (De Jaegher and Di Paolo 2007, p. 485).

Crucially, without being embodied in a social interaction—without the return of the gaze of the other—we would be hard-pressed to recognize them *spontaneously* as moral agents. The spontaneity of recognizing the other as moral agent is inherent in the embodied interaction, whereas disembodied interaction requires the agent to exercise more mental efforts in attributing moral agency to the other (Marin 2024, p. 227). Similarly, for the good of friendship, disembodied interactions make it harder to form meaningful friendship because technologically mediated virtue friendships require more discipline for a diminished experience (Grasso 2025).

Therefore, being embodied is critical for becoming skilled insofar as the process of ‘shared valuing’ of the skill requires some degree of *doing* something with the mentor(s), creating a shared history and a participatory sense-making regarding both the product of the skill and the skill itself. While shared valuing might be achieved in the absence of

any embodied interaction (say if the only form of interaction is mediated by online platforms), it would be much more difficult to achieve, compared to interactions that involve a significant level of embodiment.

Putting the different characteristics of skill together, I propose the following framework for conceptualizing capacity cultivation as skilling:

Being capacity-skilled is having agential control over the human capacities, obtained in a process of intersubjective, embodied habituation, manifested in competent exercise of the human capacities, as well as in societal shared valuing of the capacities.

So far I have reflected on the constituents of capacity-as-skill. To show how agents may become capacity-deskilled, I now introduce the structural dimension.

4 Structural dimension of capacity impoverishment

On the capacity-skilling framework above, deskilling would manifest in the impoverishment of the capacities: limited agential control over the capacities, as well as a diminished process of habituation and a devaluation of the capacities. Prima facie, capacity-deskilling may look like an individual vice: on a virtue ethics approach, I have a duty to myself to cultivate my human capacities, and if I fail to do so, I am failing in the arts of personhood. However, framing the problem of technology-induced capacity-deskilling as an individual ethics problem does not capture the full picture of deskilling. An important dimension of Aristotelian virtues ethics is that a well-ordered polis is a necessary part of human flourishing. Framing the response to the threats to our humanness as an individual responsibility problem therefore misses the structural dimensions of the deskilling problem.

The structural dimension of capacity-deskilling can be captured by the concept of affordances. The environment one inhabits is a ‘world of affordances’ (Zimmermann et al. 2025). Affordances are action possibilities formed by a relationship between an agent and its environment offers to persons (Gibson 1979; Nye and Silverman 2012, p. 179), that are not mere opportunities but invitations that can have a severe influence on the behavior agents will exhibit in that environment (Withagen et al. 2012). The concept of affordances helps explain how the environment is constitutive of flourishing: the proper cultivation of the core capacities is dependent on inhabiting an environment that affords action possibilities for activities, in turn triggering the capacities associated with these activities. For example, a library affords the opportunity for reading books, thereby triggering the epistemic and social capacities involved in reading.

In the context of digital technologies, the integration of AI into healthcare, for example, affords the possibility of offloading tasks to the machine, in turn likely leading to automation bias, where the user attributes authority or excessive trust to an AI tool over advice from other sources. This risks the loss of the user's critical skills necessary for unexpected scenarios such as a system malfunction (Natali et al. 2025). Another example is the smartphone: it provides possibilities for interaction that are disconnected from other persons, such as swiping, clicking and scrolling. This in turn invites determinate, separate, and repeatable tasks. This fractures the user's affordance space into disconnected fields, ultimately leading to alienation, rather than offering opportunities for the development of new cognitive and bodily skills that are mutually informing and enriching (Butler 2024). The problem depicted in these examples is that the affordance space is an impoverished environment: an environment that provides a narrow field of affordances (more on that later), limiting the opportunities to engage in activities that require the competent exercise of our capacities.

As I stressed earlier, learning to value a capacity is part of the process of its habituation. When a person develops a capacity in a shallow way, they cannot learn to value the more masterful levels of the capacity, since learning to value the capacity is a result of an ongoing, gradual process of habituation, which they haven't gone through. The danger is that when a person does not value the state of being competently skilled in the human capacities, they will not be inclined to *become* properly skilled, thereby entrenching a shallow level of the capacities. On a societal scale, when enough persons do not value excellence or mastery of a human capacity (e.g., critical thinking), the capacity is in danger of atrophying. This then provides another justification for the structural approach: we cannot expect agents to be what Mitcham (2024) calls 'virtuous superheroes' and properly cultivate their human capacities if their social environment does not value the full maturation of these capacities. First, the agent would not have the proper mentoring of the 'know-how', and second, they would not obtain the shared understanding regarding the value of the capacities.

4.1 Capacity-hostile environments

To distinguish between the types of environments according to how favorable they are to capacity cultivation, I propose to distinguish between *capacity-hostile* environments and *capacity-conducive* environments. *Capacity-hostile* environments restrict, limit or create a narrow field of affordances for capacity development and exercise. These environments will afford opportunities for shallow skilling in the capacities, and restrict opportunities for a fuller, well-rounded development and exercise of capacities. *Capacity-conducive* environments, on the other hand, are environments

that encourage the development and exercise of the human capacities. Importantly, capacity-conducive environments will afford opportunities for cultivating the capacity *to will*, by encouraging agents to develop and exercise their other capacities (recall that developing the other capacities is difficult and requires willpower). In this way, capacity-conducive environments encourage the volitional capacity without expecting agents to be virtuous superheroes.

An environment that provides inadequate conditions for the development and exercise of the capacities constitutes a narrow field of affordances: it will not provide sufficient invitations for valuable activities. The absence of invitation for activities means that the capacities involved in these activities will not be triggered (Ferdman 2024).⁶ A narrow affordance field is normatively problematic, from a structural point of view, when it leads to 'affordance shrinkage': when the reduction in action possibilities is systematic, enduring and deeply entrenched (Dokumacı 2023; Krueger 2023), such that discourages users from finding ways to compensate for the shrinkage.

Disembodied environments might constitute a narrow field of affordances for capacity cultivation. Recall the point above about the inherent difficulty of cultivating the moral and social capacities in disembodied environments since it requires more effort and self-discipline for a diminished achievement. On an individual responsibility approach, overcoming the difficulties of cultivating capacities in disembodied environments would require of individuals to become virtuous superheroes and will themselves into overcoming the difficulties. The structural dimension of skilling, on the other hand, highlights that a disembodied environment does not provide the right kind of affordances that encourage agents to develop and exercise the social capacities necessary for morality and friendship. Take social media platforms, for example. Even without malicious intent on the part of designers of these disembodied environments, agents still might find themselves as users in a position where they are being herded to behave in certain ways, because these technologies are *designed* in specific ways to keep users *engaged* with them. This is not because the users are ignorant in a sense in which they are blameworthy, but rather that users' ignorance, maintained through things like targeted adverts, filtered newsfeeds, and information

⁶ Richard Kraut mentions that the human capacities rarely operate in isolation from each other (Kraut 2007, p. 166). For example, the capacity to know involves theoretical rationality as well as practical rationality (Hurka 1993), the capacity to will ourselves to acquire knowledge (Bradford 2015), and the physical capacities, insofar as knowledge is embodied. Without developing this point in depth, to the extent that we accept Kraut's characterization, it follows that capacity-hostile environments are typically hostile to the capacities in that they undermine the opportunity to engage in activities that involve the exercise a combination of the capacities at a time.

bubbles, is utterly predictable, and, indeed, often a *feature* of their engagement with social media platforms. As such, the context in which a user's behavior occurs can better explain the behavior than appealing to the user's dispositions (Tolton 2024).

This is precisely what the structural dimension captures: a technology's affordances make certain action possibilities more likely than others, given how the environment is designed. To the extent that embodied interaction is crucial for properly developing and exercising the capacities, the environment should not only afford but encourage embodied interaction. The emphasis on encouraging rather than affording embodied interaction is, as I argue elsewhere, that merely having the opportunity to ϕ does not guarantee that the agent will actually ϕ and engage the capacities involved in ϕ -ing (Ferdman 2024). To ensure the activation of the capacities, the environment has to encourage the activity of ϕ -ing, otherwise the capacities might not be triggered. This is why the environment conducive to the capacities has to include public goods like community centers, quality public spaces, public libraries, etc. that encourage structured and unstructured embodied interactions (Sypnowich 2016). Embodied engagement in these settings would be more conducive to creating shared experiences, which are necessary for shared valuing of the social capacities.

I propose that AI tools and systems can be analyzed as components in capacity-conducive or capacity-hostile environments along the following axes: the field of affordances; the level of embodied user engagement and the level of intersubjectivity.

4.1.1 AI automation and the field of affordances

As Bainbridge noted in “Ironies of Automation” (1983), in order to be skilled in a domain, one has to have efficient retrieval of knowledge from long-term memory, which depends on frequency of use. This type of knowledge develops only through use and feedback about its effectiveness. Automating routine tasks and leaving only exception-handling to the human user, “deprives the user of the routine opportunities to practice their judgment and strengthen their cognitive musculature, leaving them atrophied and unprepared when the exceptions do arise” (Lee et al. 2025). In other words, indiscriminately automating human activity may create capacity-hostile environments, insofar as users will be deprived of the habituation process necessary for properly developing the capacities involved in routine activities.

Additionally, recall that to develop and exercise the capacities requires that the person exercise their capacity to will, since cultivating the capacities requires effort and determination. In capacity-hostile environments, the technology or tool will afford users the opportunity of *not* exercising

their willpower, for example by replacing human activity with automated activity, or by offloading human activity, such as moral reasoning, to a non-human agent such as an Artificial Moral Advisor (AMA). Automation is thus likely to entrench a narrow field of affordances when it actively discourages the user from seeking avenues to compensate for the deskilling. Insofar as the automation replaces human activity without encouraging users to seek alternative avenues to exercise their capacities, including the capacity to will, we may categorize these automation environments as capacity-hostile.

Of course, not all automation constitutes a capacity-hostile environment. Some automation may free agents from drudgery and repetition (e.g., the washing machine). However, we should note the cumulative effect of automation: automation and replacement of human activity is becoming more prevalent, across many human spheres. As such, automating multiple everyday human activities, across multiple spheres (e.g., writing, searching, planning, organizing, navigation, decision-making, entertainment, personal wellbeing), might constitute a narrow field of affordances for capacity development, where persons are discouraged from finding ways to compensate for the shrinkage, precisely because so many activities are replaced with non-human activities. The structural approach provides us with a framework to evaluate whether AI automation affords sufficient action possibilities to develop and exercise the capacities, or whether it narrows the field of affordances.

4.1.2 Embodied, intersubjective habituation

An environment conducive to the development and exercise of the capacities will afford opportunities for embodied, intersubjective habituation of the capacities. A capacity-hostile environment, on the other hand, will restrict embodied interaction or afford action possibilities that encourage persons to primarily interact through disembodied mediums (e.g., digitally mediated interaction). A capacity-hostile environment will also afford less opportunities to develop and exercise the capacities in concert with other humans. In this environment, users will be interacting with non-human agents, limiting the opportunity to develop an intersubjective valuing of the capacities. Since embodied, intersubjective habituation of the capacities is a necessary component of successfully developing and exercising the capacities, an environment with little opportunity for embodied interaction may ultimately be hostile to capacity cultivation. AI could contribute to capacity-hostile environments, or it could provide opportunities for capacity cultivation, depending on the field of affordances it creates. To demonstrate this sort of analysis, I next examine a fast-growing use case of AI: Artificial Personal Assistants.

5 Artificial personal assistants and capacity cultivation

Large language models (LLMs), such as ChatGPT, Claude, ai, Gemini, etc., are increasingly being used as artificial personal assistants (APAs). A recent report on types of uses of LLMs has found that the second and third highest uses of LLMs were for life-planning: ‘organizing my life’ and ‘finding purpose’ respectively: it found that users utilize the models to help them be more aware of their intentions (e.g., daily habits, New Year’s resolutions, introspective insights) and find small, easy ways of getting started with them. Additionally, users reported that they rely on APAs to determine and define their values, get past roadblocks, and take steps to self-develop: seeking advice on what they should do next, reframing a problem and helping them to stay focused (Zao-Sanders 2025). Another study shows that heavy users utilize LLMs for social validation, self-regulation, and interpersonal guidance (Kim et al. 2025a). Sam Altman, CEO of OpenAI, envisions a future where “eventually we can each have a personal AI team, full of virtual experts in different areas, working together to create almost anything we can imagine” (Altman 2024).

APAs are attractive because they adapt to the individual user’s needs, tailor their advice based on the user’s responses, habits and preferences, and are willing to listen to the user without judgment (Ma et al. 2023). APAs are constantly available, constituting “an always-on sounding board with whatever subject-matter expertise you need, configured with your preferred tone of interaction—anything from strict taskmaster to encouraging cheerleader” (Samuel 2025).

While proponents of using LLMs as personal assistant might argue that these tools provide benefit to the users, there is emerging evidence that becoming overly reliant on such tools might be problematic in terms of capacity-skilling. One of the features of LLMs is that they are ‘trained to do everything for everyone at every environment’ (Gebru and Torres 2023). This ostensibly gives LLMs the ability to do everything for the user, from menial tasks to planning one’s life. If users become too dependent on LLMs for things like organizing their life and finding meaning, the APA might contribute to a capacity-hostile environment. There is already tentative evidence that young people growing up with the smartphone struggle with “everyday but essential” skills like empathy, time management, speaking to other people, problem-solving and critical thinking (Halliday 2025). LLMs could potentially exacerbate this deskilling, as follows.

5.1 Life-planning

To determine whether APAs might contribute to capacity-hostile or capacity-conducive environments, we first need to unpack the meaning of ‘life-planning’. On an instrumentalist view, ‘life-planning’ is a tool to help a person achieve their goals and aims. A different approach, however, holds that life-planning is a self-constituting activity. Heyd and Miller (2010), for example, argue that the identity of the person is, at least partly, *constituted* by their life plans, and hence that their overall long-term flourishing cannot be even ideally considered as relating to an independently identifiable agent who tries to maximize their own good. Additionally, the activity of life-planning relates to the kind of person one decides they want to be, rather than to the aims or goals one wants to maximize. As such, planning one’s life is about developing a particular character or dispositions, rather than following an overall program of predetermined activities (p. 27). On this non-instrumentalist view, life-planning should not be outsourced to a tool.

On a different vein, in theorizing about agency, Michael Bratman argues we are essentially planning creatures, and that planning is a manifestation of practical rationality (Bratman 1983). Agency includes future-directed intentions, which play two important agential roles in a person’s deliberation: intentions have a settling function and a coordination function. The settling function manages the finite cognitive resources that humans have when facing a choice between options. The coordination function forms and organizes different types of plans, policies, and more specific future-directed intentions, imposing coherence constraints on one’s intentions in order to allow for the incompleteness of information and enable flexibility. What is more, plans facilitate coordination both of a person’s activities over time, and of the person’s activities with other persons, enabling us to achieve complex personal and social goals we would not otherwise be able to achieve (p. 272). Finally, planning involves responding to unexpected events and surprises, that could be sources of value and a sense of growth in one’s life (Heyd and Miller 2010, p. 32).

From the perspective of capacity-as-skill, planning and organizing one’s life requires a combination of human capacities. First, following Bratman, it involves the epistemic capacities. This manifests in two ways: in exercising theoretical rationality manifested in self-deliberation on one’s values and goals, and in exercising practical wisdom regarding coordinating one’s goals and intentions with one’s actions. Practical wisdom is especially relevant for planning one’s life. A wise person knows the effective means to one’s ends and possesses an accurate understanding of their situation and circumstances, especially in assessing the testimony of others; they also possess discernment: correct understanding of how the good and the bad apply

in particular circumstances. To be well motivated is not enough, to have practical wisdom is to be an effective and savvy agent in the real world, equipped with valuable life experience (Hursthouse 2022). Moreover, practical wisdom is necessary for critically reflecting on one's conceptions of a flourishing life, in order to set appropriate goals and strive to accomplish them in the right way (Stichter 2021). In short, life-planning requires the cultivation and exercise of the epistemic capacities.

Life-planning also involves the social capacities. First, planning involves coordinating with others, toward achieving complex goals. Second, recall that on the intersubjective approach, our reasons for action are (at least in part) drawn from our social relationships. As such, planning one's life involves the moral and social capacities, insofar as our life goals (e.g., being a good parent, help combat climate change) are similarly drawn from our social relationships. Planning one's life also involves the capacity for creativity, insofar as planning is always incomplete, and executing the plan requires the ability to reconsider and respond to novel situations, surprises and unexpected events. Finally, life-planning requires the capacity to will, that is, to develop and exercise the aforementioned capacities, as well as to resist the temptation to outsource the activity of planning to an external entity or tool.

5.2 APAs as capacity-hostile environments

From the perspective of capacities as skills, APAs might entrench a capacity-hostile environment, if persons become overly dependent on it for planning and organizing their lives. I readily acknowledge that using APAs to figure out one's daily schedule or using the APA as a sounding board for life goals does not in itself constitute a capacity-hostile environment. If a person offloads planning activities to the APA in a way that enables them to develop and exercise their capacities in other activities, then APAs would not constitute a capacity-hostile environment.

However, an APA affords much more than using it as a sophisticated calendar or occasional recommender system. APAs are a use-case of LLMs that can be integrated with other AI tools because LLMs are ultimately 'trained to do everything for everyone at every environment' (Geburu and Torres 2023). This makes them not only social companion, moral advisor, life coach and personal butler, but all of them combined in one. As such, APAs afford the replacement of many of the activities involved in both planning *and* executing one's plans, that would otherwise require the proper cultivation and competent exercise of our epistemic, social, moral and creative capacities. It is the scale and extent of the APA as all-encompassing agent that is the cause for worry: the leveraging of AI capabilities, such as dynamic reasoning, real-time adaptability, and multi-scale operational

coordination (Floridi 2025) to merge synergistically discrete uses of the APA in ways that greatly magnify their scope and power to alter lives (Vallor 2016, p. 27). The more the APA does things for the user, the more it could constitute a capacity-hostile environment, as follows.

Let us first examine APAs as a field of affordances. Recall that a narrow field of affordances not only restricts the number of action possibilities, but also discourages the user from seeking alternative avenues to compensate for the shrinkage in capacity exercise. APAs are constantly available, designed to be helpful, always prompting the user to see if they could be of further assistance. As such, they might discourage the user from seeking out new opportunities for growth and exploration (Gabriel et al. 2024). When the APA replaces most of the user's planning activities, it reduces the user's need to engage in routine, small-scale decision-making about organizing daily activities, or small-scale decisions that over time amount to broader life goals. As mentioned above, automation of routine tasks might be problematic since it deprives the user of the opportunity to "practice their judgment and strengthen their cognitive musculature" for activating in exceptional circumstances (Lee et al. 2025). Overreliance on APAs could lead to a reduced sense of personal competence (Gabriel et al. 2024), which could be explained by the offloading of small-scale decisions to the model, such as "help me plan my day" and "how should I fill an unscheduled hour between meetings",⁷ undermining the habituation of the capacities necessary for trusting one's decisions on life-planning.

Additionally, over-reliance on APAs, may, over time, acclimate the user to merely responding to triggers in the environment (i.e., the APAs outputs) rather than exercising agential control over the planning. One of the affordances that APAs provide is the invitation to become a plan-follower rather than a plan-former. Take Sunsama, a scheduling app: "Each morning, Sunsama displays tasks synced from a variety of software, alongside your calendar and unfinished tasks from yesterday. It prompts you to prioritize today's work and to link tasks to your objectives—your largest goals for the week. Sunsama then schedules tasks logically into your next available time slot, intelligently rescheduling them when other tasks run into overtime. The app reminds you throughout the day when it's time to go to a meeting or when you should take a quick break. And when you're done for the

⁷ This might strike the reader as a caricature, but anecdotal evidence shows that LLMs are being used in this way (for example, bloggers reporting on their 'quest for ultimate productivity' <https://forem.juliang.org/svilupp/the-quest-for-ultimate-productivity-building-an-llm-powered-assistant-21nk> and <https://python.plainenglish.io/how-i-built-a-personal-ai-agent-to-automate-my-everyday-tasks-367c512a7558>).

day, it prompts you to write a journal entry about what you accomplished” (Guay 2025).

One of the features of being skilled (Sect. 3) is having agential control over initiating a task, whereas fulfilling the task in response to a trigger from the environment is more like a habit than a skill (Pavese 2024). Now let us consider the distinction between forming plans that involves self-deliberation, and following plans constructed by APAs. Forming a plan by self-deliberation requires exercising control over all the steps of the activity of forming the plan, including over initiating the two functions of planning: settling (deciding between options) and coordinating. As such, forming a plan and executing it by settling and coordinating, requires exercising the epistemic, moral, social and creative capacities involved in planning and executing the plan, as well as the capacity to will oneself into initiating the activity of forming a plan to begin with.

Relying on an APA to form our plans for us is more like a habit that is developed in response to a trigger in the environment. To illustrate, a user might begin with asking the APA to help them form a plan for finding a soulmate, find a career, resolve a conflict with a friend, or structure their daily schedule. The APA will propose a plan and help the user refine it, according to their personal preferences or past behavior, and the user will then follow the plan. In this sequence of events, the user is not initiating the two functions of planning: settling and coordinating, but merely responding to the APAs recommendations. Moreover, in plan-forming, the capacity to will is involved in initiating the plan, whereas in plan-following, it is absent, because the user is responding to a trigger in the environment—the APA output. When much of the planning process is offloaded to the APA, it arguably requires a shallower exercise of the capacities that are ordinarily involved in making coherent plans and executing them. Moreover, making plans involves taking responsibility if the plan fails, affording an opportunity to be self-critical and to improve accordingly. Yet if an APA makes a plan that ends up failing, the user could more readily shift the responsibility to the APA, forgoing the opportunity for self-growth.

On a similar vein, Schuster and Lazar (2024) argue that responding to some content because it has popped up in my feed is in itself morally problematic, as it is not as good as attending to something for my own moral reasons. Likewise, following a plan because it was suggested by an automated tool, even if it is a good plan and fully coheres with my values and preferences, is not as good, in developmental-perfectionist terms, as forming my own plan (see also Zelny 2025 on this point). Plan-forming requires a more in-depth utilization of my capacities, including my capacity to will myself into making a plan, following it and knowing when to revise it. Recall that having practical wisdom, as part of the epistemic capacities, is to be an effective and savvy agent

in the real world, equipped with valuable life experience, whereas merely being well motivated is not sufficient for practical wisdom (Hursthouse 2022). Constantly turning to APAs for advice, even if the advice is good, and the user is well motivated to form valuable goals in their life, restricts the user’s real-world experience and undermines the opportunity to develop practical wisdom.

Furthermore, APAs are always there for the user, inviting the user to constantly return to the APA for answers. Long-term reliance on the APAs as a life coach/personal assistant would be problematic from a developmental-perfectionist standpoint, as its convenience and constant availability may limit agents’ ability to transition to develop the capacities on their own. Unlike a good human moral exemplar, the APA, by its purpose of giving advice, would not know to remove itself as the training wheels and resist being called upon to provide a model for every decision, large or small (Liu et al. (2022)). In this way, the APA is narrowing the field of affordances for gaining experience in becoming an independent planner.

APAs could also contribute to capacity-hostile environments by discouraging intersubjective habituation of the capacities. As APAs become more integrated into everyday systems, and replace more and more human activity, they afford an environment where the user increasingly interacts with non-human entities for the planning and execution of their daily activities. AI mediation might come at the expense of human interaction, decreasing the opportunity to develop our capacities in concert with other humans, and create the shared valuing that is intrinsic to being capacity-skilled. While the development and exercise of human capacities depend on embodied, intersubjective interactions with others, APAs decrease embodied, intersubjective interaction involved in life-planning activities.

In embodied interaction, all the constant “living together” which goes on in between and around physical, embodied activities with other persons (Grasso 2025) better allow us to test and grow our human capacities, compared to virtual environments. Preferring the advice of an APA over that of a friend or colleague may contribute to undermining the shared experience between persons. Ideally, when I seek the advice of a friend, we are creating a space of shared experiences, where the friend provides good advice *because they care*. Further, our relationship is deepened because we share the experience of knowing one another and coming to trust one another by seeking and giving advice. Thus the shared experience, exercised via the social capacities, contributes to the intrinsic good of friendship. In the context of APAs, however, even if after prolonged use the APA gets to “know” the user and will respond with outputs that *sound* caring, it nevertheless cannot have shared experiences that are necessary for mutual care, thus undermining the opportunity

to exercise the social capacities, as well as the epistemic capacities involved in learning from others.

In sum, APAs might constitute a capacity-hostile environment insofar as they afford users the possibility of offloading life-planning to the LLM in ways that discourage the user from finding ways to compensate for the loss of capacities ordinarily involved in planning.

5.3 APAs in capacity-conducive environments

APAs could be an interim step on the path to capacity cultivation, so long as they provide opportunities to develop and exercise human capacities in embodied, intersubjective interactions with others, within an adequate field of affordances. For example, APAs could incorporate serendipity and agential control such that encourage the development and exercise of the capacities, as follows:

5.3.1 Serendipity

One feature of APAs is that they can be used as a sophisticated recommender system: the APA recommends a course of action in response to the user's prompts, preferences and past behavior. Existing recommender systems are like vending machines: they prioritize efficiency and predictability of the outputs, pigeonholing users to contents they are already familiar with. In contrast, recommender systems could be designed more like a tool box, rather than vending machine. A tool box recommender system could support exploration and unexpected discoveries, by deliberately introducing elements of surprise and challenge into the user experience through productive friction. One of the main differences between the vending machine and tool box models is that the latter explicitly designs for serendipity: the art of discovering new, advantageous things through chance encounters with unexpected information (Danzico 2010; Reviglio 2019; Campbell 2023; Kim et al. 2025b). As such, serendipitous recommender systems are designed for novelty and unexpectedness (Kotkov 2024), emphasizing exploration over efficiency and predictability (Chen 2022).

Because serendipity encourages exploration, serendipity-incorporated APAs could contribute to capacity-conducive environments. First, let us consider the user's state of mind: to take advantage of the serendipitous event, the agent needs to actively put themselves in a position of "investigative ignorance"—recognizing and exploiting the state of ignorance to come up with new and more satisfying information (Arfini et al. 2020). Putting oneself in a position of investigative ignorance involves exercising the capacity to will—to be willing to step outside of one's comfort zone by considering unexpected recommendations. Additionally, to investigate, according to the Oxford English Dictionary, is to seek to understand something through examination, study,

or analysis; to examine, research, inspect. These activities require action, in contrast to merely receiving and accepting or rejecting a recommendation. Therefore, to put oneself in an *investigative* position is to take an agential, rather than passive, role.

Second, serendipity-incorporated recommender systems might encourage the volitional capacity by encouraging *cognitive patience*. Cognitive patience is a willingness to linger in difficulty. We risk losing cognitive patience when we replace difficult tasks with simpler ones. For example, replacing deep reading with skim reading on screens undermines cognitive patience by undercutting the conditions under which deep thought becomes possible (Hendrick 2025). Similarly, vending machine recommender systems suppress unfamiliar, surprising or even uncomfortable recommendations, undermining cognitive patience. With serendipitous recommendations, however, the user's capacity to will would be exercised: when we are open to put ourselves in a position of investigative ignorance, we are exercising our will to linger in difficulty by accepting to consider something unpredictable, unexpected, surprising or uncomfortable. In this way, we are also cultivating the tolerance for the conditions under which complexity—unpredictable, unexpected, surprising or uncomfortable recommendations—becomes possible.

Finally, recall that successful skilling requires exercising the skill in routine, small-scale tasks, in order to competently exercise judgment when larger decisions are necessary. Serendipity-incorporated recommender systems that includes surprising results might provide the appropriate environment for dealing with small-scale unexpected events, in preparation for the event of having to make life-changing decisions, as well being able to plan one's life without depending on the APA.

5.3.2 Cultivating agential control

As I have argued throughout, APAs used in excess are problematic from the point of view of Developmental Perfectionism, as they could narrow the field of affordances for capacity cultivation. One way to avoid excess is to ensure that the APAs are interim tools, that help the user become an independent planner. For example, the 'always on, always there for you' feature could be replaced with a design that explicitly and transparently guides the user to disengage from the tool, instead encouraging them to engage in embodied, intersubjective interactions. In other words, designing APAs to help persons cultivate agential control over planning their life, including the ability and motivation to initiate the activity of planning itself. Learning to become independent from the APA would enable the user to develop and exercise their epistemic, creative and social capacities, as well as their capacity to will, manifest in initiating the activity of

planning and exercising cognitive patience in responding to unexpected, unforeseen and uncomfortable events.

Interestingly, LLMs users are free to ask the model to provide serendipitous recommendations, or help them to become less dependent on the APA. In the language of affordances, these options are action possibilities that exist in the user's field of affordances. However, only virtuous superheroes are likely using APAs this way, while ordinary persons, inexperienced in the activity of life-planning (e.g., people growing up with LLMs) might end up over-relying on them for life-planning. From a structural point of view, if we want users to engage with serendipitous options and learn to become independent life planners, designers should purposefully work to avoid LLMs potentially becoming replacements for life-planning, and so it is important that these models are not marketed as omnipotent personal assistants. Instead, they should be framed as exercises—tools to practice planning—from mundane tasks to life goals. Clear boundaries should be established at the outset, with the user informed that the purpose is to cultivate their capacities. Features like having time limits on interactions, a focus on improving the user's real-world embodied experience and the provision of feedback and practical tips that can be applied outside the virtual space, can further reinforce this goal.

Importantly, prompting an LLM for serendipitous recommendations may not necessarily guarantee serendipitous recommendations. Whether LLMs can reliably follow prompts for serendipity depends on their training and fine-tuning. Given that these models are currently optimized for overall user satisfaction (rather than serendipity), it is unclear whether over repeated use they will be able to overcome their training and actually prioritize serendipity over 'vending-machine'/popular recommendations. More generally, this provides reason to explicitly design and train recommender systems to optimize for serendipity rather than user satisfaction and the popularity of the suggested output.

6 Conclusions

According to the Developmental Perfectionism view, capacity impoverishment is robustly bad as it leads to impoverished lives. As such, protecting humans against capacity-deskilling is a moral obligation we have to others, to ensure capacity-conducive environments for all. Currently, the more AI and automation uncritically replace human activity and interaction, the greater the risk that social interactions will be primarily mediated by a disembodied, intersubjective agent, undermining capacity cultivation. One of the implications is that to ensure proper cultivation of the capacities, environments must encourage (not merely provide)

intersubjective, embodied interactions to compensate for AI-mediated, disembodied environments.

More broadly, the capacity-skill framework provides a useful starting point from which to evaluate the goodness or badness of AI tools and systems, according to their conduciveness or hostility to capacity cultivation. Given that AI applications are not merely tools but socio-technical systems (Kudina and van de Poel 2024), fulfilling the moral obligation to ensure capacity-conducive environments in the age of AI requires paying attention not only to the technical or design aspects of AI applications, but to the institutional, cultural, and governance contexts in which they are embedded. Insofar as socio-technical environments are value-laden, a Developmental Perfectionist approach provides a useful framework for critically reflecting on the values pertaining to capacity cultivation embedded in AI applications and systems. This is because we owe it to others to ensure that as humans, we invest in the shared valuing of the human capacities, otherwise we risk losing the 'human' in us.

Acknowledgements I am grateful to participants of the Jülich Speaker Series in the Philosophy of Technology for helpful comments on an earlier version of this paper, as well as to two anonymous reviewers for helpful comments.

Author contributions Avigail Ferdman has authored all the manuscript sections.

Funding Open access funding provided by Technion - Israel Institute of Technology. The author did not receive funding from any organization for the submitted work.

Data availability No datasets were generated or analyzed during the current study.

Declarations

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Altman S (2024) The Intelligence Age. <https://ia.samaltman.com/>. Accessed 24 Apr 2025

- Arfini S, Bertolotti T, Magnani L (2020) The antinomies of serendipity how to cognitively frame serendipity for scientific discoveries. *Topoi* 39:939–948. <https://doi.org/10.1007/s11245-018-9571-3>
- Aylsworth T, Castro C (2024) Should i use ChatGPT to write my papers? *Philos Technol* 37:117. <https://doi.org/10.1007/s13347-024-00809-w>
- Bainbridge L (1983) Ironies of automation. *Automatica* 19:775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- Bonicalzi S, De Caro M, Giovanola B (2023) Artificial intelligence and autonomy: on the ethical dimension of recommender systems. *Topoi* 42:819–832. <https://doi.org/10.1007/s11245-023-09922-5>
- Bradford G (2015) Achievement. Oxford University Press, Oxford
- Bradford G (2016) Perfectionism. In: Fletcher G (ed) *The Routledge Handbook of Philosophy of Well-being*. Routledge, London, pp 124–134
- Bradford G (2021) Perfectionist bads. *Philos Q* 71:586–604. <https://doi.org/10.1093/pq/pqaa055>
- Bradford G (2025) Failure. In: Rossi M, Tappolet C (eds) *Ill-Being: Philosophical Perspectives*. Oxford University Press, Oxford
- Bratman M (1983) Taking plans seriously. *Soc Theory Pract* 9:271–287
- Budzyń K, Romańczyk M, Kitala D et al (2025) Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study. *Lancet Gastroenterol Hepatol* 10:896–903. [https://doi.org/10.1016/S2468-1253\(25\)00133-5](https://doi.org/10.1016/S2468-1253(25)00133-5)
- Butler M (2024) Fracturing the affordance space: an account of digitalized alienation. *Front Psychiatry*. <https://doi.org/10.3389/fpsy.2024.1407586>
- Boden M (2004) *The Creative Mind: Myths and Mechanisms*, 2nd edn. Routledge, London
- Campbell DR (2023) In defense of (some) online echo chambers. *Ethics Inf Technol* 25:45. <https://doi.org/10.1007/s10676-023-09715-9>
- Chen M (2022) Exploration in Recommender Systems. In: *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. Association for Computing Machinery, New York, NY, USA, p 1642
- Chiang T (2023) ChatGPT Is a Blurry JPEG of the Web. *The New Yorker*. <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>. Accessed 10 Feb 2023
- Coeckelbergh M (2019) Skillful coping with and through technologies. *AI & Soc* 34:269–287. <https://doi.org/10.1007/s00146-018-0810-3>
- Danzico L (2010) The design of serendipity is not by chance. *Interact* 17:16–18. <https://doi.org/10.1145/1836216.1836220>
- De Jaegher H, Di Paolo E (2007) Participatory sense-making. *Phenomenol Cogn Sci* 6:485–507. <https://doi.org/10.1007/s11097-007-9076-9>
- Di Paolo E, Thompson E (2014) *The enactive approach*. The Routledge handbook of embodied cognition. Routledge/Taylor & Francis Group, New York, NY, USA, pp 68–78
- Dokumacı A (2023) *Activist Affordances: How Disabled People Improve More Habitable Worlds*. Duke University Press
- Dreyfus HL, Dreyfus SE (1986) *Mind over machine : the power of human intuition and expertise in the era of the computer*. Free Press, New York
- Ferdman A (2019) A Perfectionist Basic Structure. *Philosophy & Social Criticism* 45:862–882. <https://doi.org/10.1177/0191453718820891>
- Ferdman A (2024) Human Flourishing and Technology Affordances. *Philosophy & Technology* 37:1. <https://doi.org/10.1007/s13347-023-00686-9>
- Fletcher G (2016) Objective list theories. In: Fletcher G (ed) *The Routledge Handbook of Philosophy of Well-Being*. Routledge, London, pp 148–160
- Floridi L (2025) AI as agency without intelligence: on artificial intelligence as a new form of artificial agency and the multiple realisability of agency thesis. *Philos Technol* 38:30. <https://doi.org/10.1007/s13347-025-00858-9>
- Fricker E (2021) Should we worry about silicone chip technology deskilling us? *R Inst Philos Suppl* 89:131–152. <https://doi.org/10.1017/S1358246121000138>
- Gaut B (2010) The Philosophy of Creativity. *Philosophy Compass* 5(12):1034–1046. <https://doi.org/10.1111/j.1747-9991.2010.00351.x>
- Gabriel I, Manzini A, Keeling G, et al (2024) *The Ethics of Advanced AI Assistants*. <https://arxiv.org/abs/2404.16244>
- Gallagher S, Zahavi D (2013) *The Phenomenological Mind*, 2nd edn. Routledge, London
- Gebru T, Torres ÉP (2023) Eugenics and the Promise of Utopia through AGI. *IEEE SaTML 2023 - 1st Conference on Secure and Trustworthy Machine Learning*
- Gerlich M (2025) Ai tools in society: impacts on cognitive offloading and the future of critical thinking. *Soc*. <https://doi.org/10.3390/soc15010006>
- Gibson JJ (1979) *The ecological approach to visual perception*. Houghton Mifflin, Boston
- Grasso D (2025) Disembodied friendship: virtual friends and the tendencies of technologically mediated friendship. *Ethics Inf Technol* 27:13. <https://doi.org/10.1007/s10676-025-09822-9>
- Guay M (2025) AI Can't Plan Your Whole Day—Not Yet—But We Found One App That Gets Close. *Wirecutter*. <https://www.nytimes.com/wirecutter/reviews/best-ai-scheduling-apps/>. Accessed 2 May 2025
- Hacker-Wright J (2009) Human nature, personhood, and ethical naturalism. *Philosophy* 84:413–427
- Halliday J (2025) Gen Z students in Manchester to learn 'soft skills' such as empathy and time management. *The Guardian*. <https://www.theguardian.com/society/2025/apr/21/gen-z-students-in-manchester-to-learn-soft-skills-such-as-empathy-and-time-management>. Accessed 25 April 2025
- Hendrick C (2025) Ultra-Processed Minds: The End of Deep Reading and What It Costs Us. In: *The Learning Dispatch*. <https://carlhendrick.substack.com/p/ultra-processed-minds-the-end-of>. Accessed 29 Apr 2025
- Heyd D, Miller FG (2010) Life plans: do they give meaning to our lives? *Monist* 93:17–37
- Hurka T (1993) *Perfectionism*. Oxford University Press, New York
- Hursthouse R (2022) *Practical Wisdom: A Mundane Account*. In: Annas J, Reid J (eds) *Virtue and Action: Selected Papers*. Oxford University Press, Oxford
- Kauppinen A (2025) Telic Perfectionism and the Badness of Pain. In: Rossi M, Tappolet C (eds) *Ill-Being: Philosophical Perspectives*. Oxford University Press, Oxford
- Kim Y, Oh S, Noh C et al (2025b) Design of a serendipity-incorporated recommender system. *Electronics*. <https://doi.org/10.3390/electronics14040821>
- Kim E, Choe K, Yoo M, et al (2025a) Beyond Tools: Understanding How Heavy Users Integrate LLMs into Everyday Tasks and Decision-Making. <http://arxiv.org/abs/2502.15395>
- Kotkov D (2024) Overview of Serendipity in Recommender Systems. In: Stefanidis K, Systä K, Matera M et al (eds) *Web Engineering*. Springer Nature Switzerland, Cham, pp 453–457
- Kraut R (2007) *What Is Good and Why: The Ethics of Well-Being*. Harvard University Press, Cambridge, Mass
- Kraut R (2013) Human diversity and the nature of well-being: reflections on Sumner's methodology. *Res Philos* 90:307–322. <https://doi.org/10.11612/resphil.2013.90.3.1>
- Kremer M (2020) Gilbert Ryle on Skill as Knowledge-how. In: Fridland E, Pavese C (eds) *The Routledge Handbook of Philosophy of Skill and Expertise*. pp 100–112
- Krueger J (2023) Affordances and spatial agency in psychopathology. *Philos Psychol*. <https://doi.org/10.1080/09515089.2023.2243975>

- Kudina O, de Boer B (2024) Large language models, politics, and the functionalization of language. *AI Ethics*. <https://doi.org/10.1007/s43681-024-00564-w>
- Kudina O, van de Poel I (2024) A sociotechnical system perspective on AI. *Minds Mach* 34:21. <https://doi.org/10.1007/s11023-024-09680-2>
- Lee H-P (Hank), Sarkar A, Tankelevitch L, et al (2025) The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of KnowledgeWorkers. ACM, New York, NY, Yokohama, Japan
- Liu Y, Moore A, Webb J, Vallor S (2022) Artificial moral advisors: A new perspective from moral psychology. In: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society. pp 436–445
- Ma Z, Mei Y, Su Z (2023) Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. *AMIA Annu Symp Proc* 2023:1105–1114
- MacIntyre A (2017) *Common Goods, Frequent Evils*. The Common Good as Common Project, University of Notre Dame. <https://www.youtube.com/watch?v=9nx0Kvb5U04>.
- Manne K (2013) On Being Social in Metaethics. In: Shafer-Landau R (ed) *Oxford Studies in Metaethics*, vol 8. Oxford University Press, Oxford
- Marin L (2024) Attending to the Online Other: A Phenomenology of Attention on Social Media Platforms (1st edition). In: de Boer B, Zwier J (eds) *Phenomenology and the Philosophy of Technology*. openbook publishers, pp 215–240
- Marin L, Steinert S (2025) Ctrl+ Ethics: Large Language Models and Moral Deskillling in Professional Ethics Education. In: Philipp H (ed) *Oxford Intersections: AI in Society*. Oxford Academic
- Mitcham C (2024) Brief for political philosophy of engineering and technology. *NanoEthics*. <https://doi.org/10.1007/s11569-024-00463-8>
- Natali C, Marconi L, Dias Duran LD, Cabitza F (2025) AI-induced deskillling in medicine: a mixed-method review and research agenda for healthcare and beyond. *Artif Intell Rev* 58:356. <https://doi.org/10.1007/s10462-025-11352-1>
- NHS England (2023) The risk of deskillling | Workforce, training and education. In: NHS England | Workforce, training and education | Digital Transformation. <https://digital-transformation.hee.nhs.uk/building-a-digital-workforce/dart-ed/horizon-scanning/developing-healthcare-workers-confidence-in-ai/chapter-4-workforce-transformation/the-risk-of-deskillling>. Accessed 16 Sept 2025
- Nye BD, Silverman BG (2012) Affordance. In: Seel NM (ed) *Encyclopedia of the sciences of learning*. Springer, US, Boston, MA, pp 179–183
- Pavese C (2022) Knowledge How. In: Zalta EN, Nodelman U (eds) *The Stanford Encyclopedia of Philosophy*, Fall 2022 Metaphysics Research Lab. Stanford University
- Pavese C (2023) The shape of agency, by Joshua Shepherd. *Mind* 132:586–594. <https://doi.org/10.1093/mind/fzab090>
- Pavese C (2024) Knowledge, skills, and creditability. *Philos Stud*. <https://doi.org/10.1007/s11098-024-02261-0>
- Reviglio U (2019) Serendipity as an emerging design principle of the infosphere: challenges and opportunities. *Ethics Inf Technol* 21:151–166. <https://doi.org/10.1007/s10676-018-9496-y>
- Rinta-Kahila T, Penttinen E, Salovaara A et al (2023) The vicious circles of skill erosion: a case study of cognitive automation. *J Assoc Inf Syst* 24:1378–1412. <https://doi.org/10.17705/1jais.00829>
- Rosen C (2024) *The Extinction of Experience*. W.W. Norton
- Ryle G (2009) *Knowing How and Knowing That*. In: *Collected Essays 1929 - 1968: Collected Papers Volume 2*. Routledge, pp 222–235
- Samuel A (2025) *How to Build Your Own AI Assistant*. Harvard Business Review. <https://hbr.org/2025/03/how-to-build-your-own-ai-assistant>. Accessed 23 April 2025
- Schuster N, Lazar S (2024) Attention, moral skill, and algorithmic recommendation. *Philos Stud*. <https://doi.org/10.1007/s11098-023-02083-6>
- Shepherd J (2021) *The Shape of Agency: Control, Action, Skill*. Knowledge. Oxford University Press
- Sherman N (1991) *The Habituation of Character*. In: *The Fabric of Character: Aristotle's Theory of Virtue*. Oxford University Press
- Stanley J, Williamson T (2017) Skill. *Noûs* 51:713–726
- Stichter M (2021) Differentiating the Skills of Practical Wisdom. In: Vaccarezza MS, De Caro M (eds) *Practical Wisdom: Philosophical and Psychological Perspectives*. Routledge, pp 96–113
- Stichter M (2020) Virtue as Skill: Self-regulation and social psychology. In: Fridland E, Pavese C (eds) *The Routledge Handbook of Philosophy of Skill and Expertise*. pp 487–502
- Sypnowich C (2016) *Equality Renewed: Justice Flourishing and the Egalitarian Ideal*. Routledge Taylor & Francis Group, New York
- Tacca A, Gilbert F (2024) Just copy-paste me! Assessing the risks of epistemic dependence on Large Language Models. In: Strasser A (ed) *Anna's AI Anthology. How to Live with Smart Machines?* Xenomoi Verlag, Berlin, pp 31–52
- Tollon F (2024) Technology and the situationist challenge to virtue ethics. *Sci Eng Ethics* 30:10. <https://doi.org/10.1007/s11948-024-00474-4>
- Vallor S (2015) Moral deskillling and upskilling in a new machine age: reflections on the ambiguous future of character. *Philos Technol* 28:107–124. <https://doi.org/10.1007/s13347-014-0156-9>
- Vallor S (2016) *Technology and the Virtues*. Oxford University Press, New York
- Vigani D (2021) Virtue and embodied skill: refining the virtue-skill analogy. *J Value Inq* 55:251–268. <https://doi.org/10.1007/s10790-021-09814-0>
- Withagen R, de Poel HJ, Araújo D, Pepping G-J (2012) Affordances can invite behavior: reconsidering the relationship between affordances and agency. *New Ideas Psychol* 30:250–258. <https://doi.org/10.1016/j.newideapsych.2011.12.003>
- Wong P-H (2019) Rituals and machines: a Confucian response to technology-driven moral deskillling. *Philosophies* 4:59–0. <https://doi.org/10.3390/philosophies4040059>
- Zao-Sanders M (2025) *How People Are Really Using Gen AI in 2025*. Harvard Business Review. <https://hbr.org/2025/04/how-people-are-really-using-gen-ai-in-2025>. Accessed 22 April 2025
- Zelny A (2025) Offloading wisdom: four technological relations that mediate phronesis. *Philos Technol* 38:56. <https://doi.org/10.1007/s13347-025-00889-2>
- Zhai C, Wibowo S, Li LD (2024) The effects of over-reliance on AI dialogue systems on students' cognitive abilities: a systematic review. *Smart Learn Environ* 11:28. <https://doi.org/10.1186/s40561-024-00316-7>
- Zimmermann C, Poirier P, Catala A (2025) Performing arts and affordances: moving toward epistemic justice through embodied learning. *Br J Aesthet*. <https://doi.org/10.1093/aesthj/ayaf004>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.