Citation: Smart, P. R. (2026) Building Extended Minds. *Social Epistemology*, 40(1), 35–44. [https://doi.org/10.1080/02691728.2025.2559279]

Building Extended Minds

Paul Smart

Electronics and Computer Science, University of Southampton, Highfield, Southampton, SO17 1BJ, UK (email: ps02v@ecs.soton.ac.uk)

The mind–technology problem refers to issues that lie at the intersection of technology development and the philosophy of mind. In the present paper, I explore one aspect of the mind–technology problem, namely, the role of technologies in supporting the emergence of extended minds. I approach this issue from an engineering perspective, suggesting that the project to build extended minds yields insights into a number of philosophical problems. These include our understanding of the criteria for cognitive extension and the way the borders/boundaries of extended cognitive mechanisms are delineated.

Keywords: engineering; extended mind; extended cognition; mechanism; constitutive relevance; maker's knowledge

Introduction

According to proponents of the extended mind,[1] technological devices are not just tools that can be used to perform a cognitive task; they are also, on occasion, resources that can be factored into the physical machinery of the mind. In this sense, a bio-external resource, such as a smartphone, becomes part of the realization base for human mental states and cognitive processes—part of the physical mechanisms that realize human

[1] A quick note on terminology. For the purposes of this paper, I will use the term "cognitive extension" to refer to both "extended cognition" and the "extended mind." The difference between extended cognition and the extended mind has been understood in different ways, although the primary contrast is between cognitive scientific kinds (e.g., attention, memory, problem-solving) and folk psychological kinds (e.g., dispositional beliefs). For present purposes, we can ignore this distinction; it will have little bearing on the ensuing discussion.

cognitive/mental phenomena. This rather radical-sounding claim is clearly relevant to the way we understand mind–technology relations. If, for example, we accept the basic idea of cognitive and mental mechanisms as reaching beyond the borders of skin and skull, then technological devices are poised to effect profound shifts in our cognitive and mental architectures. Once such resources have been incorporated into cognitive and mental mechanisms, they acquire the same metaphysical significance as the more usual neurological resources. Just as a neural circuit may be the (partial) realizer of a cognitive process, so too a technological device may be the (partial) realizer of a cognitive process. In this sense, technological devices are the potential building blocks of new minds, or they are, at any rate, the potential constituents of new cognitive and mental mechanisms. This, I suggest, speaks to an important (although often overlooked) aspect of the extended mind debate, namely, the way in which an engineering (or synthetic) approach informs our understanding of prominent philosophical problems. Engineering is, of course, relevant to the practical project of building extended minds; but, as we shall see, the project of building extended minds also yields insight into a number of important philosophical problems.

Building Extended Minds: An Engineering Perspective

Let us start, then, by considering how we might go about building an extended mind. In particular, let us consider a form of cognitive extension that centers on an individual's capacity to identify flower species. The focus of the relevant engineering effort will be a system that includes the following technological elements:

(1) **Augmented Reality (AR) Glasses:** These will be used to take images of a flower whenever that flower is within the field of view of the human user. The images will be taken using an onboard (forward-facing) camera. The user will

trigger the camera by uttering a verbal command (e.g., “IDENTIFY”) or by performing a gesture.

(2) **Smartphone:** The smartphone will receive images from the AR glasses via a Bluetooth connection. When an image is received, the smartphone will post the image to an online AI service.

(3) **Online AI Service:** The online AI service will identify flower species from uploaded images. It will post the result to the smartphone, which will transfer the result to the AR glasses. The AR glasses will then render the received information within the user’s field of view (via an augmented reality display).

For present purposes, we do not need to worry too much about the details of these technological elements. Suffice to say, there is nothing here that is beyond the current state-of-the-art. There are, for example, a number of systems that support the visual identification of plant species from smartphone images (e.g., Bilyk et al. 2022). At this point, then, we have the basic structure of a system that will enable users to identify flower species. Call this system the Flower Identification System (FIS). The FIS implements a process that identifies the species of flower based on the image taken by the AR glasses. Call this process the Flower Identification Process (FIP).

Now comes the tricky part. The aim is to implement the FIS in such a way that it yields a genuine form of cognitive extension (at least when coupled to a human user). The problem is that it is not entirely clear what makes the FIS a genuine mind-extending technology, as opposed to a mere ‘tool’ that is used to perform a particular task. There is, to be sure, a definite difference here. Not every form of interaction with a technological device counts as a *bona fide* form of cognitive extension. There is, as such, a difference between mere tools and mind-extending technologies (see Carter, Clark, and Palermos 2018). But what exactly is the nature of this difference?

The answer, it seems, must have something to do with the interaction between the human user and the FIS. In particular, the FIS must function in such a way as to support claims of cognitive extension. Some insight into the nature of these requirements is revealed by philosophical discussions of the extended mind, specifically those relating to the criterial evaluation of extended cognitive systems (or, more simply, the criteria for cognitive extension) (see Heersmink 2015). At this point, however, we have a further problem. The criteria for cognitive extension are many and varied, and it is not entirely clear that all of these criteria are relevant to the design of the FIS. Ideally, what we require is some more general understanding of these criteria—an understanding that reveals (in a general sense) what these criteria are trying to do.

Here, then, is how I understand the criteria for cognitive extension, or at least most of the criteria. To my mind, the criteria for cognitive extension are best understood as what I will call *establishing criteria* (or establishing conditions). They are, in short, criteria that establish the conditions under which we credit ourselves or others with the possession of a certain dispositional property, where the notion of a dispositional property refers to things like abilities, capacities, capabilities, dispositional beliefs, and so on.

As a means of making this clearer, let us consider the classic extended mind case—the one involving Otto and the notebook (Clark and Chalmers 1998). According to Clark and Chalmers (1998), it is the nature of Otto's interaction with the notebook that underwrites claims regarding the extended nature of (Otto's) dispositional beliefs. In particular, Otto is to be credited with dispositional beliefs pertaining to the location of the Museum of Modern Art (MoMA), just so long as the following conditions obtain (see Clark 2010, 46):

- **Availability:** The external resource (notebook) should be reliably available and typically invoked.
- **Automatic Endorsement:** Information retrieved from the resource should be more or less automatically endorsed.
- **Accessibility:** Information contained in the resource should be easily accessible as and when required.

These criteria, commonly referred to as the criteria of trust and glue, are (I suggest) best understood as establishing the conditions under which it makes social (and perhaps subjective) sense to credit a given individual (including, perhaps, ourselves) with the possession of a particular belief. The wrong way to think about these criteria is to regard them as criteria that (e.g.) determine the cognitive status of an external resource or that resolve issues of constitutive relevance. In respect of constitutive relevance, for example, it would make no sense to insist that an entity (e.g., a notebook) only counts as part of a mechanism if it (*inter alia*) delivers information that is automatically endorsed by another entity. The reason for this is that the problem of constitutive relevance (the problem of determining when some entity ought to be regarded as a *bona fide* constituent [or component] of a mechanism) is applicable to *all* mechanisms, not just those featuring an entity that is (in principle) capable of endorsing or rejecting the informational deliverances of some other entity. The piston and crankshaft are, for example, both components of a common (propulsion) mechanism, but neither of these components is in a position to do much in the way of endorsing or rejecting information.

Rather than helping to resolve the problem of constitutive relevance (or the problem of cognitive status), the trust and glue criteria are, I think, best understood as simply establishing the conditions under which it makes sense to credit some agent (in this case, Otto) with the possession of a given dispositional belief. Consider the notion of automatic endorsement. If Otto believes that MoMA is on 53rd Street, then this is where Otto should go whenever he desires to visit MoMA. It would make no sense to say that Otto believes that MoMA is on 53rd Street if, whenever Otto desires to visit MoMA, he goes to some location other than 53rd Street. This is why the notion of automatic endorsement is important. If the notebook contents are to serve as a reliable guide as to what Otto believes, then Otto's behavior must be aligned with the contents of the notebook. If the notebook reads that MoMA is on 53rd Street, then Otto's museum-going behaviors must be aligned with this information. In short, if the notebook is to serve as a reliable guide as to what Otto will do (when he wishes to visit MoMA), then Otto must endorse the information pertaining to MoMA's whereabouts. Such will be the case if Otto *automatically* endorses the information. If, by contrast, Otto should fail to automatically endorse the information, then Otto might decide to go somewhere other than 53rd Street. But if that should be the case, then it would make little sense to credit Otto with the dispositional belief that MoMA is on 53rd Street.

The appeal to establishing conditions is perfectly compatible with what has been dubbed *the dispositional hypothesis*—an approach that sees dispositional properties (and dispositional ascriptions) as central to our understanding of extended cognition and the extended mind (Smart 2024). According to the dispositional hypothesis, we encounter a case of cognitive extension when

(1) A cognitive/mental dispositional property[2] (*D*) is ascribed to an entity[3] (*E*).

(2) The exercise/manifestation of *D* is a constitutive mechanistic phenomenon (*P*) (e.g., a process).[4]

(3) *P* is realized/constituted[5] by a mechanism (*M*), where *M* comprises a collection of components (*C*).

(4) Some (at least one) of the members of *C* are located external to the borders/boundaries of *E*.

As noted by Smart (2024), this scheme is applicable to many of the cases that have been discussed in the philosophical (active externalist) literature.[6] In respect of the

---

[2] I use the term "dispositional property" to refer to things such as abilities, capacities, powers, tendencies, propensities, capabilities, and so on. Examples of cognitive/mental dispositional properties include cognitive abilities, cognitive capacities, dispositional beliefs, dispositional knowledge, and so on.

[3] Note that *E* could be oneself, as when we credit ourselves with the possession of (dispositional) beliefs and knowledge. In dispositional philosophy, *E* is sometimes referred to as the disposition bearer or disposition carrier.

[4] For more on the notion of constitutive mechanistic phenomena, see Kaiser and Krickel (2017). In general, constitutive mechanistic phenomena are occurrent entities (states, events, processes) that depend on mechanisms for their realization. Accordingly, the dispositional properties targeted by the dispositional hypothesis are what might be dubbed mechanism-dependent dispositional properties.

[5] To be a little more precise, the relation between *P* and *M* is one of mechanistic realization (Wilson and Craver 2007) or mechanistic constitution (Baumgartner and Wilutzky 2017).

[6] As noted by an anonymous reviewer, the dispositional hypothesis also aligns with the views of Wilson (2004), especially in regard to the notion of wide realization. This alignment is deliberate, for the dispositional hypothesis is not intended to be a radically new approach to how we understand extended cognition and the extended mind. It is more an attempt to situate existing views within a theoretical framework that speaks to recent work in both dispositional and neo-mechanical philosophy. The distinctive feature of the dispositional hypothesis is the emphasis placed on disposition ascription. Wilson (2004) is a proponent

Otto notebook case, for example, *D* is a dispositional belief (pertaining to the location of MoMA), *E* is the biological individual known as Otto, *P* is the process of retrieving information from the externally-situated notebook, and *M* is the mechanism that realizes the retrieval process. The notebook is then understood as one of the components of *M*. Given that this resource lies external to the entity that is credited with the possession of *D* (i.e., Otto), then *D* is what we might call an 'extended' dispositional property (in this case, an extended dispositional belief).

Much the same applies to other forms of cognitive extension, such as those involving the use of pen and paper resources to solve long multiplication problems. In this case, we credit a given human individual (*E*) with a capacity to solve long multiplication problems (*D*), even though the exercise/manifestation of that capacity is a process (*P*) involving the manipulation of resources that lie external to the individual.

In both these cases, the central problem is not (I think) the constitutive relevance (or cognitive status) of the external resources; it is more the degree to which we are justified in crediting the relevant individual (*E*) with the possession of *D*. Many (although by no means all) of the criteria for cognitive extension are related to this sort

---

of what has been dubbed the *narrow subjects, extended systems view* (see also Wilson 2014). According to this view, dispositional properties ought to be understood as the properties of individual subjects (e.g., individual humans) as opposed to the larger systems comprising both the individual subject and the extended mechanism (cf. Miyazono 2017). The dispositional hypothesis is neutral on this point. What matters according to the dispositional hypothesis is simply the fact that dispositional properties are ascribed to a particular entity (e.g., a human individual). It makes no claim as to whether these properties actually belong to the entity to which they are ascribed. This allows for the possibility that extension-related claims may stem from a sort of 'error' in our ascriptive practices. Specifically, in the Otto case, it may be correct to view the Otto + notebook system as the proper bearer of a dispositional belief (see Miyazono 2017), but, for whatever reason, we persist in ascribing the belief to Otto.

of problem. They are, in short, an attempt to specify the conditions under which a given individual might be plausibly credited with dispositional kinds—kinds such as believing that X, knowing that Y, or having the capacity to Z. Such a view tallies with the way that some proponents of the extended mind have understood the criteria for cognitive extension. In respect of the trust and glue criteria, for example, Wilson and Clark (2009, 67) suggest that such criteria are "meant to ensure that the capacities of the hybrid system—the biological organism plus augmentation—are plausibly seen as the capacities of a specific individual (e.g. Otto)."

From an engineering standpoint, the dispositional hypothesis is important because it yields a general understanding of what it is we are trying to do when we build extended minds. In the present case, recall, the aim is to implement a system that enables human individuals to identify flower species. If this system is to qualify as a mind-extending technology, then it should (à la the dispositional hypothesis) work in such a way as to support the ascription of certain cognitive/mental dispositional properties to the human user. In particular, we want our human user to be credited with the capacity to identify flower species, even though the relevant capacity (call it the flower identification capacity [FIC]) is one that relies on the coordinated interoperation of resources that lie external to the (biological borders of the) user. This gives rise to a set of requirements that must be met by the FIS. To help us understand these requirements, we can turn our attention to a non-extended variant of the FIC. In particular, we can ask ourselves: "What would need to be the case if a typical human individual were to be credited with a capacity to identify flower species?" Responses to this question might include the following:

(1) **Reliability:** The individual should be able to identify flower species without error (i.e., the capacity should be reliable).

(2) **Efficiency:** The individual should be able to identify flower species within a given timeframe—no more than a few seconds, perhaps.
(3) **Consistency:** The individual should respond to the same flower species in the same way over time.
(4) **Stability:** The individual should be able to identify flower species whenever they are required to do so (i.e., whenever they are required to evidence their capacity).

This is not intended as an exhaustive list of requirements, but it does give us some insight into how the FIS should work if it is to qualify as a mind-extending technology.[7] At this point, of course, we may discover that certain requirements cannot be met. If, for example, there is a significant delay (e.g., 30 seconds) in delivering a

---

[7] Just to be clear, the foregoing requirements are not intended to be understood as the criteria by which we evaluate the extended status of a process. The requirements are intended to act as a guide to the relevant design and development effort. In an engineering context, these are what would typically be referred to as *non-functional requirements*. Their purpose is to ensure that the to-be-built technology works in such a way to support the ascription of dispositional properties (e.g., abilities/capacities) to a given human individual. The core idea is that we ask ourselves what would need to be the case if we were to credit an individual with the possession of a particular capacity (or other dispositional property). We then treat the results of this effort as a set of constraints that must be fulfilled by the relevant technology. In the present case, for example, the FIS must operate in such a way that is reliable, consistent, efficient, and so on. These requirements are informed by our commonsense understanding of the conditions under which we credit individuals with the possession of dispositional properties (such as dispositional beliefs, cognitive capacities, and so on). There is, of course, no reason why philosophical criteria (e.g., those associated with specific epistemological theories) could not be included in the list of requirements. For the most part, however, the requirements are intended to reflect our quotidian ascriptive practices (i.e., the circumstances under which we credit a particular entity as having [or being the bearer of] a particular dispositional property).

result to the user, then we may need to rethink our original design. For present purposes, however, let us assume that there is no issue with the requirements. We implement the FIS and deploy it as a mobile app. Once installed, the FIS works as expected. The human user is now in a position to identify flower species. Whenever the user looks at a flower, they trigger the FIP by issuing a vocal (or gestural) command. The species name is then displayed within the user's field of view (as an augmented reality overlay to the visual scene).

Assuming the requirements pertaining to reliability, efficiency, consistency, and so on have been met, then we will have established the conditions under which the human user could be said to possess a capacity to identify flower species. The FIC will thus be ascribed to the human individual (as opposed to the glasses, the smartphone, or the online services) but the exercise/manifestation of the FIC is a process (i.e., the FIP) that reaches beyond the borders of the relevant individual. The FIP is realized by a mechanism that includes the AR glasses, the smartphone, and the online AI services, all of which lie external to the human individual. Accordingly, insofar as the human user is credited with a capacity to identify flowers, then this capacity must be one that is subject to wide (or extended) realization. It is, we might say, an 'extended capacity'—a capacity that is ascribed to a particular entity (e.g., a human individual), but which nevertheless relies on the instantiation of a mechanism that reaches beyond the borders of that entity. Such mechanisms are what are sometimes referred to as *extended mechanisms* (see Smart 2022).

If what I have said is correct, then the FIS amounts to a *bona fide* form of cognitive extension. It is, in short, the (partial) technological realizer of a recognizably cognitive (or epistemic) capacity.[8] It is, moreover, a system that we have built from the ground up as part of a deliberate attempt at cognitive systems engineering. Now that the engineering (or synthetic) effort is complete, we can turn our attention to issues of a more 'analytic' nature. In particular, let us consider one of the objections to the claim that the FIS ought to be understood as a *bona fide* form of cognitive extension.

Constitutive Relevance and the Extended Mind Debate

According to the dispositional hypothesis, the manifestation/exercise of the FIC is a process that is subject to wide (or extended) realization. That is to say, the FIP is realized/constituted by a mechanism that includes components that lie external to the disposition bearer. In the case of the FIS, the relevant mechanism is one that is constituted by the AR glasses, the smartphone, and the online AI services. In the language of neo-mechanical philosophy, these resources are to be understood as *components* (i.e., the constituents or parts of a particular mechanism). What underwrites the status of these resources as components is the fact that they are *constitutively relevant* to the FIP. There is an important distinction, here, between the notions of constitutive relevance and causal relevance (see Craver 2007). While these are both understood to be forms of explanatory relevance, they are not the same. If an entity is causally relevant to a process (*P*), then it cannot be part of *P* (or the mechanism that

---

[8] The cognitive status of the FIC ought not to be in any doubt. If, for example, we were to encounter an individual who was able to manifest the same capacity using only the resources of the biological brain, then we would (I assume) have little problem in recognizing this capacity as one of the cognitive (or epistemic) variety. But if this is the case for the non-extended capacity, then why should the extended variant be regarded any differently? To be sure, the way in which these capacities are realized (when manifest) is not the same, but it is hard to see why that would have any bearing on the way we regard the capacities themselves.

realizes *P*). By contrast, if an entity is constitutively relevant to a process (*P*), then it must be part of *P* (or the mechanism that realizes *P*). This distinction is important, for it marks the contrast between two views of cognition, namely, extended and embedded cognition. According to the proponents of extended cognition, extra-organismic resources are constitutively relevant to cognitive processes, and they are thus part of the mechanisms that realize those processes. Embedded theorists, by contrast, reject this claim. They insist that extra-organismic resources should be understood in purely causal terms (i.e., as causally—but not constitutively—relevant to cognitive processes).

Clearly, what is required here is a means of resolving constitutive relevance—a means by which we can determine the constitutive relevance of some entity (e.g., a smartphone) relative to some phenomenon (e.g., a process). This is arguably crucial if we are to defend the claim that a cognitive process ought to be understood as a specifically *extended* cognitive process. In the case of the FIS, the candidate cognitive process is the FIP. I have suggested (see Footnote 8) that this process ought to be understood as a cognitive process on the grounds that it represents the exercise/manifestation of a discernibly cognitive/mental dispositional property. I have also suggested that the process is realized by a mechanism that is constituted (*inter alia*) by the AR glasses, the smartphone, and the online AI services. As yet, however, I have done little to substantiate this claim (i.e., the claim regarding the componential status of the AR glasses, the smartphone, and the online AI services). In particular, I have said nothing about why we should regard the (e.g.) smartphone as constitutively relevant to the FIP.

There is a good reason for this. The reason relates to the way in which an engineering approach alters our view of constitutive relevance. In fact, as far as I can tell, engineers do not confront the problem of constitutive relevance, or, at any rate, they do not confront it in quite the same way as their scientific cousins. This is important, for while there has been much research into the notion of constitutive relevance, the bulk of this research is situated in the philosophy of *science*. From a scientific perspective, constitutive relevance features as part of attempts to formulate mechanistic explanations of phenomena. Consider a target (explanandum) phenomenon, such as a disease process. In attempting to explain (or understand) this phenomenon, scientists strive to discover the mechanism that is responsible for the phenomenon. This, in turn, requires the discovery of the mechanism's components. This, it should be clear, is an *epistemic* problem. In particular, the problem is to ensure that one's beliefs (regarding componential status) are suitably aligned with the factive structure of reality. The reason scientists resort to experimental (and other) techniques is because such techniques are useful in addressing this problem. In short, scientific techniques help to reveal the inner workings of mechanisms, and it is by revealing the inner workings of mechanisms that scientific techniques support the epistemic determination of constitutive relevance relations.

Now let us consider constitutive relevance from an engineering perspective. In particular, let us consider constitutive relevance in relation to the FIP. In a scientific setting, this process would be the target (explanandum) phenomenon—i.e., the phenomenon to be explained. Scientists, recall, seek to explain phenomena by discovering (and describing) the mechanisms responsible for those phenomena. So, the explanation of the FIP requires us to describe the mechanism that realizes/constitutes the FIP. In the case of the FIS, however, there is nothing to be done, for we already have a description of the mechanism. That description is simply the earlier characterization of how the FIS would work, if it were to be implemented. In short, the description of the mechanism is the specification (the design plan or 'blueprint') that is used to guide the engineering effort. It is the characterization of how the AR glasses, the smartphone, and the online AI services must interact so as to bring the target phenomenon (the FIP) into existence. As long as the engineering effort conforms to this specification, then we already know how the FIP is realized. The FIP is realized by the mechanism that was built from the design specification, and the design specification describes the mechanism's components and how they must interact so as to make the FIP materially possible. The upshot is that there is no need to resort to scientific/empirical techniques as a means of resolving relations of constitutive relevance, for we already know what these relations are.[9] In an engineering context, these relations are not something to be discovered (as is the case with naturally-occurring phenomena); rather, the aim of the engineering effort is to establish (or create) these relations. The scientist seeks to align their (component-related) beliefs with the factive structure of reality; the engineer, by contrast, seeks to change the factive structure of reality to conform with their (component-related) beliefs.

[9] As noted by an anonymous reviewer, the sort of knowledge that is being referred to here is best understood in terms of the maker's knowledge (see Floridi 2018). The maker's knowledge is a particular form of knowledge that appeals to the role of action in underwriting epistemic claims. An oft-cited example is Alice knowing that Bob's coffee has been sweetened because she herself put two spoons of sugar in it (Floridi 2018). In this case, the truth of Alice's belief is established courtesy of the way that Alice has intervened in the world so as to make her belief true. Much the same can be said of the sort of knowledge that is held by an engineer in regard to the componential structure of a mechanism. In building a mechanism, an engineer establishes a causal structure (a set of causally-interacting components) that gives rise to a particular process. When the process is observed, the engineer knows how this process is realized: the process is realized by the interactions between the elements of the causal structure that was developed as part of the engineering effort.

This difference is important when it comes to recent attempts to substantiate claims of cognitive extension via scientific (e.g., experimental) means (e.g., Gillett et al. 2022). The success (or failure) of such efforts turns on the degree to which scientific techniques support the epistemic resolution of constitutive relevance relations, and some philosophers have expressed doubts about the extent to which such techniques can fully distinguish between constitutive and causal claims (see Baumgartner and Wilutsky 2017). Such concerns are clearly relevant to the effort to understand (extended) cognitive systems from a *scientific* standpoint; but they have little bearing on the effort to actually build such systems. Again, from an engineering standpoint, we are not trying to discover relations of constitutive relevance, for prior to the implementation effort these relations do not exist. It is the goal of the engineering effort to establish relations of constitutive relevance by designing and building mechanisms. But once such mechanisms have been built, we do not then need to, in effect, 're-discover' these mechanisms. Assuming the mechanisms were built in accordance with their respective design specifications, then each design specification will describe the mechanism and thus explain the phenomenon that is realized/constituted by the mechanism. It is just a standard part of engineering practice that such specifications outline the causal and componential structure of mechanisms.

Perhaps, however, there is more to be said about the notion of constitutive relevance, for it remains unclear whether the engineering notion of constitutive relevance is the same as that seen in scientific contexts. As a means of addressing this concern, it will help to consider a recent account of constitutive relevance dubbed the Matched Interlevel Experiments (MIE) account (Craver, Glennan, and Povich 2021; see also Gillett et al. 2022). According to this account, issues of constitutive relevance are resolved by determining the causal path between phenomenon-defining endpoints. In the case of the FIS, the relevant phenomenon is the FIP, and its 'endpoints' are the events that mark the beginning and ending of the process. The FIP begins when the user issues a vocal or gestural command, and it concludes when information is presented to the user.[10] Between these events lies a causal path that winds its way through the AR glasses, the smartphone, and the online AI services. Accordingly, *all* these entities lie on the causal path that connects the beginning of the process with the end of the process. The upshot is that the AR glasses, smartphone, and online AI services are (according to the MIE account) constitutively relevant to the FIP, and they thus qualify as components of the mechanism that realizes the FIP.

There is, as such, no reason to think that engineered mechanisms trade on a radically different conception of constitutive relevance. The only difference is the way in which the componential structure is *disclosed* by the use of scientific techniques (e.g., experimental interventions) or, in the engineering case, *established* by design and development efforts. Note that a key feature of the MIE account relates to the causal path that connects one component (e.g., the AR glasses) to another (e.g., the smartphone). In the case of naturally-occurring phenomena, these links are established by (e.g.) evolutionary processes, and thus their presence needs to be resolved by (*inter alia*) the use of experimental interventions. For the FIS, however, we have no need for such experiments. We already know that the AR glasses exert a causal influence on the smartphone because it was the entire point of the engineering effort to bring this causal link into existence. Accordingly, if the FIP occurs, then there must be a causally-potent interactive exchange between the AR glasses and the smartphone. This must be the case, for it is only via this interactive exchange that the FIP is able to occur.

---

[10] Or perhaps when the delivered information is used to inform a user response.

Conclusion

The mind–technology problem refers to issues that lie at the intersection of technology development and the philosophy of mind (Clowes, Gärtner, and Hipólito 2021). In the present paper, I sought to approach the mind–technology problem from an engineering perspective, focusing on the effort to build extended cognitive systems and extended minds. Such efforts are perhaps best understood as a form of "mind design"—an attempt to create new cognitive and mental organizations from an ever-expanding array of technological assets. But the virtues of an engineering approach are not limited to the practical project of building extended minds; they also promise to inform our understanding of prominent philosophical issues, including (perhaps) our basic understanding of what it means for human mental states and processes to qualify as extended. In this sense, a synthetically-oriented shift in our thinking about the extended mind is important, not just because it helps us build extended minds; it may also hold the key to understanding what it is we are trying to build.

References

Baumgartner, Michael, and Wendy Wilutzky. 2017. "Is it possible to experimentally determine the extension of cognition?" *Philosophical Psychology* 30 (8):1104–1125. https://doi.org/10.1080/09515089.2017.1355453.

Bilyk, Zhanna I, Yevhenii B Shapovalov, Viktor B Shapovalov, Anna P Megalinska, Sergey O Zhadan, Fabian Andruszkiewicz, Agnieszka Dołhańczuk-Śródka, and Pavlo D Antonenko. 2022. "Comparison of Google Lens recognition performance with other plant recognition systems." *Educational Technology Quarterly* 2022 (4):328–346. https://doi.org/10.55056/etq.433.

Carter, J Adam, Andy Clark, and S Orestis Palermos. 2018. "New Humans? Ethics, Trust and the Extended Mind." In *Extended Epistemology*, edited by Adam J Carter, Andy Clark, Jesper Kallestrup, Orestis Spyridon Palermos and Duncan Pritchard, 331–351. Oxford, UK: Oxford University Press.

Clark, Andy. 2010. "Memento's revenge: The extended mind, extended." In *The Extended Mind*, edited by R Menary, 43–66. Cambridge, Massachusetts, USA: MIT Press.

Clark, A., and D. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1):7–19. https://doi.org/10.1093/analys/58.1.7.

Clowes, Robert W, Klaus Gärtner, and Inês Hipólito. 2021. "The Mind Technology Problem and the Deep History of Mind Design." In *The Mind-Technology Problem: Investigating Minds, Selves and 21st Century Artefacts*, edited by Robert W Clowes, Klaus Gärtner and Inês Hipólito, 1–45. Cham, Switzerland: Springer.

Craver, Carl. 2007. "Constitutive explanatory relevance." *Journal of Philosophical Research* 32:3–20. https://doi.org/10.5840/jpr20073241.

Craver, Carl F, Stuart Glennan, and Mark Povich. 2021. "Constitutive relevance & mutual manipulability revisited." *Synthese* 199 (3):8807–8828. https://doi.org/10.1007/s11229-021-03183-8.

Floridi, Luciano. 2018. "What a maker's knowledge could be." *Synthese* 195 (1):465–481. https://doi.org/10.1007/s11229-016-1232-8.

Gillett, Alexander James, Christopher Jack Whyte, Christopher Louis Hewitson, and David Michael Kaplan. 2022. "Defending the use of the mutual manipulability criterion in the extended cognition debate." *Frontiers in Psychology* 13 (Article 1043747):1–9. https://doi.org/10.3389/fpsyg.2022.1043747.

Heersmink, Richard. 2015. "Dimensions of integration in embedded and extended cognitive systems." *Phenomenology and the Cognitive Sciences* 14 (3):577–598. https://doi.org/10.1007/s11097-014-9355-1.

Kaiser, Marie I, and Beate Krickel. 2017. "The metaphysics of constitutive mechanistic phenomena." *The British Journal for the Philosophy of Science* 68 (3):745–779. https://doi.org/10.1093/bjps/axv058.

Miyazono, Kengo. 2017. "Does functionalism entail extended mind?" *Synthese* 194 (9):3523–3541. https://doi.org/10.1007/s11229-015-0971-2.

Smart, Paul R. 2022. "Toward a Mechanistic Account of Extended Cognition." *Philosophical Psychology* 35 (8):1107–1135. https://doi.org/10.1080/09515089.2021.2023123.

———. 2024. "Extended X: Extending the Reach of Active Externalism." *Cognitive Systems Research* 84 (Article 101202):1–12. https://doi.org/10.1016/j.cogsys.2023.101202.

Wilson, Robert A. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York, New York, USA: Cambridge University Press.

———. 2014. "Ten questions concerning extended cognition." *Philosophical Psychology* 27 (1):19–33. https://doi.org/10.1080/09515089.2013.828568.

Wilson, Robert A, and Andy Clark. 2009. "Situated Cognition: Letting Nature Take its Course." In *The Cambridge Handbook of Situated Cognition*, edited by Phillip Robbins and Murat Aydede, 55–77. Cambridge, UK: Cambridge University Press.

Wilson, Robert A, and Carl F Craver. 2007. "Realization: Metaphysical and Scientific Perspectives." In *Philosophy of Psychology and Cognitive Science*, edited by Paul Thagard, 81–104. Oxford, England, UK: North-Holland.