Forthcoming in *Synthese* special issue "New Directions in the Extended Mind", eds. Marco Facchin and Luke Kersten

Home as mind: AI extenders and affective ecologies in dementia care

Joel Krueger
University of Exeter

Abstract

I consider applications of "AI extenders" (Vold & Hernández-Orallo 2021) to dementia care. AI extenders are AI-powered technologies that extend minds in ways interestingly different from old-school tech like notebooks, sketch pads, models, and microscopes. I focus on AI extenders as *ambiance*: so thoroughly embedded into things and spaces that they fade from view and become part of a subject's taken-for-granted background. Using dementia care as a case study, I argue that ambient AI extenders are promising because they afford richer and more durable forms of multidimensional integration than do old-school extenders like Otto's notebook. They can be tailored, in fine-grained ways along multiple timescales, to a user's particular needs, values, and preferences — and crucially, they can do much of this self-optimizing on their own. I discuss why this is so, why it matters, and its potential impact on affect and agency. I conclude with some worries in need of further discussion.

Key words

Extended mind; AI; AI extenders, dementia; affect

Introduction

In 1998, Andy Clark and David Chalmers introduced us to Otto and his trusty notebook (Clark & Chalmers 1998). This paper, and this notebook, convinced many that minds extend. The subsequent twenty-five years or so have seen vigorous debate about this paper's central idea — the extended mind thesis (hereafter, EMT) — and fruitful applications of EMT to other areas of philosophical work (e.g., emotions, epistemology, psychiatry, aesthetics, the Internet and technology studies) that go well beyond the original paper's narrow focus on memory and belief.

I here consider EMT and Generative AI: AI-powered technologies that can both analyze and classify existing data as well as generate new content *from* this data: text, images, videos, code, and more. Consumer-facing Generative AI (hereafter, just "AI") is now ubiquitous, from large language models powering chat apps and virtual assistants (OpenAI's ChatGPT, Google's Gemini, Anthropic's Claude) to devices, wearables, apps, and services many of us use every day. And while some are rightly worried about this AI invasion — there are good reasons to be critical of, for instance, the motives of tech companies relentlessly pushing it onto consumers and potential harms it may disproportionately inflict upon certain groups (Birhane 2022; Buolamwini 2023; Kerr 2020; Noble 2018; Ruane et al. 2019) — there are also ways it may be used for good. One potential area is preventative care and wellbeing.

In what follows, I consider potential applications of "AI extenders" (Vold & Hernández-Orallo 2021) to dementia care. AI extenders are AI-powered technologies that extend minds in ways interestingly different from old-school tech like notebooks, sketch pads, models, and microscopes. I'll here focus on AI extenders as *ambiance*: so thoroughly embedded into things and spaces that they fade from view and become part of a subject's taken-for-granted background. Using dementia care as a case study, I argue that ambient AI extenders are promising because they afford richer and more durable forms of multidimensional integration than do old-school extenders like Otto's notebook. They can be tailored, in fine-grained ways along multiple timescales, to a user's particular needs, values, and preferences — and crucially, they can do much of this self-optimizing *on their own*. I'll discuss why this is so, why it matters, and its potential impact on affect and agency. And I'll conclude with some substantive worries in need of further discussion.

EMT and the rise of AI

When Clark and Chalmers (1998) first motivated EMT by introducing us to Otto and his notebook, they were keen to stress the *dynamic* and *relational* aspects of this partnership — even if these qualities weren't immediately clear.[1] One reason, perhaps, may have been the tech in this original thought experiment. When fiddling with his notebook, Otto provided all the inputs, all the action. He'd write things down and retrieve them when needed. It was easy to see his notebook as a passive receptacle; there didn't seem to be much reciprocity between them. Moreover, when in Otto's pocket or nightstand drawer, the notebook didn't do any work. It lacked cognitive agency and autonomy.

But this picture of EMT is evolving. How so becomes clear when we move from old-school tech like notebooks and sketch pads to more contemporary examples. If our Internet-enabled smartphones, for instance, are part of our extended minds, they do cognitively relevant work even when we're not engaging with them (Clowes 2019; Smart 2017). Different context-aware apps monitor a user's location and help them navigate a city or surface prompts, reminders, and perform automated tasks (unlocking the door, turning on the lights) at certain locations. They may even collaborate in the background with other apps to perform complex tasks, such as health-tracking apps that use data from other apps and wearable devices to identify subtle deviations in vital signs or behavioral patterns. Unlike Otto's notebook, these contextual apps work without user supervision. They have kinds of lives of their own (Krueger & Roberts 2024).

So, it's not just that portable hardware like smartphones — with processing power several thousand times greater than desktop PCs in 1998 (when the original article was published) — are now ubiquitous. Again, it's what's *on* this hardware that matters: namely, the rapid rise of AI-powered apps driving much of the tech we rely on for work and play. These apps connect to cloud services that, using neural networks, large language models, and other machine-learning AI algorithms, allow them to perceive, navigate, make complex decisions, recognize and produce language, plan, identify emotions, and many other things, all in complex and dynamic situations (Vold & Hernández-Orallo 2021, p. 182). Beyond quantitative advances in raw hardware processing power, these *qualitative* developments have already had a significant impact on how

[1] See Gallagher (2018), Kirchhoff & Kiverstein (2019), and Wheeler (2019) for helpful overviews of different "waves" of EMT debates.

we relate to and depend upon our tech. Vold and Hernández-Orallo's (2021) discussion of "AI extenders" helps clarify how so.

AI extenders, agency, and autonomy

What are "AI extenders" and how do they differ from old-school "cognitive extenders" like Otto's notebook? Consider an example. Imagine that Otto has a neighbor, Helen.[2] Like Otto, Helen has dementia. And like Otto, she uses a notebook — along with strategies like putting labels and pictures on cabinets and doors — to remember things. But Helen has an additional piece of tech that's quickly replaced her notebook: smart (i.e., augmented reality) glasses. Helen's glasses perform many tasks that help her function in everyday life. They label objects within her visual range; keep her agenda and surface contextual reminders when needed; recognize faces of family and friends; detect fall hazards and potential dangers when manipulating objects; provide answers to spoken questions (via linked smart speakers distributed throughout the house); and many other useful functions. Helen's glasses are connected to the cloud and gain new abilities with regular updates. They can do these fancy things because they're integrated with AI-powered capabilities — image recognition, face detection, natural language processing, event recognition, speech recognition, prediction — that help Helen live a relatively independent life, despite her dementia diagnosis.

Helen's smart glasses are an "AI extender" (Hernández-Orallo & Vold 2019; Vold & Hernández-Orallo 2021). AI extenders are a subkind of cognitive extenders like notebooks, sketch pads, maps, models, GPS systems, and calculators. The latter are "an external physical or virtual element that is coupled to enable, aid, enhance, or improve cognition such that all — or more than — its positive effect is lost when the element is not present" (Vold & Hernández-Orallo 2021, p. 183). Cognitive extenders are powerful; as Otto demonstrates, they allow agents to do things they couldn't otherwise do, cognitively speaking. However, AI extenders are potentially even more powerful — potentially more epistemically and affectively transformative — than cognitive extenders. This is because AI extenders are fueled by AI-powered capabilities, and therefore "some AI technology *is directly responsible* for the cognitive capability that the

[2] This example is taken from Vold & Hernández-Orallo (2021).

extender is able to deploy, in conjunction with its user" (Vold & Hernández-Orallo 2021, p. 183, my emphasis). Unlike cognitive extenders, AI extenders can do cognitive work on their own.

AI extenders take different forms: *devices* like smartphones, tablets, and speakers; *wearables* like smartwatches, glasses, or pins; and *apps* and *services* available across different hardware and software platforms. Despite this variety, however, the key feature of AI extenders — again, what differentiates them from cognitive extenders like notebooks or sketch pads — is that "they themselves implement kinds of cognitive processes; they are not mere static or interactive tools" (Hernández-Orallo & Vold 2019, p. 510).

For present purposes, there are two key points — both of which I'll discuss in more detail shortly. First, in contrast to cognitive extenders, AI extenders have a significant degree of cognitive agency and autonomy.[3] Not only can AI extenders implement kinds of cognitive processes in a way notebooks and sketch pads cannot. Additionally, they can do this work without constant monitoring from human agents in the control loop (Merat 2019). They are capable of "self-supervised learning": able to go beyond their initial training data and learn new things without the supervision of a human caretaker (Rani 2023). As we'll see, this ability to self-optimize, and the downstream flexibility that comes with it, means that how we live with AI extenders will be different than with (non-AI) cognitive extenders. And these new relations will, in turn, potentially impact our own sense of agency.

Second, I'll argue that AI extenders potentially afford richer and deeper forms of integration with users than do (non-AI) cognitive extenders. Early debates about EMT stressed the *functional parity* of internal and external processes when motivating cases of extension. If Otto's trusty notebook, say, plays the same cognitive role that his bio-memory does – and it's reliably "coupled" (i.e., causally connected) to him throughout the day – there's no principled reason why it shouldn't be considered part of his extended mind (Clark & Chalmers, p. 8). In response to criticisms of parity-based approaches to EMT, later-developing *integrationist* approaches stressed the *interdependence* and *interweaving* of internal and external resources (e.g., Arnau et al. 2014; Menary 2007; Heersmink 2015; Rowlands 2010; Sutton 2010). What's important, according to these approaches, is not that Otto's notebook or smartphone work roughly the same way as do his brain-based mental states. Rather, what matters is how they work together – that is,

[3] See Andrada et al. (2022), Dung (2024), and Swanepoel (2021) for thoughtful discussions of artificial agency that go well beyond the simple characterization I give here.

how they *complement* and *coordinate* their respective cognitive contributions, and the unique formats and functions by which they make these contributions, to form a dynamic and cohesive (extended) system. So, it's precisely because Otto's notebook *isn't* vulnerable to bio-effects like recency, chunking, and decay that it can be such a reliable memory aid and shore up his cognitive deficiencies.

I won't offer a separate defense of an integrationist approach to EMT here. Others have done this (persuasively, in my view) elsewhere.[4] Moreover, my primary aim is to put EMT to work in the context of AI extenders and preventative care – not to defend EMT more generally. But an integrationist approach is what I'll adopt throughout this discussion. Again, I'll argue that AI extenders potentially afford richer and deeper forms of integration with users than do (non-AI) cognitive extenders. This is because they can be tailored, in fine-grained ways along multiple timescales, to a user's particular needs and preferences. They are richly "self-referential" (Krueger & Roberts 2024).

Ottos' notebook has this quality, too, at least to a degree. It reflects particularities of their shared history: Otto's handwriting, the specific things he chooses to remember, and the unique way he organizes this information. But AI extenders offer richer *multi-dimensional* forms of self-referentiality. Users can customize their AI extenders (i.e., increase their self-referentiality). But crucially, their self-optimized AI extenders can do much of this customizing *on their own*, without the ongoing input of the user. And this richer multidimensional self-referentiality, I'll suggest, affords more complex and potentially durable forms of integration and extension than do many old-school scaffolds like Otto's notebook. It does so by fading into the background and working not so much as a tool but rather as a persistent regulative ecology or *ambiance*. I turn to these considerations now.

AI extenders and ambient smart environments (ASEs): dementia as a case study

As a way into these ideas, let's revisit Otto's neighbor, Helen. Recall that like Otto, Helen is one of nearly 7 million Americans 65 and older — and one of an estimated 55 million people around the world — with some form of dementia (2024 Alzheimer's disease facts and figures. 2024). Dementia is a syndrome (i.e., cluster of related symptoms) associated with a severe decline in

[4] See, e.g., Andrada (2019), Arnau et al. (2014), Farina & Levin (2021), Heersmink 2015, Kirchhoff & Kiverstein (2019), Menary (2007), Sutton (2006, 2010), and Rowlands (2010).

mental ability and brain function. While it generally involves memory loss, other symptoms include diminished mental acuity (quickness, sharpness, clarity, etc.); difficulty speaking and following conversations; becoming confused or lost, even in familiar places; losing track of time; difficulty making decisions or performing familiar tasks; and changes in mood and social behavior (feeling increasingly anxious, sad, or angry; personality changes; withdrawal from work or social activities) (NHS 2023; World Health Organization, n.d.).[5]

There is currently no cure for dementia (Mayo Clinic 2024; NHS 2023). Management and support generally consist of a combination of physical and social activities, medication, and – for those no longer able to care for themselves – residential long-term care homes (e.g., nursing homes). In recent years, person- and relationship-centered approaches to dementia care have looked to "fully capture the interdependencies and reciprocities that underpin caring relationships" (Nolan et al. 2002, p. 203), and address ways that issues of agency and power differentials have traditionally been ignored (e.g., excluding the person with dementia from the decision-making process) (Bartlett & O'Connor 2010; Kontos et al. 2017). While the role digital technologies might play in mitigating cognitive decline and supporting dementia care is recognized as worthy of further research, this remains an emerging topic (Benge et al 2020; Blott 2021; Wolff et al. 2021).

To return to Helen, her dementia has a profound impact on how she thinks, feels, speaks, and behaves. Helen sometimes forgets to turn the oven off when she's done cooking, leaves the bathtub faucet running, or discovers her missing keys dangling from the front door. Helen increasingly finds conversations difficult to follow. She routinely loses track of time and, although upbeat by nature, now experiences increasing bouts of anxiety, depression, and apathy. Helen's children see these changes, too; they're increasingly concerned about her ability to live on her own. During visits, they try to tell her this — but their conversations are emotionally fraught and usually end badly. Understandably, Helen doesn't want to give up her independence. Like most people with dementia, she wants to live at home as long as possible (Carter 2016).

In short, Helen's dementia makes her feel increasingly less grounded, less at home in the world. First-person accounts show that it's not unusual for people with some form of dementia to describe feeling like their grip on the world is weakening and that they are somehow less real

[5] Although its exact cause is not known, Alzheimer's disease is a complex neurodegenerative disease that is the most common cause of dementia and sixth leading cause of death in the world (Ramachandran et al. 2021).

(Hughes 2001, p. 86; see also Murphy 2024). But Helen is also grateful that she has a useful bit of tech — her smart glasses — to help counteract this feeling and support her independence. Having them with her makes Helen feel calmer and more rooted in the world than she would without them.

Otto knows this feeling, too. His notebook not only helps him remember things. It also makes him feel connected to the world, and more empowered to go out and do things. So, in addition to augmenting his diminished memory, Otto's notebook is a tool for *affect-regulation* (Candiotto & Stapleton 2024; Krueger & Colombetti 2015; Krueger & Osler 2019; Colombetti & Roberts 2015). Like Helen, he feels more confident, secure, and happy knowing it's with him, providing reliable information he needs to successfully navigate his world.[6]

Yet, as we've seen, Helen's tech is much fancier than Otto's trusty notebook. The latter is a (mere) cognitive extender. However, Helen's smart glasses are an AI extender that give her an additional amount of independence and confidence. Helen can use her smart glasses to take pictures via vocal commands and upload that information to her photos app. So, like Otto's notebook, it has a memory-augmenting function. But Helen's smart glasses can also perform a much broader range of complex cognitive tasks on their own — image recognition, face detection, natural language processing, event recognition, speech recognition, prediction — without her active intervention. This is what makes them an AI extender and why her glasses give Helen greater functionality and security beyond that of a notebook.

From smart things to smart(er) spaces

Let's now imagine that Helen's entire *house* is an AI-extender, a "smart" environment powered by the latest AI-driven software and hardware. Again, this is more than the claim that curating one's home environment can extend cognition and, in the case of dementia, support unassisted living. There's evidence that people with dementia already do this (McCabe et al. 2018). A brief discussion of how so will help clarify what new integrative resources "ambient" AI-extenders (e.g., smart spaces) offer.

[6] Early discussions of EMT largely ignored the affective dimensions of our mind-extending practices, which prompted an "affective turn" in EMT debates some years later (see Krueger & Szanto 2016 for an overview). More on this later.

Clark (2003, pp. 139–140; see also Drayson & Clark 2020) describes a powerful example of homes as cognitive extenders. He recounts discussions with Carolyn Baum, then head of occupational therapy at the Washington University School of Medicine in St Louis, Missouri. Baum and colleagues were puzzled by a sub-population of inner city Alzheimer's sufferers who, despite faring badly on standard psychological tests (e.g., CERAD protocol), could nevertheless live successfully on their own — even though by standard measurements they should not have been able to cope with the demands of everyday life (Edwards et al 1994). A closer look found that they adopted a range of extended mind-style strategies to thwart their cognitive decline.

For example, centralized message centers stored notes about what to do and when to do it; "memory books" housed new events, meetings, and plans; doors and cabinets were tagged with labels and pictures; photos of family and friends were similarly tagged; and open-plan storage strategies held important items (cooking items, medicine, check books) in plain view. As Drayson and Clark tell us, their home environments were "carefully calibrated to scaffold their biological brains" and implement "alternative but perfectly genuine forms of control and recall" (Drayson & Clark 2020, p. 588; see also Dennett 1996, pp. 138–139).

Following Vold and Hernández-Orallo, we can see how AI advances mean that smart homes can (or will soon be able to) do even more. Again, they'll be *AI extenders*, able to do significant chunks of cognitive and affective work on their own. And crucially, they'll flexibly adapt, over multiple timescales, both what information they gather and what they do with it without the active intervention of the user.

To make this clearer, imagine that Helen's children, respecting her desire for autonomy, have now set up a smart home to support her independence. At some expense, Helen has access to the latest AI-powered home tech.[7] Her home is an *ambient smart environment* (ASE) — a space infused with a rich array of networked smart technologies that communicate with one another in

[7] The cost of this technology will be an important issue insofar as it impacts *accessibility*: i.e., those who can afford this technology and the infrastructure needed to support it, and those who cannot. This is one part of a larger "digital divide" that will likely become even more pronounced as AI systems are increasingly embedded within everyday life. In addition to AI-related access, other emerging issues include *capability* (i.e., developing the knowledge and literacy needed to understand and use AI systems) and *outcome-related* issues (e.g. its potential to perpetuate structural inequalities among users with different backgrounds). See Wang et al. (2024) and Carter et al. (2020) for further discussion. My thanks to a reviewer for pressing this point.

the background, transparently supporting Helen's day-to-day life (White & Hipolito 2024; White & Miller 2024).

Safety and wellbeing flows from sensors monitoring Helen's movement and activities. These sensors activate smart locks once Helen is safely inside, and on a regular schedule. A soft alarm sounds if Helen leaves her keys in the front door or if they're detected in an unusual place (e.g., the edge of the driveway or back garden), a faucet is left on too long, or Helen forgets to turn off the oven after cooking. AI algorithms analyze data from motion sensors and other monitoring devices (e.g., built into chairs, mirrors, toilet seats, toothbrushes, forks and spoons, scales, mattresses, etc.) and wearables to learn typical patterns of behavior, detect or predict falls and alert emergency services if needed, and flag unusual deviations or anomalies (e.g., weight loss, skin conditions, disturbed sleep patterns, unusually long periods of immobility) that may signal a health or safety concern.

Cognitive support comes from voice assistants and visual displays distributed throughout the house. These things surface reminders for tasks, daily and weekly medication schedules, appointments and social events, and summaries of the day's weather. These reminders show up on Helen's phone or smart display in the kitchen as she gets her morning coffee (set up to automatically start brewing when her smart mattress indicates she's awakened); or they might be conveyed via gentle prompts, other visual cues, or friendly voices. A smart pillbox tracks Helen's medication, offers chirpy reminders for missed meds, and texts caregivers about missed doses that go beyond a day or two. When not in use, many of Helen's visual displays cycle through family photos (which can be updated remotely by her children); these photos are automatically tagged with family names, relationships, and descriptions of where and when the photo was taken.

Environmental control and comfort come from the fact that most of Helen's personal tech (lights, TVs, speakers, phone, cameras, appliances, etc.) can be activated both by voice and a simple visual interface. Lights and temperatures are adjusted automatically throughout the day, relative to Helen's preferences and the weather; smart artificial candles and air fresheners provide additional ambient texture when needed. Playlists with Helen's favorite music start playing when she's cooking dinner or doing her morning stretches; these playlists respond to data from wearables (heart rate, skin conductance) by selecting peppier tunes if Helen seems sluggish and needs a boost during an exercise session, or calmer music if she seems agitated.

Social support comes not only from the easy accessibility of family and friends via voice-activated digital assistants and displays in every room but also from digital companions and social robots. Advances in natural language processing and text-to-speech mean that Helen has regular chats with her digital assistant, "Courtney". Courtney adapts her responses to the patterns detected via the array of other devices throughout the house monitoring Helen's current state (e.g., "I notice you're a bit slow this morning. Is everything okay?"). Helen also has access to a sophisticated chatbot modeled on her late husband, Ted. Chatbot Ted mimics his style of speaking, humor, catchphrases, characteristic references, etc. Although Helen doesn't use a visual avatar of Ted (an optional upgrade she finds unsettling), occasionally hearing his voice or seeking his advice brings comfort. Finally, Helen's fancy humanoid robot (she hasn't given it a name yet) helps with household chores, undertakes simple nursing tasks when needed, and provides emotional and social support by recognizing changes in Helen's facial expressions, posture, and behavior, and adapting its responses accordingly.

This list isn't exhaustive. And some of this tech isn't yet available, of course. Moreover, interoperability (i.e., all these systems and platforms communicating seamlessly with one another) remains a challenge, as does the fact that most of these devices are not designed with the needs of dementia patients in mind (Addae et al. 2024; Moyle et al. 2021). But much of this AI-powered tech is here now — many of us use the devices, wearables, and apps described above throughout our day – and much more is on the way (Blott 2021; Ienca 2017). ASEs like Helen's will likely have a significant impact on in-home preventative care (Ienca et al. 2017; see also White and Hipolito 2024).

This may be especially true of dementia care. Many people with dementia are cared for at home by informal carers (i.e., family and friends) (WHO 2012). So, the character of the home environment plays a significant role in dementia management and quality of life for both patient *and* carer (Van Hoof & Kurt 2009). Since home is, for many with dementia, associated with positive feelings — belonging; autonomy and control; familiarity and security; a repository of memories and life histories (Aminzadeh 2010; Forsund et al 2018) — *leaving* home for assisted living housing can feel like a death sentence (de Witt et al. 2009; see also Phenice & Griffore 2013). Emerging "materialities of care" perspectives in healthcare, including dementia care, acknowledge these important links between material culture (i.e., everyday things and spaces) and wellbeing, and adapt their caregiving strategies accordingly (Buse et al. 2018; see also

Araujo et al. 2020; Malafouris & Rohricht 2024). Some forward-thinking care homes and hospices, for example, now allow residents to bring their own furniture and mementos as they transition into an assisted living space.

ASEs, as AI extenders, may help delay these transitions and support more effective dementia care in several ways. They might (1) decrease the burden on public finances by delaying the need for institutional care; (2) lighten caregiving burdens for both formal and informal (i.e., family, friends, neighbors) carers, reducing stress and anxiety; (3) compensate for decreasing numbers of human caregivers available to provide these resources while enhancing and optimizing quality of care; and (4) empower those living with dementia and improving their quality of life (Ienca et al. 2017, p. 1302).

However, Soilemezi and colleagues (Soilemezi et al 2019) note several important challenges that must be considered when adopting a more broadly externalist approach to dementia care. Two will help us now circle back to some philosophical themes discussed previously. First, they note the need to adopt *individualized* care strategies tailored specifically to the changing needs of the person with dementia (ibid., p. 1262). Second (and relatedly), environment-centered approaches should involve adapting *the physical space, objects, and behavior* of the individual's home to better meet their unique needs. As we'll now see, an important feature of AI extenders and ASEs is their ability to do this adapting, in precise ways along multiple timescales, without the user needing to be in the control loop.

Self-referential and self-supervised ASEs

So, what is the philosophical significance of all this as it relates to EMT? And what might AI extenders and ASEs tell us about links between cognition and affect, and how we might think about some implications of future forms of AI-driven extension?

First, note an important contrast between the home-as-cognitive-extender example Drayson and Clark (2020) describe and the ASE case. The former involves what Slaby (2016) calls a "user/resource" model of extension. In this example — as with Otto — individual users actively curate their home environment to extend their cognitive capacities. They assess their individual needs (e.g., to compensate for their diminishing short- and long-term memory) and structure their local resources accordingly. Crucially, these off-loading practices (Kirsh 2010; Risk & Gilbert

2016) pushing cognitive work onto bits of the environment involve a subject-to-world, or *inside-out*, dynamic. However, AI extenders — and more specifically, ASEs — flip this relation to something more like a world-to-subject, or *outside-in* (i.e., resource/user) dynamic. ASEs, as local worlds, actively wrap themselves around users (Aydin et al. 2019). This quality of "self-referentiality" is one reason this ambient tech operates so transparently in the background and is so effective.[8]

By "self-referential", I mean that ASEs are shaped by a history of prior interactions with the individual and therefore mirror the individual back to themselves (see Krueger & Roberts 2014 for a longer discussion of this idea). Consider immersive video games. Video games — particularly open world games like *Cyberpunk 2077*, *Red Dead Redemption 2*, *Baldur's Gate 3*, *Zelda: Breath of The Wild*, or *Elden Ring* — are what Robson and Meskin (2016) call "self-involving interactive fictions". These immersive game worlds are pliable and responsive in a way works of canonical fiction like novels and films are not (Wildman & Woodward 2018). No matter how intensely I identify with a favorite character in a book or movie, I have no influence over what they do. Their fictional world does not respond to mine.

But rich open-world video games are different. This is part of what makes them so immersive, so affectively arresting. If I betray a character early in the game's story, say, that interaction will have downstream consequences. It will influence future interactions with that character and others; it may make interacting with their allies more difficult but also present opportunities for new alliances with other non-player characters similarly hostile to them. What I say and do in-game therefore has both (synchronic) consequences for that initial interaction, as well as (diachronic) consequences for how the larger narrative and character arcs unfold as the story progresses. I play an active role in world-building. And the world I help shape gradually takes on my profile via the contours of this interactive history.[9]

Helen's ASE is similarly self-referential. Her network of interconnected AI-powered apps and hardware continually monitor Helen's state and, working together, adapt their responses accordingly. They collectively establish a space specifically curated to Helen's distinct needs,

[8] Although see Wheeler (2019) for a discussion of why the drive for transparency in the design of smart technology is not an unqualified good.
[9] For example, the massive open-world role-playing game *Baldur's Gate 3* has, according to the studio behind it, 17,000 possible variations on the ending based upon a player's unique dialogue choices throughout the 100 or so hours it takes to complete the main story.

preferences, and behavioral tendencies. And crucially, they do this — in increasingly fine-grained ways — in the background, without her active input. They can flexibly adapt this way thanks to a second important feature that will soon be a key feature of many (if not most) AI-extenders: *self-supervised learning* (SSL).

SSL is a machine learning paradigm that enables AI models to learn new things from unlabeled data without the supervision of a human caretaker. SSL makes the process of training AI models much more efficient than previous (i.e., supervised) methods since fewer initial input labels — traditionally supplied by human caretakers — and smaller samples are used to learn more things in a faster way (Rani et al 2023). For instance, SSL-powered computer vision apps can learn to identify objects, classify images, graph these classifications, and answer visual questions like "Is this a vegetarian pizza?" or "What kind of dog is this?" (Manmadhan and Kovoor 2020). SSL algorithms simplify their development by training models to recognize previously unforeseen objects or features based upon a few initial input labels ("pepperoni", "mushroom", "Shiba Inu", etc.). These models use "transformation strategies" — rotating, colorizing, blurring, cropping, filling in, or predicting missing bits of an image — to extract further information and formulate their own self-generated "pseudo-labels" for downstream tasks like classifying future objects and scenes (Rani et al. 2023, p. 2762). SSL works in other domains, too. In the context of audio, SSL can train models to predict masked portions of audio signals or reconstruct audio from corrupted samples (Wu et al. 2023); or in robotics, by training models to predict the next video frame of its own movements, build a continuously updated model of its own kinematics, and support autonomous movement (Ebert et al. 2017).

Helen's ASE employs SSL. Its meshed array of monitors, sensors, and AI-powered apps track things like Helen's movements, energy consumption, and appliance usage. By sharing data and working together in the background, these apps learn Helen's daily routines, construct a baseline model of expected behavior, and then monitor for anomalies — and if needed, intervene without her prompting. Again, falls, wanderings, or deterioration in hygiene habits can be tracked and caregivers alerted. SSL algorithms monitor appliances and gadgets, anticipating malfunctions (based on energy usage or deviations in operating patterns) and initiate repairs or alert remote support. Integrated cameras and microphones track facial expressions, body language, movements, and speech. By learning Helen's unique emotional and expressive repertoire, SSL can provide tailored emotional support, or prevent potentially dangerous or escalating behavior, by adjusting the lighting, playing music, initiating a video call with a family member or caregiver, or nudging Helen toward a mood-lifting activity like gardening or exercise. Finally, SSL might power personalized adaptive interfaces (e.g., in screens or appliances) that respond in real-time both to Helen's preferences (e.g., automatically surfacing preferred channels or automated commands) and, over longer temporal stretches, her gradual cognitive decline (e.g., slowly simplifying the TV remote control layout or appliance control panel).

Again, not all this tech exists yet. But much of it does — and much more is coming (Blott 2021). And a key lesson, I suggest, is that SSL algorithms will create richly self-referential ASEs keenly adapted to individual users' distinct needs. Unlike Otto, users like Helen won't have to do the work of conforming the environment to themselves. Rather, the latter will do much of this work for them in at least two ways.[10]

First, by developing *multi-modal self-supervision models* (i.e., integrating data from audio, visual, and movement sensors, along with other AI-powered apps) to develop a rich dynamic profile of both Helen and her environment. And second, by realizing *unsupervised adaptation* in that the home as AI-extender will continuously update models of both Helen and itself, adapting to changing environmental conditions and supporting a more personalized experience. Again, this will run both synchronically (i.e., detecting and responding to moment-to-moment needs and preferences) and diachronically (i.e., creating baseline models that predict and adapt to longer-term habits, behavior, changes in habitual patterns, etc.).

As a result, Helen's home ASE will quickly become, I suggest, something like what Piredda calls "an affective exoskeleton" (Piredda 2020): a shifting and dynamic assemblage of meshed AI-powered artifacts and digital models that move and adapt with and to Helen in real-time, regulating her behavior, decision-making, and affect, and supporting a range of cognitive tasks she'd otherwise be incapable of doing on her own. As an AI extender, it will do much of this integrating and regulative work on its own. White and Miller (2024) helpfully characterize this anticipatory regulative capacity as "upstream functionality" shaping an agent's present and future field of affordances, enabling agents to act more skillfully within environments self-

[10] As Aydin et al. (2019) put it, "By connecting to and exchanging data with other devices and technologies…ATEs [Active Technological Environments] are able to increasingly act autonomously. The extent to which ATEs will shape the course of their own cognitive development, will probably depend on their capacity and success to "self-structure" their own learning" (p. 325). See also Settles (2012), Smart (2017), and White & Hipolito (2024).

referentially tailored to their individual values, preferences, and moment-to-moment and long term needs.

This regulative control might emerge, for instance, via the way ASEs *simplify* Helen's home environments to facilitate smoother interactions (automatically surfacing prompts and reminders; decluttering user control interfaces; automatically placing grocery orders when supplies get low). But as White and Miller (2024) note, they may also sometimes increase the *complexity* or *friction* of the environment, too. This gentle friction might prompt a "shift into more exploratory modes of engagement" (ibid., p. 47) with an individual's environment (e.g., suggesting that Helen try a new recipe tonight for dinner or prompting her to get up and stretch after she's been watching TV for several hours) and, in so doing, help ward off or disrupt unhealthy habits (see also Wheeler 2019).

In this way, Helen's AI-powered smart home will directly modulate her sense of agency, enabling her to do things she couldn't otherwise do without the regulative scaffolding of this affective exoskeleton. It can do so because as Nguyen (2019) observes, human agency is not fixed. Rather, it is "modular and moderately fluid. We have the capacity to set up temporary agencies, layered within our larger agency, and submerge ourselves in them" (ibid., p. 426). For Nguyen, immersive game worlds invite this sort of "layering" by specifying modes of agency players can adopt within that game world (many of which are unavailable to us outside it). This is part of their affective pull. Helen's ASE, I suggest, functions in a similar way. Helen maintains a degree of independence and autonomy without it. But particularly as her cognitive condition declines, her ASE will assume an increasing amount of cognitive and self-regulatory work that will allow Helen to live well and meet many of her everyday needs.

ASEs, affective place memory, and benign invasions

There is one more point I want to make, one that brings us back to something raised at the start of this discussion. It concerns the affective dimension of these AI-powered regulative dynamics, that is, the felt vibes they set in the spaces they structure. More precisely, ASE scenarios like these, I suggest, bring into sharp relief the tight link between *memory* and *affect*. As we've already seen, Helen's ASE isn't just an epistemic ecology that helps her remember things. It's an affective ecology, too, one saturated with an array of AI-powered regulative dynamics that elicit and shape a repertoire of feelings like intimacy, trust, security, and at-home-ness. Helen's ASE is therefore an example of what Sutton (2024) terms "affective place memory": a distributed process of feeling and remembering partially constituted by the processes, activities, and resources unique to a given place.

These feelings of intimacy, trust, security, and at-home-ness — alongside the epistemic and memory-enhancing benefits of her ASE — are what help Helen feel empowered in her own home (Calkins 2018; Gitlin et al. 2006).[11] Moreover, by shaping her affect, Helen's ASE also strengthens her "spatial agency": her felt sense of both *action-possibilities* within her domestic space along with her feeling of how her body *fits into* that space (Awan et al. 2011; Krueger 2023; Kukla 2022; Schneider and Till 2009). As a self-referential "affective exoskeleton", Helen's ASE becomes what Jacobson (2009) terms a home as "second body", which helps clarify why Helen (and other dementia patients like her) not only realize forms of functional independence at home but also *feel* a particular security, empowerment, and freedom as they move and act within it.

Helen's case is therefore helpful because it highlights how these themes of extended cognition, affect, memory, and agency intermingle in rich and complex ways. However, by focusing exclusively on the epistemic benefits of AI-extenders and ASEs, these rich entanglements and affective dimensions are potentially obscured or lost (Krueger & Osler 2019; Sutton 2024).

To bring these ideas into sharper relief, let's return to Slaby's (2016) discussion of EMT and what he calls affective "mind-invasion". Slaby's work is continuous with a wave of literature that, in the wake of early EMT debates, argued that similar extensions also occur within our *emotional* life, understood broadly as encompassing a range of affective phenomena both occurrent and dispositional, including momentary emotion episodes or evaluative appraisals to temperaments, character traits, and enduring moods (Candiotto 2016; Colombetti & Roberts 2015; Coninx & Stephan 2021; Krueger & Colombetti 2015; Maiese 2016; Slaby 2014). These approaches — in different ways — argue that some emotions incorporate external resources so deeply that they can be said to extend beyond brain and body.

While Slaby (2016) is sympathetic to this idea, he adopts a critical perspective. Instead of focusing on happier individual user/resource cases like using music to positively self-regulate

[11] Although the seamless support and security Helen's ASE provides might make her less inclined to leave her home and venture into an unpredictable world beyond. I discuss this worry about social isolation and other concerns in the conclusion.

occurrent emotions (Krueger 2014), he is interested in cases where individuals and groups are vulnerable to being affectively "invaded" or "hacked" from the outside, often without their full awareness or consent. Drawing on Melissa Gregg's (2011) work on "presence bleed", Slaby argues that contemporary corporate workplace culture is the site of one such invasion. Slaby's analysis is rich and subtle; I can't do justice to it here. But the basic idea is that for many, the corporate world is "one of those expanded zones of contemporary life in which human affectivity is profoundly framed and modulated so that the affective and emotional dispositions of an individual squarely fall in line with the interaction routines prevalent in these domains" (Slaby 2016, p.2).

For example, increased emphasis on the affective value of "intimacy" in white-collar work relationships feeds into a willingness to open one's home and personal life to work colleagues and professional responsibilities (Gregg 2011, p.85). Consequently, employees habitually check always-on communication technologies (smartphones, email, chat apps) that further reinforce this felt intimacy: responding immediately to email, even on weekends; feeling guilty when they don't, or becoming anxious when they briefly check in and see decisions being made without them. Presence bleed generates a local instantiation of a workplace arrangement — including its attendant affective qualities and norms — at home. Employees enact many of the same bodily and affective habits that characterize their office routines, weakening work/life boundaries and "invading" their domestic space with dominant corporate values.

I'm sympathetic to Slaby's critique and find the notion of affective "mind invasion" useful. It highlights some individualistic limitations of a user/resource framework for EMT, which can easily overlook the sort of outside-in, or resource/user dynamics Slaby discusses — and which, I suggest, ASEs also provide. Even more importantly, Slaby's critique highlights crucial *political* and *normative* dimensions of EMT, which draw attention to ways things and spaces can function as "callous", "hostile", or "oppressive" forms of regulative and affective scaffolding, pushing on the individual in ways detrimental to their wellbeing (see, e.g., Krueger 2023; Liao & Huebner 2021; Maiese & Hanna 2019; Osler forthcoming; Protevi 2009; Rosenberger 2017).

For present purposes, Slaby's framework also helpfully offers additional conceptual resources for thinking about how ASEs and other AI extenders may increasingly work in outside-in, or bottom-up ways, "invading" users' habitual thought, feeling, and behavior. Following Vold and Hernández-Orallo, a recurring theme has been that AI-extenders do much of their extending work without constant user supervision; they have a significant degree of cognitive agency and autonomy. So, instead of requiring users to actively offload cognitive work onto them (think of Otto faithfully writing new things down in his notebook each day), ASEs and AI extenders are instead "set up to set off" a range of self-referential upstream functions (White & Miller 2024).

However, while Slaby's example of affective mind invasion shows how external resources may work *against* individuals, Helens' ASE — and the possibility of therapy-oriented ASEs more generally (Blott 2021; White & Hipolito 2024) — present us with cases of what we might term "benign invasions": AI-saturated environments that, working in an outside-in or bottom-up manner, positively *enhance* users' spatial agency by empowering their functional independence (Carter 2022, p. 473). These benign invasions are, I suggest, one pathway for developing what Carter (2022, 2024) helpfully terms "non-dominating" care strategies for dementia. A non-dominating (or "indirect-first") approach recognizes that despite their cognitive decline, people with dementia remain persons capable of generating their own authentic values and interests — both of which must therefore be respected, to the extent this remains possible, when making care decisions.[12]

For example, a *direct* approach like hiring a live-in nurse or putting Helen in a care home may be in Helen's children's best interests (e.g., alleviating their anxieties by ensuring that she has round-the-clock care). But like many with dementia, Helen's preference is to remain at home. Autonomy, self-sufficiency, and independence are some of her core values. So, Helen's ASE helps meet her vital needs while maintaining these core values. Of course, this might change as her dementia progresses. Every patient is different, and an indirect approach might not always be feasible at all stages. However, a virtue of an ASE-supported indirect approach of the sort sketched here is that the *durability* and fluid *self-referentiality* provided by AI-extenders may significantly delay the need to adopt a more direct approach that potentially subverts Helen's core values. It may also prolong Helen's ability to engage in independent decision-making – again, not just by transparently supporting Helen's short- and long-term needs but also by providing moments of gentle friction of the sort mentioned earlier that nudge Helen out of periods of stasis and prompt more attentive, exploratory engagement with the world – including her assistive tech (White & Miller 2024). Instead of designing these systems to completely

[12] Several authors have made similar arguments using phenomenological and embodied approaches to dementia and the self. See, e.g., Hughes (2011); Hydén et al. (2014); Käll (2017); Kontos (2005); Krueger (2024); Petherbridge (2019); Tewes (2021).

disappear into the background, then, perhaps a more nuanced approach is for some aspects of ASEs to support patterns of activities in which individuals "build up their spaces through "conversations" with the environment, where the history of interactions builds up new possibilities for sharing goals and sharing outcomes" (Haque 2006, p. 44; see also Wheeler 2019b). Again, the goal of this conversational approach is to provide resources for supporting the user's capacity for independent decision-making as long as possible instead of taking over all decision-making responsibilities for them.[13]

There's much more to be said about the ethics of these strategies and the decisions they involve, of course. Carter (2022, 2024) helpfully elucidates some of these complexities. Nevertheless, Helen's example — and therapy-oriented ASEs more generally — suggest the possibility of AI-driven *benign* invasions that may, *contra* Slaby, work in the best interests (i.e., by preserving their core values, autonomy, independence, and dignity) of the "invaded". Slaby's EMT-inspired framework can help us think through the dynamics of these invasions and, in so doing, bring some of these ethical considerations into sharper relief.

Final thoughts

I've put Vold and Hernández-Orallo's (2021) helpful notion of "AI extenders" to work in the context of dementia care. I've argued that AI extenders and ASEs offer promising therapeutic pathways because they potentially afford richer and more durable forms of multidimensional integration than do old-school cognitive extenders like Otto's notebook. Because these artificial systems are both *self-referential* and can employ *self-supervised learning*, I've argued that they can be dynamically tailored, in fine-grained ways along multiple timescales, to support a user's unique needs, values, and preferences. And they can do much of this work on their own, leaving users to live more empowered and independent lives.

However, as Slaby (2016) and Carter (2022, 2024) show us, both EMT and dementia care involve important ethical considerations in need of more attention. So does AI. While this discussion has adopted a relatively optimistic outlook on the potential of AI and ASEs in health care, there are many reasons to be skeptical of this technology and its potential impact across a range of domains and human agency more generally (see, e.g., Ienca 2023; Müller 2015; Vallor

[13] I'm grateful for a reviewer pressing this point.

2024). One worry concerns what Figà-Talamanca (2024) calls "digitally scaffolded vulnerability". This notion refers to "unforeseen ways digital technologies can harm, exploit, or intrumentalize their users in ways that are engendered or aggravated by the technology's design and functioning" (p.632).

In the case of ASEs and dementia care, I suggest, digital vulnerability can fall into at least two broad categories: *functional* and *psychological*. The former encompasses *security and safety risks*, such as unauthorized access, hacking, or data breaches involving sensitive health information — as well as the sudden disorientation that may occur when an individual component breaks or the entire ASE goes down (e.g., due to a power outage, hardware failure, or bad software update). Additionally, AI extenders remain imperfect. Large language models like ChatGPT and Gemini routinely "hallucinate" (i.e., provide false or misleading information). So, vigilance against the possibility of *misinterpretations and errors* is needed — e.g., a fall detection alert might activate an alarm when a person bends over or does yoga, causing anxiety for both individuals and carers. Additionally, *bias mitigation* should be a priority. AI systems not only hallucinate but also reflect the biases and prejudices that are part of their training data. And since implicit bias has already been found to lead to disparities in access to health care, quality of care received, and health outcomes (Hall et al. 2015), we should assume that similar biases will be built into many of the algorithms powering ASEs in dementia care. Measures should therefore be taken to make sure all patients receive fair treatment, regardless of background or cognitive ability. Finally, issues surrounding *data ownership and access* are critical. Patients and caregivers should be clear about what sort of data is being collected, and how they can access, modify, or delete it.

When it comes to psychological vulnerability, *loss of privacy and autonomy* is a significant concern. Constant surveillance through sensors, cameras, and AI-powered monitoring apps may, understandably, erode a sense of privacy and lead some to feel that they have less autonomy than before. *Social isolation* is also a threat. Individuals might increasingly feel disinclined to leave the security of their ASEs, reducing access to novel stimulation and embodied human contact – crucial for dementia care (Kontos 2005; Zeiler 2014) – which can lead to further isolation. Human interaction and emotional support remain crucial, and therefore AI should be thought of as a supplement – and not a replacement – for human contact in contexts of care. Another form of psychological vulnerability is something briefly mentioned earlier: *digital literacy*. These systems are complicated; few of us have the technical knowledge to understand them in a deep way. And the progressive decline of cognitive function and decision-making capacities in dementia provides further barriers to understanding. So, discussions about the ethics of informed consent and transparency (e.g., explaining to patients and caregivers what these systems do, how they work, how they'll be implemented within individual cases, etc.) are critical.

Yet, despite these worries, this tech holds promise. This is particularly so when it comes to its therapeutic potential. By offering a promising pathway for achieving "just care relationships" (Carter 2022) for people with dementia, it therefore warrants ongoing research. Those involved in EMT debates should be part of this work.[14]

[14] I am grateful for the feedback I received from three anonymous reviewers. Their critical comments greatly improved the main argument. Many thanks also to Marco Facchin for his excellent editorial support. Finally, I'm grateful for feedback I received when discussing this material with audiences in Lisbon, Mexico City, and Oxford.

References

2024 Alzheimer's disease facts and figures. (2024). *Alzheimer's & Dementia: The Journal of the Alzheimer's Association*, *20*(5), 3708–3821.

Addae, S., Kim, J., Smith, A., Rajana, P., & Kang, M. (2024). Smart solutions for detecting, predicting, monitoring, and managing dementia in the elderly: A survey. *IEEE Access: Practical Innovations, Open Solutions*, 12, 100026–100056.

Andrada, G. (2019). Cognition as an Enculturated and Extended Social Skill. *Australasian Philosophical Review*, *3*(1), 71–75.

Andrada, G., Clowes, R. W., & Smart, P. R. (2022). Varieties of transparency: exploring agency within AI systems. *AI & Society*, *38*(4), 1–11.

Araujo, M. T., Velloso, I. S. C., Ceci, C., & Purkis, M. E. (2020). The significance of overlooked objects: Materiality and care at home for people with dementia. *Nursing Inquiry*, *27*(1), e12306.

Arnau, E., Estany, A., González del Solar, R., & Sturm, T. (2014). The extended cognition thesis: Its significance for the philosophy of (cognitive) science. *Philosophical Psychology*, *27*(1), 1–18.

Awan, N., Schneider, T., & Till, J. (2011). *Spatial agency: Other ways of doing architecture* (1st ed.). Routledge.

Aydin, C., González Woge, M., & Verbeek, P.-P. (2019). Technological environmentality: Conceptualizing technology as a mediating milieu. *Philosophy & Technology*, *32*(2), 321–338.

Bartlett, R., & O'Connor, D. (2010). *Broadening the dementia debate: Towards social citizenship*. Bristol, UK: Policy Press.

Benge, J. F., & Scullin, M. K. (2020). Implications for technological reserve development in advancing age, cognitive impairment, and dementia. *The Behavioral and Brain Sciences*, *43*(e157), e157.

Birhane, A. (2022). The unseen Black faces of AI algorithms. *Nature*, *610*(7932), 451–452.

Blott, J. (2021). Smart homes for the future of dementia care. *Lancet Neurology*, *20*(4), 264.

Buolamwini, J. (2023). *Unmasking AI: My mission to protect what is human in a world of machines*. Random House.

Buse, C., Martin, D., & Nettleton, S. (2018). Conceptualising "materialities of care": making visible mundane material culture in health and social care contexts. *Sociology of Health & Illness*, *40*(2), 243–255.

Calkins, M. P. (2018). From research to application: Supportive and therapeutic environments for people living with dementia. *The Gerontologist*, *58*(suppl_1), S114–S128.

Candiotto, L. (2016). Extended Affectivity as the Cognition of Primary Intersubjectivity. *Phenomenology and Mind*, *0*(11), 232–241.

Candiotto, L., & Stapleton, M. (2024). Assistive Technology as affective scaffolding. *Topoi: An International Review of Philosophy*, 1–10.

Carter, D. (2016). *Fix dementia care: homecare*. Alzheimer's Society.

Carter, M. (2022). The indirect approach: Towards non-dominating dementia care. *Res Publica (Liverpool, England)*, *28*(3), 467–480.

Carter, M. (2024). The person as environmentally integrated. *Journal of Ethics and Social Philosophy*, *28*(1). https://doi.org/10.26556/jesp.v28i1.3185

Carter, L., University of New South Wales, Liu, D., Cantrell, C., University of New South Wales, & Virginia Commonwealth University. (2020). Exploring the intersection of the digital divide and artificial intelligence: A hermeneutic literature review. *AIS Transactions on Human-Computer Interaction*, *12*(4), 253–275.

Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford University Press.

Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, *58*(1), 7–19.

Clowes, R. W. (2019). Immaterial engagement: human agency and the cognitive ecology of the internet. *Phenomenology and the Cognitive Sciences*, *18*(1), 259–279.

Colombetti, G., & Krueger, J. (2015). Scaffoldings of the affective mind. *Philosophical Psychology*, *28*(8), 1157–1176.

Colombetti, G., & Roberts, T. (2015). Extending the extended mind: the case for extended affectivity. *Philosophical Studies*, *172*(5), 1243–1263.

Coninx, S., & Stephan, A. (2021). A Taxonomy of Environmentally Scaffolded Affectivity. *Danish Yearbook of Philosophy*, *54*(1), 38–64.

De Witt, L., Ploeg, J., & Black, M. (2009). Living on the threshold: The spatial experience of living alone with dementia. *Dementia (London, England)*, *8*(2), 263–291.

Dementia: a public health priority. (2012, August 24). World Health Organization. https://www.who.int/publications/i/item/dementia-a-public-health-priority

Dennett, D. C. (1996). *Kinds of Minds*. Basic Books.

Drayson, Z., & Clark, A. (2020). Cognitive disability and embodied, extended minds. In A. Cureton & D. T. Wasserman (Eds.), *The Oxford Handbook of Philosophy and Disability* (pp. 579–597). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190622879.013.10

Dung, L. (2024). Understanding artificial agency. *The Philosophical Quarterly*, qae010.

Ebert, F., Finn, C., Lee, A. X., & Levine, S. (2017). Self-Supervised Visual Planning with Temporal Skip Connections. *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, *78*, 344–356.

Edwards, D., Baum, C., & Morrow-Howell, N. (1994). Home environments of inner city elderly: Do they facilitate or inhibit function. *The Gerontologist: Program Abstracts of 47th Annual Scientific Meeting*, *34*(1), 64.

Farina, M., & Levin, S. (2021). The extended mind thesis and its applications. In *Handbook of Embodied Psychology* (pp. 127–147). Springer International Publishing.

Figà-Talamanca, G. (2024). Digitally scaffolded vulnerability: Facebook's recommender system as an affective scaffold and a tool for mind invasion. *Topoi: An International Review of Philosophy*, *43*(3), 631–643.

Førsund, L. H., Grov, E. K., Helvik, A.-S., Juvet, L. K., Skovdahl, K., & Eriksen, S. (2018). The experience of lived space in persons with dementia: a systematic meta-synthesis. *BMC Geriatrics*, *18*(1), 33.

Gallagher, S. (2018). The extended mind: State of the question. *The Southern Journal of Philosophy*, *56*(4), 421–447.

Gitlin, L. N., Winter, L., Dennis, M. P., Corcoran, M., Schinfeld, S., & Hauck, W. W. (2006). A randomized trial of a multicomponent home intervention to reduce functional difficulties in older adults: Multicomponent home intervention to reduce functional difficulties. *Journal of the American Geriatrics Society*, *54*(5), 809–816.

Gregg, M. (2011). *Work's Intimacy*. Polity.

Hall, W. J., Chapman, M. V., Lee, K. M., Merino, Y. M., Thomas, T. W., Payne, B. K., Eng, E., Day, S. H., & Coyne-Beasley, T. (2015). Implicit racial/ethnic bias among health care professionals and its influence on health care outcomes: A systematic review. *American Journal of Public Health*, *105*(12), e60-76.

Haque, U. (2023). Architecture, Interaction, Systems. In B. Farahi & N. Leach (Eds.), *Interactive Design* (pp. 42–47). De Gruyter.

Heersmink, R. (2015). Dimensions of integration in embedded and extended cognitive systems. *Phenomenology and the Cognitive Sciences*, *14*(3), 577–598.

Hernández-Orallo, J., & Vold, K. (2019). AI extenders: The ethical and societal implications of humans cognitively extended by AI. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 507–513.

Hughes, J. C. (2001). Views of the person with dementia. *Journal of Medical Ethics*, *27*(2), 86–91.

Hughes, Julian C. (2011). *Thinking through dementia*. Oxford University Press.

Hydén, L.-C., Lindemann, H., & Brockmeier, J. (2014). *Beyond Loss: Dementia, Identity, Personhood*. Oxford University Press.

Ienca, M. (2023). On Artificial Intelligence and manipulation. *Topoi: An International Review of Philosophy*, *42*(3), 833–842.

Ienca, M., Fabrice, J., Elger, B., Caon, M., Scoccia Pappagallo, A., Kressig, R. W., & Wangmo, T. (2017). Intelligent assistive technology for Alzheimer's disease and other dementias: A systematic review. *Journal of Alzheimer's Disease: JAD*, *56*(4), 1301–1340.

Jacobson, K. (2009). A developed nature: a phenomenological account of the experience of home. *Continental Philosophy Review*, *42*(3), 355–373.

Käll, L. (2017). Intercorporeal expression and the subjectivity of dementia. In L. Dolezal & D. Petherbridge (Eds.), *Body/Self/Other: The Phenomenology of Social Encounters* (pp. 359–385). SUNY Press.

Kerr, A. D. (2020). Artificial Intelligence, Gender, and Oppression. In *Encyclopedia of the UN Sustainable Development Goals* (pp. 1–11). Springer International Publishing.

Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended Consciousness and Predictive Processing: A Third Wave View*. Routledge.

Kirsh, D. (2010). Thinking with External Representations. *AI & Society*, *25*(4), 441–454.

Kontos, P. C. (2005). Embodied selfhood in Alzheimer's disease: Rethinking person-centred care. *Dementia*, *4*(4), 553–570.

Kontos, P., Miller, K.-L., & Kontos, A. P. (2017). Relational citizenship: supporting embodied selfhood and relationality in dementia care. *Sociology of Health & Illness*, *39*(2), 182–198.

Krueger, J. (2023). An ecological approach to affective injustice. *Philosophical Topics*, *51*(1), 85–111.

Krueger, J. (2024). Selves beyond the skin: Watsuji, "betweenness", and self-loss in solitary confinement and dementia. *Journal of Consciousness Studies: Controversies in Science & the Humanities*, *31*(5), 127–150.

Krueger, J., & Osler, L. (2019). Engineering Affect: Emotion Regulation, the Internet, and the Techno-Social Niche. *Philosophical Topics*, *47*(2), 205–231.

Krueger, J., & Roberts, T. (2024). Real feeling and fictional time in human-AI interactions. *Topoi: An International Review of Philosophy*, 1–12.

Krueger, J., & Szanto, T. (2016). Extended Emotions. *Philosophy Compass*, *11*(12), 863–878.

Kukla, Q. R. (2022). *City living: How urban spaces and urban dwellers make one another*. Oxford University Press.

Liao, S.-Y., & Huebner, B. (2021). Oppressive things. *Philosophy and Phenomenological Research*, *103*(1), 92–113.

Maiese, M. (2016). Affective Scaffolds, Expressive Arts, and Cognition. *Frontiers in Psychology*, *7*, 359.

Maiese, M., & Hanna, R. (2019). *The Mind-Body Politic*. Palgrave Macmillan.

Malafouris, L., & Röhricht, F. (2024). Re-thinging embodied and enactive psychiatry: A material engagement approach. *Culture, Medicine and Psychiatry*, 1–24.

Manmadhan, S., & Kovoor, B. C. (2020). Visual question answering: a state-of-the-art review. *Artificial Intelligence Review*, *53*(8), 5705–5745.

Mayo Clinic (2024). Dementia. Retrieved January 5, 2025, from https://www.mayoclinic.org/diseases-conditions/dementia/diagnosis-treatment/drc-20352019.

McCabe, L., Robertson, J., & Kelly, F. (2018). Scaffolding and working together: a qualitative exploration of strategies for everyday life with dementia. *Age and Ageing*, *47*(2), 303–310.

Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. New York: Palgrave Macmillan.

Merat, N., Seppelt, B., Louw, T., Engström, J., Lee, J. D., Johansson, E., Green, C. A., Katazaki, S., Monk, C., Itoh, M., McGehee, D., Sunda, T., Unoura, K., Victor, T., Schieben, A., & Keinath, A. (2019). The "Out-of-the-Loop" concept in automated driving: proposed definition, measures and implications. *Cognition, Technology & Work*, *21*(1), 87–98.

Moyle, W., Murfield, J., & Lion, K. (2021). The effectiveness of smart home technologies to support the health outcomes of community-dwelling older adults living with dementia: A scoping review. *International Journal of Medical Informatics*, 153(104513), 104513.

Müller, V. C. (2015). *Risks of artificial intelligence* (V. C. Müller, Ed.; 1st Edition). Chapman & Hall/CRC. https://doi.org/10.1201/b19187

Murphy, A. (2024). The spirited interworld: Caregiving and the liminal phenomenology of dementia. *Journal of Critical Phenomenology*, *7*(1), 57–68.

NHS. (2023). What is dementia. Retrieved December 8, 2024, from https://www.nhs.uk/conditions/dementia/about-dementia/what-is-dementia/

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.

Nolan, M., Ryan, T., Enderby, P., & Reid, D. (2002). Towards a more inclusive vision of dementia care practice and research. *Dementia (London, England)*, *1*(2), 193–211.

Nguyen, C. T. (2019). Games and the Art of Agency. *The Philosophical Review*, *128*(4), 423–462.

Osler, L. (forthcoming). Narrative Railroading. *Topoi*.

Petherbridge, D. (2019). Beyond Empathy: Vulnerability, Relationality and Dementia. *International Journal of Philosophical Studies*, *27*(2), 307–326.

Phenice, L. A., & Griffore, R. J. (2013). The Importance of Object Memories for Older Adults. *Educational Gerontology*, *39*(10), 741–749.

Piredda, G. (2020). What is an affective artifact? A further development in situated affectivity. *Phenomenology and the Cognitive Sciences*, *19*(3), 549–567.

Protevi, J. (2009). *Political Affect: Connecting the Social and the Somatic*. University of Minnesota Press.

Ramachandran, A. K., Das, S., Joseph, A., Shenoy, G. G., Alex, A. T., & Mudgal, J. (2021). Neurodegenerative pathways in Alzheimer's disease: A review. *Current Neuropharmacology*, *19*(5), 679–692.

Rani, V., Nabi, S. T., Kumar, M., Mittal, A., & Kumar, K. (2023). Self-supervised learning: A succinct review. *Archives of Computational Methods in Engineering. State of the Art Reviews*, *30*(4), 2761–2775.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive Offloading. *Trends in Cognitive Sciences*, *20*(9), 676–688.

Robson, J., & Meskin, A. (2016). Video games as self-involving interactive fictions. *Journal of Aesthetics and Art Criticism*, *74*(2), 165–177.

Rosenberger, R. (2017). *Callous Objects*. University of Minnesota Press.

Rowlands, M. (2010). *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, MA: MIT Press.

Ruane, E., Birhane, A., & Ventresque, A. (2019). *Conversational AI: Social and ethical considerations*. AICS. http://ceur-ws.org/Vol-2563/aics_12.pdf

Schneider, T., & Till, J. (2009). Beyond discourse: Notes on spatial agency. *Footprint*, 97–112.

Settles, B. (2012). *Active Learning*. Springer International Publishing.

Slaby, J. (2014). Emotions and the extended mind. In M. Salmela & C. Von Scheve (Eds.), *Collective emotions* (pp. 32–46). Oxford University Press.

Slaby, J. (2016). Mind Invasion: Situated Affectivity and the Corporate Life Hack. *Frontiers in Psychology*, *7*(266), 1–13.

Smart, P. (2017). Extended Cognition and the Internet. *Philosophy & Technology*, *30*(3), 357–390.

Soilemezi, D., Drahota, A., Crossland, J., & Stores, R. (2019). The role of the home environment in dementia care and support: Systematic review of qualitative research. *Dementia (London, England)*, *18*(4), 1237–1272.

Sutton, J. (2006). Distributed cognition: Domains and dimensions. *Pragmatics & Cognition*, *14*(2), 235–247.

Sutton, J. (2010). Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process. In R. Menary (Ed.), *The Extended Mind* (pp. 189–225). Cambridge, MA: MIT Press.

Sutton, J. (2024). Situated affects and place memory. *Topoi: An International Review of Philosophy*, *43*(3), 593–606.

Swanepoel, D. (2021). Does artificial intelligence have agency? In *The Mind-Technology Problem* (pp. 83–104). Springer International Publishing.

Tewes, C. (2021). Embodied selfhood and personal identity in dementia. In C. Tewes & G. Stanghellini (Eds.), *Time and Body. Phenomenological and psychopathological approaches* (pp. 367–389). Cambridge University Press.

Vallor, S. (2024). *The AI mirror: How to reclaim our humanity in an age of machine thinking*. Oxford: Oxford University Press.

Van Hoof, J., & Kort, H. S. M. (2009). Supportive living environments: A first concept of a dwelling designed for older adults with dementia. *Dementia (London, England)*, *8*(2), 293–316.

Vold, K., & Hernández-Orallo, J. (2021). AI extenders and the ethics of mental health. In F. Jotterand & M. Ienca (Eds.), *Artificial Intelligence in Brain and Mental Health: Philosophical, Ethical & Policy Issues* (pp. 177–202). Springer International Publishing.

Wang, C., Boerman, S. C., Kroon, A. C., Möller, J., & H de Vreese, C. (2024). The artificial intelligence divide: Who is the most vulnerable? *New Media & Society*.

Wheeler, M. (2019a). Breaking the waves: Beyond parity and complementarity in the arguments for extended cognition. In *Andy Clark and His Critics* (pp. 81–96). Oxford University Press.

Wheeler, M. (2019b). The reappearing tool: transparency, smart technology, and the extended mind. *AI & Society*, *34*(4), 857–866.

White, B., & Hipólito, I. (2024). Preventive mental health care: A complex systems framework for ambient smart environments. *Cognitive Systems Research*, *84*(101199), 101199.

White, B., & Miller, M. (2024). Ambient smart environments: affordances, allostasis, and wellbeing. *Synthese*, *204*(2), 1–24.

Wildman, N., & Woodward, R. (2018). Interactivity, Fictionality, and Incompleteness. In *The Aesthetics of Videogames* (1st Edition, pp. 112–127). Routledge.

Benge, J. F., & Scullin, M. K. (2020). Implications for technological reserve development in advancing age, cognitive impairment, and dementia. *The Behavioral and Brain Sciences*, *43*(e157), e157.

World Health Organization. (n.d.). Dementi*a*. Retrieved December 8, 2024, from https://www.who.int/news-room/fact-sheets/detail/dementia

Wu, J., Li, Q., Yang, G., Li, L., Senhadji, L., & Shu, H. (2023). Self-supervised speech denoising using only noisy audio signals. *Speech Communication*, *149*, 63–73.

Zeiler, K. (2014). A philosophical defense of the idea that we can hold each other in personhood: intercorporeal personhood in dementia care. *Medicine, Health Care, and Philosophy*, *17*(1), 131–141.