



Misplacing memories? An enactive approach to the virtual memory palace

Anco Peeters*, Miguel Segundo-Ortin

Faculty of Law, Humanities and the Arts, University of Wollongong, Australia

ARTICLE INFO

2010 MSC:

00-01
99-00

Keywords:

Method of loci
Enactivism
Functionalism
Extended mind
Virtual reality
Embodied cognition
Anti-representationalism

ABSTRACT

In this paper, we evaluate the pragmatic turn towards embodied, enactive thinking in cognitive science, in the context of recent empirical research on the memory palace technique. The memory palace is a powerful method for remembering yet it faces two problems. First, cognitive scientists are currently unable to clarify its efficacy. Second, the technique faces significant practical challenges to its users. Virtual reality devices are sometimes presented as a way to solve these practical challenges, but currently fall short of delivering on that promise. We address both issues in this paper. First, we argue that an embodied, enactive approach to memory can better help us understand the effectiveness of the memory palace. Second, we present design recommendations for a virtual memory palace. Our theoretical proposal and design recommendations contribute to solving both problems and provide reasons for preferring an embodied, enactive account over an information-processing treatment of the memory palace.

1. Introduction

Though the memory palace technique, a mnemonic making clever use of places and images, is enjoying newfound attention by researchers on virtual reality (VR), its use goes back centuries. According to one famous story, Giordano Bruno, a Napolitan philosopher and influential memory palace master, earned himself an accusation of plagiarism while presenting at Oxford in 1583. Apparently, one attentive Oxford don did not appreciate that Bruno, in a top-off-the-head lecture, recited long text passages from a contemporary scholar without a reference (Rowland, 2008, p. 146). Bruno's mnemonic use was careless, yet his memory feats remain impressive. Cognitive scientists have been trying to make use of the memory palace more accessible through visualising the technique's places and images in VR, but their efforts have so far yielded underwhelming results. In this paper, we address the issues surrounding recent attempts at operationalizing the memory palace through VR and we present a new and improved way of understanding the technique. Our proposal is inspired both by going back to the technique's roots and by insights from embodied, enactive cognitive science and should help towards solving the issues mentioned.

The main problem with mastering the memory palace technique is the time and effort involved. The technique takes long-term practice, in a suitable environment, and requires creative imagination. This explains why, given the strength of the technique, its use in education and training practices is not more prevalent. To increase accessibility of the memory palace, researchers have attempted to operationalize its use through VR devices. So far, it has proved hard to gain similar levels of remembering with the use of such devices when compared to traditional mnemonics.

* Corresponding author at: Faculty of Law, Humanities and the Arts (Bld. 19), University of Wollongong, 2522 NSW, Australia.
E-mail address: mail@ancopeeters.com (A. Peeters)

To make steps towards solving this issue, we propose to consider the difficulties in the translation of the memory palace into VR against the background of the so-called ‘pragmatic turn’ in cognitive science. The pragmatic turn signals a move towards conceiving of cognition as dynamic, embodied and enactive and away from cognition as information-processing ((Engel, 2010; Engel, Maye, Kurthen, & König, 2013)). Reframing how we think about the cognitive underpinnings of memory will help in the design of the virtual memory palace.

What is the advantage of examining the memory palace from the perspective of embodied, enacted cognition? We provide two related incentives. The first stems from the observation that current cognitivist investigations into the workings of the technique, which are based on the information-processing paradigm, have not shed sufficient light on why it is so powerful, as we will elaborate in the next section.¹ This opens the door to the consideration of an alternative paradigm. The second and related reason is that the memory palace, because it leans heavily on memory scaffolding through environmental resources, calls for a cognitive framework which places the role of the body in the environment front and centre.

Keeping in mind the pragmatic turn, our paper develops as follows. In Section 2, we will examine current cognitivist approaches to the memory palace technique and show how they are unable to explain its dynamics, concluding that there is, as we call it, an Explanation Problem. Following this, we will argue in Section 3 that current attempts to operationalize the memory palace in virtual reality fall short, because they depend on cognitivist understandings of the technique. Call this the Operationalization Problem. Because addressing the Operationalization Problem first requires addressing the Explanation Problem, we turn to the latter in Section 4, where we argue that an enactive account of the memory palace captures the technique better than its cognitivist rivals. This sets the stage for Section 5, in which we address the Operationalization Problem by presenting design recommendations for designers of virtual memory palaces based on our proposed enactive account. In doing so, we will rely on influential theories in embodied cognition, such as ecological psychology (Gibson, 1979; Chemero, 2009). We conclude with some considerations on the application of the virtual memory palace in educational settings and for future lines of research.

2. The memory palace in cognitive science

Much of our understanding of the memory palace is derived from historical sources. In her titular and seminal book on the art of memory, historian Frances Yates (1966) develops a now classic account of the memory palace. Drawing on instructions by Roman rhetoricians like Cicero and their further development by Bruno, she explains that the memory palace strategy rests on two pillars: *loci* (places) and images.²

A *locus* is characterised as part of a spacious environment with distinct features. Classic examples of such environments include large and varied buildings with decorations inside, such as churches and cathedrals. Environmental parts which qualify as *loci* are usually those that stand out when one would take a familiar route through the environment, such as a gargoye statue at the entrance, or a niche under a window. *Loci* and images play a role during both the learning and the recalling phase of the technique. In the learning phase, one moves through the building (preferably physically) and has to imagine placing images of that which has to be remembered at specific locations in and around the building. Then, during the recalling phase, one imagines moving through the building and gets triggered by the images positioned there to reconstruct the memory. It is advised to use vivid and personally resonating images for maximum recall-effect.

To illustrate the use of the technique and draw out some important aspects, let us imagine the following. While applying the technique to a talk on robot ethics, I choose the Sydney Opera House as my locus of choice. During the learning phase, I physically move around the Sydney Opera House. Initially I imagine a porter at the entrance who holds a copy of Isaac Asimov’s *I, Robot*. Moving on, I approach the Opera House’s wardrobe, where I imagine Aristotle arguing with Immanuel Kant while Jeremy Bentham hands his head to a robot attending the cloakroom. I continue to move around and create and place images for every part of my talk. When I am ready to present the talk I enter the recalling phase. During that phase, Asimov’s book serves to remind me that I need to start my talk by presenting the three laws of robotics, both as an introductory ‘hook’ and to mark them as a starting point in robot ethics. The image of Aristotle and Kant arguing, trigger me to say that virtue ethics and deontology might have something to say on robotics, though both theories are not dominant in current discussions. This is where the image of Bentham comes in, as it cues me to say that utilitarianism is currently the dominant theory in debates on robot ethics. The vividness and personal quality of the images will help me remember, and placing them at specific positions in the locus will help me to order my recollection. The use of personal imagery in combination with the scaffolding of memories through environmental cues are the defining features of the memory palace.

An impressive study by Maguire, Valentine, Wilding, and Kapur (2002), on the functional and neurological differences between normal and high-performing memorizers, shows that the memory palace technique is much alive today. Of the high-performing memorizers, drawn from a pool of participants of the World Memory Championships, 90% report using the technique for some or even all of their tasks. The goal of the study was to capture the possible causes that could differentiate superior memorizers from normal ones. As expected, the superior memorizers performed significantly better in tests on both working and long-term

¹ We take inspiration from recent critique on symbolic interfacing with augmented reality devices. Raja and Calvo (2017) argue that instead of programming augmented reality glasses (like Google Glass) to navigate spaces uses symbols and icons like arrows and text (cf. Clark, 2003, p. 52), such devices would instead function better if they leverage their user’s sensorimotor capacities through changes in brightness. Froese (2014) provides a similar, generalized critique of symbolic interfaces.

² In fact, the technique is often called *method of loci* (MOL), though this is a bit of a misnomer as it puts undue focus on the first of the two pillars.

verbal memory. No differences in terms of general intellect or brain structure between the two groups were found. However, functional brain-imaging showed that the superior memorizers, in contrast with the controls, had consistent higher activation levels in the medial parietal cortex, retrosplenial cortex, and the right posterior hippocampus. These regions are “known to be important for memory, and are implicated in spatial memory and navigation” (p. 93). Unsurprisingly, these brain areas showed increased activity during the learning phase of the task. Thus, Maguire and colleagues conclude that mnemonics like the memory palace, which they defined as “strategies for encoding information with the sole purpose of making it more memorable” (p. 93), constitute the main explanatory cause for the performative difference between superior and normal memorizers. The memory palace technique provides the “top participants of the annual World Memory Championships ...the ability to memorize hundreds of words, digits, or other abstract information units” and is therefore called the “most prominent mnemonic technique” (Dresler et al., 2017, p. 1227).

As of yet, there is no single explanation for why the memory palace technique is so effective. There is nonetheless a suspicion that the “additional motor imagery aspect is likely the reason the method of loci has been found to be particularly effective—a connection that has not been previously made” (Madan & Singhal, 2012, p. 220). This in contrast to other memory strategies which often solely depend on visual imagery. However, it is unclear exactly why motor imagery in combination with visual imagery would explain the effectiveness of the memory palace as a cognitive technique.

Moving further down these lines of thought, Dresler et al. (2017) hypothesize that with the memory palace technique “abstract and unrelated information units are transformed into concrete and related information patterns that can more easily be processed by memory-related brain structures, such as the hippocampus” (p. 1232). But what does it mean to say that “concrete and related information patterns” are more easily processed by brain structures? What does the memory palace technique do which transforms a random deck of playing cards from “abstract and unrelated information units” into “concrete and related information patterns”? This transformation seems to presuppose two types of information: abstract and concrete. Are there such different kinds, and, if so, why is concrete information more easily digested? We will take a closer look at this issue in Section 4.

The relation between, on the one hand, Yates’ account of the memory palace as deeply dependent on both the environment for structure and the individual for creating images, and, on the other, the information-processing paradigm of the previously discussed experiments, remains underdeveloped. The support of the environment is, in this paradigm, defined as the ordering of information units which are processed by a cognizer’s brain. But that this reordering allows for more efficient information processing is, at best, in need of further explanation, or, as we will argue, a fundamentally flawed approach to the understanding of the technique. Let us call this issue the Explanation Problem.

3. The virtual memory palace

The Explanation Problem, as we argue in the next section, lies at the root of why efforts at making the memory palace accessible through VR devices, are not yielding results comparable to traditional memory palace practice. Why is there a need for a ‘virtual memory palace’? Memory theorists have observed that the “primary flaw of mnemonics is that effective use often requires extensive practice” (Madan, 2014, p. 3). And, specifically in the case of the memory palace, not only does it take practice but it also takes time to familiarize oneself with a large and spacious building and to translate what one wants to remember into images which can then be placed in and around that building. Moreover, the learning phase can be extra problematic for someone who may not always have ready access to a locus that fits the described purpose. Large, easily accessible buildings fit for practice are after all not always available when one wants to, for instance, practise and memorize a talk. Furthermore, the creating-and-placing-the-images phase of the memory palace technique depends on having a creative imagination to come up with evocative pictures which translate to whatever it is one would like to remember. So while the memory palace is acknowledged as a powerful mnemonic technique, potential users are often hesitant to go through the effort of learning it.

Virtual reality technologies might hold an answer to the previously outlined challenges. Virtual environments can be tailor-made for and readily accessible to the memorizer and, when a database of (personalizable) three-dimensional models is provided, the creation of a fitting image for a certain idea in a speech would not be so complicated. The time it takes to practice the mnemonic would also decrease when a virtual environment is available, as there is no need to physically travel to a suitable environment or spend time conjuring up an imagined one. In the words of Jund, Capobianco, and Larue (2016), given “its intrinsic spatial nature, VR seems to offer the perfect technology devices to implement ...[the memory palace]. Not only [does] it allow ...immersive exploration of any given architectural environment, but it also provides rich sensory cues (spatial contiguity, optic flow, self-directed navigation)” (p. 533). In theory, virtual reality seems to be made, as it were, for the memory palace technique.

Early research on investigating the memory palace through the lens of virtual reality aimed to establish whether virtual environments can support the memory palace technique as well as conventional, physical environments do. In an initial and exploratory study, Fassbender and Heiden (2006) found that participants who interacted with a virtual environment through the use of a personal computer and desktop monitor remembered images from that virtual environment better than words from a sheet of paper. This study is limited because different types of items were compared – images with words – in a within-subject design without randomisation, and there was no between-subject comparison that compared the virtual memory palace to a conventional one. Furthermore, more immersive interfaces than a desktop computer monitor are now available for a consumer market. Higher levels of immersion in virtual environments, specifically in terms of field of vision, improve performance on memorization (Ragan, Sowndararajan, Kopper, & Bowman, 2010). This shows that it is preferable to use, for example, a head-mounted display (HMD), rather than a desktop computer monitor to interface with a virtual environment (see also Huttner & Robra-Bissantz, 2016).

In a foundational study on the virtual memory palace, Legge, Madan, Ng, and Caplan (2012) addressed the question of whether the memory palace technique works as well with aid from a virtual environment as from a physical one. In order to test this, the experimenters assigned participants to three groups: a traditional memory palace group, a virtual memory palace group, and a control group. All participants first practised on a memory task, recalling lists of words, then moved through a virtual environment, and finally performed another memory task similar to the first. The first two groups were asked to use the memory palace on the second task, with the former imagining familiar place like their home and the latter imagining the virtual environment just before encountered. The third group were not given a specific strategy to use.

The results of Legge et al.' (2012) research confirm that a virtual environment does not perform worse than a conventional space. However, at least two critical remarks can be made about the study. First, the participants in the study were not present in the virtual environment during the learning phase of the memory task. Instead, they were shown the virtual environment for five minutes and those in the virtual memory palace group were then asked to use their memory of the virtual space for their task. Hence, the study does not speak of how effective the memory palace technique could be when the whole learning phase is performed in a virtual environment. Second, the level of immersion in the virtual environment was again quite low: the environment was shown on a desktop monitor and movement occurred by means of mouse and keyboard. This runs counter to the theory of the conventional memory palace where an active, bodily involvement from the memorizer, in terms of navigation and image placement in the loci, is supposed.

In an effort to make the virtual memory palace a more immediate and immersive experience, Jund et al. (2016) present a study in which participants engaged with a virtual environment by means of an HMD that provided a stereoscopic image. Three types of environments were presented. In the first, participants were sequentially and briefly shown items for remembering in the same frontal virtual position, without spatial cues. In the second, participants were sequentially shown items to remember, with each item briefly appearing next to the location of the previous item in the virtual environment. No further spatial cues were given. The first two conditions were categorised as 'egocentric'. In the third, participants were guided through a virtual apartment with nine different rooms. In this third condition, categorised as 'allocentric', participants used a passive navigation technique: they were moved along a preprogrammed path and could only move forward by pressing a key. Jund and colleagues were surprised to find that the egocentric conditions resulted in better memorization than the allocentric condition. In a follow-up experiment, they adjusted the third condition and found that participants performed significantly better when using a virtual environment of a familiar building. We do not think this result is surprising as per Yates' (1966) suggestion that the memorizer should use a building which is intimately familiar to them. In the next sections, we argue that an essential cognitive part of the memory palace technique is the training of a cognizer's memory in such a way that it allows for effortless re-imagining of the building in question. In a manner of speaking, such a memorizer would carry the building with them, though we emphasise this should not be understood representationally. However, even with this performance improvement on the allocentric condition, Jund and colleagues found that this condition still did worse than the egocentric ones.

We point out two likely aspects which may help explain the poorer results in the allocentric condition when compared to the egocentric ones in the study by Jund et al. (2016). Both figure in the learning phase of the memorization process. First, the participants could only indicate the moment of movement, upon which they were passively moved along a pre-set path. Second, the participants were presented with images, rather than given the opportunity to create and actively place images in the virtual environment. Both aspects signify the passive relation of the participant to the employed environment and this runs counter to the active anchoring as described by Yates (1966). Jund and colleagues seem to agree, at least on the first point, when they conclude that "the navigation technique and sensory cues associated with displacement might be of primary importance when it comes to use spatial information to support memorization" (p. 537). A new and improved experimental design would be required to determine whether our proposal holds merit, though, and we will provide a design suggestion in Section 5.

In a study designed to determine whether immersive HMD interfaces perform better in memory tasks than desktop computer monitors, Krokos, Plaisant, and Varshney (2019) take an embodied and embedded approach to the virtual memory palace. Unsurprisingly, they found that the increased immersion of an HMD allows for better memory recall than a traditional desktop monitor. Of even more interest are the peripheral observations they made regarding the manner of interaction between participants and virtual environment. About a third of the participants "mentioned that they actively used the virtual memory palace setup by associating the information relative to their own body" (p. 10). The authors further remark on the previously discussed tension between active and passive movement through an environment. They refer to Barbara Brooks (1999), who found that active movement allows for more accurate familiarisation with an environment when compared to passive movement. However, as the same study also concluded that the manner of movement, namely whether it was active or passive, had no influence on the recall of items or their positions in the environment, Krokos, Plaisant and Varshney suggest that "memory was only enhanced for those aspects of the environment that were interacted with directly – particularly the environment which was navigated" (p. 4). It should further be noted that Brooks' findings are based on a traditional desktop computer monitor interface with mouse and keyboard, and it would be of interest to redo his experiment with an HMD and direct, haptic interaction of the participants.

Until now, research on the virtual memory palace has presented the memorizer as a somewhat passive participant. We think the observations made by Krokos et al. (2019), on the role of the body in (virtual) environments, merit closer attention if we are to properly understand the memory palace technique and develop appropriate interfaces for it – like, for example, via haptic controllers. In line with Krokos and colleagues, we propose to have future experiments assign free movement to the memory palace users in VR. But we suggest departing from this experiment in two ways. First, the images used for testing in the virtual

environment were pre-given, while masters of the memory palace emphasise using personalized imagery for stronger memory evocation. Second, the order of images in the virtual environment was signalled by symbols (the numbers 1, 2, and so on). In Section 5, we present a way of using lighting to direct the user's attention in virtual environments, to move away from symbolic cues.

With this review of current developments in the field of the virtual memory palace in place, we conclude there is currently no conclusive answer to the question of whether a fully immersive approach, with head-mount display and haptic controllers, can perform as well as (or even better than) conventional memory palace techniques. This means that there is a need for research which compares memory performance of memory palace practitioners both using a conventional memory palace and a virtual one.³ It would furthermore be interesting to compare the performance of memory palace practitioners not using a virtual memory palace with ordinary subjects using a virtual memory palace, to establish whether VR operationalization of the memory palace is on par with traditional usage. Based on our interpretation of Yates (1966) in relation to our review of current scientific approaches to the virtual memory palace, we surmise that new research needs to take at least the following into account. First, such an approach needs to investigate what sensory and navigational cues can best support the memory palace. Second, the role of the body in virtual environments needs to be more pronounced than it has been, specifically in terms of how the body is virtually reproduced and whether a haptic interface to the architecture of the locus and the placement of images can enhance the technique. Third, this approach has to promote the active engagement of the memorizer to navigation, choice of loci, and choice of image. Let us call this challenge, to integrate embodied implementations of the memory palace in VR, the Operationalization Problem. It should be clear by now that addressing the Operationalization Problem requires rethinking our cognitive approach to the memory palace, in other words, it requires addressing the Explanation Problem.

4. Addressing the Explanation Problem

In addressing the Explanation Problem, we consider two different and competing frameworks which put the embodiment and embeddedness of the cognizer in a larger environment centre stage: extended functionalism and enactivism. In what follows, we connect the memory palace to broader debates on embodied, extended cognition and evaluate the two proposals just mentioned. Our conclusion is that the enactive approach offers more powerful resources to account for the effectiveness of the memory palace than its functionalist competitor.

Our examination starts from recent suggestions made in cognitive anthropology and philosophy of mind. Cognitive anthropologist Edwin Hutchins (2005, p. 1564) recounts that the memory palace makes

opportunistic use of space. The spatial relations of the landmarks do not contribute any semantic content to the problem. But the landmark themselves do provide memory cues, and the sequential relations among the landmarks, that were created by mapping a particular shape of motion onto them, is inherited by the set of items to be remembered.

This seemingly supports the idea, outlined in Section 2, that smart rearrangement of 'concrete and related information patterns' allows such patterns to be more easily processed. However, understanding Hutchins this way would skirt over a crucial difference between his description and the currently salient idea on the memory palace in neuroscience. Instead of focusing on how information patterns might be picked up by the brain, Hutchins, using terms like 'landmark' and 'motion', rightly emphasizes the role of environmental triggers to cue memories and of bodily movement to help in the ordering of them.

The relevance of environmental resources to thinking about the memory palace has also been emphasised by John Sutton (2007). Using the distinction between engrams, or biological memory, and exograms, or external memory carriers, Sutton interprets the physical environments the memory palace technique relies on – like the Sydney Opera House in our example – as "prostheses" or "internalized exograms." Such prostheses, he adds, should be seen as "structuring supplements which construct and maintain the biological processes which they simultaneously and deeply transform" (p. 27).

We will now consider the contribution of such environmental resources from the perspective of extended functionalism. Extended functionalism aligns with current information-processing accounts that we have discussed in the previous sections and can be traced back to Clark and Chalmers' (1998) classic paper on the extended mind. In this paper, they question the traditional cognitive boundaries of skin and skull and argue that mind can sometimes be constituted by parts of the environment. Clark and Chalmers argue their point by way of their famous thought experiment about Inga and Otto. Inga and Otto are both looking to visit the Museum of Modern Art while in New York. But whereas Inga uses her biological memory to recall the museum's address, Otto, a sufferer of early-onset Alzheimer's, retrieves it through his notebook. The notebook, Clark and Chalmers argue, plays the same role for Otto that biological memory plays for Inga (p. 13). In it, Otto stores the things he would like to remember and his daily routine depends structurally, not incidentally, on his writing – similar to how Inga depends on her biological memory. It is important to note that this constitution claim is stronger than the trivial claim that mind is (merely) causally affected by the environment.

³ Another way to look at virtual memory palaces is through the lens of augmented reality devices. In a study performed at the MIT Media Lab, Rosello, Exposito, and Maes (2016) present the NeverMind application. NeverMind is designed to run on spectacles or 'smart glasses' which can project images on existing physical locations in the field of vision of the user. The preliminary study found that images projected along a route with NeverMind were better remembered than a list of words on a paper. While definitely an interesting approach, NeverMind still depends on having an appropriate physical environment available. Furthermore, it suffers from the same passive involvement of participants as the studies of Legge et al. (2012) and Jund et al. (2016). As such, it falls beyond the scope of our paper.

Early extended mind theorists stressed the idea that physical boundaries do not demarcate the mental and argue for this by way of the so-called parity principle. The idea is that “[i]f, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process” (Clark & Chalmers, 1998, p. 8). The parity principle encourages us to think that restraining cognitive processes merely to, for example, the brain, would be a case of misplaced neural chauvinism.

The parity principle is the main reason the extended mind is usually seen as part of the larger cognitive programme of functionalism (Clark, 2008; Wheeler, 2010, 2015), roughly the idea that mental states are to be defined and characterized by the job they perform. Focusing on functions, instead of material realizers, opens up the way to think that some cognitive processes can be implemented, at least partly, by elements outside the skull. Therefore, theorists working on functionalism are neutral with respect to the whereabouts of cognition, thus providing a natural home for the extended mind thesis.

So how exactly does the memory palace relate to the extended mind hypothesis? Sutton (2010) proposes that, even though mnemonic devices such as the memory palace are not literal external artefacts, the structures they provide function much like Otto’s notebook. In this way, Sutton expands the reach of the initial extended mind hypothesis by arguing it can capture not only natural and biological objects, but also cultural practices. He therefore concludes that “taking EM [Extended Mind] seriously ... means that we treat such architectures, systems, and practices as both cognitive and extended whether or not they happen to be outside the skin” (p. 209).

Let us then give a tentative account of the memory palace according to an extended functionalist framework. As said previously, mental states are, for the functionalist, to be understood in terms of the job they perform.⁴ Extended functionalists cast these jobs in terms of information-processing – recall Inga and Otto and that “the information in the notebook functions just like the information constituting an ordinary non-occurrent belief” (Clark & Chalmers, 1998, p. 13). Biological memory is, on this framework, understood as a process which involves the storing and retrieving of informational content, where this content is “sitting somewhere in memory waiting to be accessed” (Clark & Chalmers, 1998, p. 12). When an event is experienced, some piece of information is stored to be later retrieved when required. It has to be noted, however, that the extended functionalist would emphasize that it “doesn’t matter whether the data are stored somewhere inside the biological organism or stored in the external world. What matters is how information is poised for retrieval and for immediate use as and when required” (Clark, 2003, p. 69). In light of this framework, we could understand the memory palace technique as a way of structuring the contents and marking them through image-association. During retrieval, the memorizer recollects the relevant contents while she imagines walking through the palace. The images are encountered, the information they encode picked up, and integrated into that which was to be remembered. On this account of extended memory, remembered contents are conceived of as accessible, objective commodities (see Loader, 2013, p. 167).

This type of canonical, “first wave” (Sutton, 2010) extended cognition thinking seems to come some way in explaining the memory palace. It helps us to think of the memory palace as a cognitive structure which supports the memorizer in placing images in a particular order. However, there are two flaws with the current functionalist explanation. First, though it putatively captures the role the environment plays in the process of encoding and retrieving information, it neglects to explain why the role of bodily movement in both learning and recall phase of the memory palace is of importance. Second, it is unclear how, on this account, the extra information the memory palace would presumably require being processed during the recall phase, actually helps with remembering.

Some extended functionalists, however, have enriched their account to accommodate the role of the body. Clark (2008), in advocating extended functionalism, proposes two different takes on the role of the body. On the one hand, there is what he dubs the ‘Larger Mechanism Story’ (LMS), while, on the other, we find the ‘Special Contribution Story’ (SC). These two stories are explanatorily competitive in that they each assign a different role to the body in the context of embodied cognition.

On LMS, the body is thought to play a specific role on the larger information-processing mechanism. To illustrate, Clark (2008) compares the mental calculation of a sum by a human with how a snake, called Adder, may slither across the keys of an electronic calculator in such a way as to achieve a similar result. He concludes that in both cases the same cognitive operation is performed. The process of the snake’s body moving over the keys is functionally equivalent to whatever activity the brain putatively performs to process the relevant information. Because the calculation of the sum is defined in terms of symbol manipulation, extended functionalists can abstract away from the specific material implementations of the calculation and, as such, consider that the body of the snake is no more special than whatever parts of the brain realize these operations. Clark associates LMS with the general (extended) functionalist agenda.

The story is, unsurprisingly, different for SC. On SC, as advocated by Lawrence Shapiro (2004, 2019), the role of the body is not that of one informational piece of the puzzle among many. Instead, as the name implies, those who adhere to SC advocate that at least some of the contributions the body makes are not reducible to mere informational processes. The implication is that some of an organism’s cognitive processes are shaped by the specific features of its body in a way that does not lend itself to an explanation in terms of information-processing. Shapiro specifies that there are at least two ways in which the body may influence cognition: “first, it might generate associations that determine certain cognitive proclivities; second, the body might, via activation of motor plans, facilitate or inhibit various cognitive processes” (p. 12). Thus, on SC, for the understanding of at least some cognitive processes the consideration of the role of the body is required.

⁴ For a current and general functionalist account of memory, see Fernández (2018).

To justify the body's role in shaping cognition, Shapiro (2019) draws on empirical sources. Illustrating the first path of the body's influence, he cites research which shows that right-handers prefer to interact with objects on their right side, and left-handers on their left. The idea is that the increased ease with which people interact with objects on their dominant side informs their concept of "good" or "preferred" (Casasanto, 2009, 2014). How would that human preference for one's dominant hand be translated to LMS with a functional description such that a handless organism would exhibit similar cognitive dispositions? Or, as Shapiro (2019) puts it, should "we expect Adder to prefer objects to its right or its left given that it has no hands?" (p. 11).

Empirical evidence supports the notion that at least certain acts of memorizing depend on a special contribution from the body, and we can divide those into the two pathways distinguished by Shapiro. In terms of the first way, that of association, research in psychology has uncovered the relevance of the context-dependence of memory (Smith & Vela, 2001). One foundational study in this regard showed that divers who memorized material while under water better recalled those materials while being under water, while material learned on dry land was better recalled on land (Godden & Baddeley, 1975; Sutton & Williamson, 2014). In terms of the second way, we can draw on the idea that the activation of motor plans are relevant in acts of memorizing, particularly those acts of memory which involve the unfolding of a sequence. I might, for example, try to remember my PIN code by, physically or imaginatively, moving my fingers in its familiar pattern, or recall the order of the alphabet by mouthing parts of it. Scientific research supports this idea, showing that a specific starting point and reenactment through bodily movements is involved in the recollection of interconnected sequences both in musical parts (Ginsborg & Sloboda, 2007; Leman & Maes, 2014; Chaffin, Demos, & Logan, 2016) and dance phrases (Kirsh, 2013; Stevens, Malloch, McKechnie, & Steven, 2003). On this account, humming a tune or moving one's foot involves the triggering of the next instance in a sequence, domino-style, by the instantiation of its predecessor.

Contextual relevance and the unfolding of familiar patterns are both distinctive aspects of the memory palace technique. Yates (1966, p. 4) stresses that the strength with which a memory is triggered depends on carefully crafted and intense images. Furthermore, the whole environment of a memory palace may contribute to the act of associative recall, as with the divers underwater. Similarly, the sequence with which the images are encountered at the different loci and, as mentioned previously, the neuroscientific evidence of brain areas normally associated with navigation activating during the technique together point towards the idea that motor plans unfold offline during the recall phase (see Section 2). Such relations between the role of the body and the memory palace do not conceive of "the body as playing an information-processing role in cognition" Shapiro (2019, p. 9) and so the LMS, as cast in its familiar functionalist garb, is unable to adequately capture the memory palace.

For these reasons, we propose to look at an enactivist theory of mind and memory that is, we argue, better able to explain the special contribution of the body in acts of remembering. Enactivism understands cognition not in terms of the processing of information, but in terms of the participation of an organism in sensorimotor loops of active engagement within the context of a larger environment (Varela, Thompson, & Rosch, 1991; Thompson, 2007). Evan Thompson (2007), one of enactivism's main architects, suggests that remembering is better understood, not as the retrieval of a mental image, but as the reproduction of a person's past experience and that it "could involve emulating earlier sensory experiences and thus reenacting them in a modified way" (p. 291).

Enactivists of a radical stripe have further developed this line of thought, casting remembering as a dynamical, re-creative act. Radical enactivists argue that basic forms of cognition do not involve mental representations (Hutto & Myin, 2013, 2017). In line with this research programme, Hutto and Peeters (2018) put forward the idea that procedural memory "can be understood as the capacity to reenact embodied procedures often prompted and supported by patterns of response that are triggered by external phenomena" (p. 105). Rather than depending on the metaphor of memory as the storage and encoding of information, a radically enactivist take on procedural memory "would focus not on access to the contents of a store but on remembering as a type of action" (Loader, 2013, p. 168). Familiar patterns of response are initiated by internal or external triggers. For example, the remembering of how to prepare a specific meal is triggered by the ingredients and tools which are available to the cook. These familiar patterns involve the activation of trained neural configurations, which, according to context and circumstance, enable specific acts (see Anderson, 2014; Anderson, 2015). Following a recipe in order to prepare a meal is, on this account, not the retrieval of the stored information on that recipe, but the re-enactment of the different steps required to make dinner according to external signposts (the onion is glazed) which direct the individual to follow a specific familiar path (lower the fire).

Procedural memory is in current debates commonly characterized as not relying on information-processing (Michaelian, Debus, & Perrin, 2018), but enactivism is not limited to accounts of procedural memory per se. Recently, a number of scholars have proposed that episodic memory centrally involves the construction and consideration of possible past episodes through simulative imagining (Gerrans & Kennett, 2010; De Brigard, 2014; Michaelian, 2016). Such proposals assume that episodic acts of remembering, because of their simulative nature, necessarily involve representational content. Memory theorist Kourken Michaelian (2016), who agrees that understanding procedural memory need not depend on positing representational content, claims, by contrast, that appealing to contents in the case of episodic memory is essential. The reason is that episodic memory is declarative: it is available to consciousness and affects behaviour (pp. 27–28). However, why not allow that episodic memories, like the remembering of a conversation last week, is an act of, perhaps imperfect, simulative reconstruction through which a proposition with the content of that conversation is formed and available to consciousness? That this is indicative of current thinking about memory is shown by Michaelian, who recently argued that radically enactive remembering aligns well with an emerging tendency in discussions of philosophical of memory which cast remembering as non-contentful (Michaelian & Sant'Anna, 2019). In following Hutto and Peeters (2018), we see no need to assume that all acts of remembering through simulative re-enactment depend on the

manipulation of informational content. We maintain that acts of memory, such as using the memory palace, can be explained in a non-representational way.⁵

Applying an enactive account of memory to the memory palace then leads us to the following theory. In the remembering phase, the memorizer would either walk or imagine walking through an appropriate environment, such as the Sydney Opera House, with which she has become intimately familiar through active, bodily exploration. The order of the loci in the environment ensures that they are sequentially triggered during the recall phase, but it is up to the memorizer to ensure that the loci are then associated with the images to be remembered. During the recall phase, the memorizer will use her imagination to sequentially reconstruct the environment through the familiar triggers. For example, in the case of the Sydney Opera House, she would not remember the Opera House as a whole. Instead, she would reconstruct the relevant features while she images walking through it, letting the triggers guide her. Because of the learned association with the images, these images will spring to mind and can then be used by the memorizer to reconstruct whatever it is she would like to remember. The previously discussed findings by Dresler et al. (2017), on the structural rearrangement of neural networks for users of the memory palace, can then be reinterpreted as the construction of a network which enables the triggering sequence – in essence, a well-practised user of the memory palace carries the triggers of its loci with her. The user of the memory palace is, on the enactive account, not picking up information but reconstructing something resembling that which she was supposed to originally remember.

Though enactive remembering seems well suited to explain the role of embodiment in the memory palace while those bodily engagements are not straightforwardly intelligible in information-processing terms, extended functionalists may counter with an adjustment to their theory. In a striking experiment, Mackay, Fayard, Probert, and Médini (1998) investigated the adaptation of new electronic air-strips at an airtraffic station. In the late 1990s, traditional paper-made strips contained information about speed and direction of incoming airplanes and were used as an integral tool in the safe control of air traffic around Paris. Researchers were tasked to investigate how the use of such strips could be improved or even replaced with electronic devices. Initial trials with replacing the paper-based system with a computer-based one met with resistance by the traffic controllers. Advocating extended functionalism, Michael Wheeler (2010) observes that, from the perspective of an engineer, “one is inclined to focus, naturally enough, on the information carried by these strips. But this is not the only contribution of the strips.” (p. 33). It turns out that the strips were used in ways beyond merely carrying information. For example, they may be held in the hand as a reminder, placed at an angle to indicate two planes on a potential collision course, or, supported by the use of a strip-holding board, afford the signaling of important flight movements through body language. Wheeler’s analysis is worth quoting in full:

From a practical perspective, this recognition of the non-informational contribution of the flight strips is far from idle. The testimonial evidence suggests that a number of previous attempts to introduce new computer technology into air-traffic control may ultimately have been rejected as unworkable by the controllers precisely because *the proposed replacement systems attempted to reproduce the straightforwardly informational aspects of the flight strips while ignoring the extra factors.* (Wheeler, 2010, p. 33, emphasis added.)

Wheeler concludes that “nothing about this story undermines the extended functionalist line” (p. 33). This implies that the extended functionalist’s story either needs elaboration on the differences between ‘straightforward informational aspects’ (like the writing on the strips) and material informational factors (like the orientation of the strips), or that it need not be an information-processing story exclusively. Extended functionalism, as advanced by Wheeler, can thus allow for the materiality of artefacts, such as flight strips, to implement cognitive states as well, because it is neutral with respect to what cognitive states are made of.

Allowing extended functionalism to go beyond merely information-processing by recognizing the material roles artefacts play, looks like a promising move to give a functionalist account of the memory palace. As Wheeler admits, though, his proposal needs further analysis. We see two paths which the extended functionalist could take. The first one is to develop an account which explains the interplay between the informational processing of memories and the role the body plays when walking, imaginatively or not, through the memory palace. Recall that cognitive scientists currently explain the memory palace technique as somehow transforming abstract information units into concrete information patterns. The functionalist needs to provide an explanation of these types of information, explaining whether or not these are different kinds of information, and how transformations between the two take shape. While perhaps not logically impossible, this path seems to lead to conceptually murky waters (Hutto & Myin, 2013, Ch. 4).

A second path for the extended functionalist is to get rid of informational talk altogether and lean on an embodied approach to the memory palace which is entirely non-representationalist. This might seem like a radical move to some philosophers, but it

⁵ One might rightly ask how reconstructive or simulative processes of enactive remembering unfold if they are not based on information storage. While this is an important issue that deserves further elaboration, it is also an open question that needs to be addressed by enactive approaches to memory in general. A proper discussion of this unfortunately falls outside the scope of the current paper. As a tentative proposal, we suggest that enactive remembering involving the previously mentioned processes depend on the sensorimotor activation of familiar patterns. To illustrate, we refer to how artificial neural networks can be trained to generate images (Goodfellow et al., 2014). Such networks do not store specific pixels, but depend on adjusting the signalling strength between nodes during training. After training they may then activate areas on a pre-given (digital) canvas and thus generate an image. Similarly, a person, with an adult, developed brain, may be triggered to think about the Sydney Opera House because of a word read or a sound heard. This trigger may generate, through many intermediary steps, partial images of white, rounded domes against the background of water. It may even be that this person will use her consciousness to help herself generating the memory, for instance, by asking herself “Are the distinctive white shells of the Sydney Opera House spread across two or three separate parts of the building?” Naturally, this is a gross simplification, but it serves as an initial step towards developing a robust enactive account of remembering.

looks like Wheeler is opening the door to that possibility. And a brief look at the history of functionalism provides ground for supporting this move. As Gualtiero Piccinini (2010) argues, functionalism in its purest form is merely the metaphysical claim that cognitive processes are to be understood as structural organizations with input and output relations (see also Putnam, 1967).⁶ It seems that an extended functionalist account of the memory palace based on bodily engagement and not on information processing, is a possibility.

Yet, if functionalists surrender their commitment to the information-processing framework, then what difference is left between extended functionalist and enactivist approaches when it comes to explaining the memory palace? It seems the functionalist's metaphysical account would, to the extent to which they could explain techniques such as the memory palace in terms of bodily engagement, collapse into their competitor theories on enactivism (see Hutto, Peeters, & Segundo-Ortin, 2017). Elucidating the implications of this collapse lies beyond our current argument, but we would be interested to hear what an adapted extended functionalist story would offer that our enactivist story does not.

The extended functionalist, then, has two options. Either develop an information processing account that is not only able to explain how the body plays its role in the memory palace, but also the transformation of abstract into concrete information (whatever that may be). Or, she could surrender her commitment to information-processing altogether and adopt a fully embodied and non-representational account of the memory palace which basically collapses into an enactive account. In any case, the functionalist is currently not in a position to explain the memory palace while the enactivist is not trapped in a similar dilemma. We conclude that thinking about the memory palace from an enactivist perspective is therefore the better option.

We submit that a radically enactive account of memory, which depends on cues and triggers for re-enactment, may act as a clarifying lens through which to look at mnemonic techniques that centrally involve interaction between a person and their environment, such as the memory palace, whether virtual, imagined or otherwise. As we have seen in the previous section, cognitive scientists currently explain the memory palace in terms of information encoding and retrieval, which leads to virtual memory palaces in which the memorizer is a passive participant with only a superficially strong connection to the used locus. Such operationalizations are better served by an enactivist approach which explains why a multimodal memorization technique that heavily involves visualisation, active involvement of a body with an environment, and the reconstruction of memories is more efficient than learning words from a list. The latter mnemonic after all, provides less triggers and cues with which to rebuild memorized items, while the former builds upon such resources and abilities for reconstruction which are already in place. Our next step, then, is to determine which resources and abilities a virtual memory palace needs to work on.

5. Addressing the Operationalization Problem

How can VR technologies support the practice of the memory palace technique? We propose that VR can support practice of the memory palace in at least two ways. First, by supporting the user with a virtual memory palace inspired by recent discussions in cognitive science, thus both relieving the user of the need to go to a familiar, physical building to practise and making sure that the virtual environment evokes those sensorimotor interactions which resemble traditional memory palace usage. Second, by enhancing the memory palace technique by actually going beyond that which is feasible through traditional methods, for example by sharing virtual memory palaces with other users or by supplying the user with visual cues to improve memorisation. These two notions form the inspiration for the following operationalization proposal.

As said earlier, deciding on how best to support the memory palace technique in VR depends on one's answer to the Explanation Problem. In contrast to existing operationalizations of the memory palace we argue for an enactive and re-creative account of remembering. If this argument strikes true, it has implications for the operationalization of the memory palace in VR. Specifically, it means that such operationalizations need to be rethought through the perspective of an embodied cognizer which takes the movement within and active engagement with her (virtual) environment seriously and moves away from the idea that using the memory palace is merely a way of reordering and picking up information. In what follows, we propose that adopting an enactive take on memory will support the practice of the virtual memory palace and that it may help to solve the Operationalization Problem of current designs. We do so by giving concrete design recommendations based on this enactive approach.

To move away from the information processing model of the virtual memory palace, the role of the memorizer needs to be recast from passive observer to active participant. In order to do so, we will single out two aspects of current memory palace operationalizations and translate them into active, body-engaging modes of interaction: movement of the user, and the creation and placement of images. As discussed in the previous section, this translation has to keep in mind the unfolding of sequences through the activation of motor plans in acts of memory. This requires active participation of the body.

Regarding the first aspect, instead of the user being moved through a virtual space passively, we propose that any VR operationalization of the memory palace ought to depart from the idea that the user is actively moving herself through an environment – say a virtual apartment or cathedral. This is not only in line with Yates' (1966) account, which posits the individual moving through the space and engaging with sensori-navigational cues as an essential part of the technique, but also with the two main insights gleaned from current memory research as discussed in Section 2. The first is that, at least in some situations, the activation of motor plans supports remembering – recall the examples on PIN codes, music, and dance from the previous section. We have

⁶ Not all functionalists might agree with the claim that pure functionalism is merely a metaphysical claim. However, my aim here is not to present some kind of essential feature of functionalism, but to trace the genealogy of the extended functionalist line back to its most general shape, like Piccinini (2010) does.

argued that the memory palace is of a similar kind to those examples and thus involves motor activation. Second, neuroscientific evidence supports the idea that brain areas associated with spatial navigation are involved in the use of the memory palace. As such, we think approaches where users are either passively moved or there is no movement at all, do not support the optimal unfolding of a memory sequence.

Moving on to the second point, the placement of images, we present a similar line of reasoning. Active participation of the body in the placement of images in a virtual space would mean that the user should be able to do two things. First, she should be able to either choose or, preferably, create personalised images which may represent parts of that which she wants to memorize. A database in the virtual space, where images can be stored and retrieved, can support the user friendliness and re-use and easy adjustment of images. Second, the user should then be able to place those images in distinct locations in the virtual memory palace. Virtual reality devices with hand-held controllers that can mimic regular hand movements seem especially suited for these use-cases.

Now that we have discussed how the memory palace technique could be translated to VR, by using insights from an enactive approach to cognition to improve movement and image placement, we will present ways of potentially enhancing the virtual memory palace. Is it possible to go beyond the technique's traditional limitations? And if so, how?

One way in which to take advantage of computer technology is to highlight features of the virtual environment in such a way as to support the user's needs. In this, we take inspiration from work done by Raja and Calvo (2017), who propose a way of looking at augmented reality based on ecological psychology (Gibson, 1979). In discussing navigational apps, such as Google Maps, they argue that instead of overloading a user by presenting yet more symbolic information on a screen, for example, by showing a top down map with arrows and numbers, certain pathways might be emphasised more subtly. For instance, one can imagine a user wearing smart glasses which brighten those areas that the user should go, and darken areas the user should avoid. This nudges a user into the destination she wants to go to. Similarly, we suggest, parts of one's virtual memory palace can be highlighted during the learning phase if they offer a memorable location to carry an image associated with part of what one wants to remember (see Fig. 1). Or, also during learning, when unfolding the sequence of the memory the next part of the sequence in the virtual space that a user needs to go to can be brightened, visually, as the next space to move to. So instead of overloading the user with symbolic information, a virtual environment might support memory performance by highlighting the relevant affordances this environment offers to the user (Stoffregen, Bardy, & Mantel, 2006).

A second way of enhancing the virtual memory palace concerns what we dub 'sensorimotor realism.' Note that realism here should not be understood in its common, digitalized meaning: as the photo-realistic replication of images and textures. Contrary to this, perhaps intuitive, idea, there is empirical evidence which suggests that familiar sensorimotor interaction in virtual environments contributes more to the immersion of the memorizer in her memory palace than high-resolution imagery (Fink, Foo, & Warren, 2009). Sensorimotor interaction in VR further seems to improve one's sense of agency, in the sense of experiencing control over one's actions and their consequences (Kong, He, & Wei, 2017), which ties in nicely with and supports the previously discussed active bodily participation.

By *sensorimotor* realism we mean that a VR device involving movement needs to replicate the kind of sensory patterns we experience when we move in real life. To illustrate, think of what occurs when you approach a wall. As you approach the wall, you see how the texture gradients of the wall radiate from the centre of your visual field, causing the wall to expand from the perspective



Fig. 1. Imaginary virtual memory palace with a suggested locus highlighted.

of the perceiver (see Fig. 2). This is commonly described by saying that optic flow is centrifugal in the direction of locomotion (Chemero, 2009, p. 124). The rate at which optic flow expands is lawfully correlated to the speed to which we move towards the object – the wall in this case. By saying that a virtual environment must be sensorimotor realistic we mean that it must echo the sensorimotor experience we are used to in real life. The optic flow generated while moving towards an object in the virtual environment ought to be the same as the one we get when we do so in real life. Otherwise, our experience of moving through the virtual space will feel odd and unpleasant (Bubka, Bonato, & Palmisano, 2008), and it will require us to take extra effort to get attuned to the sensorimotor contingencies of the virtual environment. Ensuring sensorimotor realism will thus add to the immersiveness of the virtual memory palace.

Incorporating active bodily participation lies at the heart of our proposal for operationalizing the virtual memory palace. For the translation of the memory palace to VR, we argued that this requires the user to take control in the virtual environment. For potentially enhancing the virtual memory palace, we proposed to make use of sensorimotor guidance that makes optimal use of the type of interactions the user is already familiar with.

6. Conclusion

Considering the memory palace from an embodied, enactive perspective, in line with the pragmatic turn in cognitive science, helps in understanding why current operationalizations of the technique in VR leave much to be desired. Such operationalizations focus on supporting the picking up of information by the user, but we have argued that this does not capture what is at the core of the technique.

Instead, we presented design recommendations for improving the virtual memory palace, focusing on embodied cognition and affordances. Smart use of VR devices could make the learning of the memory palace more accessible and increase the usage of one of the most powerful methods of remembering on offer. Our design recommendations are ready for implementation. If their adaptation yields better results than current operationalizations, this will have both practical and philosophical implications. To start with the latter: if virtual memory palaces based on our enactive proposal work well *outside* of the head, it would provide a good reason, by way of abduction, to re-evaluate what is going *inside* the head. By way of a reversed parity principle, the enactivist research programme would have provided an impressive case in point in terms of understanding the underpinnings of memory, placing the ball squarely in the functionalist park.

The practical implications, if our proposal holds true, lie in making the power of the memory palace more accessible and their advantages are obvious. Special attention should be given to its potential use in educational settings (Putnam, 2015). We predict that using VR devices to support learning through the memory palace can greatly enhance learning experiences (in line with: Mäkelä & Löytönen, 2017; Heersmink & Knight, 2018). Not only that, but activities which are traditionally seen as boring, like the rote learning of words from a foreign language, would potentially become a lot more fun because of the engaged, bodily interaction. Furthermore, in classroom settings, both teachers and students can benefit from the shared experience which VR will allow. Unlike in the traditional technique, teachers would be able to participate in and give feedback on how their students utilize the memory palace.

Our proposal, though grounded on available empirical data, requires more experimentation. Not only to test whether the hypothesized design recommendations will improve the use of the memory palace, but also to investigate aspects of the techniques that were hereto hard or impossible to investigate. The sharing of the same loci, as described in the previous paragraph is one aspect, but this could be generalized to the investigation of loci which are not necessarily environmental landmarks as traditionally imagined. For example, how will moving objects like animals or other persons affect the technique? What about videos? Virtual realities allow for plenty of creative freedom and the memory palace is a worthy candidate for testing the limits of that freedom with respect to successful memory strategies.

Acknowledgments

Thanks to Patrick McGivern for his detailed comments on a previous draft of this paper and to Andy Clark, Daniel D. Hutto, Michael D. Kirchhoff, Mog Stapleton, Tillmann Vierkant and Michael Wheeler for discussing the paper on various occasions. Thanks also to the graduate students in Brisbane, Edinburgh and Sydney who attended presentations of this paper at different stages of its evolution and gave helpful feedback. Special thanks to Lies Bruines for her keen eye and artistic hand in helping design Figs. 1 and 2. We are grateful to an anonymous reviewer for their helpful feedback on a previous draft of this paper.

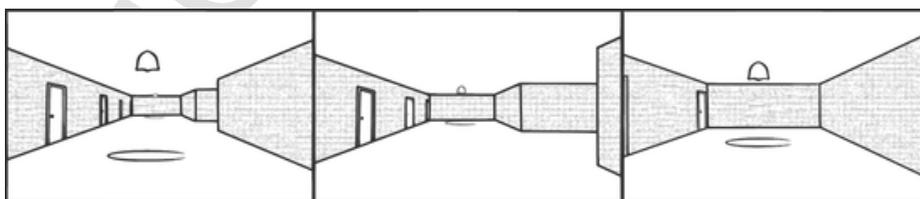


Fig. 2. The walls in these consecutive images expand in a process of optic expansion.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.concog.2019.102834>.

References

- Anderson, M. L. (2014). *After phrenology: Neural reuse and the interactive brain*. Cambridge, MA: MIT Press.
- Anderson, M. L. (2015). Mining the brain for a new taxonomy of the mind. *Philosophy Compass*, 10, 68–77. doi:10.1111/phc3.12155.
- Brooks, B. M. (1999). The specificity of memory enhancement during interaction with a virtual environment. *Memory*, 7, 65–78.
- Bubka, A., Bonato, F., & Palmisano, S. (2008). Expanding and contracting optic-flow patterns and vection. *Perception*, 37, 704–711. doi:10.1068/p5781.
- Casasanto, D. (2009). Embodiment of abstract concepts: Good and bad in right- and left-handers. *Journal of Experimental Psychology: General*, 138, 351–367.
- Casasanto, D. (2014). Body relativity. In Shapiro, L. (Ed.), *The Routledge handbook of embodied cognition* (pp. 108–117). Oxford: Routledge.
- Chaffin, R., Demos, A. P., & Logan, T. (2016). Performing from memory. In Hallam, S., Cross, I., & Thaut, M. (Eds.), *Oxford handbook of music psychology* (pp. 559–571). Oxford: Oxford University Press. doi:10.1093/oxfordhb/9780198722946.013.34.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford: Oxford University Press.
- Clark, A. (2008). Pressing the flesh: A tension in the study of the embodied, embedded mind? *Philosophy and Phenomenological Research*, 76, 37–59. doi:10.1111/j.1933-1592.2007.00114.x.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19. doi:10.1111/1467-8284.00096.
- De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191, 155–185. doi:10.1007/s11229-013-0247-7.
- Dresler, M., Shirer, W. R., Konrad, B. N., Müller, N. C. J., Wagner, I. C., Fernández, G., ... Greicius, M. D. (2017). Mnemonic training reshapes brain networks to support superior memory. *Neuron*, 93, 1227–1235. doi:10.1016/j.neuron.2017.02.003.
- Engel, A. K. (2010). Directive minds: How dynamics shapes cognition. In Stewart, J., Gapenne, O., & Di Paolo, E. A. (Eds.), *Enaction: Toward a new paradigm in cognitive science* (pp. 219–243). Cambridge, MA: MIT Press.
- Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences*, 17, 202–209. doi:10.1016/j.tics.2013.03.006.
- Fassbender, E., & Heiden, W. (2006). The virtual memory palace. *Journal of Computational Information Systems*, 2, 457–464.
- Fernández, J. (2018). The functional character of memory. In Michaelian, K., Debus, D., & Perrin, D. (Eds.), *New Directions in Philosophy of Memory* (pp. 52–72). New York: Routledge.
- Fink, P. W., Foo, P. S., & Warren, W. H. (2009). Catching fly balls in virtual reality: A critical test of the outfielder problem. *Journal of Vision*, 9, 14. doi:10.1167/9.13.14.
- Froese, T. (2014). Bio-machine hybrid technology: A theoretical assessment and some suggestions for improved future design. *Philosophy & Technology*, 27, 539–560. doi:10.1007/s13347-013-0130-y.
- Gerrans, P., & Kennett, J. (2010). Neurosentimentalism and moral agency. *Mind*, 119, 585–614.
- Gibson, J. J. (1977). *The ecological approach to visual perception*. New York: Psychology Press.
- Ginsborg, J., & Sloboda, J. A. (2007). Singers' recall for the words and melody of a new, unaccompanied song. *Psychology of Music*, 35, 421–440. doi:10.1177/0305735607072654.
- Godden, D. R., & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and underwater. *British Journal of Psychology*, 66, 325–331.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems* (pp. 2672–2680). volume 27.
- Heersmink, R., & Knight, S. (2018). Distributed learning: Educating and assessing extended cognitive systems. *Philosophical Psychology*, 31, 969–990. doi:10.1080/09515089.2018.1469122.
- Hutchins, E. (2005). Material anchors for conceptual blends. *Journal of Pragmatics*, 37, 1555–1577. doi:10.1016/j.pragma.2004.06.008.
- Huttner, J.-P., & Robra-Bissantz, S. (2016). A design science approach to high immersive mnemonic e-learning. *MCIS 2016 Proceedings* (p. 28).
- Hutto, D. D., & Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*. Cambridge, MA: MIT Press.
- Hutto, D. D., & Myin, E. (2017). *Evolving enactivism: Basic minds meet content*. Cambridge, MA: MIT Press.
- Hutto, D. D., & Peeters, A. (2018). The roots of remembering: Radically enactive recollection. In Michaelian, K., Debus, D., & Perrin, D. (Eds.), *New directions in philosophy of memory* (pp. 97–118). New York: Routledge.
- Hutto, D. D., Peeters, A., & Segundo-Ortin, M. (2017). Cognitive ontology in flux: The possibility of Protean brains. *Philosophical Explorations*, 20, 209–223. doi:10.1080/13869795.2017.1312502.
- Jund, T., Capobianco, A., & Larue, F. (2016). Impact of frame of reference on memorization in virtual environments. 2016 IEEE 16th International conference on advanced learning technologies (pp. 533–537). doi:10.1109/ICALT.2016.77.
- Kirsh, D. (2013). Embodied cognition and the magical future of interaction design. *ACM Transactions on Computer-Human Interaction*, 20, 1–30. doi:10.1145/2442106.2442109.
- Kong, G., He, K., & Wei, K. (2017). Sensorimotor experience in virtual reality enhances sense of agency associated with an avatar. *Consciousness and Cognition*, 52, 115–124. doi:10.1016/j.concog.2017.04.018.
- Krokos, E., Plaisant, C., & Varshney, A. (2019). Virtual memory palaces: Immersion aids recall. *Virtual Reality*, 23, 1–15. doi:10.1007/s10055-018-0346-3.
- Legge, E. L. G., Madan, C. R., Ng, E. T., & Caplan, J. B. (2012). Building a memory palace in minutes: Equivalent memory performance using virtual versus conventional environments with the method of loci. *Acta Psychologica*, 141, 380–390. doi:10.1016/j.actpsy.2012.09.002.
- Leman, M., & Maes, P.-J. (2014). Music perception and embodied music cognition. In Shapiro, L. (Ed.), *The Routledge handbook of embodied cognition* (pp. 81–89). Oxford: Routledge.
- Loader, P. (2013). Is my memory an extended notebook? *Review of Philosophy and Psychology*, 4, 167–184. doi:10.1007/s13164-012-0123-2.
- Mackay, W. E., Payard, A.-L., Probert, L., & Médini, L. (1998). Reinventing the familiar: Exploring an augmented reality design space for air traffic control. *Conference Proceedings on Human Factors in Computing Systems 1998* (pp. 558–565). New York: ACM Press/Addison-Wesley. doi:10.1145/274644.274719.
- Madan, C. R. (2014). Augmented memory: A survey of the approaches to remembering more. *Frontiers in Systems Neuroscience*, 8, 30. doi:10.3389/fnsys.2014.00030.
- Madan, C. R., & Singhal, A. (2012). Motor imagery and higher-level cognition: Four hurdles before research can sprint forward. *Cognitive Processing*, 13, 211–229. doi:10.1007/s10339-012-0438-z.
- Maguire, E. A., Valentine, E. R., Wilding, J. M., & Kapur, N. (2002). Routes to remembering: The brains behind superior memory. *Nature Neuroscience*, 6, 90–95. doi:10.1038/nn988.
- Mäkelä, M., & Löytönen, T. (2017). Rethinking materialities in higher education. *Art, Design & Communication in Higher Education*, 16, 241–258. doi:10.1386/adch.16.2.241_1.
- Michaelian, K. (2016). *Mental time travel: Episodic memory and our knowledge of the personal past*. Cambridge, MA: MIT Press.
- Michaelian, K., Debus, D., & Perrin, D. (Eds.). (2018). *New Directions in Philosophy of Memory*. New York: Routledge.
- Piccinini, G. (2010). The mind as neural software? Understanding functionalism, computationalism, and computational functionalism. *Philosophy and Phenomenological Research*, 81, 269–311. doi:10.1111/j.1933-1592.2010.00356.x.
- Putnam, A. L. (2015). Mnemonics in education: Current research and applications. *Translational Issues in Psychological Science*, 1, 130–139. doi:10.1037/tps0000223.

- Putnam, H. (1967). Psychological predicates. In Capitan, W. H., & Merrill, D. D. (Eds.), *Art, mind, and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.
- Ragan, E. D., Sowndararajan, A., Kopper, R., & Bowman, D. A. (2010). The effects of higher levels of immersion on procedure memorization performance and implications for educational virtual environments. *Presence: Teleoperators and Virtual Environments*, 19, 527–543. doi:10.1162/pres_a.00016.
- Raja, V., & Calvo, P. (2017). Augmented reality: An ecological blend. *Cognitive Systems Research*, 42, 58–72. doi:10.1016/j.cogsys.2016.11.009.
- Rosello, O., Exposito, M., & Maes, P. (2016). Nevermind: Using augmented reality for memorization. *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (pp. 215–216). ACM Press. doi:10.1145/2984751.2984776.
- Rowland, I. D. (2008). *Giordano Bruno: Philosopher/heretic*. Chicago: University of Chicago Press.
- Shapiro, L. A. (2004). *The mind incarnate*. Cambridge, MA: MIT Press.
- Shapiro, L. A. (2019). Flesh matters: The body in cognition. *Mind & Language*, 34, 3–20. doi:10.1111/mila.12203.
- Smith, S. M., & Vela, E. (2001). Environmental context-dependent memory: A review and meta-analysis. *Psychonomic Bulletin & Review*, 8, 203–220.
- Stevens, C., Malloch, S., McKechnie, S., & Steven, N. (2003). Choreographic cognition: The time-course and phenomenology of creating a dance. *Pragmatics & Cognition*, 11, 297–326. doi:10.1075/pc.11.2.06ste.
- Stoffregen, T. A., Bardy, B. G., & Mantel, B. (2006). Affordances in the design of enactive systems. *Virtual Reality*, 10, 4–10. doi:10.1007/s10055-006-0025-7.
- Sutton, J. (2007). Spongy brains and material memories. In Floyd-Wilson, M., & Sullivan, G. (Eds.), *Embodiment and environment in early modern Europe* (pp. 14–34). London: Palgrave.
- Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In Menary, R. (Ed.), *The extended mind* (pp. 189–225). Cambridge, MA: MIT Press.
- Sutton, J., & Williamson, K. (2014). Embodied remembering. In Shapiro, L. A. (Ed.), *The Routledge handbook of embodied cognition* (pp. 315–325). London: Routledge.
- Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Wheeler, M. (2010). Minds, things and materiality. In Malafouris, L., & Renfrew, C. (Eds.), *The cognitive life of things: Recasting the boundaries of the mind* (pp. 29–37). Cambridge: McDonald Institute for Archeological Research.
- Wheeler, M. (2015). Extended consciousness: An interim report. *The Southern Journal of Philosophy*, 53, 155–175. doi:10.1111/sjp.12124.
- Yates, F. A. (1966). *The art of memory* (ARK 1984 ed.). London: Routledge.
- Michaelian, K., & Sant'Anna, A. (2019). Memory without content? Radical enactivism and (post)causal theories of memory. *Synthese*. Online first publication. doi:10.1007/s11229-019-02119-7.