Human Memory and the Limits of Technology in Education

Puddifoot, Katherine; O'Donnell, Cian

Human Memory and the Limits of Technology in Education

Abstract Human memory systems perform various functions beyond simple storage and retrieval of information. They link together information about events, build abstractions, and perform memory updating. In contrast, typical information storage and access technologies, such as note-taking applications and Wikipedia, tend to store information verbatim. In this article, we use results from cognitive psychology, neuroscience and machine learning to argue that the increased dependence on such technologies in education may come at a price: the missed opportunity for memory systems of student learners to form abstractions and insights from newly learned information. This conclusion has important implications for how technologies should be adopted in education.

1. Introduction

Numerous technologies are now used within educational settings, with the aim of improving the learning experience and student outcomes. Examples include: distance/online/virtual learning; the use of analytics to gather and utilise data about student learning habits; interactive learning applications/tools; audio-visual teaching aids; and information storage and access technologies.[i] The increasing use of technologies within educational settings raises important questions about the extent to which traditional methods of teaching and learning should be supplanted by new methods involving the use of technology. In this paper, we highlight a potential danger to supplanting some teaching methods with alternatives that involve using technologies.

We focus on the use of what we call *information storage and access technologies*. These technologies store information that can easily and rapidly be accessed by anyone with an understanding of how to use them. They can be contrasted to technologies that directly act as a means of support to learning activities rather than providing access to information (Tondeur, van Braak and Valcke 2007). Included in the relevant category are: (i) personal devices such as flash drives, cloud storage, and note-taking applications, in which students can store information that they have been taught; (ii) open access resources such as Wikipedia or Google that contain information that other people have made available; (iii) restricted access resources such as digital textbooks or the online learning environments for specific courses of study, which include course-related documents and resources; and (iv) social media resources, in which information shared by other people can be accessed and used by a student in their studies. For some purposes, people may distinguish between devices that store self-generated information (e.g. note-taking applications) and devices that store other-generated information (e.g. Wikipedia), but for our purposes we treat both equally.

We recognise that these technologies perform numerous important roles within contemporary education settings, ameliorating the student experience and student outcomes in a variety of important ways. However, we highlight a danger associated with the adoption of these technologies within an educational setting. Because of the benefits of the technologies, some people have argued that educational methods should be overhauled, so that significantly less emphasis is placed on students engaging in learning that involves storing information to memory systems in the human brain. We argue that there are important functions performed by human memory systems— the linking together of information found in different sources, the production of abstract representations, and the updating of learnt information over time— which are unlikely to be performed if educational methods are overhauled in thus way. We argue that these functions are essential to the achievement of one of the central goals of education, i.e.: the transference of learning. Consequently, a move away from storing information internally in our brains has the potential to have a detrimental effect on educational outcomes because it can prevent students from achieving transference. Our argument draws on findings from cognitive psychology, neuroscience and machine learning.

The structure of our argument is as follows. In section 2, we show why information storage and access technologies are attractive to those working within an educational setting. In section 3, we highlight the functions performed by internal memory that will be the focus of discussion. In section 4, we show how these functions are important and valuable within an educational setting, improving the student experience and learning outcomes by facilitating transference of learning—a central goal of education. Then, in section 5, we show that these functions might not be performed, and transference not achieved, if there is an increased focus within education on using information storage and access technologies.

2. The appeal of technology in education

Information storage and access technologies have several features that make them attractive for use in education. Contemporary digital technologies have large storage capacities. Information stored to these devices is easily compressed and therefore a large amount can be stored on small physical devices (Selwyn 2016). The technologies are highly reliable at storing accurate, verbatim representations of information that are then available for retrieval. The information stored in these technologies can be easily edited (ibid.), searched through, copied, and shared. In contrast, human memory systems have only limited storage capacities (see, e.g. Cherniak 1983, Brady et al., 2011). There is a huge psychological literature suggesting that people are not only susceptible to forgetting, we are also susceptible to *misremembering* (Robins 2016), recalling details of an event inaccurately (e.g. Loftus and Palmer 1974; Roediger and McDermott 1995; Schacter and Addis 2007). Human memory systems are therefore fallible with respect to the goal of storing accurate verbatim representations of information. Finally, human memory systems are largely private. If a person wishes to access information stored in another person's internal memory systems, their success depends on the ability to identify and communicate their need, and on the other person's willingness, as well as their ability, to access and to provide this information. Meanwhile, technologies such as Wikipedia, Google and social media provide easy, public access to information.

On a view according to which memory works as a storehouse (Sutton 1998), only functioning to store and provide access to information, the types of technologies that we are discussing would seem to overwhelmingly, if not only, bring benefits, inside and outside of education. They provide better storage capacities, more accurate records of information, and more ready access to a wider set of information than internal memory systems. However, the memory-as-storehouse picture has been widely rejected within philosophy of memory (see e.g. De Brigarde 2014; Sutton 2009, 2010; Michaelian 2011; Robins 2016) and cognitive psychology (see, e.g. Schachter and Addis 2007; Schachter 2011) and is increasingly being criticised in neuroscience (see, e.g. Stickgold and Walker, 2013; Eichenbaum and Cohen, 2014; Richards and Franklin 2017). It is now widely accepted that memory systems perform numerous important functions other than storage and retrieval. Our claim is that these functions are both important to education and unlikely to be performed if students reduce the extent to which they internalise information to memory because of the adoption of information storage and access technologies.

We therefore highlight how discussions within the cognitive sciences put pressure on positions like that of connectivism within educational theory, according to which it is not the learning that occurs within a person that is important but instead the networks that they form, with computer networks, social networks, etc. as it is through these that people can acquire accurate, up-to-date information (Siemens 2005, Thota 2015). We argue that the learning that occurs within the person, in their internal memory systems, can be vital to supporting important functions of learning. Our view also highlights shortcomings of some educational practices that involve the use of information storage and access technologies to perform functions traditionally performed by internal human memory systems, e.g. allowing students to use information storage and access technologies inside the examination room (see, e.g. Wheeler 2011). We show that there are important functions of human memory systems that are less likely to be performed if such practices are adopted**.**

To be clear, our aim is not to advocate the use of technology-free examinations or any other specific traditional methods of teaching and learning. It is consistent with the claims made in the current paper that, for example, the *constructivist* view of education is correct. According to constructivism, students should be active learners, using existing knowledge to engage in activities that lead to the acquisition of further knowledge (Bruner 1996). Adaptive learning technologies may help accelerate this process by, for example, tailoring feedback or lines of instruction to suit each particular student (Desmarais and Baker, 2012). It is also consistent with the view defended in this paper that both informal and formal methods of learning are valuable. Informal methods of learning tend to take place outside of a structured learning environment and do not tend to involve rigorous testing (Marswick and Watkins 1990). One can learn informally outside the classroom in everyday life (Mills and Kraftl 2014). It is consistent with our view that active and informal learning are highly valuable. What our argument emphasises that if learning, either utilising these methods or not, does not involve the internalisation of information, student learning can be negatively affected.

3. Human Memory and Its Functions

The aim of this section is to spell out in more detail the functions other than those suggested by the memory-as-storehouse view that are performed by human memory systems. It outlines results from the fields of neuroscience, cognitive psychology and machine learning that show how the brain mechanisms underlying memory are responsible for: (i) linking together information about different events, (ii) building abstract representations, (iii) updating memories in light of most recent information.

Let us begin with considering how biological memory systems link together information about different events. In a learning setting, this will usually involve linking the newly learned information with existing, older memories for other events and concepts. In neuroscience, this linking is generally thought to occur slowly over days and nights as part of a larger process termed *systems consolidation*. Consolidation here is defined as a process that crystallises new memories so that they become less malleable and more locked-in over time (McGaugh, 2000). Neuroscientists believe that systems consolidation involves a two-stage process: first, new memories are encoded in the hippocampus directly following the experience. Second, over the following days and weeks the newly learned information is tranferred from the hippocampus to the neocortex where it is linked to existing memories and stored long-term (Alvarez and Squire, 1994; Buszaki 1989; Marr 1971). Mechanistically, the basic transfer of information from hippocampus to neocortex is believed to happen via repeated *replay* of episodic memories by the hippocampus during sleep (Buszaki 1989;

Wilson and McNaughton 1994, O'Neill et al., 2010). The idea is that memory replay by hippocampus during sleep drives neocortical brain networks, activating both some representation of the new memory, plus related older memories. This co-activation triggers strengthening of interconnections between the sets of active neocortical neurons, and so links the new memory with existing knowledge.

Of the three memory processes we describe, this standard account of systems consolidation model accounts only for the first (linking together of information). Importantly however, this consolidation process also seems to parallel the abstraction and generalisation of memories into simpler representations (Winocur et al. 2007; Stickgold & Walker, 2013). This is a process that is taken by cognitive psychologists to explain a large range of memories, as it is thought that representations of the gist of events remain as verbatim details fade (Brainerd & Reyna 2002). Although it is not known how memory generalisation works at the neural level, two recent theoretical models have been put forward. Lewis and Durant (2011) suggest that if multiple memories were to be replayed concurrently by the brain, then "the overlapping replay of related memories selectively strengthens shared elements". As the non-overlapping elements of these memories will not be reinforced, they are more likely to be forgotten. This model fits with the common-sense view that an abstraction should be built out of the common elements of different items, while the ignoring their differentiating details. O'Donnell and Sejnowski (2014) propose a different model for memory generalisation, where memory replay during non-REM sleep results in a biochemical template being laid down in the neocortical neurons that were activated by the replayed memory. This template is then used during the subsequent REM phase of the sleep cycle (when most vivid dreams occur) to selectively strengthen connections from neurons outside the template to those neurons inside the template. This should have the effect of broadening the original memory representation to incorporate a wider set of neurons. In contrast to the Lewis and Durant model, the O'Donnell and Sejnowski model proposes that generalisations are not formed by finding commonalities between pairs of memory items, but by taking a single memory and meshing it with existing prior knowledge about the world, so generalising its contents.

A related insight into how human memory generalisation might work comes from the field of machine learning. Machine learning researchers seek to build computer programs that learn how to perform a task by encountering example ‘training’ data points and updating their algorithms accordingly, mimicking how humans learn cumulatively. In this field, it is well appreciated that the ability of a computer program to generalise from specific training examples to perform well on a broader set of tasks can be impaired by *overfitting*. Overfitting happens when a statistical model learns to capture too well every detail of the specific examples that it happened to see during training. If these details are irrelevant for the broader problem, then they may impair generalisation on future tasks. Machine learning researchers have discovered several methods for reducing the effects of overfitting. One powerful solution is to build in prior knowledge that the computer programmer may have of the structure of the problem. Ideally this prior structural information will bias the computer program towards solutions that lead to better performance on data or related tasks. For example, a program that is pre-programmed to know that everyday objects, like bicycles, tend to consist of multiple parts can learn to understand new object categories from as few as one or two examples, to the same level of performance as humans (Lake et al., 2015). In contrast, otherwise identical programs that do not understand that objects can be decomposed into parts tend to generalise poorly. Hence, prior knowledge is essential for robust generalisation.

The third type of memory processing performed by the brain, memory updating, is thought to be mediated by the mechanism of *reconsolidation*. In something of a surprise to the neuroscience field, it was shown that previously consolidated memories could be made re-labile simply by appropriately cueing their recall (Nader et al, 2000). This finding implies that each recall of a memory opens a temporal window of opportunity for the brain to alter it, and potentially incorporate new information into it. Memory updating can explain findings from cognitive psychology showing that memories of specific events can be updated to reflect information, including false information, encountered after the event, in what has become known as the *misinformation effect* (e.g. Loftus and Palmer 1974). Findings from neuroscience suggest that the incorporation of information provided after the event is possible during the window of opportunity that occurs at each recall of the memory. For our purposes, the key lesson is that that new learned information is not simply linked an original memory, but in fact the original memory itself is altered. This may result in aspects of the old memory potentially being lost in the process. In this sense, the brain performs a true updating, not simply an accumulation of information.

4. Limited functioning of information storage and access technologies

In this section we show that the functions of memory outlined in section 3 facilitate the achievement of one of the most important goals of education: the transference of learning.

The transference of learning involves information being used outside of the context in which it was initially learnt (Thorndike and Woodward 1901). It is widely accepted among educators that transference is a crucial component of education (Bransford and Schwartz 1999; Perkins and Salomon 1992).[ii] Educators aim for the information that they convey to their students to be utilised under a variety of different conditions, inside and outside of the classroom (Bransford and Schwartz 1999) rather than the benefits of their learning to be confined to the context of learning. For example, a student might study population growth in ecology, learn that exponential growth of groups occurs when there are no barriers to slow growth, then transfer this learning to the consideration of infectious disease in epidemiology, concluding that diseases will spread exponentially if barriers are not in place (Kaminski et al 2013). In this case, what is learnt about one case could be applied in another educational context (in another class) or outside of the classroom. Another example would be a student of history who learns about the dangers of populism by learning about one case in history and applies their learning to other historical cases, and then compares each of these cases to what is currently occurring within her home country. In the final case, the student could develop a general picture of the dangers of populism on the basis of considering two or more historical cases, and the general picture could then be used to understand current events.

What is most important for our purposes is that all cases of transfer of learning are highly dependent on the functions of human memory outlined in section 3. This point can be understood by considering the various stages of information processing involved in the transference of learning.

For transference to occur, information about different cases must *be linked together*, e.g. information about two cases of populism. If information about case of populism A is never connected to case of populism B then learning about case A will not be transferred to case B. In order to ensure that the information is linked where appropriate, an *abstract representation* must be formed that reflects the commonalities between the cases, e.g. the feature of the cases that are due to populism rather than something else. This requires abstracting away from the details that differ between the cases and detecting a common core (Kaminski et al 2009). Where people fail to engage in transfer of learning this can be because they attend too heavily to details of a specific case, failing to see the commonalities between cases (ibid.). Once an abstract representation has been formed, there is the potential to identify numerous different learning contexts in which previously learned knowledge could be applied. By applying learning across new contexts it will be possible to refine the abstract representation to reflect the new information that becomes linked together through the process of transfer. It will also be possible to change one's view of the information learnt in the initial stages of learning. If one's initial learning about populism is challenged by comparing it to other cases of populism, through the process of transfer of learning, the earlier *learning can be updated*.[iii]

We have seen that each of these processes—linking together of information, formation of abstract representation, and updating of learned information—are facilitated by the nature of human memory systems. It therefore follows that the functions of human memory systems, other than the storage and retrieval of information—facilitate the core goal of education that is transference of learning.

It is worth noting that not all types of learning can be transferred. For some information, there is no way to find commonalities between examples to build an abstraction. For example, there is no way of predicting someone's phone number from their name. This means that it would be impossible to study a list of name/phone number pairings to discover some underlying structure with which to build an abstract model, that could later help you to predict a new person's phone number based on name alone. In situations like these it makes sense to offload the information to external devices, both because they are good at storing a large amount of information, and because the student will not miss out on any abstractions by doing so. This observation is consistent with the claims made in the current discussion, however, because our aim is to show that there is a *significant subset of* information that can usefully be transferred by forming abstractions, and that this transference is likely to be missed with increased dependence on information storage and access technologies. Our claim is *not* that *all* information can usefully be transferred in this way.

5. Educational Technologies and a Failure of Functioning

So far we have argued that human memory systems function in ways other than simply storing and retrieving information, and that these functions are important to the achievement of transference of learning—a central goal of education. This section shows that the same functions are unlikely to be performed as students increasingly depend on information storage and access technologies.

The argument outlined so far provides good *prima facie* reasons for accepting this conclusion. It has identified advantages for learning which are the result of the systems operating in ways that differ from how information storage and access technologies operate, i.e. solely providing a facility for storage and retrieval of information. This suggests that if people who increase their usage of information storage and access technologies also reduce the extent to which they internalise information to memory, they will miss out on advantages for learning.

It is, of course, important for us to show that these phenomena—i.e. (i) the increase in use of information storage and access technologies and (ii) the reduction of internalisation of information to memory—reliably co-occur. Why should we think this? First of all, there may be cases where the co-occurrence would be intentional. For example, teachers who adopt the connnectivist viewpoint might decide that it is no longer important to teach students information for storing in their internal memories, due to the existence of new technologies. This might be reflected in their teaching practice. Alternatively, widespread use of information storage and access technologies might lead to an unintentional reduction in learners' internally stored information. For example, it has been found that people tend to forget information that they think will be stored externally (Sparrow et al 2011). If the process of using these external technologies is not sufficient to trigger the brain functions that facilitate transfer of learning, transference may become less likely.

To see why it is that the use of technologies does not involve the performance of these functions it is first important to note that there is more than one way that information storage and access technologies could be adopted, depending on (a) the specific technologies and (b) the aims of those adopting them. In some cases, students might go through an initial learning experience in which they cognitively process information before storing it to one of the technologies. Students using note-taking applications could fit this description. In other cases, students might never go through an initial learning experience in which they cognitively process information. They might instead be merely instructed upon how to access information from the technologies, for example by searching Wikipedia. In both types of cases students are susceptible to missing the benefits of transference of learning.

Let us begin by focusing on the second case, those students who do not cognitively process the information. As they have not processed the information, they cannot have built abstract representations of the information. Without building these abstractions, the students will not be able to identify new cases in which previous learning is relevant.

It might be thought, however, that students could simply search for information when they encounter a task or problem for which information stored in these technologies could usefully be applied. For example, a person interested in understanding how diseases spread could search information storage and access technologies, such as Google or Wikipedia, to find an abundance of information relating to the topic. There are a number of problems with this response.

First of all, searches of this type will likely be too narrow to facilitate the type of transfer possible through the use of internal memory systems. This is because those engaging in the searches will not be aware of what information could usefully transfer and therefore would not search effectively for information. For instance, a search for information about how diseases spread that uses disease-specific keywords is unlikely to reveal information about the growth of populations in ecology. Although a description of the commonalities between these two processes may exist somewhere in the external storage system, successful discovery of this information would require that the person engaging in the search used the right search terms, i.e. terms that reflected the connection between the two types of information. However, students are unlikely to search in this way if they are unaware of the connection between the types of information because they have not formed an abstract representation reflecting the common core of the two types of information. And they will not have formed an abstract representation of this type if they have never internalised the information.

One way to understand this point is by considering Donald Rumsfeld's infamous distinction between *known unknowns* and *unknown unknowns*. One natural interpretation of this distinction is that known unknowns are things that we know we do not know. On the other hand, unknown unknowns are the things that we don't know that we don't know.[iv] A learner who sought and found information about the growth of populations in order to gain insights about the spread of disease would have to know that there was a connection between these cases that they did not know enough about. However, if they have not formed an abstract representation of the common core of the two cases then they will not be aware of the connection. The information about the connection between the cases will be an unknown unknown. Consequently, the student is unlikely to find the information through a search of information storage and access technologies.

Second, even if students were reliably able to access the information they needed by searching information storage and access technologies to identify information that would be relevant to their learning, the process may be impractically time-consuming, or might interfere with other cognitive tasks. Insights that we have gained from previous learning experiences are constantly informing new experiences that we have, inside and outside of the classroom. This can be achieved automatically or offline by the brain, therefore quickly and efficiently, without producing significant interference with other cognitive tasks. A switch to increased use of searching technologies may lead to these features being lost.

What about students who do engage in some cognitive processing before offloading information to information storage and access technologies? Will they too miss out on the benefits of transference of learning? Discussions from the neuroscience of memory suggest that they will.

The main difference between students who engage in some cognitive processing of information before offloading it and those that do not is that the former will temporarily engage with the information. Lessons from neuroscience suggest that this temporary access to information will not suffice to produce the abstract representations that are necessary for transference of learning. The research suggests that the process of memory abstraction is tied to the slow neocortical learning system so it takes time: minimally one night's sleep, but perhaps several months (Rasch and Born, 2013). This implies that if students only briefly process the information and then it is actively forgotten, as it is likely to be when a student is aware that the information will be stored and accessible in an external device (Sparrow et al 2011), the information is likely not to be stored for long enough to allow full abstraction to occur.

In sum, then, students who offload information to information storage and access devices, and depend on their ability to later access information from the devices rather than retrieve it from memory, are highly susceptible to missing out on the advantages of transference of learning that internal memory systems supply.

Conclusion

Human memory systems do not function like a storehouse. They perform functions other than storing and providing access to information, including linking together information about different events, forming abstract representations, and facilitating the updating of information stored to memory. Each of these functions supports the transference of learning, which is a core goal of education. These functions may be impaired if future students become increasingly dependent on technologies that function more like a storehouse, i.e. information storage and access technologies. If so, the increased use of information storage and access technologies could undermine one of the main goals of education. We suggest that future educationalists should design teaching plans that deliver the types of information that would most help student build links and abstractions across material, while simultaneously encouraging students to offload non-structured or detailed information to external storage technologies, as appropriate. Such a balanced approach could maximise the benefits of both minds and machines.

References

Alvarez, P, and L R Squire. 1994. "Memory Consolidation and the Medial Temporal Lobe: a Simple Network Model.." *Proceedings of the National Academy of Sciences of the United States of America* 91 (15): 7041–45.

Bortolotti, Lisa. "Epistemic benefits of elaborated and systematized delusions in schizophrenia." *The British journal for the philosophy of science* 67, no. 3 (2015): 879-900.

Brady, T F, T Konkle, and G A Alvarez. 2011. "A Review of Visual Memory Capacity: Beyond Individual Items and Toward Structured Representations" *Journal of Vision* 11 (4): 1–34.

Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science, 11*, 5: 164-169.

Bransford, John D., and Daniel L. Schwartz. "Chapter 3: Rethinking transfer: A simple proposal with multiple implications." *Review of research in education* 24, no. 1 (1999): 61-100.

Bruner, Jerome. *The Culture of Education*, Cambridge MA, Harvard University Press, 1996.

Buzsáki, G. 1989. "Two-Stage Model of Memory Trace Formation: a Role for 'Noisy' Brain States.." *Neuroscience* 31 (3): 551–70.

Cherniak, Christopher. "Rationality and the structure of human memory." *Synthese* 57, no. 2 (1983): 163-186.

Collins, Rupert A., and Robert H. Cruickshank. "Known knowns, known unknowns, unknown unknowns and unknown knowns in DNA barcoding: a comment on Dowton et al." *Systematic biology* 63, no. 6 (2014): 1005-1009.

De Brigard, Felipe. "Is memory for remembering? Recollection as a form of episodic hypothetical thinking." *Synthese* 191, no. 2 (2014): 155-185.

Desmarais, Michel C, and Ryan S Baker. 2012. "A Review of Recent Advances in Learner and Skill Modeling in Intelligent Learning Environments." *User Modeling and User-Adapted Interaction* 22 (1-2). Kluwer Academic Publishers: 9–38.

Eichenbaum, H, and Cohen, NJ. 2014. “Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function?.” *Neuron* 83 (4): 764–70.

Kaminski, Jennifer A., Vladimir M. Sloutsky, and Andrew F. Heckler. "The cost of concreteness: the effect of nonessential information on analogical transfer." *Journal of Experimental Psychology: Applied* 19, no. 1 (2013): 14.

Kaminski, Jennifer A., Vladimir M. Sloutsky, and Andrew Heckler. "Transfer of mathematical knowledge: The portability of generic instantiations." *Child Development Perspectives* 3, no. 3 (2009): 151-155.

Lake, Brenden M, Ruslan Salakhutdinov, and Joshua B Tenenbaum. 2015. “Human-Level Concept Learning Through Probabilistic Program Induction..” *Science* 350 (6266): 1332–38.

Lewis, Penelope A, and Simon J Durrant. 2011. “Overlapping Memory Replay During Sleep Builds Cognitive Schemata..” *Trends in Cognitive Sciences* 15 (8): 343–51.

Loftus, E.F. and Palmer, J.C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behaviour*, 13, 5: 585-589.

Marr, D. 1971. “Simple Memory: a Theory for Archicortex.” *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences* 262 (841): 23–81.

Marsick, Victoria J., and Karen Watkins. *Informal and incidental learning in the workplace*, London, Routledge, 1990

Michaelian, Kourken. "Generative memory." *Philosophical psychology* 24, no. 3 (2011): 323-342.

Mills, Sarah, and Peter Kraftl, eds. *Informal education, childhood and youth: Geographies, histories, practices*. Springer, 2014.

Nader, K, G E Schafe, and J E Le Doux. 2000. “Fear Memories Require Protein Synthesis in the Amygdala for Reconsolidation After Retrieval..” *Nature* 406 (6797): 722–26.

Norris, Christopher. *Epistemology: Key concepts in philosophy*. A&C Black, 2005.

O'Donnell, Cian, and Terrence J Sejnowski. 2014. “Selective Memory Generalization by Spatial Patterning of Protein Synthesis..” *Neuron* 82 (2): 398–412.

O'Neill, Joseph, Barty Pleydell-Bouverie, David Dupret, and Jozsef Csicsvari. 2010. “Play It Again: Reactivation of Waking Experience and Memory..” *Trends in Neurosciences* 33 (5): 220–29.

Perkins, David N., and Gavriel Salomon. "Transfer of learning." *International encyclopedia of education* 2 (1992): 6452-6457.

Puddifoot, Katherine. "Dissolving the epistemic/ethical dilemma over implicit bias." *Philosophical Explorations* 20, no. sup1 (2017): 73-93.

Rasch, Björn, and Jan Born. 2013. “About Sleep's Role in Memory..” *Physiol Rev* 93 (2): 681–766.

Richards, Blake A, and Paul W Frankland. 2017. “The Persistence and Transience of Memory..” *Neuron* 94 (6): 1071–84.

Robins, Sarah K. "Misremembering." *Philosophical Psychology* 29, no. 3 (2016): 432-447.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of experimental psychology: Learning, Memory, and Cognition*, *21*, 4, 803-814.

Schacter, Daniel L., Donna Rose Addis, and Randy L. Buckner. "Remembering the past to imagine the future: the prospective brain." *Nature reviews. Neuroscience* 8, no. 9 (2007): 657.

Schacter, Daniel L., Scott A. Guerin, and Peggy L. St Jacques. "Memory distortion: an adaptive perspective." *Trends in cognitive sciences* 15, no. 10 (2011): 467-474.

Scoville, W B, and B Milner. 1957. “Loss of Recent Memory After Bilateral Hippocampal Lesions..” *Journal of Neurology, Neurosurgery & Psychiatry* 20 (1): 11–21.

Selwyn, Neil. *Education and technology: Key issues and debates*. Bloomsbury Publishing, 2016.

Siemens, George. “Connectivism: A Learning Theory for the Digital Age”. *International Journal of Instructional Technology and Distance Learning (ITDL)* (2005). URL: http://er.dut.ac.za/handle/123456789/69 (accessed July 31st 2017).

Sparrow, Betsy, Jenny Liu, and Daniel M. Wegner. "Google effects on memory: Cognitive consequences of having information at our fingertips." *science* 333, no. 6043 (2011): 776-778.

Stickgold, E, and Walker MP. 2013. “Sleep-Dependent Memory Triage: Evolving Generalization Through Selective Processing..” *Nature Neuroscience* 16 (2): 139–45.

Sutton, John. “Adaptive misbeliefs and false memories” *Behavioral and Brain Sciences, 32*(6) (2009): 535–536.

Sutton, John. “Observer perspective and acentered memory: Some puzzles about point of view in personal memory.” *Philosophical Studies*, *148* (2010): 27–37.

Sutton, John. *Philosophy and memory traces: Descartes to connectionism*. Cambridge University Press, 1998.

Thota, Neena. "Connectivism and the use of technology/media in collaborative teaching and learning." *New Directions for Teaching and Learning* 2015, no. 142 (2015): 81-96.

Tondeur, Jo, Johan Van Braak, and Martin Valcke. "Towards a typology of computer use in primary education." *Journal of Computer Assisted Learning* 23, no. 3 (2007): 197-206.

Wilson, M A, and B L McNaughton. 1994. “Reactivation of Hippocampal Ensemble Memories During Sleep..” *Science* 265 (5172): 676–79.

Winocur, Gordon, Morris Moscovitch, and Melanie Sekeres. 2007. “Memory Consolidation or Transformation: Context Manipulation and Hippocampal Representations of Memory..” *Nature Neuroscience* 10 (5): 555–57.

Woodworth, Robert S., and E. L. Thorndike. "The influence of improvement in one mental function upon the efficiency of other functions.(I)." *Psychological review* 8, no. 3 (1901): 247.

---

[i] It might be useful to note that one technology can fall under various categories, e.g. a tablet.

[ii] There has been some skepticism about whether transference of learning actually occurs. See, e.g., Thorndike and Woodward’s (1901) seminal paper. However, we are construing transference very broadly, to reflect the full range of activities accepted as such in the literature, and find it highly implausible that none of the examples discussed are genuine cases in which people could transfer learning.

[iii] It might be objected that updating memories about a past learning experience is costly, leading to an inaccurate representation of the initial learning experience and a lack of self-knowledge about what has been learnt. These are real costs, undermining the student's ability to represent the past accurately. However, with respect to the goal of storing learnt information so that it can later be accessed and used, for example, in further transference of learning, the updating of the memory can be highly beneficial. It enables the information learnt through the process of transference to be stored without using extra storage space, which is limited (see, e.g. Cherniak 1983). Under such circumstances, the costs of memory updating might be viewed as the lesser of two epistemic evils (Bortolotti 2015; Puddifoot 2017), where the alternative is failing to achieve the goal of engaging in successful learning.

[iv] For an alternative philosophical interpretation of Rumsfeld's meaning see Norris 2005. Rumsfeld was heavily criticised for making this distinction but he is far from alone, for example, the distinction is used in biology (see, e.g. Collins and Cruikshank 2014).