

British Psychological Society Cognitive Section Prize Lecture (2021)

Precis of Berens, Richards, and Horner (2020): Dissociating memory accessibility and precision in forgetting

Aidan J. Horner<sup>1,2</sup> & Sam C. Berens<sup>3</sup>

<sup>1</sup>Department of Psychology, University of York, UK.

<sup>2</sup>York Biomedical Research Institute, University of York, UK.

<sup>3</sup>School of Psychology, University of Sussex, UK.

Email: [aidan.horner@york.ac.uk](mailto:aidan.horner@york.ac.uk).

HomePage: <http://www.aidanhorner.org/>

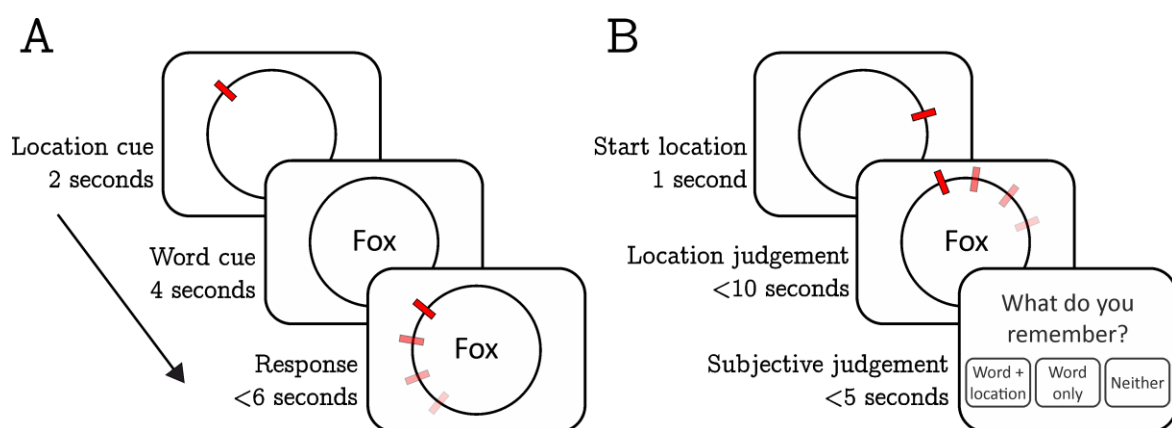
Twitter Handle: @aidanhorner

Keywords: forgetting, episodic memory, schema, generalisation

Although we know a great deal about how we remember, we know surprisingly little about how we forget. Ebbinghaus (1885) pioneered the study of forgetting, repeatedly training and testing himself on nonsense syllables over varying timescales to reveal the classic “Ebbinghaus forgetting curve” that is taught at undergraduate psychology. Decades of debate and research also centered on what causes forgetting - interference or decay (Jenkins & Dallenbach, 1924; Sadeh et al., 2014; Wixted, 2004). However, there are still many unknowns in relation to forgetting. The specific question we asked in our study was whether memories become less precise (or more ‘blurry’) over time, or whether they are simply forgotten as a whole (Berens et al., 2020).

Imagine that you discover an apple tree in an unfamiliar forest. A day later you wish to return, so you need to remember where the tree was located. One possibility is that you remember precisely where the tree was such that no forgetting has taken place. Another possibility is that you forget where the tree was located entirely; you know it must be somewhere in the forest but have no idea where to find it. A third possibility is that you remember the rough location of the tree, but not precisely. Assuming that you originally encoded the precise location, this third possibility suggests that your memory for the retrieved location has become less precise over time. In other words, your memory has become blurry. We were specifically interested in assessing whether forgetting involves a simple loss of individual memories (forgetting a location entirely) or whether forgetting involves a loss of precision (you remember the location, but less precisely than before).

How can we measure these two forms of forgetting? The recent development of ‘precision’ memory measures (Bays et al., 2009; Harlow & Donaldson, 2013; Harlow & Yonelinas, 2016; Murray et al., 2015; Richter et al., 2016) provides the ideal experimental setup to assess these two accounts. Precision memory measures typically pair an item (e.g., a word) with a continuous circular variable (e.g., a location on a circle) at encoding (Figure 1). At retrieval, participants are required to move a marker on the circle back to the remembered location associated with the cued word. For each word-location association you therefore have a



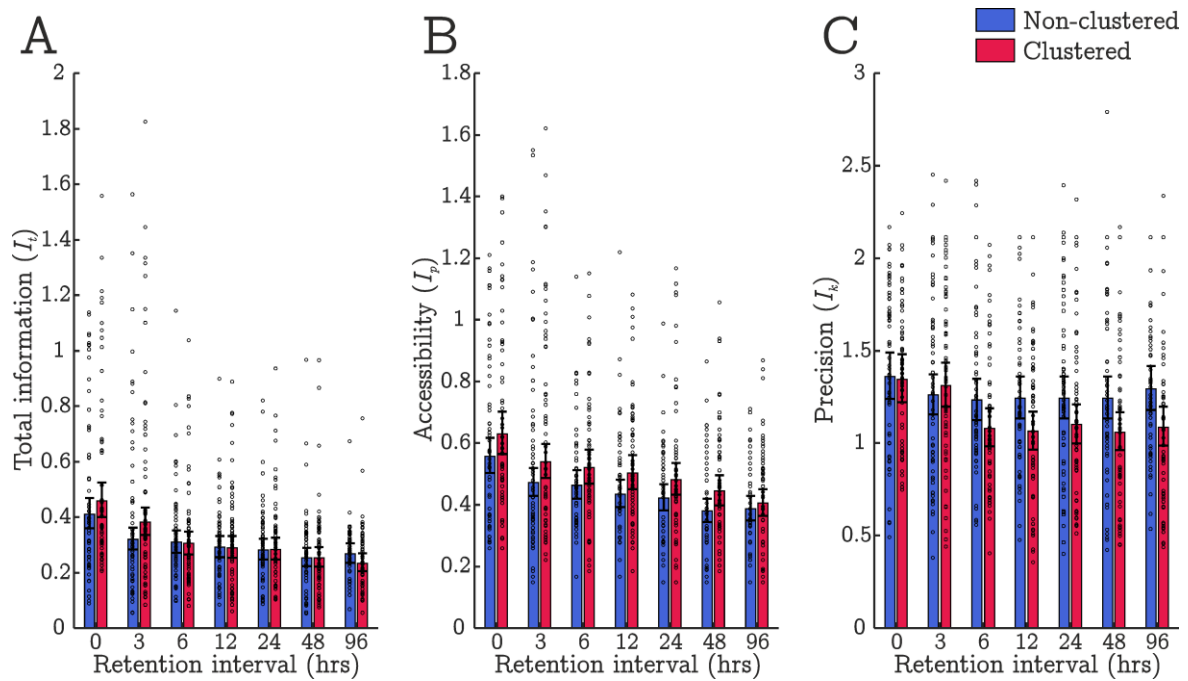
**Figure 1.** Schematic of the experimental procedure. **A.** Example encoding trial. **B.** Example retrieval trial. Taken from the [study preprint](#). © the authors, 2020. Licensed under a [Creative Commons Attribution 4.0 International Licence](#).

continuous measure of the angular error between the true location of the word and the remembered location (as opposed to a binary ‘correct/incorrect’ that you would get from a more typical memory experiment). Using a statistical approach called mixture modelling, the distribution of these angular errors across multiple word-locations can be used to extract two measures of memory: accessibility and precision. Whereas accessibility refers to the proportion of word-location associations remembered (similar to a typical memory experiment), precision refers to how precisely those accessible word-locations were remembered. Although there is theoretical work questioning whether they truly reflect independent memory processes (Schurgin et al., 2020), there is evidence to suggest they can be dissociated behaviourally (Harlow & Yonelinas, 2016) and neurally (Richter et al., 2016). You might remember few word-locations precisely (or imprecisely) or remember many word-locations precisely (or imprecisely).

Now we have two independent measures of memory, how can we use them to distinguish between the two types of forgetting outlined above? One problem with studying forgetting is that you can’t measure it directly. Take a typical forgetting experiment. Your participant learns some information, and then you test their memory for the learnt information immediately after encoding and then again after 3 days. You acquire a memory measure at two study-test delays (e.g., proportion correct). This provides you with a measure of memory performance at both tests, but no direct measure of forgetting. The inference required is that the difference in performance between the two study-test delays equates to the amount of information ‘forgotten’. All things being equal - participants paying attention during the task etc. - any change in performance over time should reflect the fact that information that was previously accessible can no longer be retrieved.

In this study, we inferred levels of forgetting by comparing behavioural performance across conditions with different study-test delays. In the case of our two accounts of forgetting, if individual word-location were forgotten in their entirety, then we should see decreases in accessibility over time (but no decreases in precision). However, if forgetting involves your memory for individual word-locations becoming ‘blurry’ we should see decreases in precision over time (but no decreases in accessibility). Thus, we can use changes in these two memory measures over time to infer what form forgetting takes.

A second problem with studying forgetting is the issue of re-testing. Ideally, you want to track memory for individual items over time. However, testing an individual item is known to have a strong effect on retention of that item (Roediger & Karpicke, 2006). There are two possible ways to avoid this problem: (1) test a subset of items at each delay in a within-participant design or (2) test all items at a single delay, but vary the delay across participants. Given the relatively large number of words needed to conduct the mixture modelling approach described above, we used a between-subjects design where each participant was tested on all word-location associations in a single session following one of 7 study-test delays (0hrs, 3hrs, 6hrs, 12hrs, 24hrs, 48hrs, 96hrs).



**Figure 2.** Results of Berens et al. (2020). Means (and 95% confidence intervals) for different measures of memory performance plotted by study-test delay and clustering condition; **A.** Overall memory performance, **B.** Memory accessibility, **C.** Memory precision. Taken and adapted from the [study preprint](#). © the authors, 2020. Licensed under a [Creative Commons Attribution 4.0 International Licence](#).

We therefore assessed accessibility and precision for word-location associations across multiple retention intervals ( $N=431$ ) (Berens et al., 2020). Given this was a registered report, we received ‘in principle’ acceptance for our design following peer review prior to data collection. Despite our hypothesis that forgetting involves losses in both accessibility and precision, we only found evidence for losses in accessibility over time. Indeed, Bayesian statistics provided strong evidence against any losses in precision. This suggests that an increasing proportion of word-location associations become simply inaccessible over time – we are unable to retrieve them – however the associations we do remember are just as precisely remembered after a delay as they are immediately after encoding. Memories do not become ‘blurry’ over time, they simply become inaccessible.

In the same study, we were also interested in seeing how the presence of a pattern across a set of words might modulate memory performance and forgetting. Returning to our apple tree example, now imagine you encounter multiple apple trees clustered in the same area of the forest. If you had no pre-existing predictions about the location of apple trees, you might be able to use this spatial clustering to extract a ‘pattern’ – that particular types of tree appear to be common in a general area. This question was inspired by theoretical (McClelland et al., 1995; Sekeres et al., 2018; Winocur & Moscovitch, 2011) and experimental (Richards et al., 2014; Tse et al., 2007) research suggesting that, when we experience a set of related events that have a generalisable pattern or structure, we can use that pattern to form schematic representations that support memory over longer retention intervals. While we might therefore forget individual apple tree locations over time, systems consolidation processes may allow for

the formation of schematic representations that can then be used to support memory decisions (i.e., to guess where another apple tree might be).

To explore this possibility, the words in our experiment belonged to two semantic categories, natural and human-made. Each category was assigned to one of two conditions. In one condition, the non-clustered condition, the locations of the words were randomly chosen such that there was no underlying structure or pattern. In the other condition, the clustered condition, word locations were clustered in a specific area of the circle. In the clustered condition participants may have been able to extract the underlying pattern and form a ‘schema’.

The extraction of an underlying pattern, or formation of a schema, is thought to support memory decisions over longer timescales (McClelland et al., 1995). This would manifest as less forgetting over time in the clustered relative to non-clustered condition. However, learning a general pattern might lead to a reduction in ‘memory specificity’ for individual items (Arpit et al., 2017). Thus, a schematic representation might support more generalised memory decisions at the expense of precision. This would lead to greater accessibility but less precision in the clustered relative to non-clustered condition. While similar experimental manipulations have been used in rodents (Richards et al., 2014), they have not been tested previously in humans in relation to forgetting (though see Graves et al., 2020; Richter et al., 2019; Tompary et al., 2020 for similar manipulations).

We therefore predicted that forgetting would be reduced in the clustered relative to non-clustered condition. We also predicted that accessibility would be greater, and precision would be reduced, in the clustered relative to non-clustered condition. While we saw strong evidence consistent with the latter prediction, we actually found strong evidence against the hypothesis that clustering decreases forgetting. The clustering of word-locations in one semantic category increased accessibility but decreased precision relative to word-locations that weren’t clustered; however, this clustering manipulation had no effect on the amount of information that was forgotten. The presence of a pattern did not lead to less forgetting over time. Interestingly, in exploratory analyses we saw a close correspondence between accessibility for word-locations and the extent to which participants placed words in the clustered condition close to the underlying pattern. This perhaps suggests that performance was primarily driven by memory for individual word-locations rather than an independent schematic representation (also see Tompary et al., 2020).

More recently, we have focussed more on the clustered and non-clustered manipulation. Theoretically, if participants are forming ‘schemas’ they should be able to use these generalised representations to make decisions about where semantically related ‘novel’ words might be located. If you know where the apple trees are located in a forest, you might be able to use this knowledge to make an accurate prediction about where a new apple tree might be located. In a series of four new preregistered experiments (both in-person and online), we used a similar design to Berens et al. (2020) but embedded ‘novel’ words at retrieval (Cockcroft et al., 2021). These words weren’t seen at encoding but did belong to the same semantic categories as the clustered and non-clustered condition. If participants are able to use the underlying pattern in

the clustered condition, they should place novel words in similar locations to the locations of semantically related words. We found participants were able to do this – they made sensible ‘guesses’ about where novel words from the same semantic category as the clustered condition might be located.

These new results suggest participants can use the clustered pattern to make sensible generalisation decisions, perhaps suggesting the presence of a schematic representation that is able to guide these decisions. Interestingly, we also saw that participants placed novel words in the non-clustered semantic category at locations opposite to the cluster (i.e., on the other side of the circle from most of the clustered words). Not only was this ‘avoidance’ effect present for novel words, but it was also seen for previously seen words in the non-clustered category. The mere presence of a pattern in one condition therefore alters both memory performance and generalisation in the other condition. While it is not clear what precise mechanism is driving this ‘avoidance’ effect, it perhaps suggests that the presence of a schema can bias decisions for unrelated items. To discover the underlying mechanisms, further experimental work is needed alongside more formal computational modelling. Ultimately, it is possible that the complex pattern of behavioural results across these two studies (and studies from other labs, e.g. Brady et al., 2015; Tompary et al., 2020) may be explicable by a combination of more basic psychological phenomena (e.g., interference or base rate neglect) without the need for the formation of independent schematic representations.

To conclude, precision memory measures offer the means to assess potentially distinct forms of memory and forgetting. We have shown that forgetting is predominantly associated with decreases in accessibility. If you forget a word-location association, then it is no longer retrievable. If you remember a word-location association you remember it with just as much precision as you did shortly after encoding. By introducing an underlying pattern (clustering semantically related words in an area of the circle), participants were able to retrieve more word-location associations (i.e., they showed greater accessibility relative to when no pattern was present). However, this increase in accessibility came at the expense of a decrease in precision. Finally, participants were able to use this underlying pattern to make sensible generalisation decisions about where novel semantically related words might be located. Although the changes in behaviour as a function of the introduction of a pattern might point towards the presence of a schematic representation, other evidence suggests this may not necessarily be the case. As such, our results place important constraints on models of forgetting and pattern learning. While further work is necessary to draw firm conclusions, this body of work demonstrates the utility of precision memory experiments in relation to understanding the psychological processes involved in memory, forgetting, and generalisation. Ultimately, the work highlights the difficult theoretical questions that are still to be addressed in the field but points towards possible solutions by combining this experimental approach with formal computational modelling.

## References:

- Arpit, D., Jastrzębski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., & Lacoste-Julien, S. (2017). A closer look at memorization in deep networks. *Proceedings of the 34th International Conference on Machine Learning, in Proceedings of Machine Learning Research*, 70, 233-242.
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10), 7-7. <https://doi.org/10.1167/9.10.7>
- Berens, S. C., Richards, B. A., & Horner, A. J. (2020). Dissociating memory accessibility and precision in forgetting. *Nature Human Behaviour*, 4(8), 866-877. <https://doi.org/10.1038/s41562-020-0888-8>
- Brady, T., Schacter, D. L., & Alvarez, G. (2018). The adaptive nature of false memories is revealed by gist-based distortion of true memories. *Journal of Vision*, 15, 948. <https://doi.org/10.1167/15.12.948>
- Cockcroft, J. P., Berens, S., Gaskell, M. G., & Horner, A. J. (2021). Schematic information influences memory and generalisation behaviour for schema-relevant and -irrelevant information. *PsyArXiv*. <https://doi.org/10.31234/osf.io/nzurq>
- Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology*. Teachers College, Columbia.
- Graves, K. N., Antony, J. W., & Turk-Browne, N. B. (2020). Finding the Pattern: On-Line Extraction of Spatial Structure During Virtual Navigation. *Psychological Science*, 31(9), 1183-1190. <https://doi.org/10.1177/0956797620948828>
- Harlow, I. M., & Yonelinas, A. P. (2016). Distinguishing between the success and precision of recollection. *Memory*, 24(1), 114–127. <https://doi.org/10.1080/09658211.2014.988162>
- Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence During Sleep and Waking. *The American Journal of Psychology*, 35, 605–612. <https://doi.org/10.2307/1414040>
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419–457. <https://doi.org/10.1037/0033-295X.102.3.419>
- Murray, J. G., Howie, C. A., & Donaldson, D. I. (2015). The neural mechanism underlying recollection is sensitive to the quality of episodic memory: Event related potentials reveal a some-or-none threshold. *NeuroImage*, 120, 298–308. <https://doi.org/10.1016/j.neuroimage.2015.06.069>
- Richards, B. A., Xia, F., Santoro, A., Husse, J., Woodin, M. A., Josselyn, S. A., & Frankland, P. W. (2014). Patterns across multiple memories are identified over time. *Nature Neuroscience*, 17(7), 981-986. DOI: [10.1038/nn.3736](https://doi.org/10.1038/nn.3736)
- Richter, F. R., Bays, P. M., Jeyarathnarajah, P., & Simons, J. S. (2019). Flexible updating of dynamic knowledge structures. *Scientific Reports*, 9(1), 1-15. <https://doi.org/10.1038/s41598-019-39468-9>

- Richter, F. R., Cooper, R. A., Bays, P. M., & Simons, J. S. (2016). Distinct neural mechanisms underlie the success, precision, and vividness of episodic memory. *elife*, 5, e18260. DOI: [10.7554/eLife.18260](https://doi.org/10.7554/eLife.18260)
- Roediger III, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249-255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
- Sadeh, T., Ozubko, J. D., Winocur, G., & Moscovitch, M. (2014). How we forget may depend on how we remember. *Trends in Cognitive Sciences*, 18(1), 26-36. DOI: [10.1016/j.tics.2013.10.008](https://doi.org/10.1016/j.tics.2013.10.008)
- Schurigin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, 4(11), 1156-1172. DOI: [10.1038/s41562-020-00938-0](https://doi.org/10.1038/s41562-020-00938-0)
- Sekeres, M. J., Winocur, G., & Moscovitch, M. (2018). The hippocampus and related neocortical structures in memory transformation. *Neuroscience Letters*, 680, 39-53. DOI: [10.1016/j.neulet.2018.05.006](https://doi.org/10.1016/j.neulet.2018.05.006)
- Tompary, A., Zhou, W., & Davachi, L. (2020). Schematic memories develop quickly, but are not expressed unless necessary. *Scientific Reports*, 10(1), 1-17. DOI: [10.1038/s41598-020-73952-x](https://doi.org/10.1038/s41598-020-73952-x)
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., ... & Morris, R. G. (2007). Schemas and memory consolidation. *Science*, 316(5821), 76-82. DOI: [10.1126/science.1135935](https://doi.org/10.1126/science.1135935)
- Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5), 766-780. DOI: [10.1017/S1355617711000683](https://doi.org/10.1017/S1355617711000683)
- Wixted, J. T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology*, 55, 235-269. DOI: [10.1146/annurev.psych.55.090902.141555](https://doi.org/10.1146/annurev.psych.55.090902.141555)