

## A Framework for Computational Models of Human Memory

**Matthew A. Kelly**

The Pennsylvania State University  
University Park, PA  
matthew.kelly@psu.edu

**Robert L. West**

Carleton University  
Ottawa, Canada  
robert.west@carleton.ca

### Abstract

We present analysis of existing memory models, examining how models represent knowledge, structure memory, learn, make decisions, and predict reaction times. On the basis of this analysis, we propose a theoretical framework that characterizes memory modelling in terms of six key decisions: (1) choice of knowledge representation scheme, (2) choice of data structure, (3) choice of associative architecture, (4) choice of learning rule, (5) choice of time variant process, and (6) choice of response decision criteria. This framework is both descriptive and proscriptive: we intend to both describe the state of the literature and outline what we believe is the most fruitful space of possibilities for the development of future memory models.

### Introduction

Imagine that memory is a cardboard box filled to the brim with things that are important to you. In fact, imagine that it has *Important stuff!* scribbled on the side of it in black marker, triple underlined. But at any particular moment, only one thing in the box is important to you: the thing you need right now. How do you find it?

Some early models postulated that finding something in memory was a serial search process (Tulving 1976). When you need to find something in the box, you go through the contents of the box one item at a time until you find the thing you need. If the thing you are looking for is not in the box, you will have to search the entire box to find that out. However, people make judgements that a test item isn't in a study set too quickly for an exhaustive serial search model to be correct (Atkinson and Juola 1974).

Other early models suggested that items in memory can be accessed directly (Kintsch 1970). If the things in the box are carefully organized, you can just grab the item you need without searching. However, direct access models have trouble accounting for the effects of other items on the recognition of a given item (Clark and Gronlund 1996).

In global matching models, memory retrieves an aggregate across all experiences. To find something in the box, you turn the box upside down, dump its contents on the floor, and rifle through the jumbled mess for the thing you need. Global matching models are able to account for the speed of retrieval, like direct access models, but also the influence of other items on retrieval, like serial search models.

Mathematical analysis of the global matching models (Clark and Gronlund 1996; Hintzman 1990; Humphreys et al. 1989; Kahana, Rizzuto, and Schneider 2005; Kelly, Mewhort, and West 2017) found strong similarities between models, including MINERVA 2 (Hintzman 1984), CHARM (Eich 1982), TODAM (Murdoch 1993), the Matrix model (Humphreys et al. 1989), and SAM (Raaijmakers and Shiffrin 1981). While different models describe the structure and dynamics of human memory using different terms, similarities in the underlying mathematics of these models provide evidence of an emerging consensus.

With the goal of advancing toward a unified theory of human memory, we propose a theoretical framework for characterizing computational memory models. This framework asserts the key assumptions of global matching models, as set out by Hintzman (1986, p. 411):

1. only episode traces are stored in memory,
2. repetition produces multiple traces of an item,
3. a retrieval cue contacts all traces simultaneously,
4. each trace is activated according to similarity to the cue,
5. all traces respond in parallel, retrieved information reflects their summed output.

Hintzman's key assumptions loosely describe the behaviour of an associative neural network (Kelly, Mewhort, and West 2017). That said, not all models that realize Hintzman's framework are neural networks, including Hintzman's own MINERVA 2 (Hintzman 1984).

Some models that seem to violate these assumptions, such as the *many-to-many* models (Jones and Mewhort 2007; Kelly, Kwok, and West 2015; Rutledge-Taylor et al. 2014) and REM (Shiffrin and Steyvers 1997) can be understood as higher level descriptions that can be translated into a model within this framework (Kelly, Mewhort, and West 2017). The declarative memory of the ACT-R cognitive architecture (ACT-R DM; Anderson and Lebiere 1998) also has strong mathematical similarities to global matching models (Dimov 2016) and *many-to-many* models (Kelly, Kwok, and West 2015). As such, ACT-R DM can also be understood as a higher level description of a model within this framework.

Thus we take Hintzman's assumptions as a description of a largely unspoken mathematical consensus across computa-

tional models of human memory: Namely, that human memory behaves like a type of associative neural network.

However, this framework of assumptions describes a vast space of possible memory models. In particular, current models in the literature differ on six dimensions:

1. Choice of representation scheme: How representations are constructed and combined.
2. Choice of data structure: *many-to-one*, *one-to-one*, or *many-to-many* model.
3. Choice of association: *Hetero-* or *auto-*association.
4. Choice of learning rule: *Error-less* Hebbian learning, or prediction *error-driven* learning.
5. Choice of time variant process: Is it the echo or the probe that determines time?
6. Choice of response decision criteria: When does the model respond?

Some of these choices are largely a matter of implementation with little consequence for theory, whereas others are grey areas where more theoretical work needs to be done. We discuss each of these six choices in what follows.

### Choice of representation scheme

The symbol processing models are the earliest mechanistic models of cognition. The LISP (list processing) language, developed in 1958, became popular in the budding fields of artificial intelligence and cognitive modelling. LISP programs operate on lists of discrete symbols. LISP heavily influenced the nature of early cognitive models, such as the cognitive architecture ACT-R (Anderson and Lebiere 1998). A number of memory models, such as SAM (Raaijmakers and Shiffrin 1981) and ACT-R DM are described in terms of the storage and retrieval of discrete symbols.

Symbolic models are systems most appropriately described at the level of expressions of symbols and the manipulation of those symbols (Clark 2001). Symbolic models take inspiration from linguistics and logic. Typically, in a symbolic architecture, the atomic units of the model are concepts (e.g., *dog*, *glowing*, *noun*, or *anger*) and complex units arise from characterizing the relations between those concepts. Arguably, Locke (1690) proposed one of the first symbolic architectures, well before there were computers, analyzing the mind in terms of a set of processes that manipulate expressions of elementary ideas.

With the advent of parallel distributed processing in the 1980s, neurally-inspired connectionist models became a popular alternative to traditional symbolic models. By contrast to symbolic models, the atomic units of a connectionist system are artificial neurons or simple neuron-like nodes, whereas a concept will have a complex implementation as a pattern of activation across those units. Connectionist models transform vectors of activation values using matrices of connection weights. While the patterns of activation at the input or output of a connectionist model may be interpretable as a known concept, the patterns of activation on the hidden layers of neurons or nodes between the input and output are often difficult to interpret.

To achieve a full, theoretical understanding of a cognitive process and how it relates to the physical world, explanations of the process need to be provided at both symbolic (i.e., representational) and sub-symbolic levels of description. The classic symbolic approaches to modelling do not account for how the symbol manipulations described in the model could arise from neural tissue, or account for how the symbols themselves come into existence. Classic connectionist approaches are more concerned with neural plausibility, but are notoriously opaque, doing little to aid our understanding of the cognitive processes modelled.

By contrast, what is known as the vector-symbolic (Gayler 2003; Plate 1995) or conceptual spaces (Lieto, Chella, and Frixione 2017) approach to modelling explicitly provides an account at both symbolic and sub-symbolic levels of description. In computational, vector-based models of memory, a *to-be-remembered item* is represented as a vector. Vector representation allows items to be treated as complex entities with internal features and also allows for errors and fuzzy composites of items to be constructed. Vectors are also amenable to implementation in neural models. Thus, there are good reasons to prefer to use vectors as a means of representing items in memory.

In vector symbolic architectures, the vectors represent symbolic information. These vectors, or symbols, can be combined and manipulated using a small number of mathematical operations (e.g., convolution, multiplication, addition, permutation), which can be understood as sub-symbolic processes (see Kelly, Blostein, and Mewhort 2013 for review). However, the information processing models built from these operations are best characterized at a symbolic level of description (Kelly and West 2012). Importantly, the modelling decisions made at the sub-symbolic level are to some degree independent of those made at the symbolic level (i.e., decisions about cognitive processes).

At the symbolic level, we have to choose how these vectors are combined to create more complex representations (e.g., how is a percept or concept structured?). The representation scheme used greatly influences the behaviour of the model. Unfortunately, there is a great deal of variability in the representations across models, indicating a lack of a consistent theory of representation.

The HEM (Jamieson and Mewhort 2011) and BEAGLE (Jones and Mewhort 2007) models differ in that HEM uses an instance-based memory whereas BEAGLE uses a concept-based memory, as discussed in the next section. However, HEM and BEAGLE use the same basic encoding scheme: Given a set of co-occurring items within a context, add to memory all combinations of them (all  $n$ -grams or skip-grams) or a random subset of all combinations to model degraded performance.

Of course, the relationships that humans learn and remember are not a random subset of all relationships observed in the environment. Integration with an attention model would allow memory models to better capture which associations between items are attended to and retained by the mind.

Vectors representing items are typically randomly generated by the models, but ought to be generated according to semantic (Kelly, Reitter, and West 2017) or perceptual (Cox

et al. 2011; Kelly, Blostein, and Mewhort 2013) features, which would require integrating memory and perception.

### Choice of data structure

Kelly, Kwok, and West (2015) classified vector-based memory models in the literature as belonging to one of three classes: *many-to-one*, *one-to-one*, and *many-to-many*. In *many-to-one* models, the memory model stores all experiences as a composite, superimposed on a single data structure in a fully distributed manner. In *one-to-one* models, the memory model stores each experience individually, as if it were a box on a shelf in a giant warehouse. In *many-to-many* models, experiences are decomposed into relationships between concepts. Concepts are stored individually and can be understood as points on a hyper-sphere, shifting on the surface of a globe of memory and meaning, moving closer or further from the other concepts on the basis of experience.

To give example applications of each model, a *many-to-one* model can be used as working memory buffer, a *one-to-one* model can be used as an episodic memory, and a *many-to-many* model can be used as semantic memory. However, working, episodic, and semantic memory are theoretical constructs, whereas these three classes of model are data structures with varying theoretical interpretation.

Given demonstrations of equivalence between *many-to-one*, *one-to-one*, and *many-to-many* models (Kelly, Mewhort, and West 2017), the choice of data structure for a memory model is not a theoretical distinction but rather a pragmatic one dependent on the scale of the model.

A *many-to-one* model stores all memories in a single fully distributed data structure (Franklin and Mewhort 2015; Humphreys et al. 1989; Murdock 1993). As such, the model is invariant in scale with respect to the number of instances of experiences or concepts stored by the model. These models are easily amenable to neural implementation (Eliasmith 2013) and are likely the ground truth of human memory.

*Many-to-one* models work best when only a small number of instances are stored in the memory system, as these models compress all memories into a single data structure of invariant scale. Conversely, if the data structure is made larger, it can accommodate more memories but becomes unwieldy to compute on a serial computer - though neuromorphic computers may make this possible in the future.

*One-to-one* models store one vector per instance of an experience, allowing memories to be stored and retrieved with fidelity. *One-to-one* models such as MINERVA 2 (Hintzman 1984) approximate the behavior of a large *many-to-one* model (Kelly, Mewhort, and West 2017). *One-to-one* models scale with the number of experiences and thus become unwieldy for simulating a lifetime of experiences, making them unsuitable for tasks such as language processing.

*Many-to-many* models, commonly known as distributional semantics models (Jones and Mewhort 2007; Mikolov et al. 2013), store one vector per distinct concept. These models can be understood as a higher-level description of a *one-to-one* model, as concepts can be understood as useful aggregates of individual experiences (Hintzman 1986). The advantage of these models is that they grow only with the

addition of new concepts, and as such, are suitable for language processing (Jones and Mewhort 2007) and big data (Rutledge-Taylor, Vellino, and West 2008).

The difference between a *many-to-many* model and a *many-to-one* or *one-to-one* model is also one of emphasis. A *many-to-many* model is a concept-based memory system designed to learn about items and the relationship between items. Conversely, an instance-based memory system is designed to record and retrieve experiences.

### Choice of association

In a *hetero-associative* memory (Franklin and Mewhort 2015; Humphreys et al. 1989; Kelly, Kwok, and West 2015; Murdock 1993; Raaijmakers and Shiffrin 1981), associations are formed between two different items. A cue can retrieve completely dissimilar memories that have become associated with that cue through experience.

In an *auto-associative* memory (Farrell and Lewandowsky 2002; Hintzman 1984; Anderson and Lebiere 1998), an item forms associations with itself. A cue always retrieves a similar memory, or combination of similar memories, and new information is a matter of remembering more detail.

For recall, the distinction between hetero- and auto-associative makes little difference. Given a pair of items to remember,  $x$  and  $y$ , and then cued with the item  $x$ , an auto-associative model retrieves a completion, an echo of both  $x$  and  $y$ . A hetero-associative memory given  $x$  just retrieves  $y$ .

However, modeling a recognition task with an auto-associative memory is more straightforward than with a hetero-associative memory. To test if an auto-associative memory system has studied a pair  $x y$ , one presents the auto-associative memory with the cue  $x y$ , and if that pair is present in memory, the model will retrieve  $x y$ .

How a hetero-associative architecture performs a recognition task is less well defined than for an auto-associative architecture. This is because recognition is fundamentally a task of giving the model a cue and seeing if you get something like that cue back again.

The choice between a hetero-associative model and an auto-associative model is a false one. The correct answer is both hetero-associations and auto-associations are useful tools for constructing a memory model.

At the level of representation, both auto- and hetero-associations can be encoded in a vector using linear algebra operations (concatenation, convolution, etc.). Both forms of association can and should be used to represent a stimulus as appropriate for the task.

At the level of the architecture of the model, a hetero-associative architecture is a set of connections between two neural groups that translates a signal into another signal. As such, hetero-associative architectures are pervasive in the brain, though they are arguably not memory but translation systems. Conversely, an auto-associative architecture is a recurrent neural loop that acts as an error-correcting buffer for holding a standing pattern of neural activation.



## Choice of learning rule

In error-less learning, a vector representing an experience,  $\mathbf{x}$ , is simply added to memory when the model has that experience. In error-driven (or *reinforcement*) learning, the memory system is only updated when something surprising happens. Memory makes a prediction  $\mathbf{e}$  on the basis of the current situation, and memory is updated by the difference  $\mathbf{x} - \mathbf{e}$  between the prediction and the actual experience  $\mathbf{x}$ .

Do these two kinds of learning necessitate separate memory systems? For example, ACT-R (Anderson and Lebiere 1998) has two types of memory, *declarative* and *procedural*, which respectively use *error-less* and *error-driven* learning.

Conversely, MINERVA-AL is a variant of MINERVA that uses reinforcement learning. MINERVA-AL makes and corrects predictions to capture numerous associative learning phenomena from both the animal and human learning literature (Jamieson, Crump, and Hannah 2012). MINERVA-AL demonstrates that error-driven learning does not require a distinct memory architecture, as aside from how it learns, MINERVA-AL is otherwise a standard MINERVA model.

Can human memory be modelled using error-driven learning alone? Recent research has found that even in one-shot declarative learning, human learning is modulated by prediction error (Greve et al. 2017), which suggests that only error-driven learning is necessary to model human memory.

## Choice of time variant process

Cognitive models live and die by their ability to model and predict human experimental data. Much of the experimental data collected is reaction time data. How long does it take the model to recall something? How long does it take for the model to decide if a cue is familiar or unfamiliar?

One possibility is that retrieval occurs over a series of iterations as the echo is used as a new cue, which in turn, retrieves a clearer echo (Farrell and Lewandowsky 2002).

Time to clean up the echo is not the only candidate for the time variant process that accounts for variability in response time in memory tasks. The Iterative Resonance Model (Mewhort and Johns 2005) suggests *sharpening*, a process of incrementally increasing the precision of memory retrieval. However, there is no clear neural implementation for sharpening, and it seems counter-intuitive to suggest that memory has a sharpness dial that can be adjusted at will. Why would the dial ever be set at less than perfectly sharp?

Cox and Shiffrin (2012), who use a variant of the REM (Shiffrin and Steyvers 1997), suggest sampling of features from the probe as the time variant process. Again, this seems counter-intuitive, as it suggests that we are slow to recall some things because we are slow to perceive them correctly. For these reasons, we prefer time to clean up the echo as the time variant process. Determining what the time variant process is in human memory is a matter for further research.

## Choice of response decision criteria

At some point, a decision needs to be made on the basis of retrieved information and the model produces a response. In an actual brain, this decision to act is made by whatever part of the brain needs to use the retrieved information. In

a memory model, the only part of the brain modelled is the memory system itself. As such, the decision is made by a set of criteria that act as a stand-in for those other systems. As a result, these criteria are a somewhat artificial imposition on the model, magic parameters that turn a theory of memory into a system that produces reaction time data.

In principle, integration into a larger cognitive system could help absolve memory models of their need for somewhat arbitrary criteria. Unfortunately, the ACT-R architecture (Anderson and Lebiere 1998) treats time to retrieve information from memory as internal to the memory model, rather than arising from the interaction between the memory retrieval process and an external decision process.

In a recognition task, the decision that a cue is familiar can be modelled by measuring the similarity between the cue and the echo retrieved in response to the cue. The decision that a cue is unfamiliar has, however, proven more difficult to model. Mewhort and Johns (2005) argue that decisions to reject a cue as unfamiliar is dependent not on similarity but on contradiction. Both Mewhort and Johns's IRM model and Johns, Jones, and Mewhort's (2012) RSS model use contradiction to reject foils, though the two models do not agree on how contradiction ought to be measured.

Criterion setting remains a black art in memory modelling because response times are influenced by the particulars of the task and the strategy that a participant uses to complete the task. Rutledge-Taylor et al. (2014) make an analogy between rejecting foils and cooking popcorn. When cooking popcorn, you learn roughly how much time passes between one kernel popping and the next. If more than the expected amount of time has passed since the last pop, you may judge that the popcorn is finished. Likewise, participants may learn how long it usually takes for them to recollect something in a memory task. If they have failed to recollect the cue in about that period of time, they may judge that the cue is a foil.

In short, response times in a memory model are a Gordian knot that tangles together participant strategy, the passage of time, and the quantity and quality of the information being retrieved from memory. By necessity, this aspect of memory models may remain ad hoc for some time.

## Discussion

Here we summarize each of the six key decisions for constructing a memory model in our framework, and the extent to which each of these decisions reflect theory-neutral implementation details (1, 2, 3), emerging theoretical consensus (3, 4), or areas in need of future research (1, 4, 5, 6).

1. Vector-symbolic architectures (Gayler 2003) and conceptual spaces (Lieto, Chella, and Frixione 2017) prove that symbolic systems, such as ACT-R (Anderson and Lebiere 1998), can arise from distributed connectionist networks. As such, the choice of symbolic, vector symbolic, or connectionist is a matter of a model's level of description, rather than a theoretical claim.

However, a theory of human memory ought to describe the structure of the representations stored in memory. Integrating memory with perception and attention is necessary if research is to develop a coherent theory of rep-

resentation. This integration is more easily accomplished using a vector-symbolic or connectionist model as the external world is not inherently carved up into symbols.

2. The choice of data structure is not a theoretical distinction but a pragmatic one dependent on the scale of the model.
3. A memory model may be a hybrid of both hetero- and auto-associative representations and architectures.
4. The ACT-R cognitive architecture postulates two types of memory, *declarative* and *procedural*, which correspond to two types of learning, *error-less* and *error-driven*. While there is considerable evidence that prediction error can drive learning in both humans and animals (Jamieson, Crump, and Hannah 2012), the evidence for pure, error-less learning is scant. Even in one-shot declarative learning, prediction error has an effect on human learning (Greve et al. 2017). While more research may need to be done on this question, we advocate that future memory models and cognitive architectures explore using prediction error in all memory encoding.
5. Why are people faster to remember somethings and slower to remember other things? Memory retrieval is a time variant process. But what is taking up all that time? There are three candidates in the literature: time to clean up the echo (Farrell and Lewandowsky 2002), sharpening the echo (Mewhort and Johns 2005), and sampling features from the probe (Cox and Shiffrin 2012). More research is needed to settle this question.
6. Response decision criteria are an artifact of modelling memory separate from other cognitive functions. Response times are emergent from the interactions of memory with perception, action, and participant strategy. Better integration of memory models with cognitive architectures is necessary to escape from ad hoc criterion setting.

We note that developing better memory models requires integration with models of perception, attention, action, and meta-cognitive strategy, in order to arrive at clear answers to fundamental questions about the nature of representations in memory and the time course of memory retrieval.

## Conclusion

We propose a theoretical framework for current and future models of memory. The framework consists of Hintzman's (1986) key assumptions and a set of six key decisions. On the basis of this framework, we sketch a course of development for future memory models.

The theoretical framework we propose allows for memory to be modelled at multiple levels of analysis, from the level of individual, biological neurons and messages passed between groups of those neurons (*many-to-one* models), to the level of the events processed by memory (*one-to-one* models), to the level of the concepts which emerge as aggregates across those events (*many-to-many* models), and on up, to the arbitrarily abstract concepts that emerge from aggregating across concepts (Kelly, Reitter, and West 2017). Furthermore, estimates of Bayesian probability arise from the vector algebra of *many-to-many* models (Kelly, Kwok,

and West 2015). Thus, models based on Bayesian probability, such as the ACT-R declarative memory (Anderson and Reder 1999), can be understood as higher-level descriptions within this theoretical framework.

Simon (1969) famously argued that to account for the complexity of human behaviour we need to account for the influence of two environments: the external environment of the world and the internal environment of memory. Thus we feel that the memory framework we propose is an important step towards a working model of the human mind.

## Acknowledgements

The authors gratefully acknowledge funding to M. A. Kelly from an Ontario Graduate Scholarship and National Science Foundation grants SES-1528409 and BCS-1734304 and to R. L. West from the Natural Sciences and Engineering Research Council of Canada.

## References

- Anderson, J. R., and Lebiere, C. 1998. *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., and Reder, L. M. 1999. The fan effect: New results and new theories. *Journal of Experimental Psychology: General* 128:186–197.
- Atkinson, R. C., and Juola, J. F. 1974. Search and decision processes in recognition memory. In Krantz, D. H.; Atkinson, R. C.; Luce, R. D.; and Suppes, P., eds., *Contemporary developments in mathematical psychology: Vol 1. Learning, memory, and thinking*. San Francisco: W. H. Freeman. 243–293.
- Clark, S. E., and Gronlund, S. D. 1996. Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review* 3:37–60.
- Clark, A. 2001. Symbol systems (chapter 2). *Mindware: An Introduction to the Philosophy of Cognitive Science* 27–38.
- Cox, G. E., and Shiffrin, R. M. 2012. Criterion setting and the dynamics of recognition memory. *Topics in Cognitive Science* 4:135–150.
- Cox, G. E.; Kachergis, G.; Recchia, G.; and Jones, M. N. 2011. Towards a scalable holographic representation of word form. *Behavior Research Methods* 43:602–615.
- Dimov, C. M. 2016. Connections between act-rs declarative memory system and minerva2. In Papafragou, A.; Grodner, D.; Mirman, D.; and Trueswell, J. C., eds., *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, 492–495. Austin, TX: Cognitive Science Society.
- Eich, J. M. 1982. A composite holographic associative recall model. *Psychological Review* 89:627–661.
- Eliasmith, C. 2013. *How to build a brain: A neural architecture for biological cognition*. New York, NY: Oxford University Press.
- Farrell, S., and Lewandowsky, S. 2002. An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin & Review* 9:59–79.

- Franklin, D. R. J., and Mewhort, D. J. K. 2015. Memory as a hologram: An analysis of learning and recall. *Canadian Journal of Experimental Psychology* 69:115–135.
- Gayler, R. W. 2003. Vector symbolic architectures answer jackendoff's challenges for cognitive neuroscience. In Slezak, P., ed., *Proceedings of the Joint International Conference on Cognitive Science*. Sydney, Australia: University of New South Wales. 133–138.
- Greve, A.; Cooper, E.; Kaula, A.; Anderson, M. C.; and Henson, R. 2017. Does prediction error drive one-shot declarative learning? *Journal of Memory and Language* 94:149 – 165.
- Hintzman, D. L. 1984. Minerva 2: A simulation model of human memory. *Behavior Research Methods, Instruments, and Computers* 16:96–101.
- Hintzman, D. L. 1986. "schema abstraction" in multiple-trace memory models. *Psychological Review* 93:411–428.
- Hintzman, D. L. 1990. Human learning and memory: Connections and dissociations. *Annual Review of Psychology* 41:109–139.
- Humphreys, M. S.; Pike, R.; Bain, J. D.; and Tehan, G. 1989. Global matching: A comparison of the sam, minerva ii, matrix, and todam models. *Journal of Mathematical Psychology* 33:36–67.
- Jamieson, R. K., and Mewhort, D. J. K. 2011. Grammaticality is inferred from global similarity: A reply to kinder (2010). *The Quarterly Journal of Experimental Psychology* 64:209–216.
- Jamieson, R. K.; Crump, M. J. C.; and Hannah, S. D. 2012. An instance theory of associative learning. *Learning & Behavior* 40:61–82.
- Johns, B. T.; Jones, M. N.; and Mewhort, D. J. K. 2012. A synchronization account of false recognition. *Cognitive psychology* 65:486–518.
- Jones, M. N., and Mewhort, D. J. K. 2007. Representing word meaning and order information in a composite holographic lexicon. *Psychological Review* 114:1–37.
- Kahana, M. J.; Rizzuto, D. S.; and Schneider, A. R. 2005. Theoretical correlations and measured correlations: Relating recognition and recall in four distributed memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31:933–953.
- Kelly, M. A., and West, R. L. 2012. From vectors to symbols to cognition: The symbolic and sub-symbolic aspects of vector-symbolic cognitive models. In M., N.; Peebles, D.; and Cooper, R. P., eds., *Proceedings of the 34th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society. 1768–1773.
- Kelly, M. A.; Blostein, D.; and Mewhort, D. J. K. 2013. Encoding structure in holographic reduced representations. *Canadian Journal of Experimental Psychology* 67:79–93.
- Kelly, M. A.; Kwok, K.; and West, R. L. 2015. Holographic declarative memory and the fan effect: A test case for a new memory model for act-r. In Taatgen, N. A.; van Vugt, M. K.; Borst, J. P.; and Mehlhorn, K., eds., *Proceedings of the 13th International Conference on Cognitive Modeling*. Groningen, the Netherlands: University of Groningen. 148–153.
- Kelly, M. A.; Mewhort, D. J. K.; and West, R. L. 2017. The memory tesseract: Mathematical equivalence between composite and separate storage memory models. *Journal of Mathematical Psychology* 77:142–155.
- Kelly, M. A.; Reitter, D.; and West, R. L. 2017. Degrees of separation in semantic and syntactic relationships. In *Proceedings of the 15th International Conference on Cognitive Modeling*. Groningen, the Netherlands: University of Warwick.
- Kintsch, W. 1970. Models for free recall and recognition. In Norman, D. A., ed., *Models of human memory*. New York: Academic Press.
- Lieto, A.; Chella, A.; and Frixione, M. 2017. Conceptual spaces for cognitive architectures: A lingua franca for different levels of representation. *Biologically Inspired Cognitive Architectures* 19:1–9.
- Locke, J. 1690. *An essay concerning human understanding*.
- Mewhort, D. J. K., and Johns, E. E. 2005. Sharpening the echo: An iterative-resonance model for short-term recognition memory. *Memory* 13:300–307.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems* 26. Curran Associates, Inc. 3111–3119.
- Murdock, B. B. 1993. Todam2: a model for the storage and retrieval of item, associative and serial-order information. *Psychological Review* 100:183–203.
- Plate, T. A. 1995. Holographic reduced representations. *IEEE Transactions on Neural Networks* 6:623–641.
- Raaijmakers, J. G. W., and Shiffrin, R. M. 1981. Search of associative memory. *Psychological Review* 88:93–134.
- Rutledge-Taylor, M. F.; Kelly, M. A.; West, R. L.; and Pyke, A. A. 2014. Dynamically structured holographic memory. *Biologically Inspired Cognitive Architectures* 9:9–32.
- Rutledge-Taylor, M. F.; Vellino, A.; and West, R. L. 2008. A holographic associative memory recommender system. In *Proceedings of the 3rd International Conference on Digital Information Management*, 87–92.
- Shiffrin, R. M., and Steyvers, M. 1997. A model for recognition memory: Remretrieving effectively from memory. *Psychonomic Bulletin & Review* 4:145–166.
- Simon, H. A. 1969. The psychology of thinking: Embedding artifice in nature. In *Sciences of the artificial*. MIT Press. 23–26.
- Tulving, E. 1976. Ecphoric processes in recall and recognition. In Brown, J., ed., *Recall and recognition*. London: Wiley.