

Lessons from Mapping Sigma onto the Standard Model of the Mind: Self-Monitoring, Memory/Learning, and Symbols

Paul S. Rosenbloom

Institute for Creative Technologies and Department of Computer Science, University of Southern California,
12015 Waterfront Dr., Playa Vista, CA 90094
rosenbloom@usc.edu

Abstract

Sigma was one of the three architectures explicitly factored into the recent development of the *standard model of the mind*. Here we dig deeper into the mapping of Sigma onto the standard model begun there to explore three lessons that illustrate outstanding “issues” with the current standard model while providing food for thought for its future development.

Introduction

The *standard model of the mind* is an attempt to develop a consensus across the international research community concerning the structure and functioning of human-like minds (Laird, Lebiere & Rosenbloom, 2017). As opposed to a *cognitive architecture* (Langley, Laird & Rogers, 2009), which is a concrete implementable model of the fixed structure that defines a mind, the standard model is an abstraction over cognitive architectures that is to include only those aspects for which a consensus exists. The origins of the standard model were in a summary presentation at the AAAI 2013 Fall Symposium on *Integrated Cognition*, but this then inspired a more detailed dialectic among, and attempt to synthesize across, three particular cognitive architectures – ACT-R (Anderson *et al.*, 2004), Sigma (Rosenbloom, Demski & Ustun, 2016a) and Soar (Laird, 2012) – in an attempt to build further on the initial consensus.

While taking a number of important steps, the standard model that resulted (1) was admittedly incomplete, (2) designedly reflected a consensus rather than unanimity, and (3) included initial discussion of topics that ultimately need more depth and clarity. Although each of these “issues” reflects appropriate choices in the development of the standard model, rather than mistakes per se, each also reflects a concern worth further attention. Incompleteness provides an opportunity, and in fact a necessity, of considering what else should be in the standard model. Consen-

sus raises the specter that the minority rather than the majority will ultimately be proven correct. Initial discussions of complex/subtle topics demand more detailed follow up.

Here, after a brief introduction to Sigma, three selected lessons from digging deeper into the mapping of Sigma onto the standard model – concerned with (1) self-monitoring, (2) memory/learning, and (3) symbols – are introduced, each of which illustrates the correspondingly numbered “issue” while hopefully also providing useful food for thought for the future development of the standard model.

Sigma

Sigma is a cognitive architecture that has been under development since 2008 at the USC Institute for Creative Technologies. Research on it is driven by four desiderata: *grand unification*, with the goal of incorporating not only all of the requisite cognitive capabilities but also the key sub-cognitive ones, such as perception, motor control, and affect; *generic cognition*, with the goal of supporting both useful artificial systems – starting with virtual humans, but ultimately also to include intelligent agents and robots – and models of human thought (at some appropriate level of abstraction); *functional elegance*, with the goal of providing the diversity of requisite intelligent capability and behavior from the interactions among small numbers of very general mechanisms, rather than from a large number of more specialized modules; and *sufficient efficiency*, with the goal of executing fast enough for both useful artificial systems and large-scale cognitive models.

Sigma’s overall approach to achieving these desiderata is based on the *graphical architecture hypothesis*, that the best approach at this time involves blending what has been learned from over three decades of independent work on cognitive architectures, such as ACT-R and Soar, and *graphical models* (Koller & Friedman, 2009). Graphical

models are a general technology for computing over complex multivariate functions by leveraging forms of independence to decompose them into products of factors representing simpler functions, mapping this decomposition onto graphs, and solving them for values of the function's variables, typically either via some form of sampling or message passing. They potentially support grand unification by enabling *mixed* (symbolic + probabilistic) *hybrid* (discrete + continuous) representation and processing; functional elegance by doing this via a general yet simple and theoretically elegant base; and sufficient efficiency by subsuming state-of-the-art algorithms across this span.

Sigma approaches this blending by stratifying the overall architecture into two layers – the *graphical architecture* and the *cognitive architecture* – with the former sitting below the latter and implementing it. These two architectures combine to instantiate a *cognitive cycle* that involves a single round of input, perception, memory access, reasoning, decision-making, learning, affect, attention, and output, and which is ultimately intended to run at ~50 ms. to enable real-time systems and models at human time scales. As in Soar, reflection occurs in Sigma when an *impasse* – i.e., the inability to select the next action to execute – occurs in decision-making. How the graphical architecture, through graph solution and modification, implements the cognitive architecture provides one of the major sources of functional elegance in Sigma. However, another is how the cognitive architecture then serves to implement a broad range of cognitive capabilities above it in terms of particular forms of skills and knowledge.

Self-Monitoring

One key aspect missing from the current standard model is *architectural self-monitoring* – how a cognitive architecture monitors its own status and makes the results available for architectural and/or cognitive (i.e., knowledge/skills-driven) response. Functionally, this is critical for robust autonomy. Psychologically, it is critical for affective appraisal (Moors et al., 2013) and quite possibly other cognitive capabilities. It may be conceived of as a distinct module on its own, or be implemented in a distributed fashion.

Robust autonomy in general requires reflection, comprising monitoring and correcting of both the architecture and the knowledge and skills implemented on top of it. The processing involved in monitoring and correcting can likewise itself be due to the architecture, to knowledge and skills, or to some mixture of the two. Here, however, the focus is specifically on monitoring of the architecture by the architecture. This is not intended to deny the importance of the other aspects, but simply to focus on the aspect that has the clearest implications at this time for the standard model. Achieving consensus on the extent to

which this broader notion of reflection should be incorporated, and how it should be done, remains for future work.

In Sigma, architectural self-monitoring has primarily been investigated in the context of affective appraisal; that is, in the initial stage of emotional processing that captures emotionally and behaviorally relevant assessments of situations in terms of a small set of variables, such as *relevance*, *desirability*, *likelihood*, *expectedness*, *causal attribution*, *controllability* and *changeability* (Marsella & Gratch, 2009). Appraisal can be fast and automatic, or slow and deliberate. Knowledge and reasoning are clearly implicated in the latter, and may conceivably also play a role in the former. However, in either case there must be an architectural aspect of the appraisal process if emotions are to modify how the architecture supports thought and behavior, whether via parameter setting or other means; and even more fundamentally to connect it to other essential aspects of emotion, such as the physiological ones that contribute to making emotions “hot”.

Sigma currently supports several appraisals via architectural self-monitoring. For example, *desirability* automatically compares goals with states to estimate how close the states are to the goals, and *unexpectedness/surprise* automatically compares distributions in long-term memory from just before and just after learning to estimate how poorly the current situation was anticipated by what was already known (Rosenbloom, Gratch & Ustun, 2015). More recent work in Sigma has expanded this to additional appraisals as well, including a measure of *correctness* that has proven central to implementing learning by backpropagation (Rosenbloom, Demski & Ustun, 2017).

Beyond appraisal, there are also other critical forms of architectural self-monitoring, such as how both Soar and Sigma monitor the decision-making process to determine when an impasse has occurred and thus when more general knowledge-driven reflective processing is necessary. Perhaps such impasse detection should simply be considered as yet another form of appraisal. But, either way, it is a potentially critical form of architectural self-monitoring.

The central question then raised by this first lesson is whether the standard model should be extended to include some form of architectural self-monitoring, in support of appraisal and robust autonomy.

Memory/Learning

As revealed by the initial mapping of Sigma onto the standard model in (Laird, Lebiere & Rosenbloom, 2017), the biggest disconnect at present between the two is in how they view memory and learning. A part of this is simply that Sigma is behind the other two more mature architectures in implementing some key cognitive capabilities. For example, it does not yet include a mechanism for acquiring

new procedural knowledge via procedural composition, and thus does not yet have a means for learning all of its forms of long-term memory content; and it has not yet been sufficiently optimized so as to run within ~50 ms. per cognitive cycle on complex tasks, as is required for real-time human performance. However, such shortcomings are all areas of active endeavor in Sigma.

What is of broader interest though is the disagreement between Sigma and the standard model over how much of the diversity of memory and learning capabilities that characterizes the standard model corresponds to architectural modules versus a combination of a few smaller yet more general architectural mechanisms plus appropriate knowledge/skills on top of the architecture. As mentioned earlier, one of the four desiderata driving the development of Sigma is functional elegance, where the wide range of intelligent capabilities are to arise out of the interactions among a small number of very general mechanisms. Scientifically, functional elegance can be viewed as a form of Occam's razor, favoring simpler explanations for the same phenomena, but it also seeks what Deutsch (2011) characterized as *explanations with reach*; that is, models – or, in this case, mechanisms – that are general enough to explain a broad range of phenomena. This is what underlies the power of, for example, both Newton's laws of motion and Maxwell's equations of electromagnetism in physics, and also the periodic table of the elements in chemistry.

The boundaries of the functional elegance achievable in any scientific endeavor are determined by a combination of the generality of the laws/mechanisms that can be invented/developed by the scientist/modeler and the constraints imposed by the data to be modeled. Cognitive scientists often follow the biological notion that evolution is a tinkerer that more often than not yields an arbitrarily messy outcome; yet, even in biology, evolution itself provides an explanation with enormous reach.

So, what is the story with respect to functionally elegant memory and learning in Sigma, and what implications might this have for the standard model? With respect to memory, Sigma's cognitive architecture distinguishes only between long-term and working memory; and even this distinction disappears in the graphical architecture, where both merely comprise distinct regions of a single unified graph (along with further distinct regions for perceptual memories, which don't show up explicitly in the cognitive architecture). The distinction between declarative and procedural memories that is found in the standard model arises in Sigma from using different combinations of *microvariations* within the cognitive (and graphical) architecture plus different forms of knowledge structures on top of the cognitive architecture (Rosenbloom, 2010).

Additional microvariations and knowledge differences then also yield a distinction between semantic and episodic declarative long-term memories that is not (yet?) found in

the standard model. Along with Sigma's hybrid nature, which provides continuous as well as discrete representations, this overall approach in addition enables distinct regions for imagery long-term (and working) memory (Rosenbloom, 2011a). With a mixture of microvariations and forms of knowledge/skill it further enables implementing long-term memories that blend traditional procedural and declarative varieties, such as the *trellis graphs* that underlie HMMs for perception and POMDPs for action (Rosenbloom, Demski & Ustun, 2016a).

With respect to learning, there is a similar story. Sigma's single native learning mechanism is a form of *gradient descent* on the functions stored at factor nodes in the graphical architecture (driven by the messages sent during graph solution). Yet this has proven sufficient, given the appropriate knowledge and skills, to implement classification, categorization, reinforcement and inverse reinforcement learning, episodic learning, perception and map learning, and action modeling/learning (Rosenbloom *et al.*, 2013; Rosenbloom, Demski & Ustun, 2016a).

This is possible because of a combination of the particular decision and learning mechanisms that operate in parallel across Sigma's graph each cycle plus the expressivity of the parallel language for knowledge and skills that the architecture provides (Rosenbloom, 2015). In essence, Sigma produces these memory and learning behaviors from a combination of the cognitive architecture and what can be called *reactive memory and learning idioms* that can operate in parallel during a single cycle. Together these enable exhibiting multiple forms of memory and learning simultaneously within a single cycle, making it challenging to distinguish this approach behaviorally from the standard model's approach of a purely architectural implementation that is composed of multiple distinct memory and learning modules that operate in parallel.

One major implication of this is that it remains open to debate and to further evidence whether the standard model's distinct procedural and declarative memories, along with their associated learning mechanisms, must be based on distinct architectural modules. There are, however, two qualifiers that need to be kept in mind with respect to this claim. First, as already discussed, Sigma does not at present exhibit all of the requisite forms of learning. Second, the knowledge/skills that enable its existing diverse forms of memory and learning must ultimately either be learnable or be part of the system's innate knowledge. If the latter, it becomes a rather fine distinction as to whether the requisite knowledge/skills should themselves just be considered part of the defining structure of an architectural module.

A second, less obvious, implication is that the memory and learning capabilities characteristic of human-like cognition may be more diverse and flexible than is stated in the standard model. This might involve blended forms of memory, such as the HMMs used for speech understanding

in Sigma (Joshi, Rosenbloom & Ustun, 2014), or it could involve broader notions of what may be contained within the current memories. For example, procedural memory is at its heart about action and control. The standard model describes this as typically occurring via rules that provide pattern-directed action invocation – as is in both ACT-R and Soar – but Sigma can combine actions not only with traditional rule conditions but also with other forms of reactive-and-contextual conditional processing that still is abstractly best described as pattern directed. Neural networks, and bounded forms of both POMDPs and analogy, for example, can fit this bill. All three of these have been demonstrated in Sigma (Rosenbloom, Demski & Ustun, 2016b; Chen *et al.*, 2011; Ustun *et al.*, 2014), but the lesson here is clearly not just limited to Sigma and its mapping onto the standard model, necessarily applying also to other architectures with similar capabilities.

Symbols

The standard model backed off from a strict interpretation of the classical *physical symbol systems hypothesis*, that “A physical symbol system has the necessary and sufficient means for general intelligent action.” (Newell & Simon 1976), by acknowledging that although a symbol system is logically sufficient for intelligence, it is insufficient as a basis for building a cognitive architecture that is to have a human-like cycle time of ~50 ms. In other words, although numeric processing can be constructed on top of a pure symbol system, as is done all of the time with digital computers, human-like cognition requires numeric processing be available with symbolic processing in the architecture.

The numeric data used in the architecture may take the form of *quantitative task information*, whether representing simple task quantities, such as height or weight, or more complex multi-dimensional task quantities, such as spatial extent. Or it may take the form of *quantitative metadata* that represents values such as probabilities, utilities, frequencies, and activations over symbolic data in the system.

The standard model also backed off, at least for now, on the full set of requirements assumed for traditional symbol systems, focusing instead more narrowly on just the ability to compose symbols into arbitrarily complex structures, while glossing over issues of designation, grounding, and the lack of symbol substructure.

Sigma embodies a particularly broad perspective on symbol systems and their relationship to quantitative data that, although it already has partially influenced the formulation of the standard model’s approach to symbols, is worth getting out in its full form to understand how the initial discussion in the standard model paper could eventually be fleshed out. A number of these points have already been discussed and debated in one form or another in

the literature (see, e.g., Touretzky & Pomerleau, 1994; Vera & Simon, 1994; Nilsson, 2007), but for simplicity we will start afresh here.

Let’s start with the notion of *primitive elements* with (at least) one quantity associated with each. In the standard model, these primitive elements are referred to as symbols, but here we will hold off on such labeling for now. Primitive elements can be classical substructure-free symbols (with a numeric value of 0 denoting absence and 1 presence); however, they can also be, for example, such things as the units of neural networks, with numeric values denoting their activation levels (Rosenbloom, Demski & Ustun, 2016b). There is no assumption in Sigma that primitive elements must *represent*, in the sense of designating something else. In other words, these are syntactic elements that may or may not have a clear semantics. Their meaning arises from their use, implying a form of *procedural or inferential role semantics*, the only form that really appears to make sense for an evolving/developing/learning system such as a cognitive architecture.

Next, the inclusion of *relations* over primitive elements enables combining them into (possibly complex) *structures* and yields *composite elements* that represent instances of these relations. Quantities associated with such composite elements can denote, for example, their truth-values, or their probabilities or frequencies (Rosenbloom, 2011b). These relations can be fully symbolic, such as *isa* or *member*; however, as with primitive elements, relations are in their essence syntactic structures, with any meaning deriving from their use. For example, the links in neural networks are specified in Sigma via instances of a binary relation between pairs of units in adjacent layers, with their associated quantities denoting the weights on the links.

Allowing *organizations* among the primitive elements expands the overall expressive power even further. Sigma supports three different forms of organization. The first form involves simple *grouping*, which enables multiple individual primitive elements to be treated as if they were one *compound element*; for example, representing sets of symbols or segments of the real number line.

The second form is an *ordering* among a set of primitive elements to define *vectors* over them, with associated vectors of quantities, that can for example represent the distributed vectors and word/concept embeddings that play such important roles at this point in machine learning (see, e.g., Mikolov *et al.*, 2013; Mnih & Kavukcuoglu, 2013) and cognitive science (see, e.g., Gayler, 2003; Jones & Mewhort, 2007). Such capabilities have been demonstrated in Sigma (Ustun *et al.*, 2014); and, when joined with the combinatorics of relations, can also yield tensors, as in (Smolensky & Legendre, 2006).

The third form of organization involves *metric* aspects, which enable representing, for example, 2D regions of visual images and 3D regions of physical space. A metric

organization is central, for example, to defining the long-term and working memories mentioned earlier for mental imagery. They may also potentially represent such things as metric organizations among the units of a single layer of a neural network, should such organizations prove useful.

Both relations and organizations provide structures over multiple primitive elements, but they are quite different in nature. The latter provides an implicit structuring across a set of elements that may add to the architecturally interpretable semantics across them, whether this involves word embeddings or the representation of mental images. In contrast, the former induces a cross product of two or more such sets, creating new elements for their combinations, and thus providing both the combinatorics required in symbol systems and a solution to the binding problem (Treisman, 1996) in cognitive science and neural networks that aligns well with the tensor approach mentioned above.

The standard model refers to primitive elements as symbols, while remaining agnostic as to the exact nature of these symbols. It also includes relations over primitive elements that generate composite structures, plus quantitative metadata over both primitive elements and composite structures. In Sigma, only a subset of the primitive elements – those without organization – are normally considered symbols, as they come closest to the traditional notion of bare substructure-free elements that can be used in arbitrary ways. However, Sigma does include all of the capabilities discussed in this section in some form. In the cognitive architecture, this involves *types* that determine the forms of organization that exist, and which implement the hybrid nature mentioned earlier, plus *predicates* over typed arguments that define relations among elements. These all compile uniformly into *piecewise linear functions* in the graphical architecture (Rosenbloom, 2011b), with the linearity stemming from going beyond assigning individual numbers to elements to assigning linear functions over groupings/segments of metric elements.

Let's now return to the other aspects of symbol systems mentioned earlier. As discussed above, *designation* in Sigma is a function of use; in particular, via *conditionals* in the cognitive architecture that provide a deep blending of concepts from rule systems and probabilistic networks, and which compile down to graphical models in the graphical architecture. In this manner, meaning can be assigned to individual elements, to composite elements that are instances of relations, and to organizations over elements; implying that meaning may accrue to individual elements – whether primitive or composite and whether discrete or continuous – as well as to both patterns of numbers over vectors of elements and the multidimensional combinations of metric organizations found in 2D and 3D environments.

The full determination of meaning in Sigma involves internal designation, of the form just discussed, plus the use, or *grounding*, of elements in interactions with the external

world. The long-term goal with Sigma is to incorporate the complete arc from perception through cognition to motor control within the architecture, ensuring full grounding of symbols. One partial such exploration included a connection, for example, from spectral labels to word recognition via conditionals (Joshi, Rosenbloom & Ustun, 2014).

With respect to the *lack of symbol substructure*, everything depends on what is called a symbol, which itself is more a matter of definition than of science, as discussed for example in (Vera & Simon, 1994) and elsewhere. Is it just primitive elements, or are composite elements and/or patterns of element organizations also included? If only primitive elements, then they indeed have no substructure; however, either of the latter involve significant substructure, with composites being of the type discussed by Vera & Simon and organizations being of the types found in neural networks, distributed vectors, vector symbolic architectures, and mental imagery.

One hope for the future of the standard model is that more detailed discussions such as the one here can help it move beyond agnosticism with respect to the form of the symbols that participate in relations – which was itself a step forward from just bare elements – to a more systematic understanding of the space of (necessary?) possibilities.

A second hope is that this can break down a barrier between how the cognitive system is characterized versus the perception and motor components (and any imagery component that may be coming). Rather than thinking of the former as symbols (defined as those elements over which relations can be defined) plus symbol structures and quantitative metadata, versus the latter as being purely numeric/subsymbolic, both can be seen at a minimum as comprising primitive elements with quantitative metadata.

Do relations and organizations then also make sense over both symbolic and subsymbolic elements? The most obvious answer would be that relations exist (just) in symbolic cognition and organizations (just) in the subsymbolic periphery; however, this need not be the case. For example, to the extent that either word embeddings or mental imagery are centrally implicated in cognition, then some form of organization is required there as well. Similarly, if the dimensional cross-product central to perception and motor control arises from relations, then they are required in the periphery as well. Thus, it may be that it all is primitive elements with relations, organizations and metadata.

A third hope is that by viewing things in this manner some of the confusion about the different possible applications of the term “symbol” in the context of the standard model and more broadly may be alleviated. Combinatorics applies to elements, via relations. Substructure arises from organizations (and relations). Quantities exist as metadata and metric organizations. Designation and grounding arise from how all of this is (learned to be) used, including usage in perception and motor control.

Conclusion

Digging deeper into the mapping from Sigma onto the standard model yields a number of lessons that both illustrate “issues” with the standard model and hopefully provide useful food for thought in its further development. In this article three issues and illustrative lessons have been explored, but more remains to be mined, such as extending from architectural self-monitoring to full reflection. Such mining across the full panoply of today’s leading architectures, followed up with interactions across the broader international research community concerning what is revealed, would appear to be one of the essential foundations of the continued development of the community consensus that must form the core of the standard model.

Acknowledgements

The work described in this article was sponsored by the U.S. Army and the Office of Naval Research. Statements and opinions expressed may not reflect the position or policy of the United States Government, and no official endorsement should be inferred.

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C. & Qi, Y. 2004. An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
- Chen, J., Demski, A., Han, T., Morency, L-P., Pynadath, D., Rafidi, N. & Rosenbloom, P. S. (2011). Fusing symbolic and decision-theoretic problem solving + perception in a graphical cognitive architecture. *Proc. of the 2nd International Conference on Biologically Inspired Cognitive Architectures* (pp. 64-72).
- Deutsch, D. 2011. *The Beginning of Infinity: Explanations that Transform the World*. London, UK: Penguin Books.
- Gayler, R. W. (2003) Vector symbolic architectures answer Jackendoff’s challenges for cognitive neuroscience. *Proc. of the ICCS/ASCS Intl. Conf. on Cognitive Science* (pp. 133-138).
- Jones, M. N., & Mewhort, D. J. 2007. Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, 114, 1.
- Joshi, H., Rosenbloom, P. S. & Ustun, V. 2014. Isolated word recognition in the Sigma cognitive architecture. *Biologically Inspired Cognitive Architectures*, 10, 1-9.
- Koller, D. & Friedman, N. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge.
- Laird, J. E. 2012. *The Soar Cognitive Architecture*. Cambridge, MA: MIT Press.
- Laird, J. E., Lebiere, C. & Rosenbloom, P. S. 2017. A Standard Model for the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*. In press.
- Langley, P., Laird, J. E. & Rogers, S. 2009. Cognitive architectures: Research issues & challenges. *Cog. Sys. Res.*, 10, 141-160.
- Marsella, S. & Gratch, J. 2009. EMA: A Process Model of Appraisal Dynamics. *Cognitive Systems Research*, 10, 70-90.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. 2013. Efficient estimation of word representations in vector space. In *Proc. of the International Conference on Learning Representations*.
- Mnih, A., & Kavukcuoglu, K. 2013. Learning word embeddings efficiently with noise-contrastive estimation. In *Advances in Neural Information Processing Systems* (pp. 2265-2273).
- Moors, A., Ellsworth, P. C., Scherer, K. R. & Frijda, N. H. 2013. Appraisal theories of emotion: State of the art and future development, *Emotion Review*, 5, 119-124.
- Newell, A. & Simon, H. A. 1976. Computer Science as Empirical Inquiry: Symbols and Search. *Comms. of the ACM*, 19, 113-126.
- Nilsson, N. 2007. The physical symbol system hypothesis: Status and prospects. In M. Lungarella (Ed.) *50 Years of AI*, Lecture Notes in Computer Science, 4850 (pp. 9-17). Berlin: Springer.
- Rosenbloom, P. S. 2010. Combining procedural and declarative knowledge in a graphical architecture. In D. D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 205-210).
- Rosenbloom, P. S. 2011a. Mental imagery in a graphical cognitive architecture. *Proc. of the 2nd International Conference on Biologically Inspired Cognitive Architectures* (pp. 314-323).
- Rosenbloom, P. S. 2011b. Bridging dichotomies in cognitive architectures for virtual humans. *Proceedings of the AAAI Fall Symposium on Advances in Cognitive Systems*.
- Rosenbloom, P. S. (2015). Supraarchitectural capability integration: From Soar to Sigma. *Proc. of the 13th International Conference on Cognitive Modeling* (pp. 67-68). Extended abstract.
- Rosenbloom, P. S., Demski, A., Han, T. & Ustun, V. (2013). Learning via gradient descent in Sigma. *Proceedings of the 12th International Conference on Cognitive Modeling* (pp. 35-40).
- Rosenbloom, P. S., Demski, A. & Ustun, V. 2016a. The Sigma cognitive architecture and system: Towards functionally elegant grand unification. *Journal of Artificial General Intel.*, 7, 1-103.
- Rosenbloom, P. S., Demski, A. & Ustun, V. 2016b. Rethinking Sigma’s graphical architecture: An extension to neural networks. *Proc. of the 9th Conf. on Artificial General Intell.* (pp. 84-94).
- Rosenbloom, P. S., Demski, A. & Ustun, V. (2017). Toward a neural-symbolic Sigma: Introducing neural network learning. *Proceedings of the 15th Annual Meeting of the International Conference on Cognitive Modeling*. In press.
- Rosenbloom, P. S., Gratch, J. & Ustun, V. (2015). Towards emotion in Sigma: From Appraisal to Attention. *Proceedings of the 8th Conference on Artificial General Intelligence* (pp. 142-151).
- Smolensky, P. & Legendre, G. 2006. *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*. Cambridge, MA: The MIT Press.
- Touretzky, D. S., & Pomerleau, D. A. 1994. Reconstructing physical symbol systems. *Cognitive Science*, 18, 345-353.
- Treisman, A. 1996. The Binding Problem. *Current Opinion in Neurobiology*, 6(2), 171-178.
- Ustun, V., Rosenbloom, P. S., Sagae, K. & Demski, A. 2014. Distributed vector representations of words in the Sigma cognitive architecture. *Proceedings of the 7th Annual Conference on Artificial General Intelligence* (pp. 196-207).
- Vera, A. H. & Simon, H. A. 1994. Reply to Touretzky and Pomerleau: Reconstructing physical symbol systems. *Cognitive Science*, 18, 355-360.