

Multiple Representations in Cognitive Architectures

David Peebles

Department of Psychology
University of Huddersfield
Queensgate, Huddersfield,
HD1 3DH, UK

Peter C-H. Cheng

Department of Informatics,
University of Sussex,
Falmer, Brighton,
BN1 9QH, UK

Abstract

The widely demonstrated ability of humans to deal with multiple representations of information has a number of important implications for a proposed *standard model of the mind* (SMM). In this paper we outline four and argue that a SMM must incorporate (a) multiple representational formats and (b) meta-cognitive processes that operate on them. We then describe current approaches to extend cognitive architectures with visual-spatial representations, in part to illustrate the limitations of current architectures in relation to the implications we raise but also to identify the basis upon which a consensus about the nature of these additional representations can be agreed. We believe that addressing these implications and outlining a specification for multiple representations should be a key goal for those seeking to develop a standard model of the mind.

Implications from a sample task

Cognitive architectures are operational models of the components and functions of minds (Anderson 2007; Newell 1990). The *standard model of the mind* (SMM) is intended to be a specification of those components and functions that are generally accepted as essential to all flavours of cognitive architecture (Laird, Lebiere, and Rosenbloom in press). The structure and content of the SMM is the subject of debate to which we contribute by exploring the implications for the SMM of the human use of multiple representations. We do this in two halves. The first considers implications for the SMM following from considerations of the use of multiple representations while the second considers the current approaches in cognitive architecture to modelling visuo-spatial and multiple representations.

Consider an example: Imagine a square with sides of one unit. At opposite corners of the square add circles with radii of $2/3$ of a unit centred on the corners. Do the two circles (a) overlap, (b) just touch, or (c) not touch? There are two common approaches that we have observed people use to solve this problem. The first is mathematical and an idealised verbal protocol might go something like this:

“The length of the diagonal between two opposite corners is the square root of 2—about 1.414 units—using Pythagoras’ theorem, so the centre of the square is about 0.707 units from each corner. The radius of each

circles is about 0.667 units, so neither circle’s perimeter will reach the centre of the square, and thus they do not touch”.

The second approach is to use mental imagery and might be accompanied by a verbal report something like this:

“I’m imagining a square and I’ve picked the top-left and bottom-right corners. I can see that circles with unit radius will definitely overlap; in fact they intersect each other at the other corners of the square. Now I’m imagining circles of $1/2$ unit and I can see that they clearly don’t meet, in fact the circles cross the mid point of each side of the square and curve away from the centre. I’m thinking of circles with radii of $2/3$, it’s hard to be certain how big they should be, but they seem to... overlap each other”.

“...just touch each other” or “...be separate” are also likely responses. What implications for the SMM then follow from this example, in relation to the use of multiple representations? Let us consider four.

First, the two approaches rely upon quite different representations for their solution—one declarative and mathematical, the other visuo-spatial and exploiting imagery in the mind’s eye. Obviously, cognitive architectures must encompass the use of alternative problem representations (Markman 1998; Tabachneck-Schijf and Simon 1996), but more importantly, our use of multiple representations supports the idea that the SMM should explicitly include multiple sub-modules for different modalities within each of the main SMM declarative, procedural and working memory modules. For a given task the format of information, operators, methods of indexing information, heuristics and goal structures of a task can differ considerably with alternative representations (Cheng, Lowe, and Scaife 2001; Simon and Kaplan 1989; Larkin and Simon 1987).

Without the presence of sub-modules for different modalities, cognitive models will be unrealistically homogeneous parodies. For example, the use of imagery-based reasoning is so imprecise that reasoners are likely to give the wrong answer, but this characteristic of mental imagery is absent from the proposition-mathematical representation. Markman’s detailed exposition on knowledge representation (Markman 1998, summarised in Markman 2012) usefully identifies three classes of representation used by humans,

which we believe a SMM should also encompass: *mental spaces*, including generic spatial representations (Markman 2012), geometric conceptual spaces (Gärdenfors 2000), and variable spaces (Halford et al. 2005); *featural representations* in which discrete characteristic properties distinguish objects; and *structured representations* grounded in networks that for example include semantic networks, logic-based representations, scripts and schemas.

Second, some people are able to solve the sample problems using either approach (and even other methods), so another implication for the standard model is how it deals with multi-representational phenomena. These include: (a) the choice of an initial task representation when an individual plans to use just one of several representations that they know, for example starting with imagery; (b) switching representation during problem solving, which can happen for diverse reasons, for example, we may be uncertain of the answer because we cannot precisely imagine the circles, or we may not have the skill or knowledge to complete a solution (e.g., perhaps we don't remember the square root of two); (c) simultaneously using two representations that would in other circumstance be used alone.

The existence of different patterns of usage of multiple representations raises the question of whether the cognitive architecture possesses some general processes to handle the selection of—and transitions between—representations, or whether such processes are dealt with by meta-level domain relations.

Third, use of multiple representations, either sequentially or simultaneously, and whether based on domain specific or general processes, will involve information beyond the domain level content of the representations. This higher order information concerns the individual's relation to the representation. For example, we each have some personal understanding of the limited precision of mental imagery and may make judgements about whether those limits make a solution too inaccurate to be relied upon, before we even start a solution attempt.

Beyond the precision of a representation, we grasp many other characteristics of representations and appear sensitive to how well suited they are to different classes of problems. So, an implication for the SSM is the status of the information underpinning such knowledge about representations. Is it generic architectural level meta-data, akin to process related frequencies or measures of similarity or utility (Laird, Lebiere, and Rosenbloom in press)? Does the information take the form of meta-relations about the domain, where the notion of what constitutes a domain includes conceptual content together with knowledge about the representation that encodes that content? Or does the information come in both forms, some meta-data and meta-relations of domain content?

Fourth, although our example starkly contrasts a mathematical-propositional representation with a visuo-spatial imagery-based representation, alternative representations can often occur within the same modality (Peebles and Cheng 2003; Zhang and Norman 1994). For example, in an alternative mathematical-propositional approach, we might test whether the area between the perimeters of the

two quarter circles plus the area of the two quarter circles is greater than that of the square (hence they overlap) by writing a formula for the distance between the circumferences of the circles and using integral calculus to find the area. This is a sledge-hammer-to-crack-a-nut approach for our sample problem, but it is suggestive of a potential range of representations that may be deployed on any given task.

The underpinning question for a SMM is whether it should recognise sub-modules for distinct classes of information that inherently differentiate alternative representations at a level below the gross distinction between modalities. Such classes of information might include:

- Types of quantity scales (nominal, ordinal, interval, ratio), which underpin distinct categorical, qualitative and quantitative forms of reasoning;
- Forms of symmetry, such a reflective, rotational and translational;
- Frames of reference, including egocentric and exocentric environmental views.

Such classes of information are near universal and implicated across diverse domains—Pinker (2011) even argues that the quantity scales underpin four fundamental classes of models in moral reasoning. So, we argue that such classes should be treated as intrinsic in cognitive architectures and thus should be included as fundamental sub-modules in a SMM.

To summarise, we contend that the example scenario we have described illustrates the widely demonstrated ability of human cognition to deal with multiple representations of information and highlights four implications for a proposed SMM:

1. A SMM must incorporate alternative problem representations and explicitly include multiple sub-modules for different modalities within each of the declarative, procedural and working memory modules.
2. A SMM must incorporate some form of metacognitive monitoring and control processes to handle the selection and monitoring of, transitions between, and integration of different representations.
3. A SMM must be able to incorporate meta-level information about the characteristics of different representational formats (e.g., level of precision afforded, ease of computation, suitability for a given problem etc.).
4. A SMM must also be able to incorporate—and be able to choose between—alternative representations within the same modality.

These four implications follow fairly directly from the sample task, but they are by no means unique claims, nor are they a comprehensive set of potential representationally related implications for cognitive architectures. Markman (1998) gives seven proposals for the use of representations in cognitive models that partially overlap the four we have identified and that also casts a wider net. One is the proposal that cognitive models must attend to social context in relation to representations, but this seems unlikely to

demand specific architectural commitments for the SMM. However, Markman’s proposal that cognitive models must use representations at multiple grain sizes lends weight to our second claim that general procedures are needed to handle multiple representations, because many tasks must integrate inferences about content at different spatial or conceptual grain sizes, but often the content may be held—and inferences made—in different representations.

Multiple representations in cognitive architectures

The nature of mental representations is a fundamental issue in cognitive science and one of the longest running debates in this area has been whether mental imagery is supported by abstract, amodal propositional representations or depictive representations grounded in perception (Kosslyn and Pomerantz 1977; Pylyshyn 1973; Anderson 1978).

A key distinction between these forms of representation is the degree to which they bear some structural correspondence to what they represent. In contrast to abstract propositional representations for example, imagistic visual representations depict rather than describe what they represent and retain the spatial relationships of their referents by having elements with geometric properties organised topographically (Reisberg 2013). Imagistic visual representations are not simply neutral depictions however but contain interpretive information that acts as an organisational reference frame describing such things as the level of detail, specific viewing position and angle, figure-ground relationships, orientation, configuration and occultation, suggesting that some degree of propositional encoding is included in imagistic visual representations (Kosslyn 1989; Pinker 1990; Reisberg 2013).

Visual representations are also *hierarchical* in that they are represented at multiple levels of abstraction (e.g., objects consisting of component parts) and *hybrid* as they consist of distinct qualitative and quantitative elements at each level (Lovett and Forbus 2017). The qualitative elements are symbolic descriptions of the relations between objects while the quantitative elements consist of measurable values (e.g., coordinate location, size, orientation etc.) associated with each object. The characteristics of these depictive representations require processes grounded in perception (e.g., various transformations, projection, comparison etc.) and visual routines (Ullman 1984) to operate on them to enable imagery-based reasoning in addition to symbolic representations and processes.

Cognitive architectures are rooted in the classical tradition of cognitive science and the *physical symbol system hypothesis* (Newell and Simon 1976) and so, until relatively recently, have relied predominantly on amodal symbolic representations. The need to represent diverse information formats in cognitive architectures and computational models has been recognised for many years however and a number of approaches to this problem have been proposed which we outline below.

Proposed depictive representations

Accepting the argument for modal depictive representations, the question then arises as to what and how information is represented and what processes operate upon them. Several proposals have been put forward to address the issue of depictive representations which can be differentiated in terms of where they lie on the representational spectrum—ranging from purely array based pixel maps and associated transformational processes to more abstract representations consisting of discrete objects with their visual and spatial properties¹.

An example of the former is the recent *affine* model of problem solving on the Raven’s Progressive Matrices test (Kunda, McGreggor, and Goel 2013) which exclusively employs representations consisting of two-dimensional arrays of grayscale pixels with operations restricted to the orthogonal unary transformations of rotation, reflection and translation and composition operations to implement image union, intersection, subtraction, back-subtraction, and exclusive-or transformations. Using these representations and processes alone, the affine model is able to solve between 55% and 63% of Standard Progressive Matrices problems. It remains unclear to what extent this relatively restricted set of representations and processes can be extended to other, more complex tasks.

At the other end of the spectrum are systems that eschew any lower-level pixel-based format and processes but represent visual information more abstractly. An influential cognitive model which incorporated such representations is Koedinger and Anderson’s (1990) *Diagram Configuration Model* (DCM). Developed to account for the performance of expert geometry problem solving, a crucial aspect of the DCM was that expert knowledge was represented as schemas combining a geometric image (e.g., *perpendicular-adjacent-angles*) with associated facts. Identifying these images in the diagram activated the schemas and allowed rapid access to the interpretive knowledge.

A similar approach has been taken to account for expert graph comprehension performance (Peebles 2013). Constructed within the ACT-R cognitive architecture, the model represents expert interpretive knowledge as a set of facts and rule-based processes triggered by recognition cues (typically distances between plot points) from the graph.

A more generally applicable example of this approach is the *Diagrammatic Representation System* (DRS) (Chandrasekaran 2006; Chandrasekaran et al. 2004; Chandrasekaran and Lele 2010) which represents diagrams as a data structure containing three basic types of object (point,

¹In our discussion we focus exclusively on systems and models designed for two-dimensional perception for diagrammatic reasoning and interacting with screen-based task environments. Developing cognitive architectures that are capable of functioning in fully three-dimensional environments as embodied agents and robots will no doubt require more complex representations and processes, for example to incorporate different frames of reference (McNamara, Rump, and Werner 2003). There are a number of efforts towards this goal using some of the cognitive architectures we discuss (Trafton et al. 2013; Kirk, Miner, and Laird 2016).

curve and region). DRS also contains a set of perceptual routines that identify object features (e.g., angle, length) and spatial relations between objects (e.g., left-of, inside-of, distance etc.) and a set of action routines that create and modify objects (e.g., translate, rotate etc.). DRS has been integrated with several existing cognitive architectures including ACT-R (Matessa and Brockett 2007), Soar (Chandrasekaran and Kurup 2007) and Polyscheme (Kurup and Cassimatis 2010).

The ACT-R cognitive architecture (Anderson 2007) has a somewhat similar (although less comprehensive) method of representing visual information through its vision module. Objects in the simulated external environment that ACT-R interacts with are represented by two symbolic *chunks* in declarative memory, one representing their features (e.g., type, shape, colour etc.), the other representing their coordinate location. These representations have been useful in allowing ACT-R models to interact with computer-based psychology experiments and simulated computer interfaces. ACT-R can also represent and manipulate (e.g., translate and rotate) visual information internally using its *imaginal* module. Unlike DRS however, these abilities, as with many other ACT-R models of spatial cognition (Dimperio, Gunzelmann, and Harris 2008; Peebles 2013) are only made possible by using ad hoc mechanisms which are not intrinsic to the architecture.

Two systems with currently the most well developed and comprehensive sets of representations in this category are the *Spatial/Visual System* (SVS) (Lathrop, Wintermute, and Laird 2011; Wintermute 2012), an extension of the Soar cognitive architecture (Laird 2012) and the sketch understanding system *CogSketch* (Forbus et al. 2011).

Until recently, SVS contained two layers of representation: a *visual depictive* layer (a bitmap array representation of space and the topological structure of objects), and a *quantitative spatial* layer (an amodal symbolic/numerical representation of objects and their spatial coordinates) but the visual depictive level has been omitted from the most recent (9.6.0) version of Soar. Now, SVS is just the quantitative spatial layer which consists of a *scene graph*, a hierarchical tree structure representation of the objects in the external task environment, each with associated three-dimensional information about its location, rotation and scaling. SVS also contains a number of operations to transform objects in the scene graph and a set of *filters* which transform the continuous information in the scene graph into symbolic information that can be used by Soar for reasoning. These processes allow Soar agents to perform mental imagery operations that can manipulate the scene graph and then extract spatial relationships from the modified states (Laird et al. 2017).

CogSketch takes as input a 2D sketch consisting of a set of separate objects, and generates a representation of the qualitative spatial relations between them, including relative position and topological relations such as containment and intersection. The generated representations form a three-level hierarchy of edges, shapes and groups, each with descriptions of the elements and relations at that level. These representations are then used by other (typically analogical) processes for visual problem solving (Lovett and Forbus 2011).

An early and influential cognitive model is notable in its combination of a pixel array based representation with more abstract representations containing information about the properties of the objects and their spatial relationships. The CaMeRa model of expert problem solving with multiple representations (Tabachneck-Schijf, Leonardo, and Simon 1997) includes a pictorial short-term memory (or ‘Mind’s Eye’) consisting of a *visual buffer* (a 150×150 pixel bitmap and parallel network) that represents the external diagram together with two additional data structures, one containing higher-level information about the shape, colour, and form of the objects in the bitmap, the other representing the locations of the objects in the array.

This combination of representations and processes creates what the authors call a “complete mental image” which allows complex visual reasoning with an external representation of a graph using both the parallel network and symbolic structures. The visual buffer can contain perceptual images from both the external diagram and mental images retrieved from pictorial long-term memory and mediates between them and the symbolic structures in short-term memory using feature extraction and Gestalt processes to identify objects in the external diagram. The hybrid nature of CaMeRa with its combined feed-forward parallel network and production system architecture and amalgamation of bitmap and symbolic representations presents a unique example of a computational model of cognition with multiple representations.

The above discussion illustrates the range of models that incorporate multiple representations and the diversity of approaches used—from exclusively array-based representations to exclusively symbolic/numerical representations, and some which have combined both. This demonstrates the widely held conviction that multiple representations in some form or another are crucial for explaining human cognition and are an essential element in any model of the human mind.

Implications for a standard model of the mind

In an influential paper, Newell (1973) proposed a strategy for overcoming the limitations he perceived in the approach to cognitive psychology at the time which has been a cornerstone of cognitive architecture research ever since. Newell advocated constructing complete processing models that could account for all aspects of one single complex task in great detail or to a wide range of different tasks. This has no doubt been a productive approach which has resulted in significant theoretical advances. We believe that the issue of alternative representations—both as internal components of the cognitive system and as external task environments—constitute an additional dimension to the strategy proposed by Newell which must be incorporated into any standard model of the mind.

A key issue in this regard concerns the nature of the representations that a SMM must incorporate in addition to the core symbolic system proposed in the initial SMM specification of Laird, Lebiere, and Rosenbloom (in press). For example, it is not clear whether mental imagery must be supported

by some form of bitmap representation depicting the topological structure of objects or whether these functions can be performed using more complex representations containing edges, lines, shapes and quantitative information such as spatial coordinates extracted from the perceptual system. In other words, is there something necessary about pixel array based representations and processes that means they have to be incorporated into a cognitive architecture or can another computationally equivalent way of representing and processing this information be used?

Although pixel arrays have a number of useful properties for spatial processing, none of the current cognitive architectures employ them but typically rely exclusively on symbolic/numerical representations. Further research should investigate the computational capabilities of different representations by modelling tasks that require multiple internal and external representations to provide behavioural evidence for which representations are being used.

The different models we have outlined above illustrate the limitations of current architectures in relation to the implications we have raised and it is clear that, despite recent developments, few of our proposed requirements are capable of being met by existing cognitive architectures. The examples also demonstrate however the growing conviction that theories of the human cognitive architecture must extend the set of representations and endow cognitive architectures with mechanisms for visual-spatial cognition and mental imagery and there is seems to be a consensus on two points—that (a) imagistic representations consist of hierarchically structured objects with their properties (e.g., shape, relative size and orientation) and the spatial relationships between them, and that (b) images and imagery-based reasoning are grounded in representations and processes that function to provide operations and routines utilising visual and spatial information which is not exclusively symbolic.

Beyond these general statements however there is little in the way of explicit agreement regarding the architectural commitments of a SMM and we believe that this should be a primary goal for those seeking to develop a standard model of the mind and a priority for future development in cognitive architecture research.

References

- Anderson, J. R. 1978. Arguments concerning representations for mental imagery. *Psychological Review* 85(4):249.
- Anderson, J. R. 2007. *How can the Human Mind Occur in the Physical Universe?* New York, NY: Oxford University Press.
- Chandrasekaran, B., and Kurup, U. 2007. A bimodal cognitive architecture: Explorations in architectural explanation of spatial reasoning. In *AAAI Spring Symposium: Control Mechanisms for Spatial Knowledge Processing in Cognitive/Intelligent Systems*, 17–22.
- Chandrasekaran, B., and Lele, O. 2010. Mapping descriptive models of graph comprehension into requirements for a computational architecture: Need for supporting imagery operations. In Goel, A. K.; Jamnik, M.; and Narayanan, N. H., eds., *Diagrammatic representation and inference*, LNAI 6170, 235–242. Springer.
- Chandrasekaran, B.; Kurup, U.; Banerjee, B.; Josephson, J. R.; and Winkler, R. 2004. An architecture for problem solving with diagrams. In Blackwell, A.; Marriott, K.; and Shimojima, A., eds., *Diagrammatic Representation and Inference*. Springer-Verlag. 235–256.
- Chandrasekaran, B. 2006. Multimodal cognitive architecture: Making perception more central to intelligent behavior. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, 1508–1512.
- Cheng, P. C.-H.; Lowe, R.; and Scaife, M. 2001. Cognitive science approaches to understanding diagrammatic representations. *Artificial Intelligence Review* 15:79–94.
- Dimperio, E.; Gunzelmann, G.; and Harris, J. 2008. An initial evaluation of a cognitive model of UAV reconnaissance. In *Proceedings of the Seventeenth Conference on Behavior Representation in Modeling and Simulation*, 165–173.
- Forbus, K.; Usher, J.; Lovett, A.; Lockwood, K.; and Wetzel, J. 2011. CogSketch: Sketch understanding for cognitive science research and for education. *Topics in Cognitive Science* 3(4):648–666.
- Gärdenfors, P. 2000. *Conceptual spaces: The geometry of thought*. A Bradford book. MIT Press.
- Halford, G. S.; Baker, R.; McCredden, J. E.; and Bain, J. D. 2005. How many variables can humans process? *Psychological Science* 16(1):70–76.
- Kirk, J.; Mininger, A.; and Laird, J. E. 2016. Learning task goals interactively with visual demonstrations. *Biologically Inspired Cognitive Architectures* 18:1–8.
- Koedinger, K. R., and Anderson, J. R. 1990. Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science* 14:511–550.
- Kosslyn, S. M., and Pomerantz, J. R. 1977. Imagery, propositions, and the form of internal representations. *Cognitive Psychology* 9(1):52–76.
- Kosslyn, S. M. 1989. Understanding charts and graphs. *Applied Cognitive Psychology* 3:185–226.
- Kunda, M.; McGregor, K.; and Goel, A. K. 2013. A computational model for solving problems from the Raven's Progressive Matrices intelligence test using iconic visual representations. *Cognitive Systems Research* 22:47–66.
- Kurup, U., and Cassimatis, N. L. 2010. Integrating constraint satisfaction and spatial reasoning. In *Proceedings of the twenty-fourth AAAI conference on artificial intelligence*. AAAI Press.
- Laird, J. E.; Congdon, C. B.; Assanie, M.; Derbinsky, N.; and Xu, J. 2017. *The Soar Users Manual Version 9.6.0*.
- Laird, J. E.; Lebiere, C.; and Rosenbloom, P. S. in press. A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *AI Magazine*.
- Laird, J. E. 2012. *The Soar Cognitive Architecture*. Cambridge, Mass: MIT Press.

- Larkin, J. H., and Simon, H. A. 1987. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science* 11:65–100.
- Lathrop, S. D.; Wintermute, S.; and Laird, J. E. 2011. Exploring the functional advantages of spatial and visual cognition from an architectural perspective. *Topics in Cognitive Science* 3(4):796–818.
- Lovett, A., and Forbus, K. 2011. Organizing and representing space for visual problem-solving. In *Proceedings of the 25th International Workshop on Qualitative Reasoning*.
- Lovett, A., and Forbus, K. 2017. Modeling visual problem solving as analogical reasoning. *Psychological Review* 124(1):60–90.
- Markman, A. B. 1998. *Knowledge Representation*. Psychology Press.
- Markman, A. B. 2012. Knowledge representation. In Holyoak, K. J., and Morrison, R. G., eds., *The Oxford handbook of Thinking and Reasoning*. Oxford University Press.
- Matessa, M., and Brockett, A. 2007. Using a diagram representation system with ACT-R. In *Proceedings of the sixteenth conference in behaviour representation in modeling and simulation (BRIMS '07)*.
- McNamara, T. P.; Rump, B.; and Werner, S. 2003. Egocentric and geocentric frames of reference in memory of large-scale space. *Psychonomic Bulletin & Review* 10(3):589–595.
- Newell, A., and Simon, H. A. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19(3):113–126.
- Newell, A. 1973. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In Chase, W. G., ed., *Visual Information Processing: Proceedings of the Eighth Annual Carnegie Symposium on Cognition*. New York: Academic Press.
- Newell, A. 1990. *Unified theories of cognition*. The William James lectures. Cambridge, Mass.: Harvard University Press.
- Peebles, D., and Cheng, P. C.-H. 2003. Modeling the effect of task and graphical representation on response latency in a graph reading task. *Human Factors* 45:28–45.
- Peebles, D. 2013. Strategy and pattern recognition in expert comprehension of 2×2 interaction graphs. *Cognitive Systems Research* 24:43–51.
- Pinker, S. 1990. A theory of graph comprehension. In Freedle, R., ed., *Artificial Intelligence and the Future of Testing*. Hillsdale, NJ: Lawrence Erlbaum Associates. 73–126.
- Pinker, S. 2011. *The better angels of our nature: The decline of violence in history and its causes*. Penguin.
- Pylyshyn, Z. W. 1973. What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80(1):1.
- Reisberg, D. 2013. Mental images. In Reisberg, D., ed., *The Oxford Handbook of Cognitive Psychology*. Oxford University Press. 374–388.
- Simon, H. A., and Kaplan, C. A. 1989. Foundations of cognitive science. In Posner, M., ed., *Foundations of cognitive science*. MIT Press.
- Tabachneck-Schijf, H. J. M., and Simon, H. A. 1996. Alternative representations of instructional material. In Peterson, D., ed., *Forms of representation*. Oxford, UK: Intellect Books. 28–46.
- Tabachneck-Schijf, H. J. M.; Leonardo, A. M.; and Simon, H. A. 1997. CaMeRa: A computational model of multiple representations. *Cognitive Science* 21:305–350.
- Trafton, G.; Hiatt, L.; Harrison, A.; Tamborello, F.; Khemlani, S.; and Schultz, A. 2013. ACT-R/E: An embodied cognitive architecture for human-robot interaction. *Journal of Human-Robot Interaction* 2(1):30–55.
- Ullman, S. 1984. Visual routines. *Cognition* 18(1–3):97–159.
- Wintermute, S. 2012. Imagery in cognitive architecture: Representation and control at multiple levels of abstraction. *Cognitive Systems Research* 19:1–29.
- Zhang, J., and Norman, D. A. 1994. Representations in distributed cognitive tasks. *Cognitive Science* 18(1):87–122.