Constructing a Standard Model: Lessons from CHREST

Fernand Gobet,[1] Peter C. R. Lane[2]
[1]Department of Psychological Sciences, University of Liverpool
[2]School of Computer Science, University of Hertfordshire
fgobet@liv.ac.uk, peter.lane@bcs.org.uk

Abstract

Although it might be too early for a standard model of the mind (SMM), comparison between current cognitive architectures is a useful exercise. This article highlights some of the likely difficulties facing the development of a SMM – both empirical and theoretical. In particular, it follows Newell (1990) by arguing that a viable model of the mind must be constructed taking advantage of experimental constraints, based on comparisons of the model with (human or animal) data. We then describe our proposed methodology for ensuring a tight link between psychological data and a cognitive architecture. We also discuss CHREST, a cognitive model with a particular emphasis on modelling psychological results. CHREST has been applied in several domains, such as language acquisition and expertise. The article concludes by highlighting some of the features that distinguish CHREST from architectures such as Soar and ACT-R. Some of these differences are significant, creating challenges for a SMM.

Cognitive Architectures and Standard Model

Understanding how the mind works is one of the central and most challenging questions in science. Numerous are the challenges that such an enterprise faces. But numerous are the benefits should this succeed.

In order to provide sufficient rigour to this endeavour, cognitive scientists have proposed to replace the kind of verbal theories typically used by psychologists and neuroscientists by cognitive architectures, where structures and mechanisms are embodied in a set of computer programs. Recently, Laird, Lebiere and Rosenbloom (in press) have proposed to develop a standard model of mind (SMM), which summarizes our knowledge of how the mind works, using the level of abstraction provided by cognitive architectures. While an SMM is not an architecture itself, it is a summary of the structures and mechanisms agreed – by consensus, and not necessarily unanimously – by the community to operate in the mind.

Laird et al. (in press) focus on *human-like* minds, accepting architectures that function in ways similar to humans, but not exactly so, thus including AI and robotics systems that are not meant to be theories of human cognition, but that might share important commonalities.

In this paper, we will narrow the focus, and concentrate on a "strict" SMM (SSMM): a standard model of mind that provides a theory of how the human mind works. We take this stance because we believe that the main means of developing a cognitive architecture or an SMM is to use simulations of empirical data to constrain the space of possible architectures (or SMMs). In our view, with the broader definition adopted by Laird et al., there are too many degrees of freedom for converging to a correct solution. By doing this, we use the strategy eloquently defended by Newell in his book *Unified Theories of Cognition* (Newell, 1990).

This stance is also motivated by the fact that developing cognitive architectures and standard models would be of great help to (cognitive) psychology in its efforts to understand the human mind. In particular, an SSMM would help develop the kind of sufficiently powerful formal theories that are currently lacking in psychology.

We first comment on the concept of cognitive architectures and discuss whether we need a standard model, highlighting advantages and disadvantages of this approach. Empirical data are essential for developing an SSMM, and we discuss ways to select these data. Similarly, developing cognitive architectures and SSMMs requires a sophisticated methodology, and we present a recent proposal, centred on the use of empirical data for validating an architecture, which offers considerable benefits. Finally, we discuss CHREST, a cognitive architecture that shares some commonalities with the cognitive architectures examined by Laird et al., but that also differs in important ways. Implications for the development of SSMMs will be drawn.

Cognitive Architectures

As was lamented many years ago by Newell (1973), psychology is characterized by piecemeal research that addresses specific questions but lacks a unifying theory. As a way to remedy this, Newell (1990) proposed to develop unified theories of cognition (or cognitive architectures), theories implemented as computer programs that would account for as many empirical phenomena as possible. As a possible candidate, Newell chose Soar (Laird, 2012; Laird, Newell, & Rosenbloom, 1987).

Nearly 45 years after Newell's chapter, not much has changed in psychology (Gobet, 2017). Despite the considerable amount of data that have been collected, refined methodologies, great advances in neuroscience, there is still no unifying framework in psychology, and in fact most theories are rather local, in that they are developed to account for a specific experimental paradigm. Moreover, most theories in psychology are verbal and therefore imprecise and often self-contradictory.

The presence of a few cognitive architectures (e.g., ACT-R, CHREST, Soar; see Samsonovich, 2010 for a survey) does not change this judgement much, as the number of phenomena simulated is still limited, and experimental psychology has paid little attention to these cognitive architectures. However, it is now a great time for reminding psychologists of the failings of their discipline and how cognitive architectures might help improve the quality of science. In addition to the obvious fact that unification is a goal in science, as it allows summarizing large numbers of data into central mechanisms, scientific psychology currently faces a serious crisis. This crisis is caused by the difficulty in replicating many studies and also by the occurrence of isolated but high-profiles cases of scientific fraud, where for example data were fabricated (Gobet, 2016a; 2017). Many of the non-replicable findings are surprising but also unlikely (e.g., the claim that very brief exposure to the American flag would affect participants' political views many months after). Having a cognitive architecture or a standard model covering many different subfields of psychology would make it possible to make a priori predictions, and thus call attention to predictions that are so unlikely that they require further scrutiny.

Do We Need a Standard Model?

Clearly, comparison between architectures is useful, and is not carried out as often as it could (but for examples, see Lloyd-Kelly, Gobet & Lane, 2015a; Johnson, 1997; Jones et al., 2007). But Laird et al.'s idea is more ambitious: it is to extract a meta-architecture, so to speak, from the current architectures. An SMM would identify and put together the common features of the best available cognitive architectures and would thus be the best description of the human cognitive system based on our current knowledge.

An insight into the most likely candidate subcomponents would certainly be useful. Amongst the several benefits of an SMM highlighted by Laird et al., an agreement – even a partial one – between researchers about the basic component of an SMM would be a sign that the field is heading towards the right direction.

However, this enterprise seems difficult. To begin with, how to decide which architectures to use? Presumably, the aim is to focus on those that are successful. But how is this defined? This difficulty is compounded by the fact that SMMs concern human-like minds. As noted above, this loses a main source of constraint: empirical data. Focusing on SSMMs and architectures that simulate human data, as advocated in this paper, lessens this problem, as at least we have a metric for success. Nevertheless, evaluating the goodness of fit of a simulation, and in particular comparing one or several models or architectures is fraught with difficulties (e.g., number of free parameters, overfitting, etc.), as is well known in the literature (e.g., Ritter, 1991). (We'll take up this point below, with a methodology addressing many of these problems.)

Another difficulty is that architectures differ in fundamental ways (e.g., whether they are symbolic or non-symbolic, or whether they are embodied or not). This difficulty lies in part on the kind of behaviours we expect the architecture to exhibit, and what we wish to analyse. Fine details of mistakes made during a memory task may require different techniques to a problem-solving exercise.

What data should an SMM explain? Presumably, phenomena that psychologists find important. However, there is much disagreement in the field, unlike in physics. What looks like gold to some experimenters is seen as pyrite to their neighbours. In addition, as noted above, there are issues with measurement and replication of phenomena. One possibility is to look at a number of textbooks, and identify the phenomena that are discussed in most of them. We shall say more about this option below.

A further issue relates to the difficulty of separating structure from function. Many experiments in memory research have tried to estimate the capacity of short-term memory. However, recent empirical and computational research has shown that many estimates are incorrect, because they underestimate the amount of information that is already chunked in long-term memory, making it difficult to estimate STM capacity (Jones, Gobet & Pine, 2008). Difficulties also exist in distinguishing between architectural mechanisms and strategies, which can be influenced by instructions. For example, in concept formation experiments, participants' strategies can make their behaviour consistent with classical, exemplar, or prototype theories of categorization (Gobet, Richman, Staszewski, & Simon, 1997). But these theories are assumed to address fundamental mechanisms of categorization – the kind of mechanism that would be incorporated in an SMM.

A final issue is the role of neuroscience. At the moment, neuroscience does not provide the kind of constraints that were expected for developing cognitive architectures (Gobet, 2014). Just like psychology, difficulties of replication are endemic in neuroscience. In addition, there is the issue of what level of analysis/data would be useful for cognitive architectures. So far, the cognitive architectures that have included information about neuroscience have used data from a fairly high level of analysis – e.g., ACT-R refers to brain regions such as the anterior cingulate cortex and the basal ganglia – a level at which there is still much uncertainty. However, there is vast disagreement in neuroscience about the "correct" level, if any, at which cognition should be studied (Uttal, 2011).

Classic Marble

If one starts from a given architecture to decide on the phenomena that should be simulated, one faces the problem that only a (biased) small subset of cognitive phenomena will be considered. There is also the issue that the focus will be on phenomena that naturally flow from the architecture (e.g., power law of learning with Soar). It is important of course to test the predictions of an architecture, but there is still a risk of confirmation bias.

In practice, researchers have simulated data in a rather ad hoc fashion, focusing attention to phenomena they happened to be interested in for other reasons (for example, in the case of CHREST, chess).

A practical, bottom-up approach is to use, for selecting the phenomena that an architecture should account for, the type of results typically talked about in cognitive psychology textbooks. There is high consistency about the main topics they discuss. Typically, they start with perception, then move to learning, memory, before dealing with more complex behaviours such as problem solving, decision making, concept formation and language. Many include even more complex questions such as creativity, expertise and consciousness.

There is also considerable agreement about many key experiments in cognitive psychology. For example, Tversky and Kahneman's (1974) experiments on decision making, or Chase and Simon's (1973) experiments on experts' memory can be found in nearly all textbooks. Textbooks also tend to report the standard explanations for these phenomena. However, beyond textbooks, there is sometimes little assent about the mechanisms explaining these phenomena. For example, the literature on memory is replete with alternative explanations for Murdock's (1962) classic results.

In *Foundations of Cognitive Psychology*, an introductory textbook co-authored by one of us (Gobet, Chassy, & Bilalić, 2011), the final chapter highlights the key mechanisms and related empirical results that were discussed in the book. They were captured by the acronym Classic Marble: **C**ognition is **L**imited **A**ttention, **S**elective **S**earch, **I**nformation-processing, **C**hunking, **M**emory, **A**daptation, **R**ecognition, **B**ounded rationa**L**ity and **E**motion.

This list is only indicative and informal. It would be interesting to create a database of key, well-replicated results in psychology, with description of the experimental protocols, and links to replications of the original experiments. This would be a great boost for developing, testing and comparing architectures.

How to Develop Architectures?

We have argued for the importance of psychological data to provide constraints on an SMM. Our own proposal for incorporating such constraints into the SMM development process is the *robust-testing methodology* (Lane & Gobet, 2012a). Throughout we assume that the SMM is implemented as computer code in the form of a *cognitive architecture*: this is a weak assumption as a feasible SMM must be formalised, mathematically or in code, or else we return to the problem of verbal theories discussed above. Our main interest is in constraints from experiments on humans (an SSMM), but similar considerations apply for data from any natural human-like mind.

At the heart of our proposal is the definition of an *experimental constraint*. This definition links the constraints described by Newell and the computational development of the cognitive architecture and its associated models of individual results. Lane and Gobet (2003) proposed that the following four items make up a single constraint:

1. the performance of human participants in the experiment;
2. the stimuli used in the experiment;
3. the protocol for running the experiment and assessing the performance of the participant or model; and
4. the procedure for comparing the performance of the model and the participants.

For example, in a chess recall experiment, the target performance may be an average accuracy of different levels of player; the stimuli would be the chess positions used; the protocol would capture the order and timing of presentation of the pieces; and the comparison would look at the difference between the model's average accuracy and those of the human players. Ideally, a database of key, well-replicated results in psychology, as mentioned at the end of the previous section, would be a source of constraints.

Each experimental constraint can be written, in computer code, as a *test* to be applied and run, or included in a suite of constraints which can be checked automatically; we call such constraints *canonical results*, as they define the confirmed behaviour of the model. As an architecture is used, a suite of such canonical results will be created, covering a range of confirmed behaviours.

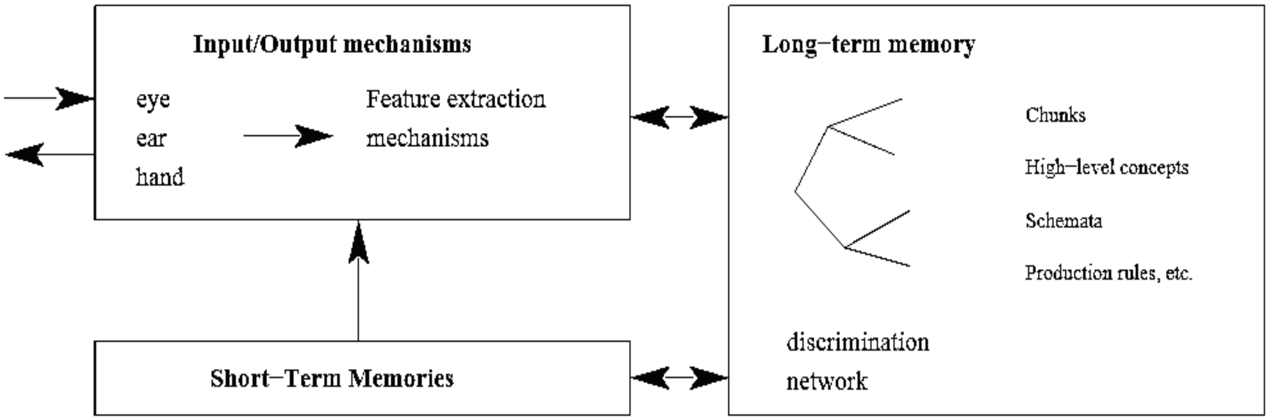

Figure 1: Overview of the CHREST architecture

Creating such a suite of canonical results, and ensuring that the results hold as the architecture evolves, provides a number of methodological benefits. First, there is a significant advantage in *documentation*. Each canonical result is a concrete example of the architecture used in an experimental study, which can be examined and analysed to enhance our understanding of the architecture and its application.

Second, a challenging issue with a number of models is the role of under-determined parameters within the model: in some cases, different values can dramatically alter the quality of fit. With the support of canonical results, parameters can be systematically altered using some form of optimisation process, and the range of viable parameter values explored. In Lane and Gobet (2013), we show how genetic algorithms can locate sets of parameters which define high quality models in the 5-4 task.

Third, the development of individual models is currently a laborious, and potentially error-prone, process of writing computer code and checking performance. A canonical result defines an acceptance test for the computer model; the model is a computer program written to meet the formal specification given by the canonical result. In certain domains, it is possible to automate the process of building the model. We have performed some pilot studies in this area (Frias-Martinez & Gobet, 2007; Lane et al., 2016), where computer models for tasks such as the Delayed Match to Sample are evolved using a genetic programming algorithm.

The fourth and final benefit, which may be one of the most important scientifically, is that our proposed methodology makes explicit many of the important properties of the developed model. This helps support and encourage replication and validation of results by other researchers. For example, an independent implementation could be written and run against the canonical results, to verify that the description of an architecture adequately defines the process of finding quantitative results (by leading to computer implementations with the same behaviour).

As an aside, the methodology we are proposing is, at heart, an adaptation of the agile development techniques currently popular in software engineering to the demands of an extended scientific project. These techniques rely on automated testing, constant rewriting of code (refactoring) to alter the program's design to include new features, regression testing to preserve behaviour over time, and continual interactions with the customer to ensure the program meets their (changing) requirements. For more on this correspondence, see Lane and Gobet (2003; 2012a).

CHREST

CHREST is a symbolic cognitive architecture, developed over 20 years to explore how experience affects how we remember, categorise and think about the world. CHREST's pattern of development is based on the principles we described above: trying to form a close fit to data taken from psychological experiments on humans. The general theory behind CHREST is the *template* theory (Gobet & Simon, 1996), and most of its mechanisms also embody Herbert Simon's theory of *bounded rationality* (e.g., Gobet, 2016b; Simon, 1982). Each principal result is compared against the performance of humans, as studied by psychologists.

A distinctive feature of CHREST is its discrimination network, used to retrieve information from LTM. CHREST models typically begin by training from naturalistic data to construct this network and associated long-term memories.

There are four main components to CHREST (see Figure 1). First, the input/output unit, which can handle visual or verbal information. Second, an STM, limited to holding a fixed number of items. Third, the LTM itself, and fourth the discrimination network, which forms an index into LTM.

The input mechanisms of CHREST are particularly sophisticated. In particular, the human eye is carefully simulated to capture data in expertise. This includes where the eye is fixating, various heuristics controlling how the eye moves, and timing parameters.

CHREST has been used to model human performance in a wide range of different applications. These include detailed models of expertise in chess (de Groot & Gobet, 1996; Gobet & Simon, 2000) and computer programming (Gobet & Oliver, 2016); diagrammatic reasoning (Lane, Cheng & Gobet, 2001); ageing (Smith et al., 2007); the role of emotions in problem gambling (Schiller & Gobet, 2014), and implicit learning (Lane & Gobet, 2012b). Related versions of CHREST's discrimination network are found in EPAM (Feigenbaum & Simon, 1984) and MOSAIC (Freudenthal et al., 2010), each of which carefully model human performance from naturalistic data.

Comparison with Other Architectures – Implications for SMM

CHREST is one of a number of cognitive architectures, such as ACT-R or Soar, which are the focus of the SMM discussion. One way in which CHREST complements these other architectures is its emphasis on how "declarative" memory is constructed and indexed, in particular with its links to perception, whether visual or verbal.

Laird (2012), after describing the requirements of Soar, states "we will still fall short of creating human-level agents until we encode, or *until the systems learn on their own*, the content required for higher-level knowledge-intensive capabilities" (p. 40, emphasis added). With CHREST, the construction of the discrimination network and associated learning of chunks[1] and their relations is a natural way to explain the origins of (some aspects of) declarative memory in knowledge-intensive tasks.

The way knowledge is acquired leads to a number of interesting conclusions about knowledge representation. First, CHREST is essentially bottom-up: more complex knowledge is incrementally and recursively built upon simpler elements, as a function of the interaction with the environment. In that sense, CHREST is self-organizing. By contrast, ACT-R and Soar architectures are more top-down: knowledge is built (by the programmer) around goals. Second, the distinction between declarative and procedural knowledge (productions) is not as clear-cut as in ACT-R and Soar, where these two types of knowledge are stored in different modules. In most CHREST applications, productions use perceptual chunks as conditions. Thus, productions are strongly anchored in perception. Third, due to their perceptual component, productions in CHREST are less powerful than productions in ACT-R and Soar; in particular, they do not offer by default the possibility of encoding variables. (Variables could be learnt by a mechanism similar to that used in the construction of templates.) Finally, because of the way knowledge is acquired, CHREST offers more constraints than ACT-R and Soar on the type of theories that it could implement.

Apart from learning, Langley et al. (2009) suggested that many cognitive architectures are overly focussed on problem-solving tasks, and that attention should be given to categorisation and understanding. The same authors suggest that architectures need to consider "visual, auditory, diagrammatic and other specialized representation schemes" used by humans, and should better reflect the limited resources available for perceiving and affecting the world. CHREST currently has this focus on categorisation and understanding, although it is relatively weak in formal problem-solving abilities, and in handling non-symbolic data. It embodies resource limitations in many areas, including a reliance on timing parameters, human-inspired limitations in perception, and short-term memory constraints; these limitations arise naturally as CHREST is built around the principle of bounded rationality.

With regard to the limitation in problem-solving ability, we are currently working on ways of improving this aspect of CHREST by combining its pattern-recognition abilities with problem-solving algorithms: see Lloyd-Kelly et al. (2015b) for one approach inspired by Dual Processing theories. Some other hard problems for cognition and how CHREST relates to them are discussed in Gobet and Lane (2005).

The time scales at which CHREST works are also very broad. From eye movements and encoding of data into template slots, CHREST can operate on the order of hundreds of milliseconds. However, its discrimination network requires a simulated time of many years to build up to expert level. As we discussed in Lane and Gobet (2012b), the same mechanisms in CHREST can be used to simulate learning which occurs over seconds and learning which requires many years to take place.

This focus on large scale knowledge acquired over long time scales is interesting theoretically because CHREST offers considerable constraints about what can be learnt. By contrast, architectures such as Soar and ACT-R make fewer ontological commitments about the nature of knowledge.

For these reasons, we believe that CHREST offers an example of how to use psychological data to drive the design of a strict cognitive architecture (or SSMM), and is also an example of a theory capable of operating across multiple timescales and with varied kinds of data.

[1] It is a very unfortunate twist of history and a source of much confusion that the term "chunk" has three different meanings in ACT-R, Soar and CHREST: a unit of declarative knowledge, a unit of procedural knowledge, and a node in long-term memory encoding perceptual information, respectively. For a discussion of these and other meanings of the term, see Gobet, Lloyd-Kelly, & Lane, 2016).

References

Chase, W. G., and Simon, H. A. 1973. Perception in chess. *Cognitive Psychology* 4:55–81.
De Groot, A. D., and Gobet, F. 1996. *Perception and memory in chess.* Assen: Van Gorcum.
Feigenbaum, E. A., and Simon, H. A. 1984. EPAM-like models of recognition and learning. *Cognitive Science* 8:305–336.
Freudenthal, D., Pine, J. M. and Gobet, F. 2010. Explaining quantitative variation in the rate of Optional Infinitive errors across languages: A comparison of MOSAIC and the Variational Learning Model. *Journal of Child Language* 37:643–669.
Frias-Martinez, E. and Gobet, F. 2007. Automatic generation of cognitive theories using genetic programming. *Minds and Machines* 17:287–309.
Gobet, F. 2014. W. R. Uttal: Mind and brain: A critical appraisal of cognitive neuroscience. *Minds and Machines* 24:221–226.
Gobet, F. 2016a. *Understanding Expertise: A Multidisciplinary Approach*. London: Palgrave.
Gobet, F. 2016b. From bounded rationality to expertise. In *Minds, Models and Milieux*, eds. R. Frantz and L. Marsh, 151–166. London: Palgrave Macmillan.
Gobet, F. 2017. Allen Newell's program of research: The video game test. *Topics in Cognitive Science* 9:522–532.
Gobet, F., Chassy, P. and Bilalić, M. 2011. *Foundations of Cognitive Psychology*. London: McGraw Hill.
Gobet, F. and Lane, P. C. R. 2005. The CHREST architecture of cognition: Listening to empirical data. In *Visions of Mind: Architectures for Cognition and Affect*, ed. D. Davis, 204–224. Hershey, PA: IPS.
Gobet, F., Lloyd-Kelly, M. and Lane, P. C. R. (2016) What's in a Name? The Multiple Meanings of "Chunk" and "Chunking". *Frontiers in Psychology*, 7:102.
Gobet, F. and Oliver, I. 2016. Memory for the random: A simulation of computer program recall. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
Gobet, F., Richman, H. B., Staszewski, J. J., and Simon, H. A. 1997. Goals, representations, and strategies in a concept attainment task: The EPAM model. *The Psychology of Learning and Motivation* 37:265–290.
Gobet, F. and Simon, H. A. 1996. Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology* 31(1):1–40.
Gobet, F. and Simon, H. A. 2000. Five seconds or sixty? Presentation time in expert memory. *Cognitive Science* 24:651–682.
Johnson, T. R. 1997. Control in ACT-R and Soar. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, 343–348. Austin, TX: Cognitive Science Society.
Jones, G., Gobet, F. and Pine, J. M. 2008. Computer simulations of developmental change: The contributions of WM capacity and long-term knowledge. *Cognitive Science* 32:1148–1176.
Jones, R. M., Lebiere, C. and Crossman, J.A. 2007. Comparing modeling idioms in ACT-R and Soar. In *Proceedings of the 8th International Conference on Cognitive Modeling*, 49–54. London: Taylor & Francis.
Laird, J. E. 2012. *The Soar Cognitive Architecture*. Boston: MIT.
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., Newell, A. and Rosenbloom, P. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence 33*:1–64.
Lane, P. C. R., Cheng, P. C-H., and Gobet, F. 2001. Learning perceptual chunks for problem decomposition. In *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 528–33. Austin, TX: Cognitive Science Society.
Lane, P. C. R. and Gobet, F. 2003. Developing reproducible and comprehensible computational models. *Artificial Intelligence* 144:251–263.
Lane, P. C. R. and Gobet, F. 2012a. A theory-driven testing methodology for developing scientific software. *Journal of Experimental & Theoretical Artificial Intelligence* 24:421–56.
Lane, P. C. R. and Gobet, F. 2012b. CHREST models of implicit learning and board game interpretation. In *Proceedings of the Fifth Conference on Artificial General Intelligence, LNAI* 7716, 148–157. New York: Springer.
Lane, P. C. R. and Gobet, F. 2013. Evolving non-dominated parameter sets for computational models from multiple experiments. *Journal of Artificial General Intelligence* 4:1–30.
Lane, P. C. R., Sozou, P. D., Gobet, F., and Addis, M. 2016. Analysing psychological data by evolving computational models. In *Analysis of Large and Complex Data*, eds. A. Wilhelm and H. Kestler, 587–597. New York: Springer.
Langley, P., Laird, J. E and Rogers, S. 2009. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research* 10:141–160.
Smith, R. Ll., Gobet, F. and Lane, P. C. R. 2007. An investigation into the effect of ageing on expert memory with CHREST. In *Proceedings of The Seventh UK Workshop on Computational Intelligence*. London, UK: Imperial College.
Lloyd-Kelly, M., Gobet, F., and Lane, P. C. R. 2015a. Piece of mind: Long-term memory structure in ACT-R and CHREST. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
Lloyd-Kelly, M., Gobet, F., and Lane, P. C. R. 2015b. A question of balance: The benefits of pattern-recognition when solving problems in a complex domain, *LNCS Transactions on Computational Collective Intelligence* XX:224–258.
Murdock, B. 1962. Serial position effect of free recall. *Journal of Experimental Psychology* 64:482–488.
Newell, A. 1973. You can't play 20 questions with nature and win. In *Visual Information Processing*, ed. W. G. Chase, 283–308. New York: Academic Press.
Newell, A. 1990. *Unified Theories of Cognition*. Cambridge: Harvard University Press.
Ritter, F. E., 1991. Towards fair comparisons of connectionist algorithms through automatically optimized parameter sets. In *Proceedings of the Annual Conference of the Cognitive Science Society* (pp. 877–881).
Samsonovich, A. V. 2010. Toward a Unified Catalog of Implemented Cognitive Architectures. *BICA*, *221*:195–244.
Schiller, M. and Gobet, F. 2014. Cognitive models of gambling and problem gambling. In *Problem gambling: Cognition, prevention and treatment,* eds. M. Schiller and F. Gobet, 74–103. London: Palgrave Macmillan.
Simon, H. A. 1982. *Models of Bounded Rationality: Behavioral Economics and Business Organization.* Cambridge: MIT Press.
Tversky, A. and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–1131.
Uttal, W. R. 2011 *Mind and Brain: A Critical Appraisal of Cognitive Neuroscience*. Boston: MIT Press.