A General Instance-Based Learning Framework for Studying Intuitive Decision-Making in a Cognitive Architecture

Robert Thomson, Christian Lebiere, John R. Anderson, and James Staszewski

Carnegie Mellon University

Corresponding Author's Contact Information:

Robert Thomson

thomsonr@andrew.cmu.edu

Abstract

Cognitive architectures such as ACT-R have not traditionally been used to understand intuitive decision making; instead, models tend to be designed with the intuitions of their modelers already hardcoded in the decision process. This is due in part to a fuzzy boundary between automatic and deliberative processes within the architecture. We argue that instance-based learning – as applied to ACT-R – satisfies the conditions for intuitive decision-making described in Kahneman and Klein (2009), separates automatic from deliberative processes, and provides a general mechanism for the study of intuitive decision making. Also, to better understand the critical role of the environment in decision-making, we propose a description of decision-making and biases as arising from three sources: the mechanisms and limitations of the human cognitive architecture, the information structure in the task environment, and the use of heuristics and strategies to adapt performance to the dual constraints of cognition and environment. Several decision-making models are described according to this framework and potential lessons on how to use cognitive models to aid intuitive decision making are discussed.

A General Instance-Based Learning Framework for Studying Intuitive Decision-Making in a Cognitive Architecture

This paper describes computational models of intuitive decision making expressed within the constraints of the ACT-R cognitive architecture (Anderson et al., 2004). The models are noteworthy for their ability to explain a variety of heuristics and biases in terms of the processes and representations that produce them. These phenomena have largely been captured and defined as results of experimental manipulations (Kahneman & Tversky, 1996) but not in terms of process models (Dimov, Marewski, & Schooler, 2013). There is a perception that modeling intuitive decision-making behavior using cognitive architectures is confounded by the explicit decisions (e.g., strategy selection; Marewski & Schooler, 2011) encoded by the modelers. This criticism can be described as: instead of modeling intuitive behavior *per se*, cognitive models make explicit the intuitions of their designers (Cooper, 2007; Lewandowsky, 1993; Schultheis, 2009; Shallice & Cooper, 2011). We address this criticism by showing that the instance-based learning mechanisms in the ACT-R cognitive architecture (Gonzalez, Lerch, & Lebiere, 2003) exhibit the characteristics of intuitive decision-making as described in Klein and Kahneman (2012), and provide a clear distinction between automatic and implicit (system 1) processes and deliberative and explicit (system 2) processes. In addition, we specifically address the strategy selection criticism by showing that the explicit strategies of the models instantiate the theories of the model designer and thus are a mechanism for model evaluation rather than a confounding factor in model development.

In making this argument we recommend adopting a tripartite explanation of decision-making and biases that illustrates the critical role of the task environment in the decision-making process. We argue that decision-making should be described in terms of: 1) the mechanisms and limitations of the architecture; 2) the information structure in the task environment; and 3) the use of heuristics and strategies to adapt performance to the dual constraints of cognition and environment. From examples of existing models, we show that simulating behavior within a cognitive architecture is a useful methodology for the study of the mechanisms, variables, and time-course in complex decision-making processes that are impossible in experimentation due to exploding combinatorics.

What is Intuitive Decision-Making?

Simon (1992) characterized intuitive decision-making skill as “nothing more and nothing less than recognition” (p. 155). In their seminal work on expertise, Chase and Simon (1973) identified that chess experts require upwards of a decade of study to retain 50,000 to 100,000 distinct and rapidly-accessible patterns of chess movements. Intuitive decision-making has been studied in both the naturalistic decision-making (NDM) and heuristics and biases (HB) literature, with the former focused on the successes of intuitive reasoning, while the latter focused on its failures (Klein & Kahneman, 2012). A distinguishing feature of intuitive decision-making is that a single plausible solution rapidly ‘comes to mind’ in its entirety without explicit or conscious awareness of the causal factors entering into the decision (i.e., not being consciously derived in a piecemeal, step-by-step, or in a ‘deliberative’ manner; Newell & Simon, 1972; Simon, 1995). As such, intuitive reasoning is considered ‘System 1’. For example, Klein, Calderwood, and Clinton-Cirocco (1986) found that fire marshals tended to make rapid decisions by generating a single alternative, mentally simulating its outcome, and either making minor revisions or adopting the next closest alternative. Effectively, fire marshals were pattern-matching based on their prior experiences. This strategy has been termed *recognition-primed decision-making*.

Conversely, deliberative decision-making is often characterized as strategic, effortful, slow, and rule-oriented (Klein, 1998), and as such is considered ‘System 2’ thinking (Kahneman & Frederick, 2004; Stanovich and West, 1999, 2000). Interestingly, the act of verifying an intuition is generally seen as optional, effortful, and thus a function of System 2 (Klein & Kahneman, 2012).

In order to gain intuitive expertise, several conditions first need be met. The first condition is that people receive extensive practice in a task environment that is sufficiently stable and provides causal or statistical cues/structures that may at least theoretically be operationalized (Brunswik, 1957). This need not be deterministic (e.g., playing poker is a probabilistic but stable environment; Klein & Kahneman, 2012). The second condition is that there must be sufficient feedback from the task environment which provides people an opportunity to learn the relevant cues/structures. In other words, feedback must be sufficient to generate a relevant internal problem space. This requirement of feedback and interaction with the task environment drove our adoption of the tripartite level of description.

Why use a Cognitive Architecture?

Before getting into the details of mechanisms, models, and results; there is an important argument to be made for the role of cognitive architectures in general, which is best characterized by Herb Simon (in 1971, no less):

> The programmability of the theories is the guarantor of their operationality, an iron-clad insurance against admitting magical entities into the head. A computer program containing magical instructions does not run, but it is asserted of these information-processing theories of thinking that they can be programmed and will run. They may be

> empirically correct theories about the nature of human thought processes or empirically invalid theories; [but] they are not magical theories. (p. 148)

In modern terms, simulations using a cognitive architecture provide a falsifiable methodology for the study of entities, mechanisms, and variables. They serve several theoretical functions including: organizing and relating a substantial number of cognitive mechanisms, making testable predictions, and explaining the cognitive processes underlying human performance. In many cases, cognitive models can perform a task too complex to analyze with traditional experimentation due to the combinatorics of possible decisions. For example, much in the HB literature is explained in very broad brush strokes, and the evidence comes from only very simple tasks. As will be described, a single ACT-R model has explained anchoring and adjustment, confirmation, and probability matching biases across a range of geospatial intelligence tasks using a common instance-based learning approach (Lebiere, Pirolli, Thomson, et al., 2013). Similarly, Marewski and Melhorn (2011) were able to specify 39 process models studied in decision-making using a smaller subset of 5 ACT-R models. In short, architecture constrains explanations to scientifically-established mechanisms and (hopefully) easily describable processes.

This is not to say that cognitive architectures are a panacea for studying decision-making (or psychology in general), but they are a valuable tool in the generation and exploration of theories (c.f., models) too complex for traditional piecemeal experimental methods. In particular, intuitive decision-making is cognitively ‘opaque’ with little observable evidence and what little evidence there is being highly fallible introspection. As such, many descriptions of intuitive decision-making are inherently qualitative or characterized as relatively simple experimental results (Dimov, Marewski, & Schooler, 2013). An advantage of cognitive architectures is not only their ability to objectively explain accuracy and response times in terms of the operation of both symbolic and sub-symbolic mechanisms (and in the case of ACT-R, links to neural structure), but also the ability to go ‘under the hood’ and actually *look at the code to explicitly examine causal processes*. Such computational cognitive models also make testable predictions of what is going on inside the mind of someone performing intuitive decision-making.

One measure of validation is to perform model tracing. Model tracing is a technique where a model is forced to respond with some or all of the same values as a human participant at the level of component processes, and then the internals of the system are examined to determine the influence of these ‘forced’ decisions. By examining the commonalities between the model’s internal states and human behavior, we are potentially able to make causal claims about the nature of mental processes within participants, that is, explain how human performance is produced by various mechanisms and their interaction. This performance includes traditional measures such as accuracy and response time, but also predictions of fMRI bold response for specific brain areas associated with functional modules of the cognitive architecture (Anderson, 2007).

We cannot get away without discussing some limitations of cognitive architectures, but we maintain that these limitations can be outweighed by the advantages of this method of theorizing. Perhaps the main criticism of cognitive architectures is the *degrees of freedom* argument (e.g., Roberts & Pashler, 2000). This argument relies on notion that the numerous parameters of the architecture offer little constraint on the kinds of models that may be generated to solve a given problem. There are two means of mitigating these biases: the first is to limit the freedom of parameters through the use of scientifically-justified default values, and the second is to develop models which are relatively insensitive to parameter values. For instance, the base-level learning rate of .5 in ACT-R has been justified in over 100 published models (see the ACT-R website: act-r.psy.cmu.edu). In addition to architectural parameters, there are ways of controlling ‘knowledge’ parameters (e.g., knowledge representation) through approaches like instance-based learning which use a common knowledge representation determined directly by the interaction with the external environment. Finally, there have been efforts to use a single, more general ACT-R model to perform a series of decision-making tasks (e.g., Lebiere et al., 2013; Marewski & Melhorn, 2011; Taatgen & Anderson, 2008). Using more general models of cognition reduces the degrees of freedom in the architecture by aligning a common set of mechanisms, parameters, and strategies into a cohesive modeling paradigm.

The benefits of cognitive architectures can be seen as a synthesis between formal mathematical theories (such as Bayesian modeling) and knowledge-level strategies (e.g., heuristics). Bayesian models belong to a broad class of abstract models that formally (i.e., mathematically) explain human behavior in terms of processes updating probabilities over a set of possible decisions. While Bayesian models do provide an explanation of behavior, it is not generally accepted to be a cognitive one and where the underlying mechanisms driving the processes are somewhat vague (Bowers & Davis, 2012). This is not a criticism specific to Bayesian models, but can also be applied to prospect theory (Kahneman & Tversky, 1979) and models of quantum probability (Busemeyer, Pothos, Franco, & Trueblood, 2011). On the other hand, explanations in the form of heuristics (i.e., knowledge-level explanations) tend to be vague as to the underlying processes leading to the biased behavior. In short, a description at either the formal or heuristic level is underconstrained. Cognitive architectures tie together the formal and heuristic levels of description (Gonzalez & Lebiere, 2011) by combining subsymbolic algorithms akin to formal theories with symbolic knowledge structures controlling those processes in a heuristic manner. As such, they provide not only explanations of existing behavior, but predictions of the mechanisms, knowledge, and behaviors of future actions based on a model's prior experience (Thomson & Lebiere, 2013).

Why Use ACT-R?

What is ACT-R?

ACT-R is a cognitive architecture defined as a set of modules which are integrated and coordinated through a centralized production system. Each module is assumed to access and deposit information into buffers associated with the module, and the production system only responds to the contents of the buffers, not the internal encapsulated processing of the modules. The declarative memory and production system modules, respectively, store and retrieve information that corresponds to *declarative knowledge* and *procedural knowledge*. Declarative knowledge is the kind of knowledge that a person can attend to, reflect upon, and usually articulate in some way (e.g., by declaring it verbally or by gesture). Procedural knowledge consists of the skills we display in our behavior, generally without conscious awareness.

Declarative knowledge in ACT-R is represented formally in terms of chunks. The information in declarative memory corresponds to episodic and semantic knowledge that promotes long-term coherence in behavior. Chunks have a type, and consist of slot-value pairs of information. Chunks are retrieved from long-term declarative memory by an activation process. Each chunk has a base-level activation that reflects its recency and frequency of occurrence. Activation spreads from the current focus of attention through associations among chunks in declarative memory. These associations are built up from experience, and they reflect how chunks co-occur in cognitive processing. Chunks are compared to the desired retrieval pattern using a partial matching mechanism that subtracts from the activation of a chunk its degree of mismatch to the desired pattern, additively for each component of the pattern and corresponding chunk value. Noise is added to chunk activations to make retrieval a probabilistic process governed by a Boltzmann (softmax) distribution. While the most active chunk is usually retrieved, a blending process (i.e., a blended retrieval) can also be applied that returns a derived output reflecting the similarity between the values of the content of all chunks, weighted by their retrieval probabilities reflecting their activations and partial-matching scores. This process enables not just the retrieval of previously encountered symbolic values but also the generation of continuous values such as probability judgments in a process akin to weighted interpolation.

Production rules are used to represent procedural knowledge in ACT-R. They specify procedures that represent and apply cognitive skill in the current context, including how to retrieve and modify information in the buffers and transfer it to other modules. In ACT-R, each production rule is a set of conditions and actions which are analogous to an IF-THEN rule. Conditions specify structures that are matched in buffers, and correspond to information from the external world or other internal modules. Actions represent requests and modifications to the contents of the buffers, including queuing perceptuo-motor responses (e.g., speaking, typing, or looking to given location). Matching production rules effectively says: if the conditions of a given production match the current state of affairs (i.e., the state of the modules and contents of the buffers) then perform the following actions.

ACT-R uses a mix of parallel and serial processing. Modules may process information in parallel with one another. However, there are two serial bottlenecks in processing. First, only one production may execute at a time. Second, each module is limited to placing a single chunk in a buffer. In general, multiple production rules can apply at any point. Production utilities, learned using a reinforcement learning scheme, are used to select the single rule that fires. As for declarative memory retrieval, production selection is a probabilistic process. Based on experience and matching certain criteria, two production rules may be automatically compiled together into a new and more-efficient rule, which accounts for proceduralization of behavior.

What does ACT-R have to do with Intuitive Decision-Making?

There are a broad range of ACT-R models studying problem solving, decision-making (including intuitive decision-making; Kennedy & Patterson, 2012), and implicit learning (see act-r.psy.cmu.edu/publications for examples of each; also, see Anderson, 2007 and Lebiere & Anderson, 2011 for an overview). Specific examples include a model of how batters predict baseball pitch speed (Lebiere, Gray, Salvucci & West, 2003), a model predicting risk aversion in a repeated binary choice task (Lebiere, Gonzales & Martin, 2007), a model of sequence learning (Lebiere & Wallach, 2001), and a model of playing Paper Rock Scissors (West & Lebiere, 2001). These models all work by storing problem-solving instances in declarative memory, then making decisions by retrieving those instances, leveraging the cognitive architecture's activation processes to extract environmental regularities.

There is a misconception that intuitive processes – being implicit – are governed using only *procedural* memory processes, while deliberative processes – being explicit – are governed by only *declarative* memory processes (this implicit/explicit distinction has been mistakenly attributed to Wallach & Lebiere, 2003a). In fact, while each declarative chunk is usually considered a piece of conscious knowledge, the sub-symbolic activations that control the retrieval process (e.g., base-level activations and strengths of associations) are consciously inaccessible and constitute the implicit knowledge of the model (Gonzales & Lebiere, 2005). In essence, the process of retrieving a chunk is the implicit part of the declarative system, while what the model does with the retrieved chunk (and the content of the chunk itself) is the explicit part of the declarative system.

An interesting interplay between system 1 and system 2 processes occurs during a retrieval request. When a production makes a retrieval request it specifies the type of chunk to retrieve, and potentially a set of slot-value pairs from which to match; which is essentially the *specification* of what to retrieve. While the production system is generally seen as an implicit (system 1) process, the constraints in matching the retrieval request come from explicitly setting which slot-values pairs to match against. Since this is something coded by the modeler, it could be argued that it is a totally explicit strategy (i.e., system 2) based on the modelers *intuition* (c.f., theory) of how the retrieval should function. This is a bit of a false dichotomy because every model is a blend of both system 1 and system 2 processes, and the retrieval request links both strategic (e.g., requesting a specific chunk) and implicit processes (e.g., spreading activation). In terms of a cognitive architecture and the *no-magic doctrine (Anderson & Lebiere, 1998)*, we argue that the retrieval specification is instead best described as an implicit heuristic, albeit still being effectively the modeler's theory of how the retrieval process should unfold.

It is possible to make retrieval driven more by implicit processes by using a technique we call making 'open' retrievals. A retrieval is considered *open* when only the type of chunk is requested and *no* (or in a relaxed case, *minimal*) slot-value pairs are used in the specification of the retrieval request. An example would be when one is given a set of indirect clues about the identity of a person and the name of the person pops up in one's mind from the convergence of the clues rather than any specific information retrieval process. Effectively, *open* retrievals are a kind of context-driven free association. By using *open* retrievals, the model is relying more on sub-symbolic activations – which are driven by experience – to control the retrieval process. For instance, performing a retrieval by specifying only the context and doing free association on the outcome allows the model to match the best outcome based on the recency and frequency of prior outcomes. This stands in contrast to specifying a particular outcome in the retrieval request, which is more analogous to the model engaging in a more explicit retrieval strategy.

A similar theme between system 1 and system 2 processes is the nature of heuristics in decision-making. Are heuristics explicit because of their symbolic nature, or implicit because the decision maker is often unaware of them? The choice of which simplifying heuristics are available to the model tends to be a conscious strategy of the modeler (as opposed to being chosen by the model; Lewandowsky, 1993). This may be considered *explicit*, although this is an uncharitable view of the modeler's strategy selection (e.g., the modeler's theory of which heuristics are available; Taatgen and Anderson, 2008).

This criticism can be mitigated by viewing the task instruction as a kind of heuristic imposed by the task environment, thus by parsing the task instructions the model derives the heuristic. Another possibility is to adopt the strategy of Gigerenzer (2002) and provide a psychologically-validated set of heuristics (i.e., an *adaptive toolbox*) from which the model may select. Of course, this leaves open to theory the strategy selection of the model (Marewski & Schooler, 2011; Marewski, Gaissmaier, Schooler, Goldstein, & Gigerenzer, 2010). One of the difficulties in providing a more complete answer is because the traditional dichotomy of automatic versus deliberative (or procedural versus declarative) is insufficient to explain the source of heuristics (mainly from the task environment), which is a key indicator of whether the heuristic should be seen as implicit, explicit, or both.

How does Instance-Based Learning tie into Intuitive Decision-Making?

Instance-based learning theory (IBL; Gonzalez, Lerch, & Lebiere, 2003; Taatgen, Lebiere, & Anderson, 2006) is the claim that implicit expertise is gained through the accumulation and recognition of experienced events or instances. Unlike instance-based machine learning algorithms (Gagliardi, 2011) that are essentially strict exemplar models of categorization applied to big data (Erickson & Kruschke, 1998), IBL allows for more generalization and bootstrapping learning with weak methods. Weak methods are relatively knowledge-free methods of action and exploration (e.g., random choice) that are used to bootstrap learning. These 'weak' methods are more procedurally-driven because there is insufficient domain knowledge (i.e., instances) at that point to make good decisions. Once enough instances are stored, these weak methods are supplanted by the retrieval of decisions based on these prior instances.

Similar to theories of intuitive expertise (Klein & Kahneman, 2012), IBL argues for the necessity of receiving effective feedback. The requirement for feedback (either externally from the environment or by an internally-generated heuristic) results in a common condition → action → outcome representational structure that reflects the requirements for effective performance. Supporting this structure, Lebiere, Gonzalez, and Warwick (2009) have shown how Klein's (2009) recognition-primed decision-making (RPD) and IBL use similar mechanisms and make similar predictions in the context of naturalistic decision-making. IBL, having been formulated within the principles and mechanisms of cognition in ACT-R, makes use of the dynamics of chunk retrieval and using blended retrievals to generalize knowledge, which provides an additional level of explanation and predictive power to complement the process specified in Klein's analysis. As such, RPD and other similar naturalistic processes can be seen as a macrocognitive substrate that naturally complements the microcognitive mechanisms of a cognitive architecture (Lebiere & Best, 2009).

The main claim of IBL is that implicit knowledge is generated through the creation of instances. These instances are represented in chunks with slots containing the conditions (e.g., a set of contextual cues), the decision made (e.g., an action), and the outcome of the decision (e.g., the utility of the decision). Before there is sufficient task-relevant knowledge, decision-makers implicitly evaluate alternatives using heuristics (e.g., random choice, minimize loss, maximize gain). Once a sufficient number of instances are learned, decision-makers retrieve and generalize from these instances to evaluate alternatives, make a decision, and execute the task. The process of feedback involves updating the outcome slot of the chunk according to the post-hoc generated utility of the decision. Thus, when decision-makers are confronted with similar situations while performing a task, they gradually abandon general heuristics in favor of improved instance-based decision-making processes (Gonzalez & Lebiere, 2005).

Comparing IBL with the necessity claims of intuitive decision-making from Klein and Kahneman (2012), both consider intuitive knowledge to be learned via instances. Also, in both cases decisions are made by pattern-matching over prior instances (and/or supplemented by heuristics) and then retrieving the best-fit. In the case of IBL, however, this best fit is computed using a generalization across the closest neighbors using partial matching or blended retrievals. Both require the task environment to be sufficiently regular to be able to implicitly learn the statistical correlations between condition, action, and – through either internal or external feedback – outcome. However, IBL offers constraints on explanation by grounding implicit learning within the mechanisms of a cognitive architecture. For instance, the dynamics of an instance’s sub-symbolic activations (e.g., frequency and recency in the base-level learning equation) provide a scientifically-justified mechanism for determining which instances are likely to be retrieved for a given situation, and also can explain *why* they were retrieved and what factors came into play. This provides a much more rigorous explanation of intuitive decision-making than case-studies and introspection of experts.

IBL – as instantiated in ACT-R – also provides for a clear distinction between automatic and deliberative processes. The act of encoding and retrieving instances is a fully automatic process, guided by either (hopefully *open*) retrieval or implicit heuristics. However, once the retrieval is completed, what the model does with the retrieved chunk (e.g., the structure of the subsequent productions) is an explicit heuristic/strategy. For instance, while the retrieved chunk might provide a recommended action, it is up to the model (through the production system) to determine whether to verify the action, discard the action, perform the action, or simulate possible other outcomes.

Models related to decision-making and problem-solving models in ACT-R over the past 10 years have seen increasing use of IBL (whether explicitly referred-to as such or otherwise; e.g., Kennedy & Patterson, 2012) to learn intuitive knowledge structures. This is unsurprising given that ACT-R's declarative memory module and chunk structure is an excellent match for the storage and retrieval of instances, which effectively guides people to some form of IBL. In other words, the design and constraints of the architecture lead people to adopt an IBL-like approach by using the architecture in the most direct and intuitive way.

Why a Tripartite Description?

An essential feature in being able to explain how a model performs decision-making is to examine not only the sources of generating expertise (e.g., the role of IBL in naturalistic decision-making), but also to examine both *where* heuristics come from and *how* they are applied (e.g., heuristics and biases). The implicit versus explicit argument ignores the question of where heuristics may come from – such as the structure of the task environment – something which is essential for the implicit learning of expertise. We argue that a tripartite description is sufficient to explain the source of both successes and failures in decision-making. These three levels include a description of the mechanisms and limitations of the architecture, the information structure in the task environment, and the use of heuristics and strategies to adapt performance to the dual constraints of cognition and environment.

The first level of description entails an understanding of the constraints imposed by the mechanisms and limitations of the cognitive architecture. In ACT-R, these include an understanding of the impact of recency and frequency of the likelihood of an instance being retrieved, which also influences the ability of the model to generalize to new situations when using blended retrievals to generate a derived output rather than a specific instance. Other sources of constraint include the serial nature of the production system, only a single chunk being in a buffer at a time, and matching human time-course of responses. A common source of biased behavior in IBL decision-making models is the use of blended retrievals, which have a tendency to retrieve values that are pulled towards the mean of all values in memory. This common mechanism can lead to both anchoring and confirmation biases based on how far the anchored value varies from the mean across all instances in memory (Lebiere et al., 2013). It is important to note that this wholly implicit process is not consciously available to the model.

The second level of description entails an understanding of the constraints imposed by the task environment. This kind of description has been somewhat neglected in discussions of the validity of cognitive models; however, it is a critical feature in understanding both the consistency of learning and the nature of biases. An understanding of the statistical and quantifiable regularities within the task environment drives the overall ability and rate of learning, and the nature of environmental feedback provides further evidence. In other words, the detection of affordances (Gibson, 1977) provided by the task environment influence the kinds of information that the model can accumulate and the actions that the model may perform. To further push this issue, Simon (1990) argued that "[h]uman rational behavior ... is shaped by a scissors whose two blades are the structure of the task environments and the computational capabilities of the actor'' (Simon, 1990, p. 7). Using another example from Simon; if you want to study the movement of an ant across the beach you need look no further than the hills and valleys in the sand to determine its path. Anderson (1990) explored how the mechanisms of the cognitive architecture are shaped by the structure of the environment in which it evolved.

The third level of description entails an understanding of how the joint constraints of architecture and task environment influence the kinds of heuristics and strategies available to the model. In ACT-R terms, this is the explanation of the nature and sequence of productions firing. This level is most important to describe as it entails most of the choices of the modeler in designing the model. Even simple heuristic structures can greatly influence the output of the model, which could overly constrain decision-making but also makes complex problems tractable. Another possible focus of explanation is to justify the nature of the model's metacognitive awareness (if any). For instance, in IBL the model transitions from reasoning via heuristics to reasoning via instances with enough experience. The implementation of metacognitive awareness is one way to determine and explain how the model moves from heuristic- to instance-based reasoning.

Todd and Gigerenzer (2003) argue for a similar position they refer to as *ecological rationality* which is composed of three focuses: a focus on the mind (the study of cognitive illusions), a focus on the world (the study of constraint optimization), and a focus on putting mind and world together (the study of ecological rationality). What we have done is adopt a similar framework for the study of cognitive architectures, and focus on explaining behavior in terms of the more transparent constraints of the architecture, abstracted task environment, and model structure.

Now that we are armed with a theory (IBL) and a means of describing model output (the tripartite description), we can delve into some examples.

Intuitive Decisions in Sensemaking

Rather than provide an overview of many examples, we would like to focus on an in-depth analysis of a single ACT-R model of sensemaking that uses IBL to perform multiple complex geospatial intelligence tasks and provides both an explanation of biases and a close fit to human data (see Lebiere et al., 2013, for a more complete description of the tasks and quantitative model fits). Sensemaking is a concept that has been used to define a class of activities and tasks in which there is an active seeking and processing of information to achieve understanding about some state of affairs in the world, which has also been applied in organizational decision-making (Weick, 1995).

Our sensemaking task is composed of three sequential components. The first is focused on learning statistical patterns of events and then generating probability distributions of category membership based on the spatial location and frequency of these events (e.g. how likely does a given event belong to each of the categories). The second requires the application of probabilistic decision rules in order to generate and revise probability distributions of category membership (e.g., if a given feature is present at an event, then that event is twice as likely to belong to category A). The third involves making decisions about the allocation of resources based on the judged probabilities of the causes of perceived events, and is effectively a metacognitive measure of confidence in one's judgment. An example of the task interface can be seen in Figure 1.

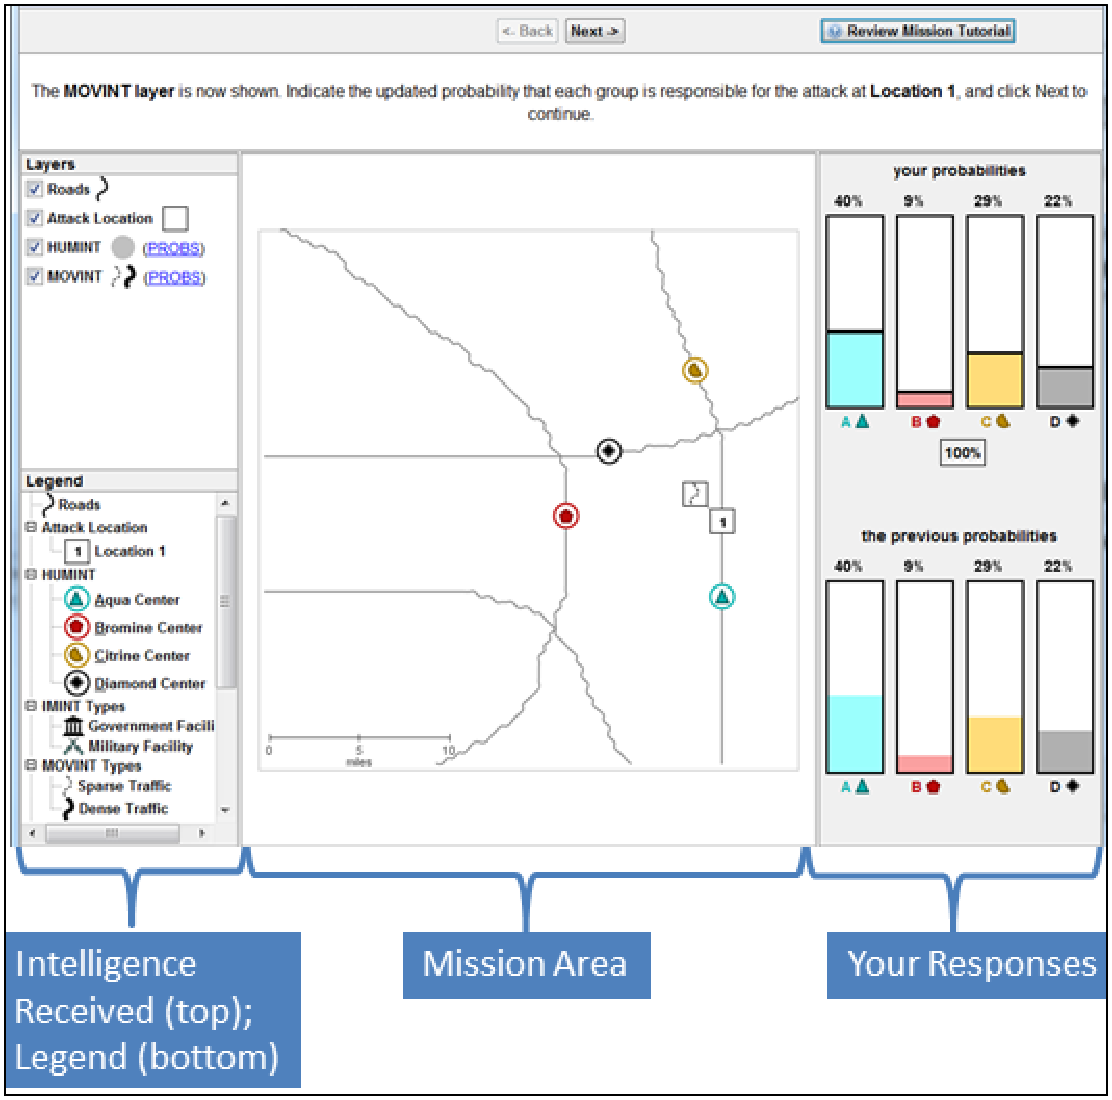

Figure 1: A sample of the task interface. To the left is a legend explaining all the symbols on the map (center). To the right are the probability distributions for the four event categories. The pane across the top provides step-by-step instructions for participants.

Biases in Group Center Generation

In the first task component, the flow of an average trial began with participants perceiving a series of events labeled according to which category the event belonged, each corresponding to a group icon on the central map, after which a probe was then displayed (see Figure 2). Participants were then required to generate a center of activity for each category's events, and generate a probability estimate for each category (summed to 100%).

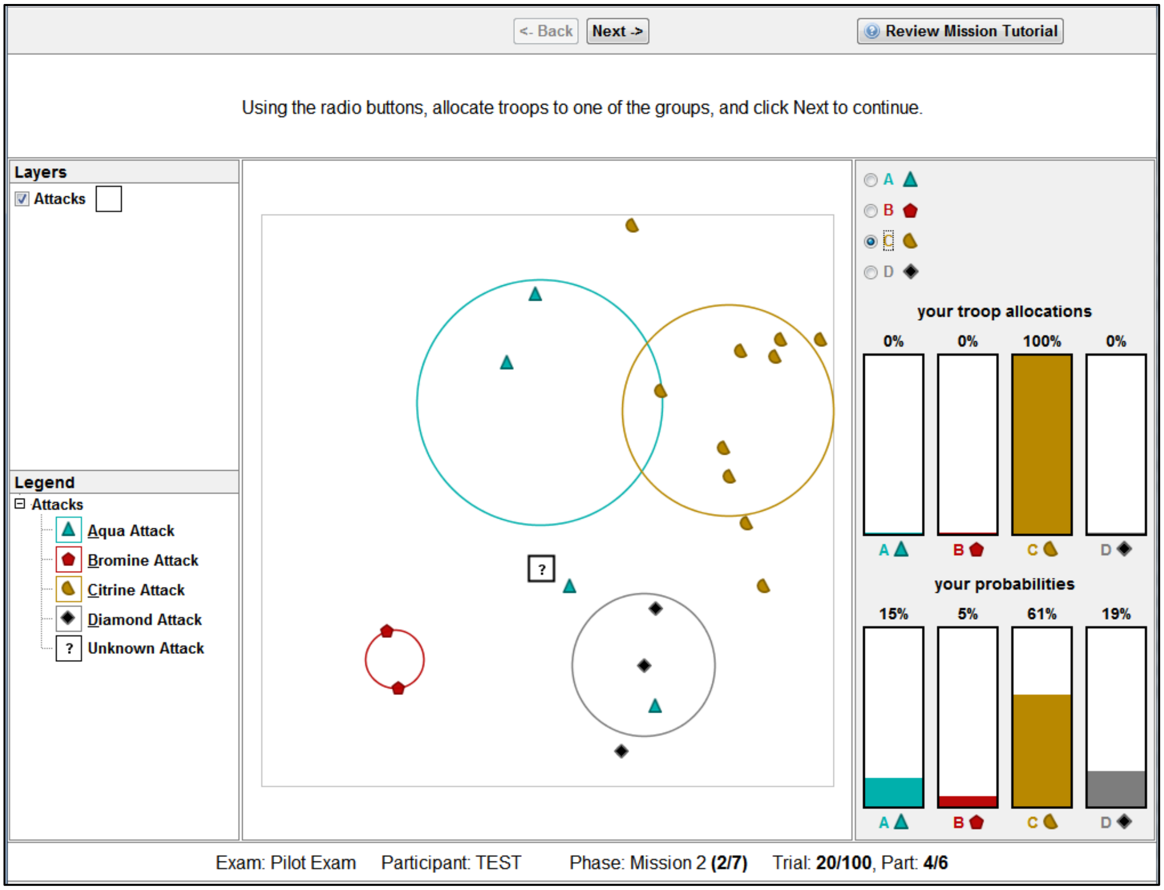

Figure 2: Participants generated likelihood that a probe event (denoted by the "?") was produced by each category, and then allocated resources to maximize their trial score. In addition, participants drew a 2-to-1 boundary for each category whose boundary encapsulates 2/3 of that category's events and whose center represented the center of activity for that category.

When group centers were generated directly from a retrieval of events represented in memory, the blended retrieval process in ACT-R reflected a disproportionate influence of the most recent events given their higher base-level activation. A strategy to combat this recency bias consisted of generating a final response by performing a blended retrieval over all the group centers (both current and past centers generated for previous trials) stored in memory, thereby giving more weight to earlier events by compounding the influence of earlier centers over the subsequent blended retrievals. This second-order blended retrieval is done for each category across their prior existing centers, which we refer to as the generation of a *centroid-of-centroids*. This effectively implements an anchoring-and-adjustment process where each new estimate is a combination of the previous ones together with the new evidence.

A fundamental difference with traditional implementation of anchoring-and-adjustment heuristics is that this process is entirely constrained by the architectural mechanisms (especially blending) and does not involve any additional degrees of freedom. Moreover, because there are an equal number of centroid-of-centroid chunks (one per category created after each trial), there is no effect of base-rate on the model's later probability judgments, even though the base-rate for each category is implicitly available in the model based on the number of recallable events. This illustrates the metacognitive nature of heuristics in our tripartite organization: given that the nature of cognitive mechanisms gives rise to a recency bias that is incompatible with the task environment (assuming a stable distribution), the centroid-of-centroid heuristic is used to give more weight to older instances and circumvent the recency bias. Note that the bias toward recency in architectural mechanisms arose because it indeed reflects the nature of many environments (Anderson & Schooler, 1991), making it well adapted to those settings. There is no such thing as suboptimal bias: just a mismatch between assumptions and environment that occasionally needs to be supplemented with the proper heuristic adjustment.

Biases in Probability Adjustment

In this task component, event features – such as the location or context of events – were presented in sequential layers on the display (see Figure 3). Initial distributions for each category were provided to participants, after which participants updated their beliefs after each feature was revealed. Beliefs were updated based on a set of provided probabilistic decision rules (e.g., if the MOVINT feature shows dense traffic, then groups A and C are four times as likely as groups B and D). When all the layers were presented, participants were required to allocate resources to each category.

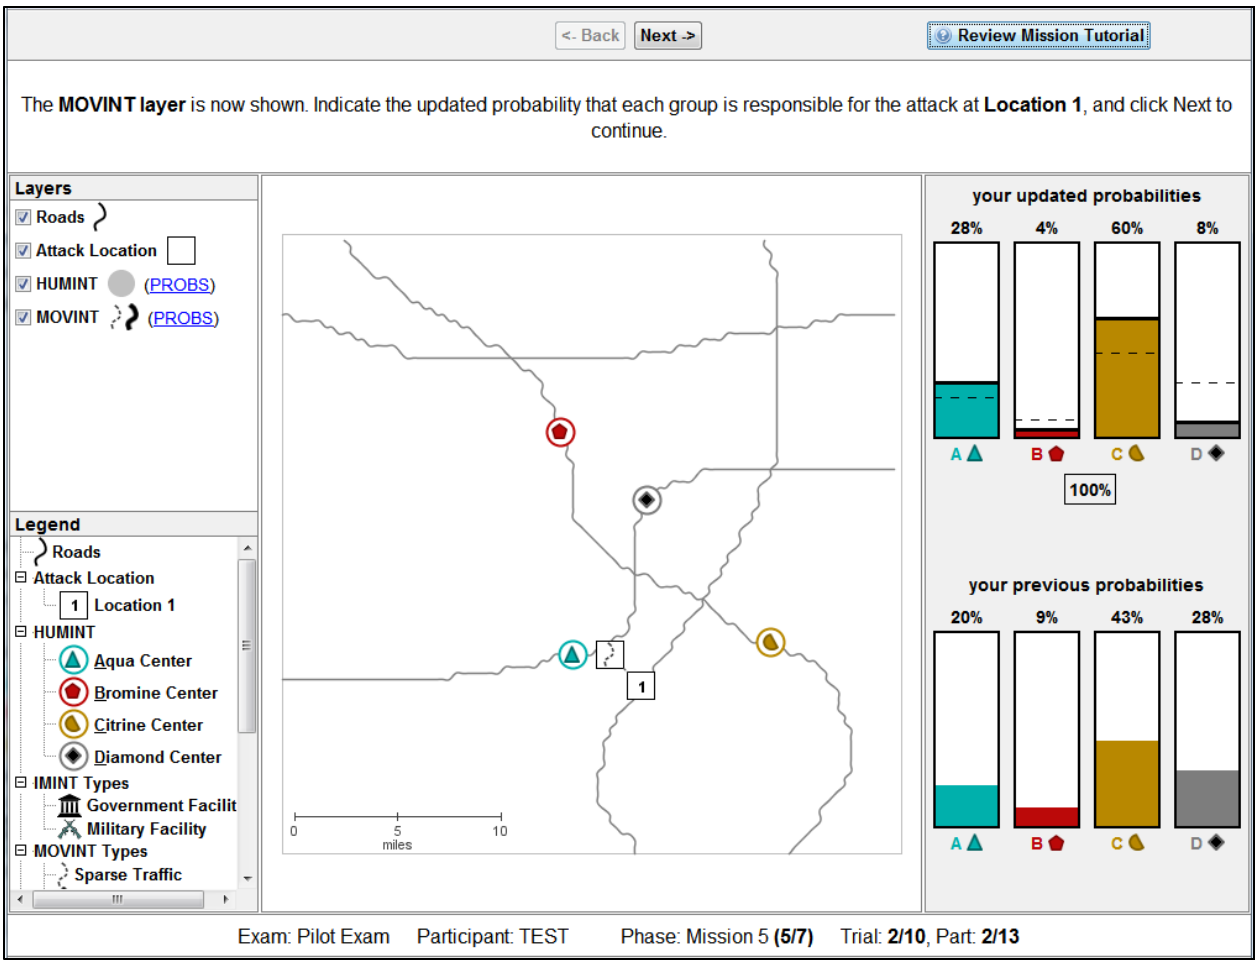

Figure 3: Participants generated the likelihood that a probe event (denoted by the probe event "1") was produced by each category. Participants updated their likelihoods as new features were revealed.

To leverage an IBL approach for probability adjustment, the ACT-R model's memory was seeded with a range of instances consisting of triplets: an initial probability, an adjustment factor, and the resulting probability. The factor is set by the explicit rules of the task. When the model is asked to estimate the resulting probability for a given prior and multiplying factor, it simply performs a blended retrieval specifying prior and factor, and then outputs the posterior probability that represents the blended consensus of the seeded chunks.

When provided with linear similarities between probabilities (and factors), the primary effect is an underestimation of the adjusted probability for much of the initial probability range (i.e., an anchoring bias), with an overestimate on the lower end of the range (i.e., confirmation bias). While the magnitude of the biases can be modulated somewhat by architectural parameters, the effects themselves are *a priori* predictions of the architecture, in particular its theoretical constraints on memory retrieval.

A simpler and more *implicit* model of probability adjustment can be produced by representing the various hypotheses as chunks in memory and using their activation as an estimate of their strength of support. When evidence is received, it is matched against patterns linking it to various hypotheses and the best matching one is retrieved, leading to a boost in activation. If contradictory evidence starts accumulating, two biases will emerge. First, new evidence will sometimes be misinterpreted because the current dominant hypothesis is most active and can overcome some degree of mismatch. Second, even if the evidence is correctly interpreted and the correct hypothesis reinforced, for the new hypothesis to attain primacy it will take some time to sufficiently build activation and for the activation of the previously dominant hypothesis to sufficiently decay over time. This process has been given a number of names, from anchoring bias to persistence of discredited evidence.

A number of structured analytic techniques have been proposed to remedy these biases emerging from the dynamics of our cognitive system (Heuer & McPherson, 2010). The most prominent one might be Analysis of Competing Hypotheses (ACH), which proposes a process by which all competing hypotheses are evaluated against each piece of evidence and the sum of their support only computed and compared at the end. This is done to prevent the early emergence of a favored hypothesis and the resulting biases emerging. An analogue to ACH has been implemented in our model and can be shown to directly affect the activation dynamics described above. Each hypothesis chunk receives a rehearsal at each step, equalizing the influence of base-rate from their activation and preventing a winner-take-all dynamic. The result is that their activation over time will simply reflect the degree of support that they have received. In this example, structured analytic techniques can also be seen as metacognitive heuristics that leverage the beneficial aspects of cognitive mechanisms while defeating or at least limiting their potential biases and thus provide external aids to our intuitive decision making.

Biases in Resource Allocation

Resource allocation makes use of the same IBL paradigm as probability adjustment. This unified mechanism has no explicit strategies, but instead learns to allocate resources according to the outcomes of past decisions. The model generates a resource allocation distribution by focusing on the leading category and determining how many resources to allocate to that category. The remaining resources are divided amongst the remaining three categories in proportion to their assigned probabilities. Representation of a trial instance consists of three parts: a decision context (i.e., the probability of the leading category), the decision itself (i.e., the resource allocation to the leading category), and the outcome of the decision (i.e., the payoff).

The model's control logic takes a hybrid approach between choice (Lebiere & Anderson, 2011) and decision models (Wallach & Lebiere, 2003b), involving two steps of access to experiences in declarative memory rather than a single one. When determining how many resources to apply to the lead category, the model initially has only the probability assigned to that category. The first step is done by performing a blended retrieval on chunks representing past resource allocation decisions using the probability as a cue. The outcome value of the retrieved chunk is the expected outcome for the trial. The second step is to generate the decision that most likely leads to that outcome given the context. Note that this process is not guaranteed to generate optimal decisions, and indeed people do not. Rather, it represents a parsimonious way to leverage our memory of past decisions in this paradigm that still provides functional behavior. A significant theoretical achievement of our approach is that it unifies control models and choice models in a single decision-making paradigm.

After feedback is received, the model learns a resource allocation decision chunk that associates the leading category probability, the quantity of resources assigned to the leading category, and the actual outcome of the trial (i.e., the resource allocation score for that trial). Additionally, up to two counterfactual chunks are committed to declarative memory. The counterfactuals represent what would have happened if a winner-take-all resource assignment had been applied, and what would have happened if a pure probability-matched resource assignment (i.e., using the same values as the final probabilities) had been applied. The actual nature of the counterfactual assignments is not important; what is essential is to give the model a broad enough set of experience representing not only the choices made but also those that could have been made. The use of a counterfactual strategy to generate a diversity of outcomes, experienced or imagined, can be seen as a very general and effective metacognitive heuristic.

The advantage of this approach is that the model is not forced to choose between a discrete set of strategies such as winner-take-all or probability matching; rather, various strategies could emerge from instance-based learning. By priming the model with the winner-take-all and probability matching strategies (essentially the boundary conditions), it is possible for the model to learn any strategy in between them, such as a tendency to more heavily weigh the leading candidate, or even suboptimal strategies such as choosing 25% for each of the four categories (assuring a score of 25 on the trial) if the model is unlucky enough to receive enough negative feedback so as to encourage risk aversion. IBL can thus be seen in this instance as a highly flexible metacognitive strategy from which a number of more limited, hardwired strategies can emerge.

Discussion

So far, we've argued that cognitive architectures aid in the study of intuitive decision-making by providing a falsifiable methodology for the study of entities, mechanisms, and variables involved in decision-making. By using a cognitive architecture, one is adopting constraints involved in managing the flow of knowledge and processes involved in these knowledge operations. Architectures expand our ability to go beyond 'just-so' explanations to describe the underlying processes and knowledge leading up to decisions. They also provide more flexibility beyond the constraints of expert systems when you break out of very constrained and/or very stable environments. In many cases, models based in a cognitive architecture can perform tasks and provide testable predictions that are too complex to analyze with traditional experimental methods due to the combinatorics of possible decisions.

The next step in the development of cognitive architectures should be mechanisms to support the generalizability of models and reduce degrees of freedom. Some preliminary thrusts include: the integration of neurally plausible associative learning to drive implicit statistical learning of regularities within the environment (Thomson & Lebiere, 2013), the development of expectation-driven cognition to cue episodic memory formation (Kurup, Stentz, Hebert, & Lebiere, 2013), and more generally the development of metacognitive awareness within the architecture to guide the selection of features used in the representation and retrieval of instances (Reitter, 2010; Reitter, Juvina, Stocco, & Lebiere, 2010; Lebiere, Gonzalez, & Warwick, 2009b).

Ideally, the strategies and heuristics implemented in the architecture should be selected (if not created) by the model itself rather than provided by the modeler. The model (driven by the architecture) should be responsible for the selection and evolution of strategies. To get started, however, several general procedures are needed to bootstrap learning until sufficient knowledge is learned, at which point processes implicated in generating expertise should lead to interesting emergent behaviors (and novel predictions) within the model. The question of which minimal set of procedures best captures human performance is an empirical one, and one that needs to be a center of focus. The adoption of general frameworks such as IBL and the adoption of a common set of heuristics across tasks appear to be the next step in the right direction.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.

Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review, 89,* 369-406.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of mind. *Psychological Review, 11*(4), 1036-1060.

Anderson, J. R., & Lebiere, C. J. (Eds.). (1998). *The atomic components of thought*. Psychology Press.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory.

Psychological Science, 2, 396-408.

Bowers, J. S., & Davis, C. J. (2012). Bayesian Just-So Stories in Psychology and Neuroscience. *Psychological Bulletin, 138 (3)* 389-414.

Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review, 118* (2), 193.

Cooper, R. P. (2007). The role of falsification in the development of cognitive architectures: Insights from a Lakatosian analysis. *Cognitive science, 31* (3), 509-533.

Dimon, C. M., Marewski, J. N., & Schooler. L. J. (2013). Constraining ACT-R Models of Decision Strategies: An Experimental Paradigm. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp .2201-2206). Austin, TX: Cognitive Science Society.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General,* 127, 107–140.

Heuer, R. J., & Pherson, R. H. (2010). *Structured analytic techniques for intelligence analysis*. SAGE.

Klein, G. (2009). *Streetlights and shadows: Searching for the keys to adaptive decision making*. The MIT Press.

Gagliardi, F. (2011). Instance-based classifiers applied to medical databases: Diagnosis and knowledge extraction. *Artificial Intelligence in Medicine*, 52 (3), 123–139.

Gibson, J. J. (1977). The concept of affordances. *Perceiving, acting, and knowing*, 67-82.

Gigerenzer, G., & Selten, R. (Eds.). (2002). *Bounded rationality: The adaptative toolbox*. MIT Press.

Gonzalez, C., & Lebiere, C. (2013). Cognitive architectures combine formal and heuristic approaches. *Behavioral and Brain Sciences*, *36* (3), 285-286.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*, 591-635.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: a failure to disagree. *American Psychologist*, *64* (6), 515.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 103 (3), 582-591.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 263-291.

Kennedy, W., & Patterson, R. (2012). Modeling intuitive decision making in ACT-R. In *Proceedings of the Eleventh International Conference on Computational Modeling* (pp. 1-6).

Lebiere, C., & Anderson, J. R. (2011). Cognitive Constraints on Decision Making under Uncertainty. *Frontiers in Cognition: Psychology*, *2*, 305-309.

Lebiere, C., & Best, B. J. (2009). From Microcognition to Macrocognition: architectural support for asymmetric adversarial behavior. *Journal of Cognitive Engineering and Decision Making* 3(2), 176-193.

Lebiere, C., Gonzalez, C., & Warwick, W. (2009). Convergence and constraints revealed in a qualitative model comparison. *Journal of Cognitive Engineering and Decision Making*, *3* (2), 131-155.

Lebiere, C., Gonzalez, C., & Martin, M. (2007). Instance-based decision making model of repeated binary choice. In *Proceedings of the 8th International Conference on Cognitive Modeling*. Ann Arbor, Michigan, USA.

Lebiere, C., Gray, R., Salvucci, D., & West, R. (2003). Choice and learning under uncertainty: A case study in baseball batting. In *Proceedings of the 25th Annual Meeting of the Cognitive Science Society* (pp. 704-709).

Lebiere, C., Pirolli, P., Thomson, R., Paik, J., Rutledge-Taylor, M., Staszewski, J., & Anderson, J. R. (2013). A Functional Model of Sensemaking in a Neurocognitive Architecture. *Computational Intelligence and Neuroscience.*

Lebiere, C., & Wallach, D. (2001). Sequence learning in the ACT-R cognitive architecture: Empirical analysis of a hybrid model. In *Sequence learning* (pp. 188-212). Springer Berlin Heidelberg.

Lebiere, C., Wallach, D., & Taatgen, N. (1998). Implicit and explicit learning in ACT-R. In *Second European Conference on Cognitive Modelling.*

Lewandowsky, S. (1993). The rewards and hazards of computer simulations. *Psychological Science*, 4. 236-243.

Marewski, J. N., Gaissmaier, W., Schooler, L. K., Goldstein, D., & Gigerenzer, G. (2010). From recognition to decisions: Extending and testing recognition-based models for multialternative inference. *Psychonomic Bulletin & Review,* 3, 287-309.

Marewski, J. N., & Melhorn, K. (2011). Using the ACT-R architecture to specify 39 quantitative process models of decision-making. *Judgment and Decision Making,* 6 (6), 439-519.

Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review, 118* (3), 393.

Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 14). Englewood Cliffs, NJ: Prentice-Hall.

Reitter, D. (2010) Metacognition and multiple strategies in a cognitive model of online control. *Journal of Artificial General Intelligence*, 2(2):20-37.

Reitter, D., Juvina, I., Stocco, A., & Lebiere, C. (2010). Resistance is Futile: Winning Lemonade Market Share through Metacognitive Reasoning in a Three-Agent Cooperative Game. In *Proceedings of the Behavior Representation In Modeling and Simulations (BRIMS 2010) Conference*. Charleston, SC.

Simon, H. A. (1995). Rationality in political behavior. *Political Psychology,* 45-61.

Simon, H. A. (1992). What is an “explanation” of behavior?. *Psychological Science*, *3*(3), 150-161.

Simon, H. A. (1990). *Reason in human affairs*. Stanford University Press.

Schultheis, H. (2009). Computation and explanatory power of cognitive architectures: The case of ACT-R. In *Proceedings of the 9th International Conference of Cognitive Modeling*.

Shallice, T., & Cooper, R. (2011). *The organisation of mind*. Oxford University Press.

Taatgen, N., & Anderson, J. R. (2010). The past, present, and future of cognitive architectures. *Topics in Cognitive Science*, *2* (4), 693-704.

Taatgen, N. A., & Anderson, J. R. (2008). Constraints in cognitive architectures. *Cambridge handbook of computational psychology*, 170-185.

Taatgen, N. A., Lebiere, C., & Anderson, J. R. (2006). Modeling paradigms in ACT-R. *Cognition and multi-agent interaction: From cognitive modeling to social simulation*, 29-52.

Thomson, R., & Lebiere, C. (2013). Constraining Bayesian inference with cognitive architectures: An updated associative learning mechanism in ACT-R. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society.* Austin, TX: Cognitive Science Society.

Todd, P. M., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, 24, 143-165.

Veksler, V. D., Gray, W. D., & Schoelles, M. J. (2009). Goal-Proximity Decision Making: Who needs reward anyway. In *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*

Wallach, D., & Lebiere, C. (2003a). Implicit and explicit learning in a unified architecture of cognition. *Advances in Consciousness Research*, *48*, 215-252.

Wallach, D. & Lebiere, C. (2003b). Conscious and unconscious knowledge: Mapping to the symbolic and subsymbolic levels of a hybrid architecture. In Jimenez, L. (Ed.) Attention and Implicit Learning. Amsterdam, Netherlands: John Benjamin Publishing Company.

Wallach, D., & Lebiere, C. (2000). Learning of event sequences: An architectural approach. In *Proceedings of the 3rd International Conference on Cognitive Modeling*.

Weick, K. (1995). *Sensemaking in organizations*. Thousand Oaks, CA: Sage.

West, R. L., & Lebiere, C. (2001). Simple games as dynamic, coupled systems: Randomness and other emergent properties. *Journal of Cognitive Systems Research*, 1(4), 221-239.