

***Knowns, Known Unknowns, and Unknown Unknowns:
Time and Uncertainty in Naturalistic Decision Making***

Marvin S. Cohen, Ph.D.
Perceptronics Solutions, Inc.

The longer I live the more I see that I am never wrong about anything, and that all the pains I have so humbly taken to verify my notions have only wasted my time.

George Bernard Shaw (Laurence, 1985)

TIME MANAGEMENT

This chapter focuses on a critical issue in NDM research – experts’ management of the time they take to gather information and verify assessments and decisions – and its implications for training and decision aiding from an NDM perspective. There has been very little research on how, or even if, proficient decisions makers in various domains and task contexts actively manage the time they take to think, and if they do, whether it is a distinctive part of their expertise. This verged on being a non-problem in early studies of expert problem solving and naturalistic decision making, because performance superiority was attributed to rapid recognition of familiar patterns [e.g., Chase & Simon, 1973; Recognition Primed Decision Making (RPD) level 1, Klein, 1993]. Time management might seem *unnecessary* because the first solutions generated by experts tended to be accurate. Time management might be *impossible* because the relevant processes were largely preconscious.

Time management becomes critical if experts excel along dimensions other than direct recognition, because when decision makers mismanage time, there can be serious consequences. Peters, Jackson, Phillips, and Ross (2008, p. 205), observing a military exercise featuring new displays, noted that “commanders commonly sacrificed the speed advantage of their lightly armored force in order to satisfy their perceived need for information.” Omodei, McLennan, Elliott, Wearing, and Clancy, (2005) found experimental support for overuse of unfamiliar information resources by relatively experienced decision makers. General George C. McClellan

exemplified this syndrome during the American Civil War while others, such as Grant and Jackson, warned against it, nicely demonstrating that the phenomenon cannot be blamed entirely on modern technology. Deliberative processes have in fact been acknowledged by NDM researchers: mental simulation (Klein, 1993, 1998), deliberative top-down processes (Endsley, 2000a), critiquing and correcting products of recognition (Cohen, Freeman, & Wolf, 1996; Cohen, Freeman, & Thompson, 1998; Cohen & Thompson, 2001), and cyclical frame-data interactions (Klein, Phillips, Rall, & Peluso, 2007). Not surprisingly, all these researchers emphasize the importance of timely action. Klein's proposed criterion for stopping deliberation refers to the *benefits* of reducing *known* uncertainties – for example, sensemaking stops when it has consistently fleshed out a frame (Klein, et al., 2007) or when affordances are recognized (Klein, 2000). Endsley (2000a) intentionally abstracts from real-world time constraints when she states that “more SA is always better.” I have emphasized balancing *costs of delay* against benefits, and have included in the latter both resolving known and discovering *unknown* uncertainties in novel situations.

The assigned mission in writing this chapter was to identify new directions for Naturalistic Decision Making research. My argument will be that the management of time and uncertainty in decision making serves this purpose. Moreover, to deal with it effectively, a number of habitual assumptions in NDM or cognitive psychology will have to be re-examined – for example, that attention is voluntarily controlled by a “central executive” (for a critique, see Monsell & Driver, 2000); that voluntary processes cannot be explained “mechanistically” (for a critique, see Bargh, 2000; Newell, 1981); and that “units” of knowledge such as schemas, frames, and mental models are well-defined in advance of their “use” by decision makers (Rumelhart, Smolensky, McClelland, & Hinton, 1986). More specific temptations include: treating macrocognitive

phenomena – such as story-building, developing mental models, or handling uncertainty (Schraagen, Klein, & Hoffman, 2008) – as primitives that explain rather than merely name the phenomena (for a critique, see Flach, 2008), treating “intuitive” decision making and critical thinking as diametric opposites (e.g., assuming that the former involves no inference while the latter is linear and logical), downplaying preference and motivation as determinants of action, and assuming that differences in thinking strategies do not contribute to expertise.

Recognition / Metacognition (R/M) theory posits three concurrent and iterative cycles of cognitive activity, as shown in Figure 1. (1) In the most basic cycle (left side of Figure 1), existing knowledge and preferences interact intimately to interpret the environment and generate actions to influence it. R/M treats recognition-primed decision making as an *anytime* process, i.e., it produces actionable results very quickly, which may then be improved by investment of additional time. (2) A second cycle (inner loop on the right side), conditional on investing time in deliberation, incorporates core recognitional conclusions into a story that explains missing or conflicting evidence, evaluates assumptions required to make the story work, and moves to another core conclusion only if the story requires too much of a stretch. Since critical thinking may be motivated either by *known* uncertainty or by the prospect of *unknown* uncertainty, critiquing and correcting not only answer questions but proactively search for questions to ask. (3) A final cycle (outer loop on the right side) determines how much time, if any, to allocate to deliberation by balancing the potential for improved outcomes against risk and foregone opportunities. The chapter’s main focus is on the third cycle. Instead of the usual non-explanatory placeholders, e.g., an “active agent,” higher-level control is executed by an automatic process called *melioration* (Herrnstein, 1997) that also applies to animal foraging. Melioration can lead to premature commitment or to inappropriate delay under predictable

conditions, but improves with the right kind of experience. The chapter describes experiments that tested and confirmed the hypothesis that experts are reliably better at regulating time than non-experts. I conclude that time allocation affords important overlooked opportunities for naturalistic aiding and training.

NATURALISTIC DECISION MAKING IS AN ANYTIME PROCESS

An anytime algorithm in artificial intelligence is a computational process that does not need to run to completion in order to provide useful information. The quality of the answer increases with additional time spent computing (in practical reasoning about complex domains, there is no stopping point beyond which improvements are demonstrably no longer possible), but the current best solution can be obtained and used at “any time.” The anytime principle plays an important role in Shruti R/M, a connectionist implementation that contends for what might be called the triple crown of simulated “automatic” or recognition-based reasoning: (1) Shruti is *relational* (Shastri & Ajjanagadde, 1993; Shastri & Mani, 1997). Unlike standard probabilistic systems (i.e., Bayes nets and influence diagrams, which represent events by propositions that lack internal structure) or associative networks (which represent events and objects by conjunctions of features), Shruti reasons with predicates or frames (e.g., *x attacks y at time t*), and uses temporal synchrony of activation to track the objects that fill roles in the relations that predicates represent. (2) Shruti is *probabilistic* (Wendelken & Shastri, 2000). Unlike standard systems for relational reasoning, it propagates uncertainty (i.e., strength of belief) across a network of causally and semantically linked predicates, inferring causal explanations from dynamically changing observations and episodic memories, and predicting future effects from combinations of causes and conditions (Shastri & Wendelken, 1998). More recent enhancements integrate reasoning about belief with reasoning about preferences and actions (Wendelken &

Shastri, 2002; Cohen et al., 2000b). (3) Shruti is extremely *fast*. Unlike both probabilistic and relational systems, Shruti draws conclusions from massive knowledge bases in a second or less. In this respect as well as others, it provides a neurally plausible representation of rapid *non-deliberative* human reasoning (Shastri & Mani, 1997; Shastri, 2001).

Shruti R/M explains the early availability of a solution (i.e., the first part of the anytime principle) as the rapid emergence of stable, self-reinforcing cycles of activation in a long-term memory associative network (Shastri & Ajjanagadde, 1993; Shastri & Mani, 1997). In the case of highly experienced decision makers, these cycles of activation tend to include key elements of both a situation model and plan, which can be acted on immediately under conditions that are highly time-constrained or routine. Figure 2 illustrates some notable features of this process. It depicts the hypothesized spread of activation through an associative network for the Tactical Action Officer (TAO) of a US Aegis cruiser during an incident in which a Libyan gunboat approached his ship in the Gulf of Sidra.¹

In this example, evidence includes recent observations, now in episodic memory, that *a gunboat emerged from a Libyan port, turned toward own ship, and increased speed* (the latter is not shown). Activation spreads up causal chains via query nodes (represented by question marks) to find explanations (*the Libyan gunboat intends to attack own ship*) and explanations of explanations (*Libyans have a motive, the opportunity and means to attack*). In the case of *motive*, it spreads even further to causal preconditions that were already primed (*the gunboat is Libyan, and own ship is in Libyan claimed waters*). When activation reaches facts in episodic memory or perceptual observations, they serve as pathways for activation to flow back to the original predicate's belief collectors (represented by plus and minus signs for truth and falsity,

¹ This example and other critical incident interviews that motivated the R/M theory are analyzed in more detail in Cohen et al. (1996), Cohen et al. (1998), and Cohen, Thompson, Adelman, Bresnick, Tolcott, et al. (1995).

respectively). Simultaneously, activation spreads down causal chains to likely future events: *own ship is within range of the gunboat's weapons*; *the gunboat fires a missile at own ship*; and *own ship is damaged*, a third step that was already primed by an active goal. Preference spreads in parallel with belief, so that when a desirable or undesirable event is predicted (such as *own ship is damaged*), activation travels through the pathway created by a *reward fact* to a positive or negative utility collector (represented by dollar signs followed by a plus or minus sign, signifying preference for the predicate's truth or falsity, respectively). Activation then travels upstream via utility collectors to causes of the predicate, seeking actions that can influence its state for the better. If facts representing feasible actions (e.g., *own ship fires at the gunboat before time t*, at which own ship will be within the gunboat's weapons range) are encountered and not inhibited, activation returns over the new pathways provided by performance of the actions, shutting down the circuit generated by the reward fact. In sum, a single very rapid process of parallel constraint satisfaction enables the TAO to explain recent and on-going observations, anticipate the future evolution of events, and find actions to deal with them. There are no distinct "processing stages" from observations to interpretations to mental simulation to decisions. They emerge together in a process of mutual interaction that does not require a fixed starting point.

Spreading activation in the Shruti R/M associative network poses questions and attempts to answer them. In Shruti R/M *questions* are predicate instances whose query nodes have been activated by goals, perceptual observations, or focal attention. *Answers* are the activation returning to the predicate instance regarding its truth/falsity or desirability/undesirability. Activation returns via directly associated perceptual, episodic, or reward facts or indirectly via causally or semantically linked predicates that are directly associated with such facts. Stable

situation pictures and plans emerge over very short periods of time if cascading activation returns over multiple fact-created pathways to the same belief collectors of the originally queried predicates. The larger the number of confirming facts, the stronger the returning activation, shutting down weaker pathways corresponding to support for competing causal hypotheses, inhibiting competing goals or actions,² and creating a stable self-reinforcing cycle of activated beliefs and intentions. *Knowns* (as in the chapter title) are questions that have received an (as yet) unchallenged answer.

With time and experience, an associative network adapts to the statistical properties of the environment (Brunswik, 1955). A natural outcome of learning is the emergence of clusters of predicates –such as the *hostile intent* “mental model” in Figure 2—whose activation levels covary because of their links in long-term memory, which reflect the experienced, reported, or inferred co-occurrences of corresponding events (Wendelken & Shastri, 2003). A second natural outcome is the emergence of causal structure among the correlated variables, which both simplifies relationships (Pearl, 1989) and enables effective action (Pearl, 2000). Clusters of this kind are often centered on a structurally definable *core conclusion*. A core conclusion is the current state of a predicate that has a unique organizing and simplifying role in an active network and a significant influence on utility and actions. *Intends to attack* satisfies both conditions in Figure 2.³ Notwithstanding internal structure, from a connectionist and naturalistic perspective large knowledge structures (e.g., *schemata*, *frames*, or *mental models*) are not fixed or well-defined entities. Each predicate in active memory is typically connected to many other concepts,

² For example, strong activation in favor of *own ship fires at gunboat* suppresses the general reluctance to fire at another ship reflected in the link of the *average utility fact* to the negative utility collector (\$-).

³ *Intends to attack* is the most central predicate in Figure 2 according to several network measures. It has more incoming ties (i.e., causal precursors like *motive*, *means*, and *opportunity*) and more outgoing ties (i.e., observable actions like *turning toward own ship* and *approaching* undertaken to implement intent). In addition, it mediates a larger proportion of the connections between other predicates. These properties reduce the number of connections required to represent the situation.

both active and inactive, which themselves belong to other loosely defined overlapping correlational and causal structures. What makes a situation seem routine or familiar is not its match to a self-contained “pattern,” but the conformity of the elements that happen to be recognitionally activated at a given time to pre-existing correlations among those elements in long-term memory, enabling activation to rapidly spread and cohere. When situation interpretation and action commitment are accomplished without intervention by deliberate attention, they qualify as *automatic*, as defined by Logan (1988): they involve *single step memory retrieval*. Because experts have more information and have organized it more effectively than non-experts, their initial single-step recognition of a situation is more likely to be correct (Feltovich, Spiro, & Coulson, 2006), but not guaranteed.

CRITIQUING AND CORRECTING

Shruti R/M accounts for the second property of anytime algorithms – that continued processing can improve results – by the use of skilled attention to *prolong* and *extend* the same types of activation processes as those described in the last section. Successive shifts of attention widen the sphere of activation and bring new knowledge to bear over multiple retrieval cycles. Attention shifting, unlike general-purpose analytical strategies, draws its strength from the same domain-specific knowledge that fuels recognition. It is effective when it stimulates activated regions in long-term memory to expand or coalesce into more coherent and inclusive solutions based on more knowledge (the inner loop on the right in Figure 1). Shruti R/M represents a nuanced view of automaticity, according to which processes like spreading activation can be both initiated and guided by deliberate attention, and deliberate attention in turn is influenced by recognized patterns of spreading activation (Bargh, 2007; Logan, 2004; Logan, Taylor, & Etherton, 1996; Wendelken & Shastri, 2005). Success, therefore, depends not only on

substantive knowledge but also on the metacognitive skills that monitor and regulate this reciprocal interaction (e.g., Dunlosky & Bjork, 2008). The purpose of the following sections is to provide a framework for understanding the skill component of deliberation. This section reviews some of the behavior that is to be explained.

During critiquing and correcting, decision makers attempt to *hold onto* initial recognitional conclusions while at the same time deepening understanding. Rather than discounting or ignoring problematic evidence or prior beliefs, they reconcile evidence and beliefs with their initial conclusions by incorporating them all within a consistent story in episodic memory, ultimately replacing initial core conclusions only if the story requires more assumptions than expected. Iterative application of a small set of critiquing and correcting skills accomplishes this:

1. The most basic critiquing skill is finding *known unknowns* in the results of recognitional processing. These are questions that have been asked but not answered, i.e., queried predicates whose returning activation is weak, conflicting, or unstable. In the Libyan gunboat example, the captain of the US Aegis cruiser noticed that a causal condition of intent to attack own ship – that own ship was an appropriate *opportunity for attack* based on its being in Libyan claimed waters – was unstable due to the unexamined assumption that a *better* target was not available.
2. The most basic correcting skill is *shifting attention* to known unknowns in order to expand the scope of activation in the associative network. The aim is to answer relevant questions by activating perceptual observations, episodic memories of observations, or well-established general beliefs. For instance the captain shifted attention directly to *opportunity to attack* and noticed that another US cruiser was closer to the gunboat. The gunboat had bypassed a better target, casting doubt on its *hostile intent*.

3. If shifting attention does not resolve known uncertainty, correcting includes *clamping* an active predicate (e.g., own ship is a suitable *opportunity for attack*) as a temporary fact in episodic memory. The aim is to bias information retrieval toward evidence supporting the clamped hypothesis, thereby reconciling the core conclusion with observations and prior beliefs. For instance, the best explanation the Captain could come up with was that the Libyans were taking care of the other US cruiser in some other way.
4. If direct evidence is not found to strengthen and stabilize conclusions and explain conflict, correcting includes the retention of clamped, reconciling hypotheses as *explicit assumptions* in episodic memory. The result is a stable and coherent story that consistently combines the core conclusion, observations, and relevant beliefs. To reconcile their assessment of hostile intent with the presence of the other US vessel, the captain and the TAO hypothesized that Libyans would take care of the other ship in some other way. A similar cycle of critiquing and correcting resulted in the assumption that the gunboat was a suitable *means of attack* against the cruiser despite the obvious difference in capabilities, because the Libyans were throwing everything they had into the fight.
5. Critiquing includes judgments about the likelihood of significant *unknown unknowns*, i.e., hidden issues which if not managed, might produce surprises down the road. Experts implicitly assess the likelihood of questions whose relevance has not yet been recognized based on the overall unfamiliarity of the situation and the number of assumptions necessary to save the core conclusion.
6. If unknown unknowns are deemed likely, critiquing expands to search for them even in the absence of known unknowns. Unknown unknowns may be exposed as by-products of shifting attention and expanded scope of activation, as in step 2. They may be

intentionally sought out by shifting attention to apparent *knowns*, i.e., observations or beliefs that apparently support the core conclusion, to find weak evidence, conflicting beliefs, or unstable implicit assumptions underlying the favorable interpretation of the evidence. For instance, as the number of required assumptions increased, the situation seemed more problematic to the Captain, and he shifted attention to evidence that seemed to support *hostile intent*, the gunboat's *turning toward own ship*. As a result, the captain recalled that the gunboat was unable to localize own ship at the relevant range. To maintain the hostile intent story, some additional assumptions, such as possible Soviet technical training, or operational assistance, were necessary.

7. Critiquing includes tracking the *overall plausibility* of the set of explicit assumptions in the story. Because residual uncertainties are deliberately translated into assumptions, their collective plausibility serves as a uniform currency for evaluating the story as a whole, and the plausibility of the story as a whole is a basis for judging the plausibility of the core conclusions. It took three assumptions, none of them very compelling, to patch up the *hostile intent* story: that an unknown Libyan plan included taking care of the other cruiser, that the Libyans were willing to sacrifice the gunboat, and that the gunboat was able to localize own ship at the relevant distance before it turned.
8. If the total set of assumptions is regarded as too implausible, correcting may release them and adopt an alternative core conclusion. The captain finally dropped the *hostile intent* story and took a serious look at another core conclusion, that the gunboat was *not hostile*. After several more cycles of critiquing and correcting, a new, *non-hostile* story emerged which also required several assumptions.
9. *Priming* (Shastri & Wendelken, 1999) complements attention shifting by prolonging

activation of information that might otherwise have been left outside the radius of activation. Newly activated information can be integrated with information maintained by priming to facilitate new results. A better story or core conclusion may emerge from the spontaneous convergence of ideas primed over repeated rounds of critiquing and correcting. For example, in the gunboat incident the captain and TAO ultimately concluded that the Libyan gunboat was probably looking for targets of opportunity. The *target of opportunity* story required only one unsupported assumption: that *Libya is willing to sacrifice the gunboat*. This became the Captain's working hypothesis, and the incident did in fact end with the gunboat's destruction.

FORAGING FOR SURPRISES

A New Model of the Quick Test

The outermost loop in the R/M model (Figure 1, right panel) regulates whether deliberative processing will take place, and if so, where it will focus and when it will stop. According to the R/M theory (Cohen et al., 1996, 1998), deliberation takes place with regard to some component of a primary task if and only if the stakes are high, the opportunity costs of deliberation (e.g., delayed action or deliberation about other topics) are relatively low, *and* the decision is either uncertain (i.e., there are known unknowns) or unfamiliar (i.e., there are possible unknown unknowns). When these conditions are met, the quick test inhibits immediate action or irreversible commitment to the relevant task component and permits the iterative process of critiquing and correcting described in the last section. When the conditions fail, critiquing and correcting halt or do not begin. Since most intentional behavior occurs in circumstance that are familiar or pose little risk, by far the most common result is action without deliberation. One should not be surprised if 80 to 90% (or even more) of "decisions" described in retrospective

critical decision interviews are rapid and recognition-primed, as reported by Klein and his colleagues. The remaining 10 to 15%, however, tend to be very consequential, and are typically the most salient parts of the incident to the decision makers themselves.

An obvious constraint is that the quick test itself not consume many cognitive resources. Unless it is truly quick, it would consume too much of the time it is meant to allocate, and the costs of delay would overwhelm any possible benefits. This is not to say that deliberation about deliberation is never appropriate or that it must lead to an infinite regress. Experts often articulate judgments about how much time they have, and effective crews take actions for the explicit purpose of buying more time before action is necessary (Orasanu & Strauch, 1994; Cohen et al., 1996; Cohen et al., 1998). But for thought to be possible at all, *the highest level of “executive control” over mental processes must be simple, automatic, and largely outside conscious awareness* (Son & Kornell, 2005).

Some insight into how a simple, automatic process can produce adaptive decisions about attention allocation comes from ethological studies of animal foraging in the wild (Stephens & Krebs, 1986; Stephens, 2007) and experimental analysis of animal choices in the laboratory (Herrnstein, 1997). Presumably without explicit deliberation, herbivores somehow “decide” how to allocate their time among patches that vary in the quantity and quality of food they contain; predators “know” how to allocate time among locations that vary in the incidence of prey. Patches may become depleted over time, are separated by distances that take time to traverse, and may expose the forager to predation. The present model of the quick test is based on an analogy between foraging in the wild and cognitive exploration of idea threads in an associative network under the risks associated with delayed action or other missed opportunities. Unlike so-called *optimal foraging theory*, the Shruti R/M *foraging for surprise* hypothesis does not purport

to guarantee an optimal result (Simon & Newell, 1958). Ironically, it provides a pathway for improvement in contexts where some “optimizing” algorithms actually do not. It utilizes implicit metacognitive assessments about knowledge, time, and uncertainty, which I discuss in turn.

Knowledge about Knowledge

While animals forage in patches to discover and consume sources of energy, decision makers forage in predicate clusters to discover and benefit from *surprises* that may lead to more successful actions and outcomes. An associative network like Figure 2 naturally decomposes into dynamically evolving “patches” of information, i.e., clusters of predicates that are both topically related to one another and whose relevance to decisions is correlated across situations. For example, if the predicate *opportunity to attack* is relevant to a course of action and its outcomes, linked information about spatio-temporal relationships with other potential targets is also likely to be relevant; if the predicate *means of attack* is relevant, so is linked information about relative platform capabilities and availability of better weapons. Patch-like organization of information is an ecological constraint to which deliberation skills may adapt through learning (Pirolli & Card, 1999; Pirolli, 2007). A patch is available for attention only if it is represented by at least one predicate in active memory, which can function as an entry point. As expertise accumulates in a domain, decisions about allocating attention to a predicate in active memory will be shaped increasingly by experience with the patches to which it belongs. This does not imply that decision makers are familiar in advance with all the predicates in a patch, or that patches are well-defined enough for such knowledge to be possible. New patch members may be identified dynamically as predicates are encountered during exploration of the associative network.

Three deliberation-related alternatives may exist any given time: implement the currently preferred action without (further) deliberation, attend to another active predicate in the currently

attended patch, or switch attention to an active predicate in a new patch. Choice is determined by the relative *values* of the available alternatives. An ideal rational model would equate value with the total expected future utility of each choice. In order to globally maximize return from deliberation, it would consider *subsequent* attention shifts, resulting discoveries, their influence on new attention shifts, and so on, to include all possible sequences and durations of deliberation across all available patches and their impact on future decisions and outcome utility. Aside from demanding far more effort than deliberation would, the “rational” approach relies entirely on the model that deliberation is meant to improve. The requirement for explicit modeling assumes the impossibility of unknown unknowns in novel situations (Cohen, Parasuraman, & Freeman, 1998; Cohen & Freeling, 1981). A simpler heuristic, the *mean value theorem* (MVT; Charnov, 1976) is the central result in optimal foraging theory and the most widely cited (e.g., Pirolli and Card, 1999; Pirolli, 2007). It also has serious limitations, especially when applied to deliberation. Stopping rules based on the MVT presuppose that the current instantaneous rate of return from a patch is known with certainty rather than learned or adjusted (McNamara, Green, & Olsson, 2006), that the rate of return always decreases with time in a patch, and that the optimal overall rate of return is already somehow known (Green, 1984).

By contrast, the foraging for surprise model utilizes a process called *melioration*, which is arguably the most psychologically natural and empirically best supported heuristic simplification (Herrnstein, 1997; Davison & McCarthy, 1988). It has very significant advantages over alternative heuristics for attention allocation. Because it bases value on the average rate of return *actually experienced over episodes in an activity*, adjusted in the light of *known unknowns* and *unfamiliarity*, it makes far more realistic assumptions about what decision makers need to know. Unlike explicit rational models and the MVT, it does not *necessarily* lead to sub-optimality when

there is no clear bound on significant information – i.e., when yet-to-be-discovered unknowns might influence decisions or returns in a patch may be positively accelerated with time.

The foraging for surprise model works as follows: The value of switching attention to an active *predicate* is proportional to its average absolute change in utility activation during complete episodes of attending to the *patch* to which the predicate provides entry. Dividing this by the average time per episode in the patch, we get a *rate of return*. When the predicate to be attended is in a new patch, the time required to switch attention is added to the denominator (time to switch attention *within* a patch is already incorporated in episode duration). The benefit of switching attention to an active predicate thus takes into account the patch of related predicates that it activates and which may as a result be further explored. (In Shruti R/M these recursive, prospective effects are captured without explicit modeling because changes in utility activation spread among the causally linked predicates in a patch.) The rate of return is updated continuously in the current episode by adding the absolute utility change thus far to the numerator and the amount of time in the patch thus far to the denominator. On-going results will have little effect if there is extensive prior experience in the patch but may be quite significant if prior experience is sparse (Olsson, 2006; Green, 1984). The updated ratio at the conclusion of the current episode becomes the new baseline. Finally, weights in Shruti R/M steer attention to predicates on the edge of the current active network because activation of average facts declines as the predicates that are averaged over become active.

Knowledge about Time

Since utility gains are seldom linear with time, a decision maker's implicit assessment of an activity's rate of gain is sensitive to the *durations* of experienced episodes over which it was implicitly measured. Some insight into how melioration interacts with selective sampling of

episode durations can be obtained by comparing the *value functions* of different activities, that is, the rate of gain that would be experienced if episodes in a patch had a particular duration. The top row of Figure 3 illustrates cumulative gains in utility as a function of time in a patch. The bottom row shows the corresponding value functions, obtained by dividing the cumulative gain (y-axis value in the top chart) by time in the patch (the corresponding x-axis value in either chart). Each chart shows two mutually incompatible activities, A and B, that divide a time interval (the length of the x-axis) between them. Because A and B are incompatible (e.g., deliberation versus immediate action, or deliberating in one patch versus deliberating in another), time spent in one equals the total time minus time spent in the other. A specific division of the interval between an episode in A and an episode in B corresponds to a point on the x-axis. Time to the left of that point was spent in activity A, after which the decision maker switched to activity B; time to the right of that point was then spent in activity B. Notice that Figure 3 plots gains in B from right to left. Since B starts when the decision maker switches out of A, this is initially counterintuitive; however, it enables us to compare the returns from both tasks for every possible allocation of time. For *any chosen time allocation between A and B*, represented as a point on the x-axis, the corresponding y-axis values for the A and B curves show how much utility will have been gained from A and from B, respectively, at the end of the time period. The dotted line is the sum, which a rational model would try to maximize.

According to R/M's melioration rule, the allocation of time among activities is governed by a simple preconscious strategy: spend more time in activities with higher rates of return and less time in activities with lower rates of return. As an example, consider a recurrent situation in which the decision maker allocates attention between patch A and patch B. On the first occasion in that situation, the decision maker switches attention out of patch A into patch B at some point

on the x-axis, and spends the rest of the time in patch B. If the result is a higher rate of return from patch A than from patch B, on the next occasion the decision maker will spend a bit more time in A and a bit less time in B; that is, according to melioration, the selected point on the x-axis will shift to the right. (Of course, it would shift left if the decision maker experienced a higher rate of return in B.) If the extra time in A does not produce a proportionate increase in the utility gained from A, its rate of return will decline. If the reduction in time does not result in a proportionate decrease in the utility gained from B, B's rate of return will increase. Repeated over occasions of decision making, these adjustments produce exclusive preference for an activity whose experienced rate of gain remains higher than other activities as more time is allocated to it, eliminate activities whose experienced rate of gain remains lower as less time is allocated to it, and in all other cases adjust time allocation until experienced rates of gain are equal across the surviving activities.

Because of melioration, decision makers do not need global knowledge of how time affects value. Value functions are ecological constraints that people adapt to by sampling specific durations (i.e., points on the horizontal axis) and learning the results. Melioration is a local adjustment process among activities rather than a rational "central executive."

The left column of Figure 3 shows that melioration works well when activities A and B both have negatively accelerated returns.⁴ The main thing to notice is the equilibrium represented by crossing lines i.e., equal rates of return, in the lower chart. The equilibrium is (1) stable, because deviations will be corrected, (2) likely to be reached by incremental adjustments regardless of the initially sampled durations of the two activities, and (3) very close to the global maximum rate of

⁴ In an associative network, the value of additional periods of deliberation in a patch ordinarily declines both as attention moves further away from decision outcomes and as other predicates in the patch become active. Declining returns may occur for other reasons as well. Marginal benefits will decrease and marginal costs rise as "low hanging fruit" is depleted. Animals become satiated with food and people with topics of thought. Additional samples have a decreasing impact on the accuracy of population estimates.

return (i.e., it gives more time to the more productive activity, B); in any case returns are flat over a wide region near the equilibrium. The same charts (left column of Figure 3) can illustrate favorable conditions for choosing among patches and for deciding whether to deliberate or act. With regard to the latter, suppose that A is deliberation that delays irreversible commitment to some action, and B represents irreversible commitment to the action at the expense of A. The value of pushing action B forward by any fixed amount of time is represented by the corresponding size of the increase in B's value from right to left. The benefit of more timeliness declines with increasing immediacy of action (i.e., proximity to the left corner), just as the benefit of more deliberation declines with time attending to patch A (i.e., proximity to the right corner). The appropriate balance will be found by incremental adjustments from virtually any starting sample along the horizontal axis. Decreasing returns are a *necessary* feature of proposed optimizing heuristics. Fu (2007) proposes that they are an "ecological invariant" in representative information seeking environments. The foraging for surprise model explains why diminishing returns promote effective performance, but it does not *assume* or *need* them.

Returns may increase over time spent deliberating when predicates in a patch depend on one another for value, like pieces of a single puzzle. This may occur because value is not realized until predicates are combined into a single story which is evaluated as a whole, or until later items have been checked to see if they overturn interpretations of earlier items, prompting the emergence of a new story. In addition, when the value of a patch is uncertain, initial results may be useful primarily as indicators of unknown unknowns, i.e., they raise previously unasked questions which can be answered by seeking other items in the patch (Olsson & Brown, 2006). In sum, positively accelerating returns distinguish foraging for ideas from foraging for food.

The rightmost column of Figure 3 shows that positively accelerating returns for an activity

(B) make adaptation more difficult. As before, there is a stable equilibrium involving both A and B which is close to the global maximum. However, there is a much narrower region of starting points from which decision makers can reach this equilibrium. Outside this region, deviations are amplified rather than corrected. As a result, decision makers' initial experiences have a large influence on the size of returns they ultimately realize. If decision makers' initial inclination favors activity A (e.g., deliberating in patch A or pushing forward some action A), episodes of attending to patch B will be short, hence unproductive; melioration will allocate even more time to A until patch B is neglected altogether (despite the fact that decision makers would do better allocating the entire interval to B than to A), The decision maker will receive a significantly lower total return in the absence of an intervention, like mentoring, that forces exposure to B for periods long enough to realize its value.⁵ If we suppose, on the other hand, that A is deliberation and B is pushing action forward in time, the opposite problem occurs. If decision makers initially deliberate too long (proximity to right corner), the costs of further delay may become vanishingly small because the damage has been done; melioration may lead them to deliberate even more on future occasions, until they are virtually incapable of acting effectively (even though they would have done better by not deliberating at all). Yet another occasion for these effects is afforded by the extra time required to switch attention from one patch to another. If deliberation starts in patch A, the interruption penalty that would be incurred by switching to B may lead decision makers to persist inordinately long in A.

All of these predictions have training implications. Decision makers can deal with interdependent evidence, in part, by learning to base foraging decisions on meaningful patches, and in part by spending enough time in a patch to realize its value. They can practice acting

⁵ See papers in Fiedler and Juslin (2006) for discussion of decision errors due to limited experiential sampling.

quickly to realize the value of timely action and practice shifting attention to realize the value of deliberation. They should learn how to factor tasks into components with varying degrees of uncertainty and varying costs of delayed commitment

Knowledge about Uncertainty

As we have seen, the value of an activity reflects previously experienced rates of return updated by current results. It can also be adjusted in the light of other cues, which indicate that rates of return will be lower or higher than average. In the case of deliberation, these cues pertain to uncertainty. Thus, the value of deliberation in the foraging for surprise model is the product of two ratios. The first (as discussed) is the average swing in utility caused by deliberating in a patch, divided by average time in the patch; the second is current uncertainty, divided by the average uncertainty for problems of the same type. The first captures the costs of the errors prevented by deliberation (i.e., stakes); the second captures the probability of errors in the absence of further deliberation. Without some uncertainty, the value of deliberation falls to zero; and unless the product of stakes and uncertainty is higher than the associated cost of delayed action, melioration will direct attention away from deliberation.⁶

Uncertainty is inferred from (1) specific *known uncertainties* in the active predicate that provides entry to the patch and (2) the *novelty* of the problem as a whole, implying the possibility of *unknown unknowns*. Known uncertainty is the sum of the measures of *incompleteness*, *conflict*, and *lack of resolution* for the active predicate. The possibility of unknown unknowns is based in part on *general cues*, e.g., associations of situation and task attributes with lack of experience, poor performance, or complexity and in part on the implausibility of the *explicit assumptions* that have had to be adopted thus far to reconcile the current core conclusion with

⁶ The new formula is thus roughly consistent with the original qualitative formulation of the quick test: deliberate only if costs of errors outweigh the costs of delay and the problem is uncertain.

other information.

The model predicts details of proficient decision making performance observed in incidents such as the Libyan gunboat example. For example, if known uncertainties are resolved by adopting explicit assumptions (as they were for the *opportunity* patch and the *means* patch in the gunboat example), the perceived *novelty* of the situation grows, increasing the chance that the decision maker will continue deliberation even in the absence of *known* uncertainty (as the Captain in the gunboat example did by shifting attention to the *localization* patch). A patch's relevance in deliberation can be affected not only by the patch, but by facts and assumptions elsewhere in active memory. In the Libyan gunboat example, the prioritization of patches based on known uncertainties changed dramatically when a new core conclusion was adopted (e.g., *opportunity* and *means* were no longer of interest given *non-hostile intent*, but *speeding up*, *turning toward own ship*, and *ignoring warnings* were of greater interest). If *known uncertainties* are resolved by facts discovered in the patch (and new uncertainties are not revealed at the same time), the value of deliberating in that patch declines. For example, given the core conclusion *non-hostile intent*, the Captain quickly found a satisfactory explanation for *turning toward own ship*: i.e., *coincidence*, supported by observations of the large number of US ships in the area. When an innovative solution emerges with significantly fewer assumptions than average (e.g., *the gunboat was looking for targets of opportunity*), the novelty of the situation declines, and the tradeoff between action and deliberation shifts in favor of action on the new solution.

Melioration can also be applied to the decision to keep or reject the current core conclusion, yielding more predictions. The value of the core conclusion is inversely related to the implausibility of the assumptions required to make it fit, divided by the time spent deliberating since it was adopted. Because no deliberation has occurred regarding any *new* core conclusion,

its value is the inverse of the average of this rate in similarly complex and unfamiliar problems. Melioration implies that the current core conclusion is rejected, and the explicit assumptions supporting it are dropped, when its value falls below this average. (The initial value of a new core conclusion is the average, which is then updated by new results.) This rule takes the novelty and complexity of the situation into account when deciding how long to work with the current solution. It predicts that the more challenging the situation, the more assumptions will be tolerated in the story supporting a given core conclusion. For example, the Captain stayed with *hostile intent* until three separate assumptions were needed, about *opportunity*, *means*, and *localization*. Moreover, the expected implausibility of assumptions for an *alternative core conclusion* is only slightly less than that required for the rejected core conclusion. In this case, assumptions required for the *non-hostile intent* story – pertaining to *speeding up*, *venturing into harm's way*, and *ignoring warnings* – were comparable to those for *hostile intent*.

Each component of the value formula corresponds to a category of factors shown empirically to influence metacognitive assessments, such as *judgments of learning* while studying material to be tested later, *feeling of knowing* after failing to recall the answer to a question, and *confidence* in answers that have been retrieved (Reder, 1996; Overschelde, 2008). First, these assessments are influenced, to varying degrees, by so-called direct access to memory traces, i.e., initial retrieval results and the effort or time required to obtain them. In the foraging for surprise model, these are represented by utility gains and time thus far in the current episode, and by known unknowns in active predicates. Second, metacognitive assessments may be inferred from factors not directly related to traces of the material, e.g., the quantity and familiarity of cues pertaining to the task. These factors are represented in the model by novelty and the concomitant possibility of unknown unknowns. Third, assessments are influenced by

stored knowledge about knowledge, represented in the present model by the baseline value of deliberating on particular topics and the average rate of required assumptions. The model generates testable predictions about the relationships among all relevant factors that influence metacognitive regulation.

Time Management Is Expert Behavior

Recent studies provide direct evidence that decision makers *learn from their experience in real-world tasks* to monitor and regulate their use of time. Khoo and Mosier (2008) found that pilots with more automation experience took more time and accessed more information about a possible engine fire when time pressure was not severe than when under time pressure. Although they were clearly aware of the time factors, less experienced pilots did not adapt the time they spent or the information they accessed to time pressure.

Cohen and colleagues found very similar interactions in two studies in which experienced pilots adapted their time to both time pressure and uncertainty while less experienced pilots, despite having the requisite information, did not. In one study (Freeman, Cohen & Thompson, 1998), the costs of delaying a diversion decision was manipulated by varying the amount of remaining fuel. Highly experienced pilots waited longer to divert when the cost of delay was lower (i.e., they had more fuel) and acted sooner when the cost of delay was higher; less experienced pilots took about the same amount of time to act regardless of how much fuel they had (Figure 4a). Yet there was no significant difference in the relevant *knowledge* possessed by the two groups (Figure 4b). A second study (Cohen, Adelman, Thompson, 2000) varied *known uncertainty*, represented by inconsistency of reports from different sources (company dispatch, air traffic control, and company airport operations center) about the estimated time of clearance for landing at the destination. More experienced pilots took more time before diverting when

reports were inconsistent than when they were consistent. Less experienced pilots took the same amount of time in both cases (Figure 4c). Uncertainty led more experienced pilots to focus inquiries on factors directly related to the known uncertainty (Figure 4d, e) and to postpone inquiries about other matters (Figure 4f). Uncertainty led less experienced pilots to delay *all* inquiries. The quick test in R/M predicts the behavior of the experienced pilots: Action is delayed, and active information acquisition and deliberative problem solving occur, when relevant factors are known to be uncertain and irreversible action can be postponed.

CONTRASTING R/M WITH ANOTHER VIEW OF TIME MANAGEMENT

The introduction of this chapter mentioned two objectives: identify future directions for NDM research and call certain habitual assumptions into question. To accomplish these, it is necessary to define more sharply than usual some of the *differences* between R/M and other NDM approaches. A recent proposal by Klein et al. (2007) provides an instructive contrast. They hypothesize, as a stopping rule, that “sensemaking usually ceases when the data and frame are brought into congruence...” (p. 126). If this were correct, time management would be a byproduct of knowledge organization and would not require regulatory skills based on implicit learning about costs and benefits, as described by R/M. A comparison of the two positions will increase understanding of both and bring hidden assumptions to the surface.

Klein et al. (2007) say that any continuation past their stopping rule is “mere grinding out of inferences.” The question they do not address is how, in a challenging and novel situation, a decision maker can *know this in advance*. Nothing guarantees that elements outside a “frame” will not turn out to be relevant, unless we define frames as fixed and isolated patterns, with *no possible connections* to significant unknown unknowns outside themselves. The data-frame congruence stopping rule assumes that decisions are made in an artificial small world: that there

is a bright line delimiting the correct and complete theory of the situation (i.e., “frame”) along with the data it declares relevant. Without that illusion, how does the decision maker know that there is no evidence, hypothesis, or assumption whose importance has not yet been recognized? Klein et al. have gotten the case for a stopping rule backwards. The problem is not that continuation would mean grinding out useless inferences. The true concern is that *it might not*: in realistically complex problems, the possibility cannot be ruled out that another inference might significantly *improve* the response by overthrowing previous assumptions. A stopping rule is needed to balance this possibility against pragmatic demands for timely action. But data-frame congruence ignores both sides of the tradeoff between knowledge and timeliness.⁷

Ironically, the data-frame congruence stopping condition contradicts the central constructivist claims of Data-Frame theory: that data are defined and selected by frames, which are put together on the spot (*just in time*) from smaller “units” (p. 132) and which may be modified and elaborated during sensemaking. If neither data nor frames are fixed in advance, but are dynamically and reciprocally redefined in the course of deliberation (like working memory and primed episodic memory in R/M), data/frame congruence can be no more than a *temporary equilibrium* between data thus far defined and knowledge thus far integrated. Such temporary equilibria can be passed over by proficient decision makers like speed bumps (specifically, interruption delays associated with changing patches in the foraging formula) rather than barriers to pragmatically justified exploration and extension of knowledge. A constructivist account

⁷ The assumption that there are context-independent criteria for good reasoning, or the sufficiency of an argument, is inherited from the syntactic definition of proof in deductive logic. Although Klein’s stopping rule is semantic, it is equally non-pragmatic. Klein et al. (2007) cite plausible reasoning and non-monotonic logic, but they miss the main point, that *real-world* (as opposed to *small world*) reasoning always remains open to defeat by unsuspected problems. See Cohen, Salas, & Reidel (2001) and Cohen et al. (2006) for discussion of why stopping rules for “plausible reasoning” must be *decision oriented* and why the quest for “deduction-lite” argument criteria fails.

requires that the stopping rule determine the “frame,” not the reverse.⁸

The other half of the data-frame congruence stopping rule – that sensemaking “continues as long as key data elements remain unexplained or key components of a frame remain ambiguous” (p. 126) – is even more puzzling in its neglect of pragmatic considerations. A decision maker might stop deliberating and act before *known* unknowns are resolved, or leave an *undepleted* patch to explore a new patch that promises higher rates of return.

CONCLUSION

NDM was spawned by the second generation of problem-solving research, when emphasis shifted from general-purpose reasoning methods (the first generation paradigm) to substantive knowledge (Chase & Simon, 1972), e.g., a rich ensemble of finely differentiated frames (Klein, et al. 2007). The second generation paradigm assumes that experts perform better because they know more than novices about their domain, but claims – as Klein et al. (2007) report in the context of one of their studies – that novices and experts show “no differences in their reasoning processes” (p. 126). In support of this, Klein et al. apply a very coarse first-generation classification of reasoning types: “both [experts and novices] tried to infer cause-effect connections...both tried to infer effects... both experts and novices were not going about the task in accordance with the precepts of deductive logic....” This is like arguing that because both expert and novice writers use nouns and verbs, there are no differences in their linguistic skills.

R/M is based on an emerging third generation view (Holyoak, 1991), that expertise is

⁸ Almost as an afterthought, Klein et al. (2007) add a significant qualification to their main account (p. 126): “Note that sensemaking may continue if the potential benefits of further exploration are sufficiently strong.” Unfortunately, they do not specify what benefits these are. Similarly, on page 136, Klein et al. say, “There are times when people try to make sense of events because they believe they can deepen their understanding, even without any surprise” (i.e., known unknowns). But they do not explain when or why “deeper understanding” would or should be sought. They describe the example of a general officer who recognizes a pattern of enemy positions as diagnostic of a particular type of unit. Instead of stopping because of the supposed “data-frame congruence,” the general proceeds to look for enemy emplacements associated with larger scale patterns of higher echelon units. For R/M, this is expected expert behavior when the costs of errors outweigh the costs of delay and the situation is novel. For Data-Frame theory, it is inexplicable.

adaptation of both knowledge and thinking skills to the requirements of the domain in question, which typically includes both routine and novel situations. The tradeoffs that determine how much time to take for information acquisition and reflection originate both in the underlying properties of memory and attention and in the ecologies to which proficient decision makers adapt. Ecologies shape implicit assessments about the potential payoff and time course of deliberation, which in turn shape attentional strategies needed to activate relevant knowledge in novel situations. R/M's foraging for surprise hypothesis predicts that thinking strategies will improve as decision makers learn to factor tasks into components that vary in uncertainty and cost of delay, acquire better understanding of how knowledge is organized into patches, sample value functions by deliberating in patches for varying amounts of time, become sensitive to known unknowns, novelty, and explicit assumptions, and learn to update their value assessments accordingly. In short, satisfactory management of time for deliberation requires not only substantive domain knowledge but also the acquisition of strategic skill.⁹

R/M provides a naturalistic basis for critical thinking training that enhances rather than replaces intuition. One variant of such training is called STEPS,¹⁰ in which trainees master the

⁹ Klein et al., (2007) downplay the importance of evidence from their own research that experts and novices differ in thinking skills: "a few of the novices...at first seemed reluctant to speculate too much without more data [i.e., to build stories that account for observations]. None of the experts exhibited this scruple." They cite other findings that "novices were less likely to notice when expectancies were violated... novices went down the garden path [held on to an inaccurate interpretation] but experts broke free" (p. 127). Klein et al. attribute the latter to experts' having "more differentiated frames than novices," i.e., more expectations to be violated rather than better strategies for handling violations. But their story is inconsistent: They imply that expectancies were violated for *both groups*. They also state (correctly) in the same paragraph that "a person will not start to question a frame simply as a result of receiving data that do not conform to the frame." In fact, there is abundant evidence that experts take more time to verify their conclusions (Larkin, McDermott, Simon, & Simon, 1980; Larkin, 1981; Chi, Glaser, and Rees, 1982) and have greater ability to correct problems and handle exceptional conditions (McLennan, Pavlou, & Omodei, , 2005; Feltovich, Spiro, & Coulson, 2006; Rudolph, 2003; Shanteau, 1992). On a longer time scale, according to Ericsson (2005), experts *become* experts by intentionally disrupting an existing level of performance in order to achieve a wider convergence. The destabilization of established equilibria in order to find better ones is a primary engine for creating new knowledge.

¹⁰ The training is called STEP in publications based on the earliest version. The "S" was added to make the time management aspect more salient, but does not represent any change in content.

following steps: *Story*, to identify known unknowns and reconcile them with the core conclusion – *Test*, to find and explain unknown unknowns by challenging supporting evidence – *Evaluate*, to assess a core conclusion in terms of the assumptions required to make the story work – *Plan*, to make plans more robust against known or unknown unknowns – *Stop*, to bypass or halt deliberation at any time when costs of delay are higher than potential benefits. Training with STEPS has utilized both tactical decision games (e.g., Cohen et al., 2000a) and live simulator sessions (e.g., Freeman & Cohen, 1996; Cohen et al., 1998) for practice and feedback. Six different experimental tests of STEPS training, mostly in counterbalanced pretest-posttest comparisons, demonstrated that it successfully influenced targeted behaviors and produced decisions that were more similar to those of experts than the decisions of untrained officers (e.g., Cohen, Thompson, Adelman, Bresnick, Shastri, et al., 2000a; Cohen & Thompson, 2001; Freeman & Cohen, 1996). The foraging for surprise model introduced in this chapter suggests some interesting avenues for expansion of this training, such as a more intensive effort to familiarize decision makers with the likely results and pitfalls of specific ways of combining intuition and deliberation, deliberation and action, and deliberation on different topics.

Acknowledgements

The work described here was supported by contracts with Cognitive Technologies, Inc., from the Office of Naval Research, the National Science Foundation, and the Army Research Institute. Thanks to Bryan Thompson and Lokendra Shastri who were indispensable partners throughout the development of Shruti R/M and to Shastri for his original development of Shruti. Finally, I thank the editors of this volume for assistance and advice above and beyond the call of duty.

REFERENCES

- Bargh, J. A. (2000). Beyond Behaviorism: On the Automaticity of Higher Mental Processes. *Psychological Bulletin*, 126(6), 925-945.

- Bargh, J. A. (Ed.). (2007). *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes*. New York: Psychology Press.
- Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, 62, 193-217. Reprinted in K. R. Hammond & T. R. Stewart (Eds.), *The Essential Brunswik* (pp. 135-56). Oxford: Oxford University Press.
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129-136.
- Chase, W., & Simon, H. (1973). The mind's eye in chess. In W. Chase (Ed.), *Visual Information Processing* (pp. 215-81). New York: Academic Press.
- Chi, M., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. Steinberg (Ed.), *Advances in the Psychology of Human Intelligence* (Vol. 1, pp. 7-75). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cohen, M. S., Adelman, L., Bresnick, T. A., Marvin, F., Salas, E., & Riedel, S. (2006). Dialogue as Medium (and Message) for Training Critical Thinking. In R. Hoffman (Ed.), *Expertise Out of Context*. Mahwah, NJ: Erlbaum. Available in <http://www.cog-tech.com/papers/Dialogue/DialogueMediumMessage.pdf>
- Cohen, M., Thompson, B.B., Adelman, L., Bresnick, T.A., Tolcott, M.A., & Freeman, J.T. (1995). *Rapid capturing of battlefield mental models*. Arlington, VA: Cognitive Technologies, Inc. Available in <http://www.cog-tech.com/papers/mentalmodels/MentalModelsPhaseI.pdf>
- Cohen, M. S., Adelman, L., & Thompson, B. T. (2000). *Experimental investigation of uncertainty, stakes, and time in pilot decision making*. Arlington: Cognitive Technologies, Inc.
- Cohen, M. S., & Freeling, A. N. S. (1981). *The impact of information on decisions: Command and control system evaluation*. Falls Church, VA: Decision Science Consortium, Inc. Available in <http://www.cog-tech.com/papers/Scanned/Bayesian/Information.pdf>
- Cohen, M., Parasuraman, R., & Freeman, J. T. (1998). Trust in decision aids: A model and its training implications. In *Proceeding of the 1998 Command and Control Research and Technology Symposium: DoD C4ISR Cooperative Research Program*. Available in http://www.cog-tech.com/papers/Trust/C2_Trust_paper_revised.pdf
- Cohen, M., Freeman, J.T., & Thompson, B.B. (1998). Critical thinking skills in tactical decision making: A model and a training method. In J. Cannon-Bowers & E. Salas (Eds.), *Decision-Making Under Stress: Implications for Training & Simulation* (pp. 155-89). Washington, DC: American Psychological Association. Available in <http://www.cog-tech.com/papers/chapters/tadmus/tadmus.pdf>

- Cohen, M., Freeman, J.T., & Wolf, S. (1996). Meta-recognition in time stressed decision making: Recognizing, critiquing, and correcting. *Human Factors*, 38(2), 206-219. Available in http://www.cog-tech.com/papers/humfac/hf_1996a.pdf
- Cohen, M. S., Salas, E., & Riedel, S. (2001). *Critical Thinking: Challenges, Possibilities, and Purpose*. Arlington VA: Cognitive Technologies, Inc. Available in http://www.cog-tech.com/papers/CriticalThinking/Phase_I_%20Report.pdf
- Cohen, M. S., & Thompson, B. B. (2001). Training teams to take initiative: Critical thinking in novel situations. In E. Salas (Ed.), *Advances in Human Performance and Cognitive Engineering Research* (Vol. 1, pp. 251-91). Amsterdam: JAI. Available in http://www.cog_tech.com/papers/CriticalThinking/Training%20Critical%20Thinking%20for%20Initiative%20HFES%202002.pdf
- Cohen, M. S., & Thompson, B. B. (2005). Metacognitive processes for uncertainty handling: Connectionist implementation of a cognitive model. In M. Anderson & T. Oates (Eds.), *Metacognition in Computation: Papers from the 2005 Symposium* (pp. 36-41). Menlo Park, CA: American Association of Artificial Intelligence.
- Cohen, M. S., Thompson, B. B., Adelman, L., Bresnick, T. A., Shastri, L., & Riedel, S. (2000a). *Training critical thinking for the battlefield. Volume II: Training system and evaluation*. Arlington VA: Cognitive Technologies, Inc. Available in http://www.cog-tech.com/papers/mentalmodels/Vol_IITraining.pdf
- Cohen, M. S., Thompson, B. B., Adelman, L., Bresnick, T. A., Shastri, L., & Riedel, S. (2000b). *Training Critical Thinking for the Battlefield: Volume III: Modeling and Simulation of Battlefield Critical Thinking*. Arlington, VA: Cognitive Technologies, Inc. Available in http://www.cogtech.com/papers/mentalmodels/Vol_IIISimulation.pdf
- Cohen, M., Thompson, B.B., Adelman, L., Bresnick, T.A., Tolcott, M.A., & Freeman, J.T. (1995). *Rapid capturing of battlefield mental models*. Arlington, VA: Cognitive Technologies, Inc. Available in <http://www.cog-tech.com/papers/mentalmodels/MentalModelsPhaseI.pdf>
- Dunlosky, J., & Bjork, R. A. (2008). *Handbook of Metamemory and Memory*. New York: Psychology Press.
- Endsley, M. (1995a). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32-64.
- Endsley, M. R. (2000a). Theoretical underpinnings of situation awareness: A critical review. In M. R. Endsley & D. J. Garland (Eds.), *Situation awareness analysis and measurement* (pp. 3-32). Mahwah, NJ: Erlbaum.

- Endsley, M. R. (2000b). Direct measurement of situation awareness. In M. R. Endsley & D. J. Garland (Eds.), *Situation awareness analysis and measurement* (pp. 147-74). Mahwah, NJ: Erlbaum.
- Entin, E., & Serfaty, D. (1990). *Information gathering and decision making under stress*. Burlington, MA: Alphatech, Inc.
- Ericsson, K. (2005). Superior decision making as an integral quality of expert performance: Insights into the mediating mechanisms and their acquisition through deliberate practice. In H. Montgomery, R. Lipshitz & B. Brehmer (Eds.), *How Professionals Make Decisions* (pp. 135-67). Mahwah, NJ: Lawrence Erlbaum Associates.
- Feltovich, P. J., Spiro, & Coulson, R. L. (2006). Issues of expert flexibility in contexts characterized by complexity and change. In P. Feltovich, K. m. Ford & R. R. Hoffman (Eds.), *Expertise in Context* (pp. 125-46). Cambridge: Cambridge University Press.
- Fiedler, K., & Juslin, P. (Eds.). (2006). *Information Sampling and Adaptive Cognition*. Cambridge, UK: Cambridge University Press.
- Flach, J. M. (2008). Mind the gap: A skeptical view of macrocognition. In J. M. Schraagen, L. G. Militello, T. Ormerod & R. Lipshitz (Eds.), *Naturalistic Decision Making and Macrocognition* (pp. 27-40). Burlington, VT: Ashgate.
- Freeman, J., & Cohen, M.S. (1996). *Training for complex decision-making: a test of instruction based on the recognition/metacognition model*. Paper presented at the Proceedings of the 1996 Command and Control Research and Technology Symposium, Monterey, CA.
- Freeman, J., Cohen, M.S., & Thompson, B. (1998). *Time-stressed decision-making in the cockpit*. Paper presented at the Association for Information Systems Americas Conference, Baltimore, MD.
- Fu, W.-T. (2007). A rational-ecological approach to the exploration/exploitation tradeoffs: Bounded rationality and suboptimal performance. In W. D. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 165-79). Oxford: Oxford University.
- Green, R. F. (1984). Stopping rules for optimal foragers. *The American Naturalist*, 123(1), 30-43.
- Herrnstein, R. J. (1982). Melioration as behavioral dynamism. In M. L. Commons, A. Kacelnik & S. J. Shettleworth (Eds.), *Quantitative Analysis of Behavior: Matching and Maximizing Accounts* (Vol. II, pp. 433-58). Cambridge, MA: Ballinger.
- Herrnstein, R. J. (1997). *The Matching Law: Papers in Psychology and Economics*. New York: Russell Sage Foundation.
- Herrnstein, R. J., & Prelec, D. (1992). Melioration. In G. Loewenstein & J. Elster (Eds.), *Choice over Time* (pp. 235-63). New York: Russell Sage Foundation.

- Holyoak, K. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise*. Cambridge UK: Cambridge University Press.
- Khoo, Y.-L., & Mosier, K. (2008). The impact of time pressure and experience on information search and decision-making processes. *Journal of Cognitive Engineering and Decision Making*, 2(4), 275-294.
- Klein, G. (1998). *Sources of Power*. Cambridge, MA.: M.I.T. Press.
- Klein, G. (1993). *Naturalistic Decision Making: Implications for Design*. Wright Patterson Air Force Base, Ohio: Crew System Ergonomics Information Analysis Center.
- Klein, G. (2000). Analysis of situation awareness from critical incident reports. In M. R. Endsley & D. J. Garland (Eds.), *Situation awareness analysis and measurement* (pp. 51-71). Mahwah, NJ: Erlbaum.
- Klein, G., Phillips, J. K., Rall, E. L., & Peluso, D. A. (2007). A data-frame theory of sensemaking. In R. R. Hoffman (Ed.), *Expertise Out of Context* (pp. 113-55). Mahwah, NJ: Erlbaum.
- Larkin, J., McDermott, J., Simon, D.P., & Simon, H.A. (1980). Expert and novice performance in solving physics problems. *Science*, 20(208), 10-335-1342.
- Larkin, J. (1981). Enriching formal knowledge: a model for learning to solve textbook physics problems. In J. Anderson (Ed.), *Cognitive Skills and Their Acquisition* (pp. 311-44). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Laurence, D. H. (Ed.). (1985). *George Bernard Shaw Collected Letters, Vol. 3: 1911-1925*. New York: Viking.
- Logan, G. (2004). Attention, automaticity, and executive control. In A. Healy (Ed.), *Experimental cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer* (pp. 129-39). Washington DC: American Psychological Association Press.
- Logan, G. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.
- Logan, G., Taylor, S. E., & Etherton, J. L. (1996). Attention in the acquisition and expression of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 620-638.
- Mani, D. (1995). *The design and implementation of massively parallel knowledge representation and reasoning systems: A connectionist approach*. Unpublished Ph.D. Dissertation, University of Pennsylvania.
- McLennan, J., Pavlou, O., & Omodei, M. M. (2005). Cognitive control processes discriminate between better versus poorer performance by fire ground commanders. In H. Montgomery, R.

- Lipshitz & B. Brehmer (Eds.), *How Professionals Make Decisions* (pp. 2009-22). Mahwah, NJ: Lawrence Erlbaum Associates.
- McNamara, J. M., Green, R. F., & Olsson, O. (2006). Bayes theorem and its applications in animal behavior. *Oikos*, *112*, 243-251.
- Monsell, S., & Driver, J. (2000). Banishing the control homunculus. In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes: Attention and Performance XVIII*. Cambridge, MA: M.I.T.
- Newell, A. (1981). The knowledge level. Pittsburg, PA: Carnegie Mellon University.
- Nisbett, R., & Ross, L. (1980). *Human inference: strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Olsson, O., & Brown, J. S. (2006). The foraging benefits of information and the penalty of ignorance. *Oikos*, *112*, 260-273.
- Omodei, M. M., McLennan, J., Elliott, G. C., Wearing, A. J., & Clancy, J. M. (2005). "More is better?": A bias toward overuse of resources in naturalistic decision-making settings. In H. Montgomery, R. Lipshitz & B. Brehmer (Eds.), *How Professionals Make Decisions* (pp. 29-42). Mahwah, NJ: Lawrence Erlbaum Associates.
- Orasanu, J. & Strauch, B. (1994). Temporal factors in aviation decision making. *Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting* (pp. 935-939). Santa Monica, CA.
- Overschelde, J. P. V. (2008). Metacognition: Knowing about knowing. In J. Dunlosky & R. A. Bjork (Eds.), *Handbook of Metamemory and Memory*. New York: Psychology Press.
- Pearl, J. (1989). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo CA: Morgan Kaufmann Publishers.
- Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge UK: Cambridge University Press.
- Peters, D. J., Jackson, L. A., Phillips, J. K., & Ross, K. G. (2008). The time to decide: How awareness and collaboration affect the command decision making. In A. Kott (Ed.), *Battle of Cognition: The Future Information-Rich Warfare and the Mind of The Commander* (pp. 194-211). Westport CT: Praeger Security.
- Pirolli, P. (2007). *Information Foraging Theory: Adaptive Interaction with Information*. Oxford: Oxford University.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, *106*(4), 643-675.
- Reder, L. M. (Ed.). (1996). *Implicit Memory and Metacognition*. Hillsdale, NJ: Erlbaum.
- Rescher, N. (2001). *Cognitive pragmatism: The theory of knowledge in pragmatic perspective*. Pittsburg, PA: University of Pittsburg Press.

- Rudolph, J. W. (2003). *Into the Big Muddy and Out Again: Error Persistence and Crisis Management in the Operating Room*. Dissertation: Boston College, The Carroll Graduate School of Management, Department of Organization Studies.
- Rumelhart, D.E., Smolensky, P., McClelland, J.L., & Hinton, G.E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition, Volume 2: Psychological and biological models* (pp. 7-57). Cambridge, MA: The MIT Press.
- Shanteau, J. (1992). The psychology of experts: An alternative view. In G. Wright & F. Bolger (Eds.), *Expertise and Decision Support*. New York: Plenum Press.
- Shastri, L. (1992). *Neurally motivated constraints on the working memory capacity of a production system for parallel processing*. Proceedings of the Fourteenth Conference of the Cognitive Science Society, Bloomington, IN.
- Shastri, L. (1999a). Types and quantifiers in Shruti: A connectionist model of rapid reasoning and relational processing. In S. Wermter & R. Sun (Eds.), *Hybrid Neural Symbolic Integration*. Berlin: Springer-Verlag.
- Shastri, L. (1999b). Advances in Shruti: A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence*, 11, 79-108.
- Shastri, L. (2001). A computational model of episodic memory formation in the hippocampal system. *Neurocomputing*, 38-40, 889-897.
- Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Brain and Behavioral Sciences*, 16, 417-494.
- Shastri, L., & Grannes, D.J. (1996). *A connectionist treatment of negation and inconsistency*. Proceedings of the Eighteenth Conference of the Cognitive Science Society, San Diego, CA.
- Shastri, L., & Mani, D.R. (1997). Massively parallel knowledge representation and reasoning: Taking a cue from the brain. In H. K. J. Geller, & C. Suttner (Ed.), *Parallel Processing for Artificial Intelligence 3*: Elsevier Science.
- Shastri, L., & Wendelken, C. (1998). *Soft computing in Shruti: A neurally plausible model of reflexive reasoning and relational information processing*. Paper presented at the Proceedings of Soft-Computing '99, Genoa, Italy.
- Shastri, L., & Wendelken, C. (1999). *Knowledge fusion in the large -- taking a cue from the brain*. Paper presented at the Proceedings of the Second International Conference on Information Fusion, Sunnyvale, CA.

- Schraagen, J. M., Klein, G., & Hoffman, R. (2008). The macrocognition framework of naturalistic decision making. In J. M. Schraagen, L. G. Militello, T. Ormerod & R. Lipshitz (Eds.), *Naturalistic Decision Making and Macrocognition* (pp. 3-25). Burlington, VT: Ashgate.
- Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6(2), 1-10.
- Simon, H. A. (1997a). Rationality as process and as product of thought. In H. A. Simon (Ed.), *Models of bounded rationality: Empirically grounded economic reason* (pp. 444-59). Cambridge MA: MIT Press.
- Simon, H. A. (1997b). From substantive to procedural rationality. In H. A. Simon (Ed.), *Models of bounded rationality: Empirically grounded economic reason* (pp. 424-43). Cambridge MA: MIT Press.
- Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over the course of decision making. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 27(5), 1250-1260.
- Son, L. K., & Kornell, N. (2005). Metacognitive judgments in rhesus macaques: Explicit versus implicit mechanisms. In H. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness*. Oxford: Oxford University Press.
- Stephens, D. W. (2007). Models of information use. In D. W. Stephens, J. S. Brown & R. C. Ydenberg (Eds.), *Foraging: Behavior and Ecology* (pp. 31-58). Chicago: University of Chicago Press.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging Theory*. Princeton NJ: Princeton University Press.
- Thompson, B. B., & Cohen, M. S. (1999). Naturalistic decision making and models of computational intelligence. In A. Jagota, T. Plate, L. Shastri & R. Sun (Eds.), *Connectionist symbol processing: Dead or alive?* (pp. 26-8).
- Wendelken, C., & Shastri, L. (2000). *Probabilistic inference and learning in a connectionist causal network*. Proceedings of the Second International Symposium on Neural Computation Berlin, Germany, Berlin, Germany.
- Wendelken, C., & Shastri, L. (2002). *Combining belief and utility in a structured connectionist agent architecture*. Proceedings of Cognitive Science 2002, Fairfax, VA.
- Wendelken, C., & Shastri, L. (2003). *Acquisition of concepts and causal rules in SHRUTI*. Proceedings of Cognitive Science 2003, Boston, MA.
- Wendelken, C., & Shastri, L. (2005). Connectionist mechanisms for cognitive control. *Neurocomputing*, 65-66, 663-672.

Figures

Figure 1. Overview of Recognition-Metacognition Model.

Figure 2. Simplified view of a Shruti R/M network depicting initial recognition of hostile intent in the Libyan gunboat example. Unshaded rectangles are predicates, each of which has a positive and negative belief collector (+ and -, respectively), a query node (?), a positive and negative utility collector (\$+ and \$-, respectively, with some not shown to save space), and a node (x, y, z) for each role played by objects in the relation represented by the predicate. Shaded rectangles are facts about objects that instantiate predicate roles. Solid lines represent cycles of self-reinforcing belief and utility activation. Numbers illustrate pathways that activation takes from the perceptual input fact. Faint broken lines represent synchronized activation of objects that instantiate roles and the role nodes that are instantiated by them.

Figure 3. Each point on the horizontal axis represents a particular division of available time between two mutually exclusive activities, A and B. Time in Activity A is measured from left to right on horizontal axis, and time in Activity B is measured from right to left on horizontal axis. The two columns of charts represent three different scenarios. A has negatively accelerated returns in both, and B has a higher asymptote in both. On left, B has negatively accelerated returns; on the right, B's returns follow an S-shaped pattern. Top row of charts represents *cumulative utility gain* as a function of time in activity A (solid curves) and in activity B (dashed curves). Dotted curve is the *total utility* gained from both activities. Vertical lines indicate equilibrium allocations, for which the two activities have equal rates of gain. (This is cumulative gain from the activity divided by time spent in the activity, corresponding to slopes of lines drawn between cumulative utility at the beginning of an activity to cumulative utility at the equilibrium time allocation.) The bottom row of charts shows the *rates of utility gain* for each activity as a function of different allocations of time between them. According to melioration, whenever A's rate of gain is higher than B's, more time will be allocated to A (thick arrow pointing right). Whenever B's rate of gain is higher than A's, more time will be allocated to B (thick arrows pointing left). Equilibria occur when rates become equal, corresponding to intersection of the two curves.

Figure 4. Results of experiments with commercial airline pilots. (a) With more time available, experienced pilots waited longer before diverting; available time had no effect on diversion by less experienced pilots. (b) Experienced and less experienced pilots did not significantly differ in accuracy of fuel estimates. (Tendency of more experienced pilots to be more accurate with more time is suggestive but not significant.) (c) Uncertainty led experienced pilots to wait longer before diverting, but did not influence time taken by less experienced pilots. (d) Uncertainty led experienced pilots to inquire earlier about progress of snow removal, but had no effect on such requests by less experienced pilots. (e) Uncertainty led experienced pilots to inquire earlier for relevant information about accident cleanup, but led less experienced pilots to postpone inquiries. (f) Uncertainty led both more and less experienced pilots to postpone requests for less relevant traffic information.







