Metacognitive Processes for Uncertainty Handling: Connectionist Implementation of a Cognitive Model

Marvin S. Cohen & Bryan B. Thompson

Cognitive Technologies, Inc.
4200 Lorcom Lane
Arlington VA 22207
mcohen@cog-tech.com

Abstract

An empirically based cognitive model of real-world decision making was implemented in Shruti, a system capable of rapid, parallel relational reasoning. The system effectively simulates a two-tiered strategy associated with proficient decisions makers: Recognitional or reflexive activation of expectations and associated responses, accompanied by an optional, recursive process of critiquing and correcting, regulated by the stakes of the problem, the time available, and the remaining uncertainty. The model and implementation are inconsistent with the conventional claim that decision makers fall back on formal analytical methods when pattern recognition fails. Instead, they learn simple metacognitive strategies to leverage reflexive knowledge in novel situations. In addition, the model suggests that the development of executive attention functions (metacognitive strategies) may be necessary for, and integral to, the development of working memory, or dynamic access to long term memory, and that strategies developed for uncertainty handling may accelerate the reflexive learning of remotely connected concepts.

Introduction

Decision makers are sometimes said to shift from an intuitive to an analytical mode of decision making in novel, unstressed situations (e.g., Hammond, 2000; Payne, 1993). Intuition involves rapid and relatively automatic processing based on preexisting knowledge, while analysis involves slow and deliberative application of learned procedures. On this viewpoint, strategy selection may be metacognitive because it involves monitoring and regulating first-order cognitive processes, but neither the intuitive nor the analytical strategy involves metacognition in an *intrinsic* way. Our hypothesis, on the contrary, is that metacognition is the fundamental means by which decision makers deal with uncertainty. More specifically, we argue that expert decision makers iteratively monitor for, identify, and attempt to correct qualitatively distinct uncertainty issues in the products of intuition, within the constraints of available time. It follows that reasoning with uncertainty in naturalistic environments involves a less abrupt discontinuity with intuitive processing than the intuitive-analytic dichotomy suggests.

What we describe here are the interim results of an effort to test and explore the implications of a cognitive model, which is based on theory and findings in research on real-world decision makers, by implementing it in a computational framework that reflects a biologically plausible model of connectionist reasoning. The immediate goal was to develop a tool that (1) can perform rapid recognitional inferences and planning (i.e., intuitive processing) within a large (expert) belief network, (2) exemplify human limitations on computational resources and attention, and (3) implement metacognitive control process that simultaneously regulate recognitional processing, help overcome computational limitations, and deal with uncertainty. Such a tool could form the basis for adaptive training and decision support systems that simulate reasoning and decision making in uncertain real-world situations.

A Cognitive Model

The Recognition / Metacognition model is motivated by research on expert problem solving and decision making that suggests that proficient real-world decision makers differ from less experienced decision makers in two ways: (1) recognitional skills, i.e., the ability to recognize familiar patterns and to rapidly retrieve associated expectations, goals, and responses, and (2) the ability to apply their experience in unfamiliar, novel, or anomalous situations which fail to closely match previous patterns. Under the latter conditions, recognition needs to be supplemented by metacognitive processes that monitor the products of recognition and, when problems are found, take steps to resolve them unless the urgency of immediate action outweighs the expected benefits. The Recognition / Metacognition model provides a systematic account of processes that decision makers use to understand and plan in such situations.

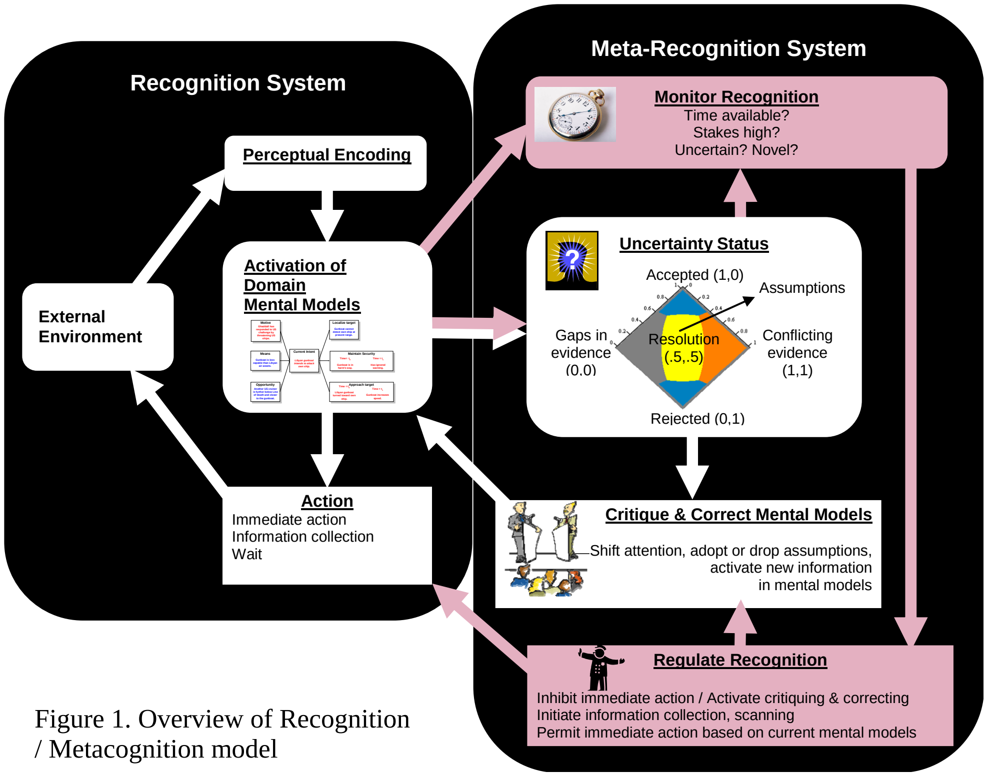

Figure 1. Overview of Recognition / Metacognition model

The R/M model (Figure 1) involves three interrelated processing cycles. The most basic is recognitional (on the left), in which external stimuli are perceptually encoded, integrated into mental models when they activate matching knowledge - goal - action complexes in long-term memory, and induce action. In routine situations, recognitional cycles proceed uninterrupted. Recognition is monitored and regulated, however, by two nested metacognitive cycles (on the right). The outer loop detects problems in the current recognitional solution and crudely but quickly balances the urgency of action against the magnitude of the problems and the stakes. If time is available, stakes are high, and the current solution is uncertain, the outer metacognitive loop inhibits the recognition-based response and activates the inner metacognitive loop.

The inner loop critiques and corrects currently active mental models. It monitors for different kinds of uncertainty, definable as patterns of positive and negative support, including gaps in knowledge, conflicting evidence or goals, and shiftable or unreliable assumptions. When problems are found, it implements corrective strategies that have been effective in the past under similar conditions for that type of uncertainty. Metacognitive correcting strategies are composed of two kinds of actions – shifting attention and/or changing assumptions – which function as queries to long-term memory. Their joint effect is to activate information not retrieved in the initial recognitional situation or in previous metacognitive cycles, but potentially relevant for resolving uncertainty. Newly activated information is automatically integrated with information primed during previous cycles, and may fill gaps, resolve conflict, or induce revision of weak assumptions. Critiquing and correcting are iterative because correcting one problem may sometimes (but not always) lead to another problem. Critiquing and correcting stop as soon as uncertainty is sufficiently reduced, the stakes change, or the costs of further delay outweigh the benefits. At any time, the decision maker is ready to act on the current best solution.

Complex behaviors can result from reinforcement learning with these simple elements. Among the behaviors which a computational implementation of the R/M model should simulate are: monitoring for and noticing conflicts between predictions and observations, going beyond pattern matching to create plausible explanations for novel situations, elaborating a story to explain a conflicting cue rather than disregarding or discounting it, sensitivity to the reliability of assumptions used to explain away conflicting data, deciding when to generate an alternative coherent story to account for data or an alternative plan to achieve goals, and a refined ability to balance the time available for decision making against the costs of errors. Primary published sources for the R / M model, including examples from critical incident interviews with Army and Navy officers, are Cohen, Freeman, & Wolf (1996); Cohen, Freeman, & Thompson (1998); Cohen, & Thompson (2001).

Training based on the R/M model has been developed and tested successfully in a number of domains. In recent work, for example, it was used to develop course materials for critical thinking in the Advanced Tactics course at the Command & General Staff College, Leavenworth, KA (Cohen et al., 2000a). In over a half dozen evaluation studies with Navy and Army officers, critical thinking training based on R / M has significantly improved the quality of trainees' assessments and decisions (as measured by agreement with domain experts), while also eliciting strong judgments of face validity from the trainees themselves (Cohen et al., 2000a; Cohen, Freeman, & Thompson, 1997). A computational implementation might significantly enhance the effectiveness and efficiency of such training.

Connectionist Implementation

Overview

According to the Recognition / Metacognition model, decision makers integrate voluminous observational and testimonial inputs about relatively familiar situation into models (typically causal) that provide coherent interpretations and plans. We simulate these rapid recognitional processes (i.e., intuition) using Shruti, a connectionist architecture developed by Lokendra Shastri of the International Computer Science Institute, Berkeley (Shastri, 1992; Shastri & Ajjanagadde, 1993; Shastri, & Grannes, 1996; Shastri & Mani, 1997; Shastri, & Wendelken, 1998; Shastri, 1999; Shastri, 1999). Shastri refers to rapid, parallel, associative activation processes as reflexive reasoning. To explore and test the R/M model, a metacognitive layer has been implemented on top of reflexive reasoning, which monitors for different types of uncertainty and learns strategies for attention shifting and assumption revision.

Limits on dynamic access to long term memory (LTM) emerge naturally from the computational structure of Shruti and the neuro-biological constraints it respects. These limits imply that all associated information known by the agent cannot be brought to bear at the same time by a single cycle of spreading activation. The existence of such limits means that inference and planning processes must be capable of (a) dynamically determining the scope of active human memory from which they draw at any given time, and (b) of remaining coherent within those limits. This requirement underlies the fluidity with which a reasoner is able to focus limited computational resources at different levels of spatial and temporal abstraction (the chunking problem in AI), and shift planning horizons from moments to years and back to moments. At the same time, this need for fluid changes in focus introduces the necessity for an adaptive dynamics of executive attention. The mechanisms of attention shifting, in turn, form a developmental basis for acquiring skilled metacognitive behaviors, which monitor and regulate recognitional processing under conditions of uncertainty. Thus, the metacognitive processes described in the R/M model could develop naturally as extensions of skilled attention-shifting behaviors within a resource-limited active memory. Key differences between experts and novices would include the knowledge structures and metacognitive strategies that enable them to manage these resource limitations and apply the appropriate information during reasoning.

Metacognition enables reasoning about highly mediated relationships in long-term memory, i.e., interdependencies that are implicit in long-term memory but which are too distant to influence decisions (because the elements are not easily co-activated). By shifting attention among currently active nodes, and/or clamping their positive or negative activation, metacognitive processes shift the sphere of activation in long-term memory to otherwise inaccessible contents. In short, uncertainty triggers the introduction of a wider perspective and brings more information to bear on the problem at hand, while the same metacognitive processes accelerate learning of initially remote connections within a computationally constrained reasoning process.

In addition to versions of Shruti that run on serial platforms, a version of Shruti has also been implemented on a parallel machine, the CM-5, (Shastri & Mani, 1997). The resulting system can encode knowledge bases with over 500,000 (randomly generated) rules and facts, and yet respond to a range of queries requiring derivations of depth five in under 250 milliseconds. Even queries with derivation depths of eight are answered in well under a second. In our implementation, the same Shruti resources are used for both rapid reflexive processes and the reflective processes that monitor and regulate them. Since resource limitations are shared, a reflective decision maker achieves less in any particular cycle of reflexive processing, but may receive a net benefit by extending the span of reflexive processing across multiple cycles of attention shifting. The implementation affords a means of studying this important tradeoff.

Simulations of a tactical Army battlefield scenario and of a Navy tactical scenario have been implemented in this system, and are suitable for use in training systems and decision aids.

Features of the Implementation

The reflexive system supports causal and semantic reasoning over relational representations, utility propagation and decision making, priming, and supervised learning. We focus here on features that bear most directly on the interaction of recognitional and metacognitive processes.

Traditional connectionist models support “associations of ideas,” but not the relational inferences that people also arrive at quickly and accurately (e.g., If Tom gave Susan a book, it is Susan, not Tom, who now has the book). Shruti uses dynamic variable binding to keep track of objects and the roles they play in relations, enabling the same entities to recur in different parts of the inference bound to different predicates. Object identity is represented by temporal synchrony of firing at nodes in different places in the network; and identity constraints on the application of rules are enforced by temporal pattern matching. Support for relational reasoning significantly enhances simulation power and is virtually unique among rapid, parallel systems.

Predictions regarding resource constraints derive naturally from the representational and computational features of the system. Because Shruti uses temporal synchrony for object identify, finite bandwidth means that (i) only a limited number of objects can be tracked at any given time, and (ii) if jitter increases with the length a signal travels in the long-term memory network, accurate inference about objects (as opposed to coarser associations of ideas) is limited in depth. These features, among others, account for limits on the amount of long-term memory that can be active in working memory at one time.

Weights in the reflexive system can be adapted to the statistical properties of the environment through experience. The Shruti simulator tunes network weights and rule-strengths via supervised learning, using a form of backpropagation. These weights reflect the co-occurrences of elements in event-action-goal complexes that define causal mental models. Activation spreads automatically from a relation, forward to generate predictions and backward to find explanations. Because Shruti also supports taxonomic / semantic reasoning, rules framed in terms of general categories can mediate inferences about instances of those categories. Moreover, activation of belief about an event propagates not only belief but also utility across causally linked events and actions. For example, the negative utility of a predicted event heightens the activation of actions that might prevent it. The network settles on appropriate actions virtually at the same time as it settles on beliefs about a situation.

Shruti aggregates information about the activation flowing through a node over time. When an active node is on the edge of the currently active network, this information is used in inferencing. It reflects the average historical effect of nodes linked to the active node but beyond the current boundary of activation. Aggregated information about belief in effect represents the prior probability of a causal explanation, while aggregated information about utility represents, in effect, the expected utility of an action or outcome. Each is an average over specific situational features: The former averages over evidence which, if known, might change the probabilities of different explanations (with possible consequences for predictions and actions). The latter averages over diverse outcomes, for which preferences might significantly vary if they were considered separately (with possible consequences for action). Instances of aggregated historical information are effectively assumptions, because the validity of inferential outcomes within the active part of the network depends on assuming implicitly that historical information fits the present situation. Acting on averages is part of what meant by habit, i.e., failing to mindful of particulars of a unique situation. In our model, assumptions of this sort are a principal target of reflective, metacognitive monitoring.

A predicate in Shruti is represented as a “supernode” whose components include distinct positive and negative collectors for belief activation (corresponding to support for truth or falsity, respectively) and for utility (corresponding to positive and negative preference, respectively), various entity roles associated with the relation, and an enabler. The enabler plays a key role in metacognitively controlled attention. Supplying activation to the enabler of a relation is equivalent to querying its truth or falsity, which may be influenced by directly encoded facts or by activation of potential causes and effects that provide relevant evidence. Metacognitive strategies focus attention by activating enablers; they revise assumptions about events, actions, and goals by clamping or unclamping positive or negative collectors.

Such strategies are effective when they bring relevant knowledge into view that was previously dormant, and permit its integration with what is already known. To make this possible, Shruti exhibits priming effects that preserve the results of reflexive reasoning across successive shifts of attention. Priming enables new information activated by an attention shift to be combined with information that was active before the shift, and to generate new inferences. For example, shifting attention to a node on the edge of the active network disaggregates the historical information stored at that node, by spreading activation into neighboring parts of the network. Priming allows the overall activation pattern to adapt to facts about the present situation that are represented in the newly activated part of the network. Attention shifting thus removes some of the reliance on habit in decision making. Determining when and where to shift attention so as to obtain a more detailed verification of current recognitional conclusions is a key function of metacognitive strategies.

The metacognitive system recognizes qualitatively different patterns of uncertainty in the active network and can adapt its uncertainty handling strategy accordingly. In real-world decision making, people do not reduce all types of uncertainty to a single measure, such as probability, nor do they respond to every type of uncertainty in the same way (e.g., Ellsberg, 1988; Lipshitz, 1994). To develop a richer representation, the metacognitive system exploits Shruti’s independent registration of evidence for and against a hypothesis. It can thereby discriminate four qualitatively distinct, synchronous uncertainty patterns at a node: (1) Gaps in information exist when here is little or no activation either for or against a hypothesis/action (i.e., the sum of + and - activation is significantly less than one); (2) conflict involves strong activation both for and against a hypothesis/action (i.e., the sum of + and - activation is greater than one); (3) acceptance or rejection occurs when activation is strong either for or against a hypothesis/action but not both; and (4) lack of resolution exists when there are moderate amounts of activation both for and against a hypotheses. Lack of resolution tends to be more prevalent in nodes at the edge of the network, where information is averaged over diverse situations. By introducing a temporal dimension, the system can identify a fifth type of uncertainty that may overlap the others, (5) dependence on assumption, i.e., activation that is subject to change as more information is considered. Assumptions may be implicit, corresponding to low resolution nodes at the edge of the network, or a result of explicit choices to clamp a hypothesis or action.

The hypothetical beliefs and actions produced by clamping may reflexively activate relevant information that would not otherwise have been considered. As a result, what-if reasoning is a powerful tool for ferreting out and resolving hidden uncertainty. The most obvious example is uncovering and filling gaps in the evidence with respect to a hypothesis by clamping it true or false – that is, by asking, What would we expect to find if the hypothesis were known to be true (or false)? This query may activate both potential explanations and potential predictions, which supply evidence for or against the hypothesis in question. If the explanations and predictions are themselves as yet unverified (i.e., more gaps), they can be further investigated or at least assessed for plausibility, eventually filling the gap in knowledge about the hypothesis. To uncover potential conflicting evidence, a decision maker can focus on beliefs or actions that are already accepted or rejected, and reverse the verdict. This is equivalent to asking: Can I find a reason to think I was wrong in arriving at that conclusion? The result may be awareness of previously unnoticed factors in the situation that lead to a different prediction or explanation.

Perhaps the most interesting what-if strategy is for ferreting out hidden assumptions in reflexive inference, in which the decision maker clamps the evidence as true and the conclusion false. For situation assessment the strategy poses questions like, (1) This situation will not produce the predicted effects – Why not, and what will it produce instead? and (2) This situation is not explained by the inferred causal factors – What else could have caused it? The analogous questions in the case of decision making are: (1) Our current plan is destined to fail (i.e., it will not achieve the desired outcomes) – how could that happen? and (2) Our goals will be achieved but not by our current plan – how could we accomplish that? In response to (1), the system reflexively searches for premises that were implicit in the falsified inference (i.e., necessary conditions for the relevant causal process to go off as expected) and which might turn out to be inaccurate. In response to (2), the system activates alternative possible explanations of the evidence or alternative action options that achieve our goals, which we may have overlooked, disregarded, or rejected too quickly. A context in which assumptions are especially important is the effort to resolve conflict in evidence. Since different lines of reasoning support both a proposition and its negation, conflict itself is a cue that assumptions somewhere in one or both lines of reasoning are false. Such assumptions may be uncovered by strategies like those just discussed.

Conclusion

We implemented a cognitive theory of real-world decision making (the Recognition / Metacognition model) in Shruti, a system capable of rapid, parallel relational reasoning. The system effectively simulates a two-tiered strategy used by proficient decisions makers: Recognitional or reflexive activation of expectations and associated responses, accompanied by an optional, recursive process of critiquing and correcting, regulated by the stakes of the problem, the time available, and the remaining uncertainty. Together, these processes generate, verify, and modify evolving mental models (often resembling stories) to account for events, anticipate the future, and develop plans in relatively novel circumstances. The system is able to identify qualitatively different types of uncertainty, i.e., gaps in the evidence, conflicts in the evidence, and changeable or unreliable assumptions, and learn strategies for addressing each particular type. The model and implementation shed light on results that are inconsistent with the conventional claim that decision makers fall back on formal analytical methods when pattern recognition fails. Instead, they learn simple metacognitive strategies to leverage reflexive knowledge in novel situations.

The hypothesis that we wished to test was that an extremely simple and small set of reflective functions can have a very significant impact on the performance of the reflexive system, especially for novel and uncertain decision making tasks. A goal in this development, therefore, was to keep the reflective layer as simple as possible and to base it on plausible psychological and neurological assumptions. In fact, the operations required for the metacognitive simulation – inhibiting recognitional responding, shifting focal attention, and clamping truth values – were readily implemented in Shruti. In addition to being simple, these operations are psychologically and biologically plausible; they add little to widely accepted mechanisms of action inhibition, attention shifting, and what-if reasoning, and to suggestions in the behavioral decision making literature that people deal differently with different types of uncertainty. Finally, in demonstrations of the system with test scenarios, the operations worked as expected. Simulated decision makers noticed more factors and made better decisions with metacognitive skills than without.

In addition, the simulation promises to shed light on the role of metacognitive reasoning in dynamic access to long term memory. According to Shruti, there are inherent, dynamic, limits on the scope of long term memory information that can be brought to bear in responding to any inputs. The key interaction between the reflexive and reflective systems is the adaptive direction of focused attention within the reflexive memory by means of learned metacognitive behaviors. Recency (or priming) effects are used to assemble such intermediate results into composite assessments, thereby tending to improve results. The model suggests that the development of executive attention functions (metacognitive strategies) may be necessary for, and integral to, the development of working memory, or dynamic access to long term memory.

Another constraint that we imposed on the implementation was to endow the metacognitive layer with no knowledge beyond what is already encoded in the reflexive knowledge base and what in addition can be learned from its own experience. The metacognitive system must learn to combine its elementary operations in response to different patterns of uncertainty by reinforcement and associative learning processes. Through such learning, the metacognitive system should acquire a repertoire of uncertainty handling strategies that enable it to make better decisions in the present, and to accelerate the acquisition of new reflexive knowledge for the future. While we have implemented backpropagation learning, we have not yet tested the ability of the system to accumulate the full measure of required skill or to regulate decision making successfully in terms of stakes, uncertainty, and time.

Acknowledgements

Work by Cognitive Technologies, Inc., has been funded by the Army Research Institute (Contract DASW01-97-C-0038), the Office of Naval Research (Contract N00014-95-C-0182; Contract N00014-00-M-0070) and the National Science Foundation (Contract DMI-9861411). We thank Lokendra Shastri for his major contributions to this work, Carter Wendelken for his continuing efforts, and Sharon Riedel of ARI for her support and encouragement.

References

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.

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. 155-189). Washington, DC: American Psychological Association.

Cohen, M. S. & Thompson, B. B. (2001). Training teams to take initiative: Critical thinking in novel situations. In E. Salas (Ed.), Advances in Cognitive Engineering and Human Performance Research (pp. 251-291). Amsterdam: JAI.

Cohen, M., Freeman, J. & Thompson, B. (1997). Integrated critical thinking training and decision support for tactical anti-air warfare. In 3d International Command and Control Research and Technology Symposium Proceedings

Cohen, M., Thompson, B.B., Adelman, L., Bresnick, T.A., Shastri, L., & Riedel, S. (2000). Training critical thinking for the battlefield. Volume II: Training system and evaluation. Arlington VA: Cognitive Technologies, Inc..

Cohen, M.S., Thompson, B.B., Adelman, L., Bresnick, T.A., Shastri, L., & Riedel, S. et al. (2000). Training Critical Thinking for the Battlefield: Volume III: Modeling and Simulation of Battlefield Critical Thinking. Arlington, VA: Cognitive Technologies, Inc..

Ellsberg, D. (1988). Risk, ambiguity, and the Savage axioms. In P. Gardenfors & N.-E. Sahlin (Eds.) Decision, probability, and utility: Selected reasoning (pp. 245-269). Cambridge, UK: Cambridge University Press.

Freeman, J., & Cohen, M.S. (1998). Critical thinking training and decision support: an evaluation of combined effects on tactical decision making. Arlington, VA: Cognitive Technologies, Inc..

Hammond, K.R. (2000). Judgment under stress. Oxford, UK: Oxford University Press.

Lipshitz, R. (1994). Decision making in three modes. Journal for the Theory of Social Behavior, (24), 47-66.

Payne, J., Bettman, J.R., & Johnson, E.J. (1993). The adaptive decision maker. New York: Cambridge University 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 (pp. 159-164). Bloomington, IN:

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. Proceedings of Soft-Computing '99 Genoa, Italy:

Shastri, L. (1999). Advances in Shruti: A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. Applied Intelligence, 11, 79-108.

Shastri, L. (1999). 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.