

Acquisition of concepts and causal rules in SHRUTI

Carter Wendelken and Lokendra Shastri
International Computer Science Institute
1947 Center Street, Suite 600
Berkeley, CA 94704
{carterw,shastri}@icsi.berkeley.edu

Abstract

The SHRUTI model demonstrates how complex cognitive functions can be realized by neural circuitry. This paper addresses how some key elements of this circuitry can be learned in a neurally plausible manner. Two basic mechanisms, causal Hebbian learning and recruitment learning, are used to learn relational concepts and causal rules.

Introduction

Although a great deal is known about neural representations in sensory, somatosensory, and motor cortices, the neural structures underlying higher cognitive processes are largely unknown. The SHRUTI model [Shastri and Ajjanagadde, 1993, Shastri, 1999] provides a promising set of proposals about what sorts of neural circuits are necessary to realize functions such as memory, reasoning, and decision-making. The biological plausibility of a connectionist architecture, however, depends not only on the fact that it is composed of abstract neurons and connections between them, but also on the assumption that its structure can be realized in the brain as a result of learning and development. This has been demonstrated for some of SHRUTI's structures, namely, episodic facts [Shastri, 2002b] and categories [Shastri and Wendelken, 2003]. In this paper we discuss how two other representational elements of SHRUTI – relational focal-clusters and causal rules – can be learned.

The SHRUTI model

SHRUTI is a connectionist model that demonstrates how a network of neuron-like elements could encode a large body of structured knowledge and rapidly perform explanatory and predictive reasoning. An expanded version, called SHRUTI-agent, also represents goals and makes decisions [Wendelken and Shastri, 2002]. The model can encode different types of conceptual knowledge including relational schemas/frames for encoding action and event types (e.g., falling); causal rules between relational schemas (e.g., if you fall you get hurt); entities, types, and the relations between them (e.g. John is a Man); and different types of facts such as episodic facts that record specific events (e.g., John fell in the Hallway today), taxon facts that record general statistical knowledge (e.g., children often fall), and utility facts

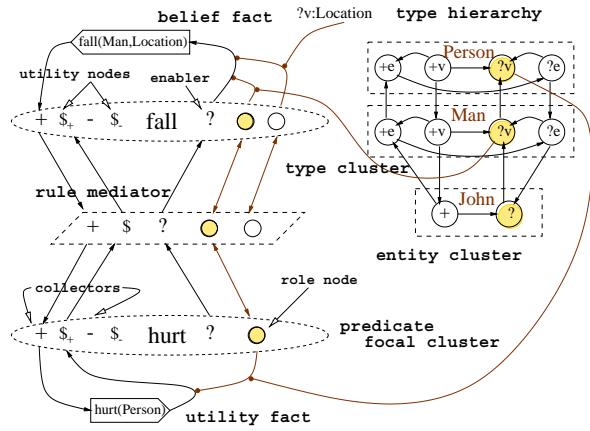


Figure 1: Diagram showing major structures of the SHRUTI architecture.

that record associations between situations and reward or punishment (e.g, being hurt is bad).

SHRUTI suggests that the encoding of relational information (e.g., event schemas) is mediated by neural circuits composed of structured cell ensembles, termed *focal-clusters* (see Figure 1). A relation focal-cluster (e.g., for the relation *fall*) consists of (i) role cells (e.g., for the role *fallees*) whose *synchronous* firing with entity cells (e.g., for the entity *John*) encodes role-entity *bindings* (e.g., the binding *fallees*=*John*) comprising the currently active relational *instance* (e.g., *John* fell in the hallway), (ii) + and - collector cells whose firing signifies belief and disbelief, respectively, in the currently active relational instance, (iii) enabler cells (?) whose firing signifies a search for support for this relational instance, and (iv) $\$_+$ and $\$_-$ utility cells that encode the desirability of this relational instance and its negation, respectively. Focal-clusters of entities and types are structured in a similar manner. The grouping of cells within a focal-cluster highlights their functional cohesiveness, but does not imply physical proximity.

The processing of relational information involves the transient propagation of *rhythmic* activity across relation and entity focal-clusters. Persistent facts in long-term memory, including episodic facts, taxon facts, and utility facts, are realized as temporal pattern-matching cir-

cuits. Causal knowledge (e.g., a rule) is encoded by links that enable the propagation of rhythmic activity between antecedent and consequent focal-clusters. Rules involve directed connections from antecedent collector to consequent collector, from consequent enabler to antecedent enabler, from consequent utility nodes to antecedent utility nodes, and between role nodes in both directions. All of these connections pass through an additional focal-cluster called the rule mediator, which serves as a locus for evidence combination.

Connectionist learning mechanisms

A variety of learning mechanisms have been applied to connectionist models. Supervised learning via error backpropagation is particularly powerful and is widely used; however, its plausibility as a brain mechanism is a matter of controversy. On the other hand, Hebbian learning, an unsupervised learning process wherein links between (source) neurons that contribute to the firing of other (target) neurons tend to have their connection strengths increased [Hebb, 1949], is well established as a neural mechanism. Hebbian learning provides the foundation for two distinct mechanisms, causal Hebbian learning and recruitment learning, that are employed by SHRUTI in the creation of relational clusters and causal rules.

Causal Hebbian learning

Hebbian learning is ideal for building associations. However, SHRUTI's causal model encodes more than just associations between components of a rule; it also encodes a directionality that is vital for correct inference. Causal Hebbian learning (CHL), has been proposed as a partial solution for the experience driven learning of simple cause-effect relationships [Wendelken and Shastri, 2000], wherein the occurrence of event *A* followed soon after by event *B* is taken as evidence of *A* being a potential cause of *B*, especially, if this sequence of events occurs repeatedly. Although it was developed independently, CHL has much in common with temporally asymmetric Hebbian learning [Abbot and Song, 2000], and is neurobiologically grounded in the phenomenon of spike-timing dependent plasticity ([Roberts and Bell, 2002]) In CHL, weight updates depend on the relative timing of pre- and post-synaptic firing (source and target activation). Furthermore, different connection types may exhibit different dependences; some may update their weights only when pre-synaptic firing precedes post-synaptic firing, while others may do so only in the reverse scenario.

The existence of a learning mechanism that allows for different connection types with different temporal dependences is key to learning the bidirectional, asymmetric links of SHRUTI's causal rules. When applied to learning SHRUTI rules, the causal Hebbian learning mechanism operates as follows: connections between collector cells are enhanced whenever the source is active before the target (within a certain window of time), whereas connections between enabler cells are enhanced whenever the

target is active before the source. Thus, when event *A* is observed preceding event *B*, both the forward and backward links for the rule $A \implies B$ are strengthened. No specific temporal relation, beyond co-activation within a certain time frame, is required for learning role to role links. The requisite initial situation consists of a set of relation focal-clusters, with weak connections linking each node to others of its type (collectors linked to other collectors, role nodes linked to other role nodes, etc.) The observation of a large number of events (or the re-observation of memorized events) leads to the formation of rules reflecting the apparent causal structure of the environment.

For a collector link, if the source has been active sufficiently long and the target then becomes active, the link weight is updated as $w_{t+1} = w_t + \alpha * (1 - w_t)$ where $\alpha = 1/\#updates$; otherwise if the source has been active sufficiently long and the target fails to fire, the weight is decremented $w_{t+1} = w_t - \alpha * w_t$. If the target becomes active before the source, then there is no change. It is easily seen that $w_{t+1} = (\#increases) / (\#updates)$ for $w_0 = 0$, correctly encoding the probability that the target of the link follows the source within the specified time parameters. A modification of this rule including a normalization term, to account for the possibility of multiple sources, reduces the weight increase on a link by a factor proportional to the number of active links impinging on the same target. This allows the link weight for $+A \rightarrow +B$ to encode the causal strength of *A* for *B* (i.e. the probability of *B* given *A* and no other cause.) Note that the update parameter α decreases with each successive update; it follows that link weights become quite difficult to affect after they have been subject to a significant amount of experience. A modification to this rule, where $\alpha \geq 1/\#updates$, may be desirable in order to allow recent events to have a greater impact on link weights than events that are further in the past.

For enabler links, the learning rule is nearly the reverse of the above. In this case, a similar weight increase occurs whenever a link target has been active for sufficiently long and a source then becomes active, and a weight decrease occurs when a target remains active for too long without activity at the source. If the source becomes active first, there is no change. It is easily seen that the enabler link correctly records the probability that the source fires after the target (within designated time parameters), and for a link $?B \rightarrow ?A$ this can be reasonably interpreted as $P(B|A)$. Note that this rule makes the counterintuitive prediction that potentiation can occur when post-synaptic spike precedes pre-synaptic spike, a process that has in fact been observed in the cerebellum of an electric fish [Bell et al., 1997].

CHL is an effective mechanism for learning the statistics of causal rules via activation of nodes in the SHRUTI network. A more complete story of rule learning in SHRUTI, however, requires recruitment of rule mediator structures in addition to the updating of causal weights.

Recruitment learning

The technique of recruitment learning [Feldman, 1982, Shastri, 1988] is a biologically plausible learning mechanism that is useful for building structured representations of concepts within a connectionist network. Recruitment learning (and also vicinal algorithms [Valiant, 1994]) can be described informally as follows: Learning occurs within a partially structured network containing a large number of richly interconnected nodes. *Recruited* nodes in the network are nodes that have acquired distinct functionality by virtue of their *strong* interconnections to other recruited and sensorimotor nodes. *Unrecruited (free)* nodes are connected via weak links to a large number of free, recruited, and sensorimotor nodes. Free nodes form a pool of nodes from which suitably connected nodes are recruited for representing new functional units. The recruitment process transforms a quasi-structured network into a collection of nodes and circuits with specific functions. Weight updates during recruitment conform to the principles of Hebbian learning [Hebb, 1949]. It has been shown that recruitment learning can be grounded in the biological process of long-term potentiation [Shastri, 2002a].

Structured recruitment learning

Learning the relatively complex structures of SHRUTI poses a significant challenge for the recruitment learning approach because learning such structures involves recruiting not just individual cells, but also structured ensembles or functional circuits. The probability of extracting a particular structure in a randomly connected structure quickly approaches zero as the complexity of the structure increases. Structures in SHRUTI are sufficiently complex that it would be unreasonable to expect them to emerge in a random network. Consequently, some level of pre-existing organization is required. We argue that the nature of the pre-existing structure required to learn the SHRUTI model, essentially a few basic circuit patterns repeated over and over again, is just what one might expect to find in the brain. For one thing, the repetition of a few simple patterns is essentially the level of organization that computational modeling has shown to be obtainable from a genetically based developmental process [Marcus, 2001]. Moreover, it is well known that different brain regions are organized differently even prior to experience, and these differences are likely to relate to demands imposed by different cognitive functions. An excellent example of this may be found in the idiosyncratic architecture and local circuitry of the hippocampal formation, which has been shown to be ideally suited for supporting the rapid encoding of episodic memories via the recruitment of a set of complex neural circuits [Shastri, 2002b].

Recruiting relational concepts

An active mind forms new conceptual representations not primarily because it has sufficient environmental stimuli to do so, but rather because the existence of a

new representation will make it more effective in some way. A key consideration, especially in the context of the SHRUTI- agent architecture, is the ability of the system to make good decisions quickly. The addition of a new concept or rule is often the key that turns a complex, difficult decision problem into one that can be solved with relative ease. In fact, expertise within a given domain can be equated with having learned the most appropriate and efficient set of representations required to solve problems within that domain.

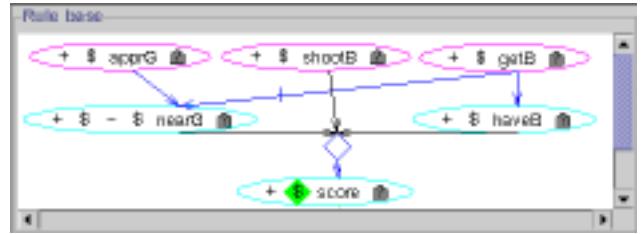


Figure 2: Soccer network before learning.

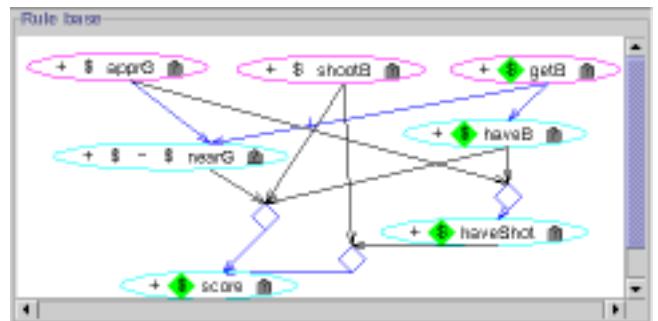


Figure 3: Soccer network after learning.

Consider for example the situation faced by a soccer player who wants to score a goal. In order to score, he must be in possession of the ball, be in range of the goal, and take a shot. A SHRUTI-agent network governing the player's behavior is shown in Figure 2. This is a sequential decision problem, since possible actions interfere with each other (e.g., going to get the ball negates the effects of approaching the goal.) It turns out that this sort of decision task is not directly amenable to solution via simple spreading activation in the SHRUTI network (for details, see [Wendelken and Shastri, 2002]). This failure is evident in Figure 2 where utility has failed to propagate to any of the three possible actions. In particular, *haveB* appears useless without *nearG* being true (inhibitory links enforce this in the network), and the utility of *nearG* depends similarly on *haveB*. However, addition of a new concept *haveShot*, that represents the result of executing the *approachGoal* action in the context of *haveBall*, facilitates the optimal action sequence and simplifies this problem. The correct sequence of actions can be obtained via spreading activation in the modified

network, as shown in Figure 3, where utility has successfully propagated to $getB$, which is the optimal first action.

The learning of a new concept such as $haveShot$ can be achieved by a process called *utile concept learning*. Utile concept learning involves the recruitment of new relational focal-clusters into an existing causal network. In order to be recruited, a focal-cluster must satisfy two conditions. First, the focal-cluster must represent some combination of existing representations; for example, such a focal-cluster could be a rule mediator that combines existing predicates as antecedents or it could be a conjunctive concept that binds together existing actions and/or predicates. Second, a recruited focal-clusters must also provide new pathways for the propagation of utility.

The central mechanism driving utile concept learning is as follows: the connection between two relational clusters A and B (specifically, the rule $A \implies B$ encoded as a specific pattern of directed links) is strengthened whenever (i) A is seen as a cause of B and (ii) utility at A appears to be derived from utility at B . This occurs whenever there is sufficient activity at both the collector and utility nodes of focal-clusters A and B . Standard Hebbian learning accomplishes much of what is required here – any link could tend to gain in efficacy when both source and target nodes are active – but an additional interaction effect between collectors and utility nodes, such that enhancement of links leading into and out of each node is significantly boosted when both node types are active, is also important. Requiring this interaction serves to restrict recruitment to situations in which a new structure is both a combination of existing beliefs and a conduit to some goal.

The creation of new rules via the recruitment of rule mediators as well as the instantiation of new relational concepts via the recruitment of predicate focal-clusters, can be accomplished via the mechanisms of utile concept learning. In both cases, a relational cluster may be considered to be recruited when it is fully linked into an existing causal network; this occurs when there are strong connections between the cluster and previously recruited relations in both the forward and backward directions. Thus, to recruit a new rule, the basic learning mechanism must be applied to more strongly link each antecedent to the new rule mediator structure, and also to more strongly link that structure to its consequents. To recruit a new predicate that is a conjunction of existing predicates and/or actions, the links between each component and the new conjunctive structure must be strengthened (this might involve recruiting a new rule structure, e.g. $A \wedge B \implies A \text{and} B$), as must the links connecting the conjunctive structure to its source of utility (and this too can require recruitment of a new rule, e.g. $A \text{and} B \implies C$).

The initial condition for utile concept learning involves a set of unrecruited structures randomly connected to existing recruited structures; specifically, a set of free relational focal-clusters are randomly connected to other free and recruited predicate clusters and rule mediators. For each random high-level directed connection

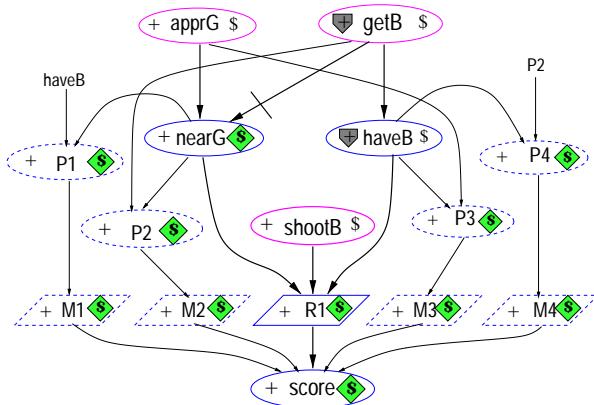
between clusters, there is a systematic pattern of links between nodes which corresponds to the connectivity described for rules.

Utile concept learning may occur as a result of an experienced or imagined sequence of actions and situations. In particular, the mechanisms involved in simulating a sequence of actions are conducive to the learning of useful new relations and rules.

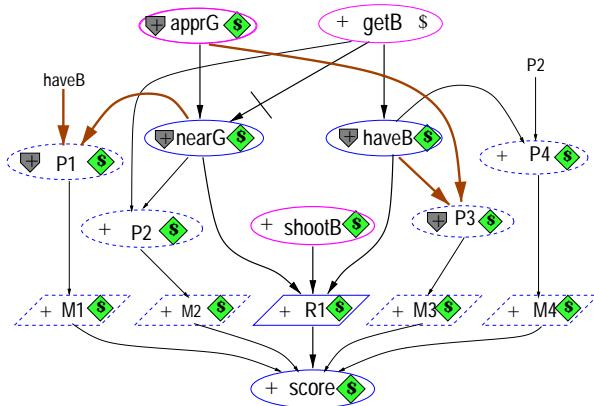
Utile concept learning example We now examine the soccer scenario pictured in Figure 4(a). This simulator snapshot illustrates the propagation of utility over a network consisting of the actions $apprG$ (approach goal), $getB$ (get ball), and $shootB$ (shoot ball at goal), and predicates $nearG$ (near goal), $haveB$ (have ball), and $score$ (scored on goal). Rules indicate that approaching the goal causes one to be near the goal, getting the ball causes one to have the ball but also takes one away from the goal, and shooting while near the goal and in possession of the ball leads to scoring. As discussed above, spreading activation is insufficient for decision-making, given this knowledge base. The task for the system is to recruit a new conjunctive concept and associated rule such that decision-making involving the represented actions is simplified. Figure 4 illustrates the causal network described here along with a set of initially unrecruited predicate focal-clusters and rule mediators (labeled $P1 - P4$ and $M1 - M4$, respectively). The unrecruited predicate clusters shown include different combinations of existing predicates, including $haveB + nearG$ ($P1$), $getB + nearG$ ($P2$), $apprG + haveB$ ($P3$), and $haveB + P2$ ($P4$). Unrecruited rule mediators shown are a subset of those which include $shootB$ as an antecedent and $score$ as a consequent.

The correct sequence of actions required to accomplish the goal of scoring is to first get the ball, then approach the goal, then shoot. When the system simulates execution of these actions in the proper sequence, recruitment of useful new concepts can occur. In Figure 4(a), simulation of the first action $getB$ is depicted. The result if this action is that $haveB$ is believed or predicted to be true. Also, utility is allowed to propagate from the goal $score$; it reaches each of the rule mediators (free and recruited) that are attached to $score$ and each of the unrecruited predicate clusters attached to these mediators. Activity at unrecruited clusters is weak since all links leading into and out of them have low weight. By virtue of belief associated with $haveB$, utility is also able to propagate to $nearG$ and $apprG$. None of the connected pairs of focal-clusters have both collectors and utility nodes active at this point, so no recruitment takes place.

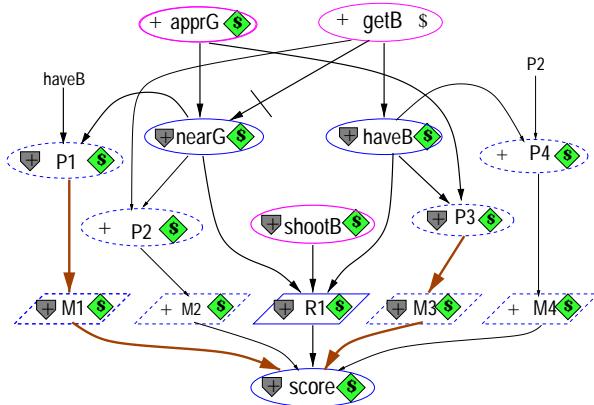
Simulation of the second action $apprG$, depicted in Figure 4(b), leads to positive collector activity at $nearG$ and also at the unrecruited clusters $P1$ and $P3$. At this point, all depicted connections leading into clusters $P1$ and $P3$ satisfy the basic condition for utile concept learning – i.e., all collectors and utility nodes are active. Thus, these links are strengthened and $P1$ and $P3$ come to rep-



(a) State after execution of `getB` (get ball).



(b) State after execution of `apprG` (approach goal).



(c) State after execution of `shootB` (shoot ball).

Figure 4: Recruited clusters that make up the soccer network along with a subset of the free clusters connected to them. Thick lines indicate strong connections.

resent, at least to a greater degree than before, the combinations `nearG + haveB` and `apprG + haveB`. Recruitment of these new concepts is completed with the simulation of the third and final action `shootB`, depicted in Figure 4(3). Instantiation of `shootB` in the network allows activity to propagate through previously recruited mediator `R1` as well as through the unrecruited `M1` and `M3`, to yield the belief or prediction that `score` is true. At this point, the utile concept learning condition is satisfied for the links leading into and out of `M1` and `M3`; thus, these two rule mediators are recruited into the causal network. The result of this process is that two new pathways leading to the goal have been constructed. One of them, through `P1` and `M1`, is essentially a restatement of the existing rule, neither harmful nor particularly helpful. The other pathway, through `P3` and `M3`, involves the introduction of a truly useful new conjunctive concept. The new rule $shootB \wedge P3 \Rightarrow score$ suggests that achieving `P3` is a useful thing. Since `P3` is achieved only when `apprG` is executed in the context of `haveB` being true, it enforces the particular sequencing of actions required for this task. Note that if the actions had been executed in a different sequence, no learning would have occurred. If `appr` had preceded `getB`, then `nearG` and `haveB` would never have been active together. Also, if `shootB` had not been executed last, then neither the predicate `score` nor any of the rule mediators could have become active.

In the example, only a small set of free focal-clusters are shown, and all of these are connected to elements of the original network. It is safe here to ignore other free focal-clusters, since only those connected to the soccer network have any chance of being recruited. However, the existence of the depicted free focal-clusters depends on the assumption that there are enough free focal-clusters in the larger network such that, with random connectivity between clusters, the probability of not finding these particular clusters is sufficiently small. While this entails that utile concept learning requires a large number of cells, the requirement is by no means biologically implausible.

Pruning recruited structures Recruitment of focal-clusters via utile concept learning can lead to recruitment of undesirable structures, and in particular, invalid rules structures may be created. Consider the recruitment of rule mediator `M1`. This could occur when execution of the action `shootB`, in the context of belief in `nearG` and `haveB`, leads to activation (due to observation) of the goal predicate `score`. At the same time that the appropriate mediator is recruited, mediators representing combinations such as `nearG + shootB` or `nearG + haveB` may also be recruited. A rule that says $nearG \wedge haveB \Rightarrow score$, however, would be causally incorrect, since it leaves out an important precondition `shootB`, and its recruitment could potentially lead to poor performance.

The solution to this problem involves adding mechanisms for statistical learning on top of the basic recruitment story. For rule learning, this can be fairly straightforward. The causal Hebbian learning mechanism de-

scribed previously provides just the sort of link weight correction that is required to preserve legitimate causal rules and eliminate others. The links corresponding to a rule that matches the environment will be strengthened (toward accurate probabilistic weights) over the course of repeated observation, while links corresponding to an invalid rule will slowly decay back to very low weight values. Recruitment of conjunctive predicate clusters can only lead to problems through the potentially invalid rules that connect these to their source of utility. Hence, here again, causal Hebbian learning can prune away harmful structure.

A version of the learning scenario presented in Figure 4, consisting of the original soccer network (Figure 2) augmented with seven free predicate clusters and seven free rule mediators clusters, was implemented and used as a preliminary test of the full learning system. Following ordered simulation of the three actions, a majority of the represented free clusters were recruited. One set of these (equivalent to $P3$ and $M3$ in Figure 4) formed a useful new causal path, one set (equivalent to $P1$ and $M1$ in Figure 4) essentially matched the existing causal structure, and three other sets yielded undesirable causal structure (e.g. $getB + haveB \wedge shootB \implies score$). Subsequently, causal Hebbian learning was allowed to operate as the system observed a series of events (sequential predicate activations) generated according to a predefined distribution. The result was significant link weight reduction along all recruited pathways except the two that represented accurate models of the environment.

Conclusion

We have demonstrated the operation of two neurally plausible learning mechanisms, causal Hebbian learning and recruitment learning, in the learning of concept focal-clusters and causal rules in the SHRUTI model. It was seen that execution or simulation of a sequence of actions that leads to reward or punishment can cause one-shot recruitment of useful new predicates and rules, in a process termed *utile* concept learning. Moreover, the co-operation of learning mechanisms that are sensitive to the statistical nature of the environment ensures that only those recruited rules which accurately reflect this will be preserved. Although the learning mechanisms discussed here are implemented at a relatively abstract level, the success of SMRITI, a highly detailed model of episodic fact learning in the hippocampus [Shastri, 2002b], provides strong guidance and well-founded hope that a similarly detailed and biologically grounded account of concept learning will emerge from this work.

Acknowledgment: This work was supported by NSF grants SBR-9720398 and ECS-9970890.

References

[Abbot and Song, 2000] Abbot, L. and Song, S. (2000). Temporally asymmetric hebbian learning, spike timing, and neuronal response variability. *Neurocomputing*, 32-33:523-528.

[Bell et al., 1997] Bell, C., Han, V., Sugawara, Y., and Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, 387:278-281.

[Feldman, 1982] Feldman, J. (1982). Dynamic connections in neural networks. *Biological Cybernetics*, 46:27-39.

[Hebb, 1949] Hebb, D. (1949). *The organization of behavior*. Wiley:New York.

[Marcus, 2001] Marcus, G. (2001). Plasticity and nativism: towards a resolution of an apparent paradox. In *Emergent Neural Computational Architectures*, pages 368-382. Springer-Verlag.

[Roberts and Bell, 2002] Roberts, P. and Bell, C. (2002). Spike timing dependent plasticity in biological systems. *Biological Cybernetics*, 87:392-403.

[Shastri, 1988] Shastri, L. (1988). *Semantic Networks: An evidential formalization and its connectionist realization*. Los Altos/Pitman Publishing Company, London.

[Shastri, 1999] Shastri, L. (1999). Advances in SHRUTI - a neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence*, 11.

[Shastri, 2002a] Shastri, L. (2002a). A computationally efficient abstraction of long-term potentiation. *Neurocomputing*, 44-46:33-41.

[Shastri, 2002b] Shastri, L. (2002b). Episodic memory and cortico-hippocampal interactions. *Trends in Cognitive Sciences*, 6:162-168.

[Shastri and Ajjanagadde, 1993] Shastri, L. and Ajjanagadde, V. (1993). From simple associations to systematic reasoning. *Behavioral and Brain Sciences*, 16(3):417-494.

[Shastri and Wendelken, 2003] Shastri, L. and Wendelken, C. (2003). Learning structured representations. *Neurocomputing* (forthcoming).

[Valiant, 1994] Valiant, L. (1994). *Circuits of the mind*. Oxford University Press.

[Wendelken and Shastri, 2000] Wendelken, C. and Shastri, L. (2000). Probabilistic inference and learning in a connectionist causal network. In *Proceedings of the Second International Symposium on Neural Computation*.

[Wendelken and Shastri, 2002] Wendelken, C. and Shastri, L. (2002). Combining belief and utility in a structured connectionist agent architecture. In *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*.