Appeared in Behavioral and Brain Sciences, 19 (2), 1996. pp 331–337

Temporal Synchrony, Dynamic Bindings, and SHRUTI: a representational but non-classical model of reflexive reasoning

Lokendra Shastri[1]
International Computer Science Institute
1947 Center Street
Berkeley, CA 94704
shastri@icsi.berkeley.edu

1 Synchrony versus signatures

As **Lange & Dyer** point out, there are a number of important structural similarities in SHRUTI and ROBIN. The two models, however, differ significantly in the way they solve the dynamic binding problem. In SHRUTI, dynamic bindings are represented by the synchronous firing of nodes. Thus the binding (John=giver) is represented by the synchronous firing of node clusters representing 'John' and 'giver'. ROBIN on the other hand associates a pre-determined and unique pattern of activation called *signature* with each concept and represents dynamic bindings by propagating appropriate signature patterns over appropriate clusters of role nodes. Thus the binding (John=giver) is represented in ROBIN by establishing the pattern of activation corresponding to the signature of 'John' over the cluster of nodes representing the role 'giver'.

Lange & Dyer argue that the signature approach is superior to the synchronous activation approach. Their argument can be summarized as follows: (i) Whereas the signature approach allows arbitrarily many entities to be co-active, the synchronous activation approach allows at most 10 distinct entities to participate in dynamic bindings. This severely limits the representational and processing power of SHRUTI. (ii) Whereas "signatures" have semantic content, "synchrony" does not. This lack of semantic content would make learning more difficult in the synchronous activation approach.

The arguments put forth by **Lange & Dyer**, however, are based on a serious misunderstanding of SHRUTI's synchronous activation approach. Lange & Dyer make two crucial errors in understanding the synchronous activation approach. First, they misunderstand what is meant by an "entity". Second, they confuse the *medium* of representing and propagating information, namely, synchrony, with the *content* that is represented using this medium. The first error leads Lange & Dyer to the erroneous conclusion that SHRUTI will rapidly run out of phases during language processing. The second error convinces them that SHRUTI's representation of bindings lacks semantic content, and therefore, it will be difficult for SHRUTI to learn high-level knowledge.

1.1 Counting phases: What does an entity make?

Lange & Dyer are wrong when they state that the representation of the sentence:

An elderly rich man bought a large red brick house from an owner known to have ties to Columbian drug lords.

will require more than ten bindings and overload the capacity of the synchronization approach.

As explained in (Shastri & Ajjanagadde 1993), all the active knowledge pertaining to an active entity, i.e., all its activated features and all the roles it currently fills in dynamic bindings, is grouped together by the synchronous firing of appropriate nodes. Consequently, all the active knowledge about an active entity occupies a single phase. A simple analysis shows that the above sentence involves only *four* distinct entities, and therefore, its representation in SHRUTI will require only *four* phases.

The four entities involved in the above sentence are: "the elderly rich man", "a large red brick house", "the owner", and "some Columbian drug lords". An approximate paraphrase of how the meaning of the above sentence may be expressed as features of, and relations between, these four entities follows:

1. *exists(u,v,w,x)*

   There are four entities u, v, w and x.

2. *elderly(u)* ∧ *rich(u)* ∧ *man(u)*

   u is an elderly rich man.

3. *large(v)* ∧ *red(v)* ∧ *made-of-brick(v)* ∧ *house(v)*

   v is a large red brick house. Note that in the given context, being large, red, and made of bricks are *attributes* of the entity house.

4. *owns(w,v)*

   w is the owner of the house (v).

5. *drug-lords(x)* ∧ *columbian(x)*

   x is a group of Columbian drug lords.

6. *buy(u,v,w)*

   buy is a multi-place relation, three of whose roles are: buyer, buy-object and seller. The elderly rich man (u) bought the large red brick house (v) from the owner (w).

7. *has-ties(w,x)*

   the owner (w) has ties to the Columbian drug lords (x).

Although one might quibble over details, the above description should make it clear that the situation described in the sentence involves only four distinct entities, and hence, its meaning can be expressed using just four phases.

Furthermore, the complete representation of the entire example text "An elderly rich man bought .... drugs were found hidden in its walls." will require only nine phases — even if *all* the information provided in the text is assumed to be active *simultaneously*.[2] Typically, this would not be the case since some of the less salient information would become inactive. In the above example, one can imagine that the entity "two months later" might become inactive and release its phase. The information associated with this entity, however, would remain in the medium term memory and become active if triggered by subsequent utterances.

Lange & Dyer might argue that while the representation of the *meaning* of the first sentence requires only four phases, its syntactic analysis will require more than ten phases. Any such claim would also be erroneous. Henderson (1994) has developed a connectionist parser based on SHRUTI's synchronous activation approach that can process a broad range of English sentences using at most *nine* phases. For the sentence in question, the parser would require only *five* phases. Furthermore, the parser will be able to process the entire example text using only *six* phases (Henderson 1995).

We believe, as strongly as before, that an important prediction about the nature of reflexive reasoning emerging from SHRUTI is that while the capacity of the working memory underlying reflexive reasoning is very large, the number of *distinct* entities that can participate as role-fillers during an episode of reasoning is quite small. This belief is further strengthened by the fact that analogous constraints on a parser's working memory capacity can help explain several properties of human parsing including our limited ability to deal with center-embedding.

1.2 Interaction between the medium-term memory and dynamic memory

As explained in (Shastri & Ajjanagadde 1993), SHRUTI assumes the existence of a medium term memory (MTM) capable of storing relational instances (i.e., facts) that remain active and are deemed interesting. In the context of language understanding, such facts consist of factual information explicitly conveyed in the input as well as facts that are inferred, or assumed, in order to establish referential and explanatory coherence. Given that facts stored in the MTM can be activated rapidly and are available for participating in reflexive reasoning, it is entirely consistent with SHRUTI that a fact mentioned seven sentences earlier can participate in the resolution of pronoun reference. So it is not clear why **Lange & Dyer** think that such data on pronoun reference contradicts SHRUTI (see para 15 Lange & Dyer).

Although SHRUTI depends on the existence of a MTM, it does not require the sort of interaction with the MTM posited by **Lange & Dyer** (cf. "bindings ... constantly fly in and out of static medium-term memory, even within sentence boundaries.", Lange & Dyer para. 12). Their erroneous understanding of how phases are used in SHRUTI leads them to believe that the working memory capacity of SHRUTI is overly limited and in turn prompts them to posit an overly intense interaction between dynamic facts and the MTM.

1.3 Propagation of bindings versus associative spreading activation

It was explained in Section 8.2.6 and R1.2 of (Shastri & Ajjanagadde 1993) that although the synchronization of nodes becomes diffuse after a few steps along a chain of inference, the activation continues to propagate along interconnected nodes. The latter, non-synchronous form of propagation, corresponds more to an associative spread of activation and less to systematic reasoning. Such an associative spread of activation has the desired properties and seems adequate for modeling priming effects and explaining the sort of phenomena discussed by **Lange & Dyer** in paras. 13 and 16.

1.4 Confusion regarding working memory

Lange & Dyer point to distinctions between the short term memory limit of 7$\pm$2 proposed by Miller and the bounded capacity of SHRUTI's working memory underlying reflexive reasoning (WMRR). The difference between these two types of working memories and their capacity was discussed explicitly in Section 8.2 of (Shastri & Ajjanagadde 1993). Based on biologically plausible values of the slowest frequency at which synchronous activity might occur ($\approx$ 30 Hz) and the coarseness of synchronization (a few milliseconds), we had proposed that the number of distinct entities participating in role-filler bindings is perhaps limited to 7 ± 2. These entities however, can participate as role-fillers in a very large number of simultaneously active facts! Thus the capacity of SHRUTI's working memory — measured in terms of the number of facts it can represent simultaneously — is huge. This working memory underlying reflexive processing should not be confused with the *overt* working memory for items present in the conscious memory and available for deliberate reflective processing. It is this latter type of working memory that seems to haves a more limited capacity of 7 ± 2 items/facts.

1.5 Alternate notions of working memory capacity

The work cited by **Lange & Dyer** notwithstanding, there does not exist any conclusive evidence that the capacity of the working memory underlying reflexive reasoning and/or language processing is limited by the *amount of activation*. The idea that the *total activation* in a system can be a limiting and determining factor during cognitive processing is quite implausible and devoid of any substantive computational or biological rationale. Too much should not be read into the fact that a limited body of psycholinguistic data can be modeled by placing a bound on total activation since it is always possible to model a limited dataset in a number of different ways. The proponents of the total activation approach must offer an independent biological or computational justification for positing a capacity limit based on total activation. While it is trivially true that there must exist an upper bound on the total activation that can be sustained by the brain, such a bound would be extremely high, and there does not exist any evidence to suggest that such a bound is ever reached when agents perform routine linguistic tasks.

The capacity limits on WMRR proposed in SHRUTI are a direct consequence of adopting a biologically plausible solution to the dynamic binding problem. These bounds were not posited to model specific psycholinguistic data. Therefore, it is significant that Henderson's parsing model which makes use of these capacity limits can explain some of the phenomenon that has been modeled by invoking limits on total activation (Just & Carpenter 1992). For example, the fact that subject-relative sentence are easier to process than object-relative ones can be easily explained by Henderson's parsing model.

Most proposals characterizing the capacity of the working memory underlying cognitive processing have not paid adequate attention to the *structure* of items active in the working memory. Total activation based models also fall in this category. In contrast, SHRUTI points to a structured view of working memory capacity wherein the working memory capacity is limited not by total activation, but rather by the maximum number of distinct entities that can fill arguments in active relation instances and the maximum number of instantiations of any relation that can be active simultaneously.

1.6 SHRUTI makes use of dynamic signatures

Lange & Dyer confuse the *medium* of representing and propagating information, namely, synchrony, with the *content* that is represented using this medium. This has lead Lange & Dyer to the erroneous conclusion that the synchronous activation approach is semantically deficient. SHRUTI supports at least as much – if not more — semantic richness as ROBIN. In fact, it is possible to interpret the functioning of SHRUTI in terms of "signatures" and demonstrate that the synchronous activation approach offers significant advantages over the signature approach adopted in ROBIN.

In SHRUTI, synchrony simply serves as the *medium* for binding together bits and pieces of active information pertaining to an active entity. The actual *content* that is encoded about an entity depends on which feature nodes and role nodes form part of a synchronously firing group of nodes. This collection of nodes firing in synchrony with the focal node of an entity can be interpreted as the *dynamic* signature of that entity! Observe that an entity's "signature" in SHRUTI is dynamic and changes from one episode of reasoning to another, and even within an episode, depending on which features are active in a given context. This dynamic notion of an entity's signature in SHRUTI can be contrasted with ROBIN's notion of a fixed and pre-defined signature assigned to each entity.

Instead of propagating the same pre-determined signature of an entity to all active predicate and rule structures, SHRUTI propagates only those components of a dynamic signature that are relevant to that structure. This can be illustrated with the help of an example. Figure 1 (fashioned after Figure 24 in Shastri & Ajjanagadde 1993) illustrates the encoding of the following rule in SHRUTI:

$x{:}animate \wedge y{:}solid \wedge y{:}heavy\ walk\text{-}into(x,y) \Rightarrow hurt(x)$

The rule states that an animate entity gets hurt if it walks into a solid and heavy object. The firing of this rule is mediated by a "semantic filter" which allows the propagation of activity as long the role fillers involved in the currently active situation possess the appropriate features/types. For the above rule, the propagation of bindings from *walk-into* to *hurt* occurs if the role-filler of the *agent* role of *walk-into* is *animate* and the role-filler of the *patient* role of *walk-into* is *solid* and *heavy*. The semantic filter uses synchrony to enforce the appropriate semantic restrictions on role-fillers by ensuring that bindings propagate only if (i) the *agent* of *walk-into* and the feature *animate* fire in synchrony (i.e., the role-filler of *agent* has the feature *animate*) and (ii) the *patient* of *walk-into* and the features *solid* and *heavy* fire in synchrony (i.e., the role-filler of *patient* has the features *solid* and *heavy*).

One can interpret the functioning of semantic filters in SHRUTI in terms of signatures as follows: All the features firing in synchrony with the *agent* role-filler constitute the dynamic signature of the *agent* role-filler. Similarly for the *patient* role-filler. The semantic filter checks the dynamic signature of the *agent* role-filler to see if it includes the feature *animate* and also checks the dynamic signature of the patient role-filler to check if it includes the features *solid* and *heavy*. Observe however, that the semantic filter associated with the rule does not examine the complete signatures of *agent* and *patient* role-fillers. It only examines those features of these role-fillers that are relevant to its semantic function.

In other words, SHRUTI maintains dynamic signatures of each active entity and propagates customized signatures to each semantic filter consisting only of those features that are relevant to that filter. Thus temporal synchrony provides an optimal way of realizing signatures. In contrast, the **Lange & Dyer** approach requires the propagation of pre-determined signatures. The same signature of a role-filler is sent to all the roles (predicates) it fills, irrespective of which features are relevant at a given predicate. Since different sets of features may be relevant for determining the applicability of different rules, it is expected that the bulk of the information propagated in ROBIN would be irrelevant.

1.7 Learning

In view of the discussion in the previous section, SHRUTI cannot have any semantic and representational limitations vis-a-vis ROBIN. Thus learning in SHRUTI cannot be more difficult than in ROBIN. In fact, it may be argued that learning would be easier in SHRUTI that in ROBIN.

Learning a rule in SHRUTI involves identifying which features are relevant to the rule and interconnecting

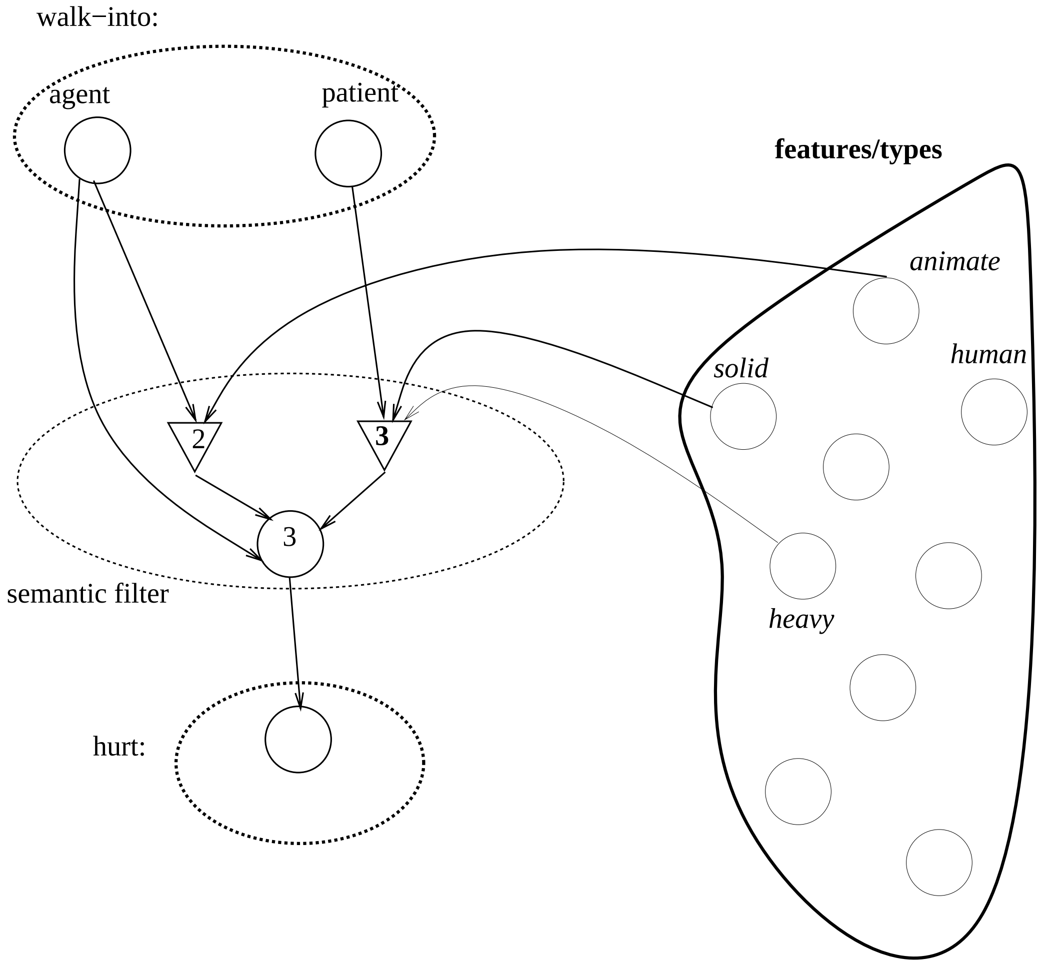

Figure 1: An example of a semantic filter for the rule: *x:animate* $\wedge$ *y:solid* $\wedge$ *y:heavy walk-into(x,y)* $\Rightarrow$ *hurt(x)*. The circular nodes are ρ-btu nodes and the triangular nodes are τ-or nodes. To make matters simple, let us assume that (i) a ρ-btu node with a threshold of k becomes active upon receiving k synchronous inputs and starts firing in synchrony with its inputs and (ii) a τ-or node with a threshold of k becomes active upon receiving k synchronous inputs and produces an uninterrupted burst of firing lasting π_{max}. The two τ-or together with the ρ-btu node in the semantic filter enforce the semantic constraint. The τ-or node on the left has a threshold of two and becomes active if the *agent* role of *walk-into* and the feature node *animate* fire in synchrony. In other words, it detects that the role-filler of *agent* is *animate*. Similarly, the τ-or node on the right detects that the role-filler of *patient* is *solid* and *heavy*. The ρ-btu node has a threshold of three and fires if both the τ-or nodes in the semantic filter become active along with the *agent* role of *walk-into*. Upon becoming active, it fires in synchrony with the *agent* role, since it receives three inputs only in the phase in which *agent* is firing.

them with appropriate role nodes to form a semantic filter. On the other hand, each role and predicate cluster in ROBIN must encode and process the complete signature of role fillers even though, many of the features occurring in a signature may be irrelevant to this particular rule. This can make learning more difficult. The use of pre-determined signatures can lead to other difficulties as well. Consider the following situation. Assume that a feature *f* is not included in the signature of some entity *E*. Subsequently, it is found that *f* is necessary for determining the applicability of a novel rule. How would a system such as ROBIN handle this situation? There seem to be two possibilities. First, the system may ignore the rule because learning it involves a feature that is not encoded in some of the signatures, and is therefore, unavailable at the predicate cluster during reasoning. Second, the system may include *f* in the signature of the entity and learn the novel rule. Doing so however, poses other problems. Since the signatures of some entities have been modified with the inclusion of *f*, other rules may have to be re-learned in order to respond correctly to the modified signatures. Furthermore, the size of the role clusters may have to be changed since these clusters must now accommodate larger signatures than before.

1.8 Other problems with signatures

Who decides which features of an object should be included in its pre-determined signature? What is the basis for including a particular feature in an entity's signature and excluding it from that of others? How does the system know in advance which set of features is going to be useful for learning rules? If we assume that signatures evolve with learning so that new features get added to existing signatures as learning progresses, how do changes in a signature induced in one domain affect the use of this signature in other domains?

Why should the same signature be used in all contexts? Why should the full signatures of role-fillers be propagated to a predicate structure even though most of the features in the signatures may be irrelevant for determining the applicability of rules attached to that predicate?

Finally, **Lange & Dyer** acknowledge that a solution based on the use of pre-determined signatures requires several times as many nodes as those required by the synchronous activation approach adopted in SHRUTI (para 8 Lange & Dyer). But they seem to underestimate the implications of increasing the size of the long-term memory network. While a complexity theorist may take an increase of a "small constant factor" in his stride, nature may not take kindly to the suggestion that the number of cells in the brain be multiplied by a factor of say, 20, in order to accommodate signatures.[3] It is unclear whether nature would prefer a scheme that increases the brain size by a factor of 20, over another scheme that offers all the advantages of the former but makes do with a network one twentieth in size.

2 SHRUTI and the representational theory of the mind

Bonatti distinguishes between three levels of commitment that a system might have to a *representational theory of the mind* (RTM). At the first level, a system can be said to be committed to RTM simply because it uses some form of representation. A second level of commitment to RTM means that the system embodies a rich representational device with its own "syntax" and rules of composition. Such a system can be viewed as using something akin to a "language of thought". Finally, a third level of commitment to RTM means that the system explicitly represents formal rules or inference schemas that are applied during inference to "sentences" in the language of thought.

Bonatti argues that SHRUTI is committed to RTM at levels one and two, and most likely, also at level three. He acknowledges that in its present form, SHRUTI is not committed to RTM at level three, but argues that if SHRUTI were extended to support a larger class of inferential behavior, it would be forced to make such a commitment.

The gist of this response is as follows: Yes, SHRUTI is committed to RTM at level one. It is also committed to RTM at level two — but in a sense that is qualitatively different from that entailed by the classical notion of such a commitment. SHRUTI, however, is *not* committed to RTM at level three.

Before elaborating further I would like to clarify that the work on SHRUTI is motivated by a desire to understand how a system of simple neuron-like elements can encode a large number of specific facts as well as systematic knowledge and perform a class of reasoning almost reflexively. We do not have a philosophical axe to grind and are comfortable being in partial agreement (and hence, partial disagreement) with competing and *seemingly* incompatible approaches to cognition. So while we are connectionists, we are not neo-behaviorists. We believe that a non-representational theory of cognitive function is untenable. But at the same time we find the classical position unacceptable in that it subscribes to a disembodied view of the mind wherein cognition can be studied without paying close attention to the computational properties of the brain.

As **Bonatti** observes, SHRUTI makes a definite commitment to representations. We view the long-term memory network (LTM) in SHRUTI as an internal representation of the agent's environment wherein long-term facts record specific situations and domain rules (i.e., meaning postulates) capture regularities in the environment. SHRUTI is also committed to RTM at level two since clearly, its encoding is systematic and has its own "syntax". Entities, domain rules, and facts are encoded in the network in specific ways. For example, each n-place relation is encoded by a cluster of nodes which in turn contains n distinct sub-clusters – one for each of the n roles of the relation. Similarly, long-term facts are composed of individual role-filler bindings, which can be attributed to specific connections between role and filler nodes. There are however, two important ways in which SHRUTI departs from the classical view outlined by Bonatti. First, SHRUTI does not share the classical dichotomy between operations (processes) and representation. In SHRUTI, the representations — i.e., the "data", and the processes that operate on the representations are indistinguishable. There is no separate interpreter that manipulates the representations according to formal rules or inference schemas. Second, nodes in SHRUTI *do not* acquire their meaning by virtue of the labels attached to them; just because we label a node "giver" does not make it the representation of the "giver". The labels on nodes are just meant to facilitate the reader in understanding the network encoding. A node (or a cluster of nodes) derives it meaning by virtue of its interconnections to other representational nodes, and eventually by being grounded in the sensorimotor, somatosensory, or emotional systems.

SHRUTI is not committed to RTM at level three. Reasoning in SHRUTI is the spontaneous and natural outcome of its behavior. SHRUTI does not apply syntactic rules of inference such as *modus-ponens*. The network encoding of LTM is best viewed as a vivid internal *model* of the agent's environment, where the interconnections between (internal) representations directly encode the dependencies between the associated (external) entities. When the nodes in this model are activated to reflect a given state of affairs in the environment, the model spontaneously simulates the behavior of the external world and in doing so makes predictions and draws inferences.

Bonatti agrees that in its current form SHRUTI does not encode abstract syntactic rules, but points to several "microproblems" that had not been modeled by SHRUTI and argue that once SHRUTI is extended to model these problems, it will have to explicitly represent abstract rules. While we have not modeled all the phenomena cited by Bonatti, we have extended SHRUTI in several ways since the appearance of (Shastri & Ajjanagadde 1993) and none of these extensions have required changes of the sort anticipated by Bonatti. For example, the SHRUTI (1993) did not deal with negated facts and domain rules involving negated antecedents and consequents. SHRUTI has now been extended to encode positive as well as negated knowledge and use such knowledge during reflexive reasoning (Shastri & Dean 1995). In addition to dealing with a richer class of knowledge and queries, the extended model explains how an agent can hold inconsistent knowledge in its memory without being "aware" that its beliefs are inconsistent, but detect a contradiction whenever it tries to make use of inconsistent knowledge. The extended model also explains how limited attentional focus or action under time pressure can lead an agent to produce an erroneous response. SHRUTI has also been extended to perform abductive inference (Ajjanagadde 1991). Neither of these two extensions have made SHRUTI more like a "classical box for reasoning."

While the result of the computations performed by SHRUTI (1993) could be explained in terms of applications of *modus ponens*, the model did not encode *modus ponens* as a formal rule. Similarly, while the computations of the extended SHRUTI system can be explained in terms of applications of *modus ponens* and *modus tolens*, the model does not encode either of these two formal rules. Furthermore, while SHRUTI supports inferences that correspond to applications of *modus ponens* and *modus tolens*, it does not do so uniformally for each domain rule. Given a particular regularity in the domain, say "if situation P occurs then situation Q also occurs", the agent's cognitive apparatus may chose to encode it in the LTM only in its direct form (i.e., activation of P would lead to an activation of Q but an activation of $\neg Q$ would not lead to the activation of $\neg P$), or only in its contrapositive form (i.e., an activation of $\neg Q$ would lead to an activation of $\neg P$, but an activation of P would not lead to the activation of Q), *or both*, depending on the sorts of situations in which the agent puts this knowledge to use. Thus an agent may be able to use some domain knowledge reflexively only in its direct form, but require *reflective* processing to use the same knowledge in its contrapositive form. In some cases the opposite might hold and the agent may be able to use some domain knowledge reflexively only in its contrapositive form.

With regards to conditional queries (**Bonatti** para. 10), simple forms of such queries can be handled by SHRUTI without requiring any significant change in its architecture. To see this, consider how SHRUTI deals with declarative and interrogative inputs (i.e., dynamic facts and queries). SHRUTI is assumed to be part of a larger cognitive system containing linguistic and perceptual processes that can communicate with SHRUTI. Thus given a declarative input such as "John bought a car", a linguistic process is expected to communicate this input to SHRUTI by activating the *collector* of *buy* and establishing the dynamic bindings (*buyer=John, buy-object=a-Car*). Similarly, when presented with a query such as "Did Susan buy a Book?", a linguistic process is expected to activate the *enabler* of *buy* and establish the dynamic bindings*(buyer=Susan, buy-object=a-Book)*. A respone to the query is generated based on the state of activation of the *collector* of *buy*. SHRUTI can deal with simple types of conditional queries in an analogous manner. We assume that a linguistic process will communicate a conditional query such as "If John bought a car, would he own it?" as follows: It will activate the *collector* of the antecedent, *buy*, and establish the bindings (*buyer=John, buy-object=a-Car*). The response process will observe the predicate *own* to ascertain if its *collector* becomes active along with the appropriate dynamic bindings of its roles.[4] Thus given a linguistic processor, SHRUTI should be able to handle conditional queries in much the same way as it handles declarative and interrogative inputs. Needless to say, the linguistic processor will have to have appropriate lexical entries and syntactic knowledge to process its input, but none of this makes SHRUTI more like a classical box of reasoning. The linguistic process only serves as an interface to SHRUTI and its presence does not change the manner in which SHRUTI performs inference.

We are not claiming that people cannot learn to apply *modus ponens* as an abstract syntactic rule. People can certainly learn to do so — along with a host of other complex skills. Our central claim is about the mechanisms by which people can perform reflexive reasoning with reference to a large body of well assimilated background knowledge. We believe that reflexive reasoning is not carried out in our mind/brain by an interpreter that scans a knowledge base of sentences "written" in some language of thought and manipulates them according to abstract and formal rules of inference. We believe instead, that our mind/brain performs reflexive reasoning in a manner analogous to that described in SHRUTI wherein reasoning is the transient propagation of a rhythmic pattern of activity over a memory network, rules are interconnection patterns that cause the propagation and transformation of rhythmic patterns of activity, and long-term facts are subnetworks that act as temporal pattern matchers that become active under suitable circumstances and create reverberatory patterns of activity.

References

Ajjanagadde, V. G. (1991) Abductive reasoning in connectionist networks: Incorporating variables, background knowledge, and structured explananda. Technical report WSI-91-7. Wilhelm-Schickard-Institute, University of Tuebingen, Germany.

Henderson, J. (1994) Connectionist Syntactic Parsing Using Temporal Variable Binding. *Journal of Psycholinguistic Research*, 23 (5) p. 353–379.

Henderson, J. (1995) Personal Communication.

Just, M.A. & Carpenter, P.A. (1992) A Capacity Theory of Comprehension: Individual Differences in Working Memory. *Psychological Review*, 99 (1), 122-149.

Shastri, L. & Ajjanagadde V. (1993) From simple associations to systematic reasoning: A connectionist encoding of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16:3 p. 417-494.

Shastri, L. & Dean, J.G. (1995) Dealing with negated knowledge and inconsistency in a neurally motivated model of memory and reflexive reasoning. TR-95-041, International Computer Science Institue, Berkeley. August 1995.

Endnotes

1. This work was supported in part by ONR grants N00014-93-1-1149 and N00014-95-C-0182.

2. In addition to the four entities referred to in the first sentence, the other five entities are: (i) "Two months later" (a temporal landmark), (ii) "(new) furniture", (iii) "police" (iv) "drugs", and (v) "walls (of the red brick house)". Note that "wall", which is a part of the entity "house", is subsequently introduced as a distinct entity because specific attributions are made about the wall.

3. This assumes that on an average, a signature would contain 20 features. The actual constant of multiplication will depend on the number of features in signatures. Since a signature is intended to include all the salient features of an entity, this assumption seems plausible. The number 20, however, is not crucial to the argument since even a factor of 2 is sufficiently high.

4. The soultion sketched out here does not distinguish between assertions and suppositions. In certain circumstances this distinction will have to made.