Baron, J. (1990) Thinking about consequences. *Journal of Moral Education* 19:77–87. [rJB]

(1994a) Nonconsequentialist decisions. *Behavioral and Brain Sciences* 17:1–42. [RB-M, MC, ROL, KES, rJB]

(1994b) *Thinking and deciding*, 2d ed. Cambridge University Press. [rJB]

(1994c) Normative, descriptive and prescriptive responses. *Behavioral and Brain Sciences* 17:32–42. [rJB]

Bereiter, C. (1984) The limitations of interpretation. *Curriculum Inquiry* 14:211–16. [KES]

Brown, J. S., Collins, A. & Duguid, P. (1989) Perspectives on socially shared cognition. *Educational Researcher* 18(1):32–42. [KES]

Cabanac, M. (1992) Pleasure: The common currency. *Journal of Theoretical Biology* 155:173–200. [MC]

Gorayska, B. & Lindsay, R. O. (1989) Metasemantics of relevance. In: *Proceedings of the first symposium on cognitive linguistics*. Linguistic Association of the University of Duisberg. [ROL]

Gorayska, B. & Lindsay, R. O. (1992) The roots of relevance. *Journal of Pragmatics* 19(4):301–23. [ROL]

Hardin, G. R. (1968) The tragedy of the commons. *Science* 162:1243–48. [rJB]

Jonas, H. (1984) *The imperative of responsibility: In search of an ethics for the technological age*. University of Chicago Press. [rJB]

Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. *Psychological Review* 80:237–51. [ROL]

Koriat, A., Lichtenstein, S. & Fischhoff, B. (1980) Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory* 6:107–18. [rJB]

Lakomski, G. (1992) Unity over diversity: Coherence and realism in education research. *Curriculum Inquiry* 22:191–203. [KES]

Langer, E. (1989) *Mindfulness*. Addison-Wesley. [KES]

Larkin, S. & McFarland, D. J. (1978) The cost of changing from one activity to another. *Animal Behaviour* 26:1237–46. [MC]

Lave, J. & Wenger, E. (1991) *Situated learning: Legitimate peripheral participation*. Cambridge University Press. [KES]

Lindsay, R. O. & Gorayska, B. (1989) On relevance: Goal dependent expressions and the control of planning processes. In: *Department of computing and mathematical sciences research report no. 16*. Oxford University Press. [ROL]

Nisbett, L. & Ross, L. (1980) *Human inference: Strategies and shortcomings of social judgment*. Prentice-Hall. [KES]

Pascal, B. (1623–1662) *Pensées*. [MC]

Posner, R. A. (1992) *Economic analysis of law*, 4th ed. Little, Brown. [rJB]

Rawls, J. (1971) *A theory of justice*. Oxford University Press. [KES]

Resnick, L. B. (1991) Shared cognition: Thinking as social practice. In: *Perspectives on socially shared cognition*, ed. L. Resnick, J. Levine & S. Teasley. American Psychological Society. [KES]

Resnick, L. B., Levine, J. & Teasley, S., eds. (1991) *Perspectives on socially shared cognition*. American Psychological Association. [KES]

Rogoff, B. & Lave, J., eds. (1984) *Everyday cognition*. Harvard University Press. [KES]

Schrag, F. (1992) In defense of positivist research paradigms. *Educational Researcher* 21(5):5–8. [KES]

Todd, E. (1991) *L'invention de l'Europe*. Editions du Seuil. [MC]

Commentary on L. Shastri and V. Ajjanagadde (1993) From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. BBS 16:417–494.

Abstract of the original article: Human agents draw a variety of inferences effortlessly, spontaneously, and with remarkable efficiency as though these inferences were a reflexive response of their cognitive apparatus. Furthermore, these inferences are drawn w[ith] reference to a large body of background knowledge. This remarkable human ability seems paradoxical given the complexity [of] reasoning reported by researchers in artificial intelligence. It also poses a challenge for cognitive science and computatio[nal] neuroscience: How can a system of simple and slow neuronlike elements represent a large body of systemic knowledge and perfor[m a] range of inferences with such speed? We describe a computational model that takes a step toward addressing the cognitive scie[nce] challenge and resolving the artificial intelligence paradox. We show how a connectionist network can encode millions of facts and ru[les] involving n-ary predicates and variables and perform a class of inferences in a few hundred milliseconds. Efficient reasoning requi[res] the rapid representation and propagation of dynamic bindings. Our model (which we refer to as SHRUTI) achieves this by represent[ing] (1) dynamic bindings as the synchronous firing of appropriate nodes, (2) rules as interconnection patterns that direct the propagation [of] rhythmic activity, and (3) long-term facts as temporal pattern-matching subnetworks. The model is consistent with recent neu[ro]physiological evidence that synchronous activity occurs in the brain and may play a representational role in neural informat[ion] processing. The model also makes specific psychologically significant predictions about the nature of reflexive reasoning. It identifi[es] constraints on the form of rules that may participate in such reasoning and relates the capacity of the working memory underly[ing] reflexive reasoning to biological parameters such as the lowest frequency at which nodes can sustain synchronous oscillations and th[e] coarseness of synchronization.

SHRUTI's ontology is representational

Luca Bonatti

Laboratoire des Sciences Cognitives et Psycholinguistique, 54, Bd. Raspail, 75006 Paris, France. **luca@afize.msh-paris.fr**

Abstract: I argue that SHRUTI's ontology is heavily committed to a representational view of mind. This is best seen when one thinks of how SHRUTI could be developed to account for psychological data on deductive reasoning.

A representational theory of mind (RTM) sees mental processes as operations over representations. In this framework, reasoning too may be seen as a special set of such operations. The view, which seems obvious at a first blush, turns out to raise many philosophical riddles. One crucial question concerns the means by which representations are internally encoded. There must be such a means, if the view is correct, but not all representations will do. For example, images are representations, but they are too poor [to] articulate thoughts and trains of thoughts (Fodor 1975). He[nce] one is led to the view that the representational device of a think[ing] being must be as rich as a full natural language with a synta[ctically] typified structure, and rules of compositions preserving consti[tu]ency relations – call it, for lack of fantasy, a language of thou[ght] (LOT). And from there, one finds harder and harder questio[ns.] How is such an internal language organized? Which kind [of] evidence could inform us about its nature? How are its r[ules] internally represented? What is its relation to observable, o[r] natural language? And what keeps in phase two things so differ[ent] as an internal language and the external world?

The temptation to cut all such riddles at their roots is strong [and] has pushed an increasing number of philosophers and scientist[s to] embrace connectionism, which has proposed an apparently a[lter]native picture, making talk of LOT, rules, and representati[ons] largely unnecessary.[1] Yet, in the domain of deductive reaso[ning] the alternative view has so far produced systems of scarce, if

...ical interest. We owe to Shastri and Ajjanagadde (1993) the first serious connectionist alternative in this area. Is their system, SHRUTI, also an alternative to LOT and RTM, as some commentators have suggested (Oaksford & Malloch 1993)? At first sight, it may seem so; one might be led to believe it when told that SHRUTI "does not apply syntactic rules of inference such as modus ponens. There is no separate interpreter or inference mechanism that manipulates and rewrites symbols" (p. 420); rather, it "spontaneously simulates the behavior of the external world, and in doing so makes predictions and draws inferences" (p. ...).

I intend to show that SHRUTI is another source of support for RTM, in that it lends support to the belief that once a system aims to model real psychological processes, it needs to make explicit use of representations. One should distinguish three different levels of representational commitment: (1) to representations of a certain sort, (2) to a LOT, and (3) to explicit representations of rules. An RTM-like system is committed at least to (1) and (2), and possibly, but not necessarily, to (3) (Fodor 1987, Ch. 1 and Appendix). I claim that SHRUTI too is committed to (1), (2), and most likely to (3). This is *not* to say that SHRUTI is a "*mere* implementation" of a classical system. There are important architectural consequences proper of SHRUTI and of no other system (see Shastri & Ajjana-gadde 1993, sect. 8). However, for what concerns representational commitment, SHRUTI is no alternative to RTM. I will proceed as follows. I will first recall the features of SHRUTI pertinent to my thesis. I will then argue for my conclusion.

SHRUTI's way with reflexive reasoning. SHRUTI intends to show how dynamically generated facts (facts the system is given "on the spot," without consulting its internal memory) can interact with long-term structures to generate quick inferences. In particular, it implements a theory of meaning postulates (Fodor et al. 1975). Such a theory was meant to account for nonlogical inferences among predicates and the solution it proposed was to list the relevant inferential relations explicitly. For example, (4) represents the derivational connection between 'give' and 'own':

4. a gives x to $b \rightarrow b$ owns x.

Thanks to (4), if you know that John gives a book to Mary, you quickly (reflexively) conclude that Mary owns the book. For a system to do so, the thematic relations between the arguments of "give" and "own" must be respected, and this is a tough problem for connectionist networks. S&A solve it by exploiting two ideas. First, they enrich the nodes for predicates with labels identifying thematic roles. So (by leaving aside other nodes not relevant to the present point) "give" is represented as in Figure 1 and (4) is represented by wiring in the connections between the thematic roles, as in Figure 2. It is very important that nodes be labeled; otherwise the system will misbehave. Likewise, to dynamically represent a fact such as

5. John gives a book to Mary,

SHRUTI has to be explicitly given the full information concerning not only the predicates and individuals involved in the fact, but their roles as well. This is achieved by activating a set of nodes that includes specific nodes for thematic roles. Second, S&A use temporal synchrony to assure that predicate arguments are correctly assigned their thematic roles: the nodes for the arguments and the nodes for their roles will fire synchronously. So, a synchronous pattern of activation such as the one generated by the set of nodes for (5) matches the antecedent of (4); this activates the consequent of (4), and SHRUTI ends up to dynamically represent another fact, (6),

6. Mary owns the book

which is the desired inference.

Representations and psychological adequacy. It is tempting to conclude that, first, SHRUTI does not need general purpose syntactic rules, because the derivation of (6) does not invoke one and second, that it may even do without a LOT, because, after all, in SHRUTI there are only nodes activating other nodes and not symbol manipulations. Two points should caution against such general conclusions, however.

The first point concerns SHRUTI's *other* representational commitments. It is true that SHRUTI does not contain explicit representations of syntactic rules such as *modus ponens*, but this only means that its representational commitment does not fully embrace level (3). As for (1) and (2), remember that node labels have to be taken seriously. They determine whether a fact is correctly expressed, or whether an inference will be triggered. So, the functionally relevant primitives in SHRUTI are nodes *with their labels*, and labeled nodes have the same roles as typified symbols in a language.[2] In fact, it is easy to match each of SHRUTI's activation states with formulas in a language: for example, the activation structure for (5) can be paraphrased thus:

7. John is the giver [the node for 'John' and the node for 'Give' fire synchronically], Mary is the receiver [the node for 'Mary' and the node for 'Receiver' fire synchronically], and book is the object [the node for 'Book' and the node for 'Given object' fire synchronically].

SHRUTI already contains temporal "and" and "or" nodes realizing logical connectives; once the roles of labels are also considered, one can see that SHRUTI implements a LOT, by means of which it explicitly represents the data structures expressing dynamic facts. So SHRUTI is committed to (1) and (2).

S&A do not run away from representational commitments. In criticizing the adequacy of distributed representations, they write that "a distributed system – at least in its pristine form – *cannot* have the necessary combination of expressiveness, inferential

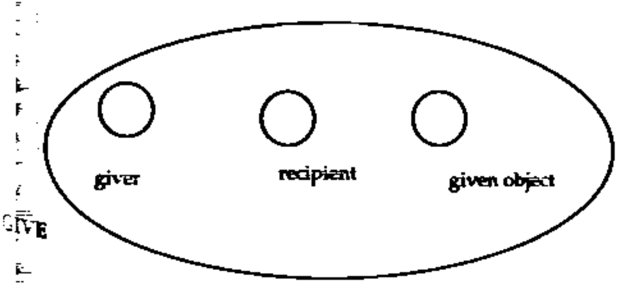

Figure 1 (Bonatti). The representation of "give" in SHRUTI.

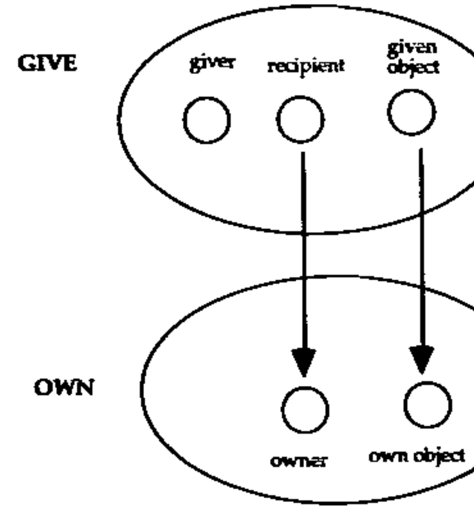

Figure 2 (Bonatti). The representation of a meaning postulate in SHRUTI.

adequacy, and scalability. . . .There is a basic tradeoff between distributed representation, systematicity, and parallelism: no amount of handwaving can make this tradeoff disappear" (p. 485). The logical conclusion is that SHRUTI is no alternative to RTM.

Nevertheless, one may argue that SHRUTI reduces the need for explicit representations, and in this sense it also reduces the role of representations in reasoning. Notably, it would show that general-purpose, syntactic rules for deriving inferences need not be explicitly represented. My second point, then, is that once SHRUTI is made psychologically more adequate for modeling reasoning, it may even be committed to type (3) representations.

S&A point out a number of extremely interesting psychological consequences of their model, but these are very broad architectural properties of a cognitive system, not specifically geared to testing theories of deductive reasoning. A plausible model for reasoning should address many microproblems that SHRUTI does not yet handle. For example, everybody masters the rule of double negation, or excluded middle, but not everyone masters unconstrained *reductio* (Braine et al. 1984). Or, to give another example, people are very good at *modus ponens*, but also fairly good at *modus tollens;* however, they do *modus tollens* better with "only if" problems than with "if then" problems (Johnson-Laird et al. 1992). Many tricky phenomena of this kind exist and call for an explanation. There may be no principled difficulties in developing SHRUTI to account for facts like these, but the system in its present form does not do it.[3]

How *could* it do it? Notice first that SHRUTI is already partially committed to level (3). Meaning postulates are explicitly represented, since for them as well one can decode labeled nodes and find unique corresponding formulas. Yet it does not represent "syntactic rules of inference such as *modus ponens*" (p. 420). This is because it does not do *modus ponens* at all: in fact, it derives certain inferences by explicitly representing meaning postulates, and *by locally building into them* the required abstract inferences. The reason for representing meaning postulates is presumably that people are good at inferences captured by them. But if it turns out that they are also good at inferences unrelated to specific contents, such as *modus ponens*, then the most natural way to incorporate such finding in SHRUTI would be presumably to treat abstract rules just like SHRUTI treats meaning postulates. This could be done by extending the notion of dynamic fact to dynamic abstract rules. We may imagine that, instead of accepting only declarative queries about facts, SHRUTI can also be given conditional queries, which create a dynamic pattern of activation such that a static fact, or a further dynamic declarative input, allows releasing their consequents. But then SHRUTI would contain a procedure that guarantees that once the antecedent of a rule/query fires, so does its consequent – that is, it would contain *modus ponens* as part of an explicit representation of the lexical entry for "if" (Braine & O'Brien 1991). Notice that even if such a modification would be a substantial enrichment of SHRUTI's reasoning abilities, it would not be a major theoretical departure. SHRUTI already contains two logical connectives (in the form of special nodes representing the truth conditions for temporal "and" and "or"); to allow conditional queries would amount to explicitly adding another connective, as abstract as logic wants it to be.

This argument applies to other possible nonreflective abstract rules for reasoning. It seems that there is a quite rich – though not too rich – set of abstract rules people use reflexively while reasoning (Braine et al. 1984). The natural way to adapt SHRUTI to accommodate these abilities would be to represent such rules explicitly just as meaning postulates are. Then the set of such rules would correspond to the content of a classical box for reasoning, realized in a parallel architecture.

In conclusion, SHRUTI cannot do without (1) and (2), but probably is committed even to (3). It is just a matter of developing it far enough to make the representational ontology emerge out of its nodes. However, this does not mean that SHRUTI *merely* implements a classical system. There are other psychological consequences that it has and classical system do not have. I[illegible] original contribution has to be found there.

ACKNOWLEDGMENT

This work was carried out with support from the EEC Human Capital a[illegible] Mobility Program, contract ERBCHB1-CT94-1325. I thank Profe[illegible] Shastri for useful comments on a first draft of this commentary.

NOTES

1. Specifically referring to reasoning, Rumelhart (1980, p. 28) wr[illegible] "some believe that the existence of inferences implies a kind of log[illegible] system similar to that employed in conventional symbolic-process[illegible] models. I have become increasingly convinced that much of what we [illegible] reasoning can better be accounted for by processes such as patt[illegible] matching and generalization, which are well carried out by PDP mode[illegible] And Bechtel and Abrahamsen (1991, p. 174) write: "if connectionists [illegible] able to provide an account on how [patterns involved in logical proofs] [illegible] recognized, as seems quite plausible, then we shall not have to try [illegible] formulate logical expertise in terms of a set of mental rules or procedur[illegible] rather, we can treat it as a quite different sort of knowledge."

2. Labeled nodes may stand for clusters of more elementary nodes, [illegible] this is not important functionally; only the labeled nodes determine [illegible] behavior of the system, and these are functionally equivalent to elem[illegible] of a language.

3. SHRUTI has many logical limitations. However, this is not necess[illegible] a shortcoming for a system meant to explain *reflexive* reasoning (co[illegible] Dawson & Berkeley 1993). The real question is whether subjects [illegible] exhibit similar limitations when reasoning reflexively. This question ca[illegible] be decided a priori: it is a matter of psychological investigation. And [illegible] point is rather that some of SHRUTI's logical limitation are excessive on [illegible] light of what we know about deductive reasoning.

Parallel reasoning in structured connectionist networks: Signatures versus temporal synchrony

Trent E. Lange and Michael G. Dyer

Artificial Intelligence Laboratory, Computer Science Department, Univers[illegible] of California, Los Angeles, CA 90024. lange@cs.ucla.edu; dyer@cs.ucla.edu

Shastri & Ajjanagadde (1993) (S&A) argue convincingly that [illegible] structured connectionist networks and parallel dynamic infer[illegible] ing are necessary for reflexive reasoning – a kind of inferen[illegible] and reasoning that occurs rapidly, spontaneously, and wit[illegible] conscious effort, and which seems necessary for everyday [illegible] such as natural language understanding. As S&A describe, re[illegible] ive reasoning requires a solution to the *dynamic binding prob[illegible]* that is, how to encode systematic and abstract knowledge [illegible] instantiate it in specific situations to draw appropriate inferen[illegible] Although symbolic artificial intelligence systems trivially solve [illegible] dynamic binding problem using computers' registers and poin[illegible] it has remained a difficult problem for connectionist systems [illegible] Fodor & Pylyshyn 1988). S&A's temporal synchrony solution t[illegible] dynamic binding problem using synchronous firing of argu[illegible] units and the entities that are bound to them illustrates one w[illegible] which connectionist networks can do this using a constraine[illegible] important class of long-term knowledge rules. Their struct[illegible] connectionist solution allows dynamic inferencing to proce[illegible] parallel and therefore has a number of advantages for refl[illegible] reasoning over most other connectionist and symbolic sys[illegible]

An alternative structured connectionist solution to the dy[illegible] binding problem that has similar advantages is to represent [illegible] concept in the network with a unique distributed patte[illegible] activation (its *signature*) and to represent each argument [illegible] ensemble of binding units (Lange & Dyer 1989). An argum[illegible] bound to a given value when the activation across its binding [illegible] is the same as the signature pattern representing that [illegible] Dynamic bindings are therefore represented as signatures a[illegible] propagated through network connections between appro[illegible] argument units much as in S&A's system. The "pattern co[illegible]