Information Networks are Better for Cognition than Symbolic Dynamics

Orlin Vakarelov*
Department of Philosophy
Duke University
West Duke Building
Box 90743
Durham, NC 27708

1 Introduction

The 1990s witnessed a number of debates in cognitive science about the proper framework for the investigation of cognition.[1] This was a result of the maturing of the connectionist program that (re)started in the '80s and the emergence of the dynamical system approach to cognition (Rumelhart & McClelland, 1986; Rumelhart, 1989). Both were offered as alternatives to the symbolic/information processing approach that has been prevalent since the '70s. With the extensive use of dynamical systems techniques in the analysis of connectionist architectures, the framework debate focused on the contrast between dynamical systems and symbol systems.[2] Since then, the framework debates left the foundational problem of cognitive science with an open problem: is there a proper unifying architectural framework for cognition? We can describe this problem as the *framework problem* for cognition.

A possible approach to the framework problem is to deny its importance by adopting a pluralistic approach. Such an approach accepts that there may be different architectural frameworks without any unifying structure.[3] Recently,

*This research was been supported by a New Faculty Fellows award from the American Council of Learned Societies, funded by the Andrew W. Mellon Foundation.

[1]See, for example, Haugeland (1997). See also Smolensky (1986); Fodor & McLaughlin (1990); Prince & Smolensky (1997); Fodor (1997).

[2]It should not be supposed that connectionism is just a version of dynamism. There is a deep connection between the dynamical approach and embodied/enactive approaches to cognition (Varela *et al.*, 1992; Riegler, 2002; Calvo & Gomila, 2008; Chemero, 2009; Shapiro, 2011), which differ in more fundamental ways from the symbolic approach. Connectionism is closer to the symbolic approach in this regard. This rift is ultimately more important, but it will not be the subject of this paper.

[3]See for example, the discussion in the special issue of the Journal of Experimental & Theoretical Artificial Intelligence dedicated to pluralism (Dale, 2008a), as well as the discussion in Dale *et al.* (2009).

there has been a wave of attempts to reconcile the dynamical approach with information processing and symbol processing, as a species of information processing. Such attempts aim to understand how notions such as computation, information processing, and symbol processing can be captured as dynamical systems phenomena. A central principle of the attempt is that the dynamical systems description is more general than a symbolic description, and captures phenomena in natural cognition (as well as offering tools for artificial cognition) that go beyond the symbolic paradigm. Symbol manipulation and digital computation are acknowledged as important tricks that many cognitive systems use, especially those that are advanced. Thus, symbol system analysis has important explanatory significance in the description of higher cognition. Such attempts may be viewed as strategies to provide a more general and encompassing architectural framework.

Reconciliation attempts must meet three conditions of non-triviality:

Novelty: Dynamical systems should not offer merely an implementation level description of otherwise symbolic systems. Even the staunchest proponents of the symbolic approach would agree that there is an underlying implementing system for every symbol system.

Non-collapse: The notion of information processing should not be so general as to apply to every possible dynamical system. Some conceptions of information, such as Shannon's, and associate conceptions of information processing, are so general that they can be used to describe any system. If such conceptions are adopted, a more fine-grained distinction must be offered. The distinction must identify the kind of information processing that is important and unique for symbol systems.

Significance: Whatever notion of information/symbol processing is defined within a dynamical system, it must have explanatory significance for understanding the function and behavior of cognitive systems, *qua* cognition.[4]

A number of authors have proposed that the mathematical theory of symbolic dynamics can be used as a general framework for bridging the dynamical and symbolic approaches. Outside of the specific debates[5], Crutchfield has used symbolic dynamics to analyze the information processing capacities of dynamical systems, and has suggested, based on a hierarchy of $\varepsilon-$machine reconstructions (see Section 2.2), how an artificial agent may engage in a creative activity (Crutchfield, 1994; Crutchfield & Young, 1990). Beim Graben (2004) has used symbolic dynamics to evaluate a debate between Smolensky and Fodor about the incompatibility between the connectionist and symbolic approaches, arguing that the claim about incompatibility is mistaken. Dale (Dale, 2008b; Dale &

[4]This condition is really a refinement of non-collapse, but it is a bit vague due to ideas such as "explanatory significance". In programmatic discussions it may be difficult to assess if it is met. (So, non-collapse may be all one may use.) But, it may be used as a guide for a more mature development of a theory.

[5]But see his comments on the BBS paper by van Gelder (1998).

Spivey, 2005), Shalizi (2005) and Edelman (2008) have offered a more systematic analysis of the ability of symbolic dynamics to provide a framework that unites the two approaches.

The enthusiasm is premature! The goal of this paper is to show that, while on the right track, symbolic dynamics is insufficient by itself to serve as the needed framework. The problem is that the different ways it can be used to connect the dynamical systems approach to information/symbol processing fail (at least) one of the non-triviality conditions. More machinery is needed to properly capture a notion of information processing for cognition. It is suggested that a related notion of information medium network (Vakarelov, 2012) offers more promise.

2 Symbolic Dynamics

2.1 Basic Concepts

The formal theory of symbolic dynamics deals with bi-infinite sequences from a fixed alphabet of "symbols" A. It studies the nature of transformations of such sequences and the nature of spaces of such sequences (shift spaces). As such, the theory has an important connection to the theory of computability. More broadly, symbolic dynamics deals with the representability of general dynamical systems, *viz.* the temporal dynamical properties of the systems, as sequences (or orbits) of symbols.

Let $D =< S, T >$ be a dynamical system with a state space S and a temporal evolution function $T(s_0, t) = s_t$, such that if the system is at state s_0 the system will be in a state s_t after time t.[6] Let $A = \{A_i | i < n)\}$ be a finite partition on the state space. Then, given an initial state $s_0 \in A_i$, the temporal evolution of the system defines the dynamics over the partition T_A such that if $T(s_0, t) = s_t$, and $s_t \in A_j$, then $T_A(A_i, t) = A_j$. Put informally, the trajectory of the underlying dynamics generates a time-sensitive stream of partition sets. If we interpret the elements of the partition as "symbols", then the dynamics generate a stream of symbols. The partition and the initial conditions define a map from the dynamical system, as a geometric object, to a set of "symbol" sequences.

As formulated, it is clear that the conception of symbolic dynamics is very general. This goes beyond the fact that the notion of a dynamical system is general. The generality comes from the arbitrariness of the partition. Consider a ball (or a cart) rolling on a hill, as a one-dimensional dynamical system of position. By choosing an appropriate partition, we can generate any arbitrary sequence of symbols.[7] Clearly then, for symbolic dynamics to be an interesting concept, one must have an interesting way of selecting the partitions.

There are several natural questions one may ask in the definition of a symbolic system. (1) Given a dynamical system and a partition, how do changes

[6] If the evolution is map from probability distributions to probability distributions over the space, then the dynamics is stochastic.

[7] If we interpret the sequence as a computation process, such simple dynamical systems can be regarded as performing arbitrary computations. This has led some (e.g. Maudlin (1989)) to question the usefulness of the computational metaphor in such cases.

of the initial conditions affect the generated symbolic sequences? (2) Given a desired symbolic behavior, are there some dynamical systems (meeting some requirements) and partitions (meeting some requirements) that generate the desired behavior? (3) Given some natural or available partition, what can we learn about the nature of the dynamical system from the symbolic sequences?

Question (1) we can always ask. So, there must be independent reasons for looking at the particular dynamical system and partition. Question (2) is the kind of question an engineer might ask. For example, the task of designing a physical computer can be interpreted as building a physical dynamical system and defining a proper partition on its possible states, so that by controlling the initial conditions (memory state) one can obtain a desired sequence of computation. The question has wider applications, however; building an ammeter can be interpreted in the same way. Question (3) is particularly interesting for a scientist investigating the dynamical system itself by looking at an observable (a reading of an instrument). The partition is defined by the observable (the instrument) and the scientist wants to know what information the sequence of outputs from the instrument — the symbol sequence — conveys about the system.

2.2 Information Processing

The interest of symbol dynamics to the framework problem of cognitive science is closely related to the possible connection to information processing. It is important to examine how some notions of information processing may be connected to symbolic dynamics, in order to assess the promise to cognitive science. There are several distinct connections that must not be confused.

Let us look first at an overly simplistic connection. One may be tempted to reason as follows: symbolic dynamics is about how symbols, as digital states, can be identified within the operation of a dynamical system. But, information processing is nothing but symbol processing. So, of course symbol dynamics is related to information processing.

There are many things wrong with such an argument. First, the fact that the word "symbolic" is used in the name of a perfectly well-defined mathematical domain is not a reason to think that the domain captures anything fundamental about symbols and symbol manipulation. There may indeed be motivating uses of symbol dynamics that relate to symbols or symbol systems (as defined by Newell & Simon (1981)), but this does not automatically reveal anything deep and interesting. (I will return to symbols later.) Second, the fact that a dynamical system with a finite partition may generate a sequence of distinct "symbols" does not automatically make it a symbol processor in any interesting way. Of course, it could be an interesting "symbol" system, but saying when that is the case (what is special about the dynamical system, the partition, or something else) is the hard question. Indeed, as noted above, by carefully choosing the partition, rather uninteresting dynamical systems can generate arbitrary "symbol" sequences. All such observations lead to is the conclusion that symbolic dynamics has the potential of defining some systems that may be interpreted as information processing, leaving the question how that may be the case unanswered.

Now, if this is all symbolic dynamics could achieve, then it would not be terribly interesting for the framework problem. Fortunately, there are better connections to information processing. Consider again question (2) above. We can ask, can a dynamical system implement a computational process that meets some functional characterization? Thus, we can ask: given an interpretation of what constitutes an input and an output, can a given dynamical system implement a computable function? Traditionally, interpretation of inputs and outputs is given as states of the system, where input is given as (a set of) initial conditions, and output as (a set of) end states. Symbolic dynamics allows for other conventions, where, for example, outputs may be given as segments of the symbol sequence generated over time. A more interesting question, for the framework problem, is whether the dynamical system simulates a local symbol-guided computational process, *viz.* its trajectory through state space.[8] Symbolic dynamics may be useful there too, as it allows an analysis of the dynamics in terms of a sequences of discrete states. Caution must be exercised, however, because the "symbols" of the sequence may not be the relevant "symbols" of the symbol-guided computation process. In the Newell/Simon conception of a symbol system, the transition rules are based on symbolic expressions, which are complex symbols with "internal structure". The "internal structure" is important for the guiding of the computation processes. In symbolic dynamics, the "symbols" are atomic. They may hide, rather then reveal, structural and causal relations in the underlying system that may be responsible for the simulation of symbol-guided processes.[9] The connection to computation in relation to question (2) does not depend on the dynamical system alone, but depends on the selection of a (potentially arbitrary) partition.

Question (3) shifts the role of the dynamical system from a system possibly implementing a computation process to a system subject to investigation by an agent. Beim Graben (2004) describes this as an epistemic (as opposed to ontological) analysis of a dynamical system.[10] The idea is that the dynamical system is observed via a non-arbitrary observable (formally, a function from the state space to some other space). Some states are indistinguishable under the observable, thus, the observable defines a partition on the state space of indistinguishable states. This need not be a finite partition, so it need not induce a symbolic dynamics yet. The observable, however, may be subject to some measurement error ε. Thus, the observable may have only finite resolution. In this case, the induced partition is finite. So far, nothing relevant for information processing. However, Crutchfield (1994) modifies the idea to analyze the

[8]A local symbol-guided process is a process where the transition rules are defined on syntactic properties of local (i.e. not total state) symbols.

[9]Examination of the sequences of symbols may reveal the structural and causal relations, but this is a different issue.

[10]The terminology is derived from analysis of quantum-mechanical systems of observables. It does not come from metaphysics. Here "ontological analysis" is an analysis of the system *as it is,* and "epistemic analysis" is analysis of the system as observable by an agent. "Epistemic analysis" is not analysis related to knowledge, as is common in philosophical discussions.

"intrinsic information processing" capacities of the dynamical system.

Crutchfield assumes the more general case of a stochastic dynamical system and considers a version of question (3) of the following form: looking at a time series up to a time t_o of an observable with some error ε, what can be predicted about the future of the time series? It is assumed that one needs only a finite length of past observations and one needs to predict only a finite length of future observations, both depending on ε. Here, a prediction is a specification of a probability distribution on the possible futures. This allows us to make further identification of states in the space, when two states lead to exactly the same predictions of the future. The partition generated by this further equivalence relation Crutchfield calls the "causal states" of the system up to ε. The symbolic dynamics generated by the original observable is converted into a symbolic dynamics of the causal states. The resulting symbolic sequence of "causal states" is called an $\varepsilon-$machine and is assumed to isolate the information processing of the dynamical system. ε-machines, which are stochastic machines, can be connected to Bernoulli-Turing machines, which are Turing machines with a random-number generating oracle. The idea is further (or primarily) used to define a measure of complexity similar to algorithmic complexity, but looking at minimally specified ε-machines instead, which are able to filter out the random part the the process, as randomness comes for free.

3 Relation to Cognitive Frameworks

Now let us assess how the tools of symbolic dynamics can be used with the aim of providing a general framework for cognition. The promise of symbolic dynamics is to offer a framework where dynamical system and symbolic approaches may coexist. The idea is that the dynamical system approach, which is very general and undiscerning, may be enhanced with tools that help us understand how information processing, and symbol-guided computational processing in particular, happens naturally within the dynamics.

First, let us look at the idea of symbols, to see how fair (or deceptive) the term "symbolic dynamics" is. There are actually two different senses of the term "symbol". One is a simple formal sense, requiring a discrete collection of objects that may be subject to recognition (identification and mutual distinction). The other is a semanticly-aimed formal sense, where a symbol is interpretable by a system and is manipulated (through symbol-sensitive rules) in order to achieve some goal. The interpretation may be used as a stand-in for an external object (as computationalists like Fodor or Pylysyn use it) or as a stand-in for another symbol, expression or process (as Newell and Simon originally used the term "semantic").[11] The semantic formal sense presupposes the simple formal sense, but not vice versa. Here is the first problem with symbolic dynamics: it only relates to the simple formal sense of symbol, which really boils down to discretization. It is the second, semanticly-aimed formal sense of symbol that is

[11]There is a third, pragmatic/semiotic sense of symbol, related to symbolizing, that extends the second. We will not discuss this sense.

relevant for cognition. Both Dale and Edelman acknowledge that the primary promise of symbolic dynamics is its ability to bridge the continuous/discrete gap in cognitive processing. Indeed, it offers nice tools for understanding how discretization in continuous processes may occur (or be identified) in dynamical systems, but the goods are not that valuable for the framework problem. This is the easy part of the problem.

However, maybe discretization, which supports the simple formal sense of symbol, is a big deal. This, after all, is required by the semanticly-aimed formal sense. Even here the achievement may be overstated. Clark (1992) has offered, in my view, one of the best analyses of the idea of a symbol in a cognitive system. He identifies two conditions: (1) *ease of recognition* of the symbol[12], and (2) *flexibility of use*. The first condition is important for the moment. It suggests that, to be a symbol for a cognitive system, a symbol must be recognizable (easily) by whatever subsystem is using the symbol. We cannot adopt the following line of thought: here is dynamical system with a partition, this provides the symbols, so the system can use them (in the semantic formal sense). The symbols must be recognized as such first, and how that happens depends on the user system. The hard problem, as suggested in the previous section, is fixing the correct partition. Without a systematic theory to do that, a theory not part of symbolic dynamics, we have no cognitive framework.

There is more to symbolic dynamics than what has been used in the argument so far, namely the connection to information processing. Maybe here lies the key. Symbolic dynamics may offer a framework where we can investigate how information/symbol processing systems may be implemented within dynamical systems. Beim Graben (2004) and beim Graben *et al.* (2004) offer several mathematical examples of simple dynamical systems that, when properly partitioned, implement various "symbol processing operations". The examples are based on a very general result from Moore (1990), showing that a single particle moving in a three-dimensional box can be interpreted as implementing a universal Turing machine. Such toy examples are very interesting. Kudos for dynamics for sure! But, such implementational observations would not move a committed computationalist one bit, for reasons we discussed above. Such technical observations fail the condition of novelty. This is not to say that such examples may not be a part of a framework that offers novelty. But such a framework has not been provided. This is the essence of the criticism. To use an architectural metaphor, we need a theory of building constructions, and we are given a theory of bricks — surely an important initial part, but far from what is needed.

Now, let us look at the connection to computation brought about by question (3). Here we are interested less in implementation, and more in analysis of existing dynamical systems. We have a dynamical system, D, analyzed through an observable (defining a partition and thus a symbolic system) with the goal of learning something about D. One thing that we may want to learn is whether it performs any information processing. Crutchfield provides a way of identify-

[12]This condition is adopted from Kirsh (1991).

ing the information processing capacity of a dynamical system with the reconstruction of an $\varepsilon-$machine from the dynamical system. This theoretically very powerful idea, unfortunately, uses an overly general conception of information processing, as general as Shannon's (1948) or Kolmogorov-Chaitin's conception of information and information complexity (Grunwald & Vitanyi, 2008). In fact, Crutchfield's measure is formally related to both. By itself, it violates the non-collapse condition. Additional conditions are needed to separate the trivial senses of information processing from the kind of information processing relevant for higher cognition.

Let us look more closely at the setup of question (3). The question of how a symbolic dynamics may be generates from a dynamical system D and an observable on D may be connected in two different ways to a given cognitive system. The first way is when the cognitive system is the system asking the question. This is actually the motivating case for Crutchfield. He is interested in how an artificial agent may produce a model of its environment (the system D) by having access to D via an observable. This is an extremely important question to ask. The question relates to cognitive architectures, but it is not a question directly about architectural frameworks of cognition. The second way is when the system D is (a part of) the cognitive system, and the agent asking the question is the investigator — the Fodor, the Smolensky, etc. This is the important question about the framework problem. Ultimately, to assess the viability of symbolic dynamics as a unifying framework for cognition, we need to understand how specifying a partition on the micro states of a brain system (e.g. a network of neurons, synaptic connections and neurotransmitters) and tracking the sequence of states in the resulting "symbol" space may tell us anything about the information/symbol processing going on.

Fodor (1997), in his response to a sequence of exchanges with Smolensky, offered a powerful criticism to the strategy of semantic analysis of distributed representations of the states of connectionist networks. The debate centered on whether compositionality and productivity can be accommodated in a connectionist framework. The gist of Smolensky's strategy was to define a map from the state of a neural network to a formal mathematical structure — a vector space. Then, the vectors, and not the network, are used to define semantics and to identify compositionality, with various fancy mathematical techniques. The problem, Fodor correctly identified, is that only the network participates in the causal goings-on of the world. The mathematical constructs cannot explain much about how cognition works, or how the system is sensitive to the semantics and structural distinctions "decoded" with the mathematical techniques. Now, Fodor was mistaken to assume that his symbolic model can do any better in this regard, but this is a separate issue.

The problem of explanatory relevance due to lack of causal significance is also applicable to symbolic dynamics, as it follows essentially the same analytic strategy: map the states of the system to a formal symbol space and use the formal space to interpret what is going on at the dynamical level — whether information processing is occurring or not. In what sense is the mathematical projection of the dynamical system into the space of "symbol sequences" explanatory about the functioning of the cognitive system? The symbols and the relations in which they enter do not have (or at least are not given by the description) any causal role. For an eliminativist of information processing this may be all one needs to offer — "symbols" are just a metaphor, a way of looking at the system. As we have formulated the framework problem, however, this move is unacceptable: we need a way to take information/symbol processing seriously. More is needed from symbols dynamics to offer a proper explanatory framework for cognition *qua* cognition; that is, not merely an analytic framework for patterns of agent behavior, but a framework for explaining the operation of the cognitive mechanisms. Symbols dynamics, as offered, fails the significance condition too. More is needed! To this *more* we turn next.

4 Information Media Networks

So far, we have identified that something is missing from symbolic dynamics to serve the role of a non-trivial framework for cognition. Here we move towards identifying what is missing. Let us start with an observation. We identified two ways an epistemic analysis of a dynamical system (through an observable) may connect to cognition: when the dynamical system is modeling an external system, and when a scientist is modeling a cognitive system. Both ways were identified as unsatisfactory. The two ways may be merged, however, in a promising way. Let D be a dynamical system. What if D is both a part of the cognitive system, and is "observed" by another part of the cognitive system D', which is also a dynamical system? What does "observe" mean? Here, it means simply that there is appropriate coupling between the two systems so that distinctions that can be made in D', define an observable that can be used to partition D. In this picture, the partition on D is not arbitrary or defined by an external observer, but is induced by another dynamical system that is "using" D.[13]

This observation suggests a somewhat different strategy for connecting symbol processing to dynamical systems. Instead of trying to get information processing and symbol processing directly from the regularity of orbits of a dynamical system, we get it by looking at relations among a large network of open dynamical systems within the agent, which mutually affect each other and interact with external systems. This strategy will allow the quality "symbols-processing" to be an extrinsic quality of a system, in that what makes a processing medium M "symbol-processing" depends on how the medium is manipulated, observed or used by something else. Yet, the strategy allows the fact that a cognitive system is using symbol processing to be intrinsic to the system (and its relation to the environment), and not only observer-dependent. The key is that what manipulates, observes or uses M need not be an agent, but another system O. Both O and M are parts of the agent, and O (and what O does in the agent)

[13]There is an interesting question of regress of definitions here, as what distinctions D' makes may depend on how it is used by yet another system D''. I will not explore this question here. I will only say that embodiment is important for the resolution of the problem.

make M into a symbol-processing system.

These ideas can be made precise with the theory of information media networks (Vakarelov, 2012, 2011).[14] An *information medium* is an (open) dynamical system with a disjointed collection of sets of states, called *information carrying states* (ICS) and a collection of functions on the dynamical states that "respect the ICS" called information preserving transformations (this could be identity only). The information preserving transformations define a natural equivalence relation on the information carrying states, which is used to define *information states*[15]. The notion of information medium is a generalization of the notion of symbolic dynamics. First, ICS need not form a partition; there may be dynamical states that are informationally inert. (Clearly, ICS may be extended uniquely to a partition.) Second, ICS need not be finite (or countable). There is no pressure to interpret the elements of ICS as symbols. Third, there is formal structure defined by the information preserving transformations, which may be, but need not be, connected to time transition of the dynamical system. Such functions will normally be defined as invariance of interaction of the dynamical system with other systems.

In addition to information preserving transformations, there may be other functions that respect the ICS states. Such functions are called information processing transformations. If the dynamical transition function (which, because the system is open, will have many control parameters) defines an information processing transformation, we call such a transformation *canonical.* In this new language, a symbolic dynamics is defined by a canonical transformation.

A function between two media that respects the ICS states is called an *information management transformation.* A collection of information media, together with a collection of information management transformations among the media, and collections of information processing transformations for each medium, define an *information media network* (IMN).

Media are open (physical) dynamical systems that may interact in various ways, thus not all interactions among media can be defined as information management transformations. Media may interact causally so that one medium M may manipulate another medium N in a way that defines (or implements) an information processing transformation on N. This allows for complex control relations in the network, including the possibility for the network structure to change and evolve.

Now I will briefly sketch a way of capturing the concept of symbols systems, and thus symbol processing, in an information media network. I will not talk specifically about semantics, as in semantic information, even though this is very important for understanding cognition.[16] Many of the details are developed in

[14]In what follows, words like "information", "information processing" etc., should be regarded as technical terms only. They are borrowed from the theory of information media networks. They have an intended meaning, but the appropriateness of the terminology cannot be justified here. It cannot be assumed that they correspond directly to the ideas discussed in previous sections.

[15]Different information carrying states may carry the same information

[16]See Vakarelov (2010) for more discussion of the role of semantic information in the picture.

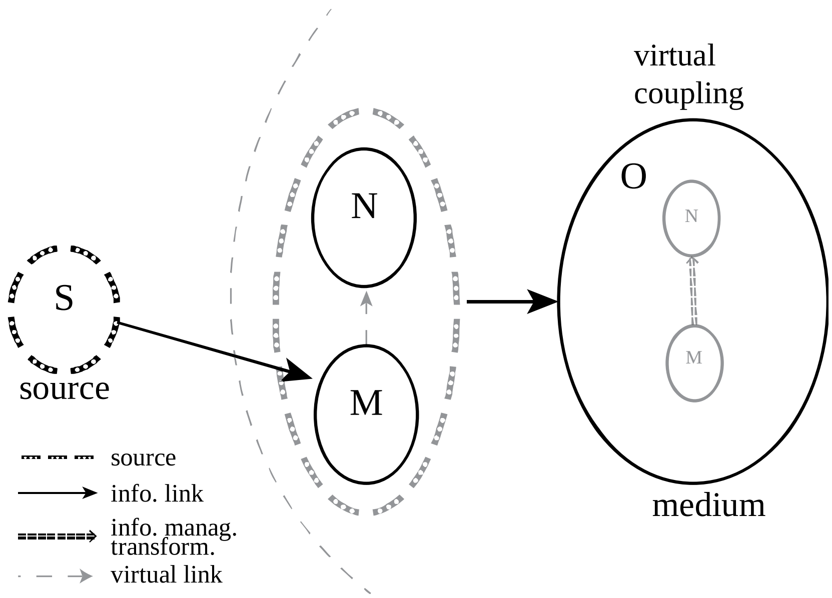

Figure 1: Virtual Coupling Diagram, from Vakarelov (forthcoming).

Vakarelov (forthcoming), except for the discussion of what makes a representational medium symbolic. Essentially, I will explain a diagram appearing in the article, containing many of the needed ideas (see Figure 1).

All solid ovals within the dashed line represent information media in the IMN of an agent. The object S is an external (or internal) information source for medium M. The medium N is another medium that is correlated (coupled) with M (in the sense that there is an information management transformation from M to N) that allows the information states of M to be "witnessed" in N. When the transformation is such that S becomes, de facto, an information source for N (the precise conditions involve semantic ideas not discussed here), we say that there is triangulation between M and N with respect to S. When another medium O has the entire complex $M \Longrightarrow N$ as an information source, and when O (with possible help of other parts of the IMN not shown here) modulates the transformation from M to N to assure better triangulation (again, semantic ideas related to goal-directed function are important here), we say that there is *triangulation monitoring*. When N is subject to triangulation monitoring, we say that N is a *representation medium for* S. (Note that according to this theory, a medium never directly represents a source. All representational relations, in the classical correspondence sense, are internal, inter-media relations.)

The figure shown here actually depicts a more complex situation. In many cases, information management transformations may be implemented by some causal process connecting two systems. For symbols systems, we do not want a fixed causal link to provide the relation. We want the symbols to be in some sense *arbitrary*. It is possible that the link between M and N exists only in the medium O (in implementations of such organization, O will be an entire sub-network). This we call *virtual coupling* between M and N. For example, a table linking words to images defines such virtual coupling. This is natural, as there is nothing specially causal that can define such a relation between words and images that lives outside the table.

We are not at symbols yet. We need three additional conditions that move beyond the discussion in Vakarelov (forthcoming) and are thus only suggestive, to be developed in future work. First, there has to be a way of actually arbitrarily assigning states in M to states in N. Second, there has to be a way of manipulating N to set it in the appropriate symbol state (some generation mechanism). Finally, and this is connected to Clark's second condition of a symbols system (see Section 3), there has to be a mechanism to channel the established virtual coupling to arbitrary control functions (flexible uses).

The first condition would require that O itself be a representational medium, having another layer of monitoring on top. In addition, there will have to be some mechanism that modulates the state of O in a way that generates different virtual coupling relations, and in a way that the virtual coupling rotations may be adapted to needed control functions (which need not be arbitrary yet).

The second condition requires that there be some mechanism in the IMN that may set the state of M (via some informational transformation) based on the demands of O or other systems modulating O. A particularly interesting case is when the information state of N may have "compositional structure", especially recursive compositional structure. Since information media are dynamical systems and we cannot assume that the compositional structure is implemented by objects with intrinsic structure (like words on paper). In the general case, such structure will be recovered through the (read/write) information management transformations to other media. (For example, M may be in a state "$A * B$" if "get-left*" transformation results in A and "get-right*" transformation result in B, etc.).

The third condition would require yet more and deeper network structure needed to represent "goals" and "actions" and virtual connections to such representational states. This is the condition that requires most complex IMN organization. The precise conditions still currently elude me but they seem graspable within the framework.

The superficial discussion presented here was meant as an illustration of the kinds of machinery needed to capture properly symbolic processing in a cognitive system. It was not intended to convince the reader that a cognitive framework based in IMNs can capture them. It should be clear that IMNs have significantly more expressive capacities for capturing the ideas needed to understand symbolic processing in a cognitive system than does the framework of symbolic dynamics. Still, since IMNs are firmly rooted in dynamical systems theory, they offer a better potential for a solution to the framework problem.

5 Conclusion

The cognitive framework supported by the machinery of information media networks allows a richer and more powerful (in the sense of more effective) language to analyze cognitive operations than the framework of symbolic dynamics. It also has the possibility (in an embodied context) to capture semantic relations (not discussed here), representations and the kinds of symbolic operations Fodor wants to have in higher cognition, but in the spirit of Clark's analysis of symbols. Yet, IMNs coexist comfortably with dynamical systems theory and may accommodate the insights to cognition, especially embodied cognition, that the dynamicist program has offered.

References

beim Graben, P. 2004. Incompatible Implementations of Physical Symbol Systems. *Mind and Matter*, **2**(2), 29–51.

beim Graben, P., Jurish, B., Saddy, D., & Frisch, S. 2004. Language processing by dynamical systems. *International Journal of Bifurcation and Chaos*, **14**(02), 599–621.

Calvo, Paco, & Gomila, Antoni. 2008. *Handbook of cognitive science an embodied approach*. Amsterdam; Boston; London: Elsevier Science.

Chemero, Anthony. 2009. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Clark, Andy. 1992. Presence of a Symbol. *Connection Science*, **4**(3-4), 193–205.

Crutchfield, James P. 1994. The calculi of emergence: computation, dynamics and induction. *Physica D: Nonlinear Phenomena*, **75**(1), 11–54.

Crutchfield, James P., & Young, Karl. 1990. Computation at the onset of chaos. *Entropy, Complexity, and the Physics of Information*, **8**, 223.

Dale, R. 2008a. Introduction to the special issue on: pluralism and the future of cognitive science. *Journal of Experimental and Theoretical Artificial Intelligence*, **20**(3), 153–153.

Dale, R., Dietrich, E., & Chemero, A. 2009. Explanatory pluralism in cognitive science. *Cognitive science*, **33**(5), 739–742.

Dale, Rick. 2008b. The possibility of a pluralist cognitive science. *Journal of Experimental and Theoretical Artificial Intelligence*, **20**(3), 155–179.

Dale, Rick, & Spivey, Michael J. 2005. From apples and oranges to symbolic dynamics: a framework for conciliating notions of cognitive representation. *Journal of Experimental & Theoretical Artificial Intelligence*, **17**(4), 317–342.

Edelman, Shimon. 2008. On the nature of minds, or: truth and consequences. *Journal of Experimental & Theoretical Artificial Intelligence*, **20**(3), 181–196.

Fodor, Jerry. 1997. Connectionism and the problem of systematicity (continued): Why Smolensky's solution still doesn't work. *Cognition*, **62**(1), 109–119.

Fodor, Jerry, & McLaughlin, Brian P. 1990. Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, **35**(2), 183–204.

Grunwald, Peter D., & Vitanyi, Paul M.B. 2008. *Handbook on the Philosophy of Information*. Amsterdam, The Netherlands: North Holland. Chap. Algorithmic Information Theory.

Haugeland, John (ed). 1997. *Mind Design II: Philosophy, Psychology, and Artificial Intelligence*. Revised edn. A Bradford Book.

Kirsh, David. 1991. *Information, Language and Cognition*. Oxford, UK: Oxford University Press. Chap. When is information explicitly represented?, pages 340–365.

Maudlin, Tim. 1989. Computation and consciousness. *The journal of Philosophy*, 407–432.

Moore, Cristopher. 1990. Unpredictability and undecidability in dynamical systems. *Physical Review Letters*, **64**(20), 2354–2357.

Newell, A., & Simon, H. 1981. Computer science as empirical enquiry. *In:* Hougeland, J. (ed), *Mind design II*. Cambridge, MA: MIT Press.

Prince, Alan, & Smolensky, Paul. 1997. Optimality: From neural networks to universal grammar. *Science*, **275**(5306), 1604–1610.

Riegler, A. 2002. When is a cognitive system embodied? *Cognitive Systems Research*, **3**(3), 339–348.

Rumelhart, D. E. 1989. The architecture of mind: A connectionist approach. *Pages 133–159 of:* Posner, M. I. (ed), *Foundations of cognitive science*. Cambridge, Mass.: MIT Press.

Rumelhart, David E., & McClelland, James L. 1986. Parallel distributed processing: explorations in the microstructure of cognition. Volume 1. Foundations.

Shalizi, C. R. 2005. Symbolic dynamics, coarse-graining, and levels of description in statistical physics and cognitive science. *Page 14–17 of: Proceedings of workshop on Symbol Grounding: Dynamical Systems Approaches to Language.*

Shannon, C. 1948. The Mathematical Theory of Communication. *Bell Systems Technical Journal.*, **27**, 379–423.

Shapiro, Lawrence. 2011. *Embodied Cognition*. Embodied Cognition: Routledge.

Smolensky, P. 1986. Information processing in dynamical systems: Foundations of harmony theory. *Chap. 6, pages 194–281 of:* Rumelhart, D. (ed), *Parallel distributed processing: Explorations in the microstructure of cognition: Foundations*, vol. 1. MIT Press.

Vakarelov, Orlin. 2010. Pre-Cognitive Semantic Information. *Knowledge, Technology & Policy*, **23**(2), 193–226.

Vakarelov, Orlin. 2011. *General Situated Cognition*. Ph.D. thesis, The University of Arizona.

Vakarelov, Orlin. 2012. The Information Medium. *Philosophy and Technology*, **25**(1), 47–65.

Vakarelov, Orlin. forthcoming. From Interface to Correspondence: Recovering Classical Representations in a Pragmatic Theory of Semantic Information. *Minds and Machines*.

van Gelder, T. 1998. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, **21**, 615–665.

Varela, Franscisco J., Thompson, Evan T., & Rosch, Eleanor. 1992. *The Embodied Mind*. Cambridge, MA: MIT Press.