Anticipation and Intentional Behaviour: Some Building Blocks

J. Barahona da Fonseca[1], I. Barahona da Fonseca[2] and J. Simões da Fonseca[3]
[1]Dept Elect Eng, Faculty of Sciences and Technology, New University of Lisbon, jbfo@fct.unl.pt, +351-214262406
[2]Faculty of Psychology, Dept of Psychophysiology, University of Lisbon, Portugal, isabelbf@fpce.ul.pt, +351-217933408
[3]Faculty of Medicine, University of Lisbon, Portugal

Abstract
We first consider the information flow and processing that is involved in perception considering the requirements necessary to explain the passage of elementary mappings to object representation and cognitive characterization of representations. The relationships between procedural implicit cognitive processes, and on the other hand, linguistic declarative processes is discussed. A reference to Marcus's results is made, taking into account experimental data and theoretical proposals, involving concepts of Quantum Chemistry made by our group.
Keywords: Valence, Quantum Chemistry, Cognition, Multilayer Perceptrons

1 Introduction

The limits on multilayer perceptron accounts can tell us something, which is how to make better models. The limits on multilayer perceptrons motivate three basic components of cognition: representations of relationships between variables, structured representations, and representations of individuals that are distinct from representations of kinds.

There is an asymmetry inherent in the research strategy of starting with simpler models as a means to discover what *fundamental* elements are really needed: while negative arguments can be decisive in ruling out particular classes of elements, positive arguments can never be decisive. At best, positive arguments can merely be what philosophers politely call *nondemonstrative*.

What this asymmetry means here is that while we can be confident that certain classes of models simply *cannot* capture a certain class of cognitive and linguistic phenomena, we can never be sure of the alternative.

We will not prove that the mind/brain implements symbol-manipulation. We will describe how symbol-manipulation *could* support the relevant cognitive phenomena. All that we can do is to provisionally *accept symbol-manipulation* as providing the best explanation.

First we need to figure out how components fundamental to make symbol-manipulation are implemented in neural hardware. It has been argued that physically localizable registers, implemented as inter or intracellular circuits with feedback, might serve as a substrate for the storage of values of variables (Marcus, 2001).

Second, even if the components of symbol-manipulation do play a real and robust role in our mental life, it is unlikely that they exhaust the set of components for cognition. It seems likely that many other basic computational elements play important roles in cognition. For instance it seems quite likely that the *representational formats for encoding images* are distinct from the sorts of *representational formats that support the encoding of propositions.*

Differences between the cognition of humans and other primates may lie not so much in the basic-level components but in how those components are interconnected, in their modular architecture. To understand human cognition, we need to understand how basic computational components are integrated into more complex devices- such as parsers, language acquisition devices, modules for recognizing objects, and so forth- and we need to understand how our knowledge is structured, what sorts of basic conceptual distinctions we represent, and so forth.

To take but one example, a system that can represent rules and structured representations has the in-principle ability to represent abstract principles that might constrain the range of variation in the world's languages, but the computational components do not by themselves tell us which among infinitely many possible linguistic constraints are actually implemented in the human mind.

If we know what the basic computational elements are, we are in a better position to understand how cognition is realized in the underlying neural substrate. It is here that connectionism has its greatest potential to help us to understand how basic computational elements work and interact with each other.

To date, progress in cognitive neuroscience has been hindered by the *enormity of the gap* between our understanding of some low-level properties of the brain on the one hand, and of some very high-level properties of the mind on the other. As we come to identify and attaining a better understanding of the intermediate-level building blocks, it may become easier to relate neuroscience to cognition.

2 Our Contribution

2.1 Introduction

There are two specific problems that we have to address when we try to make sense of the way our sensory and nervous system machinery construct a representation both of the external environment and the potentially cognitive record of events. Our formal systems for dealing with neural computation are represented by logic systems, systems of rules and inference, script grammars, case grammars, etc.

One of the major neurophysiological achievements of the XX century was the verification of the concept of *feature detectors in the frog* by Lettvin, Maturana, Pitts and McCulloch (1959) and the discovery by Hubell and Wiesel (1963) of simple, complex and hypercomplex visual cells, as well as the identification by Koch and Crick (2000) of subgroups of activated cells during the pendular alternating representations of Necker cubes or alternating gratings in area V4.

Other examples refer to the abstract representation for instance of visual responses to *doghood* or else *cathood* in the Infero Temporal Cortex and in the Frontal Cortex.

If we ask ourselves how such a representation can be attained by a sensory system like the visual system, we verify that there is a first level of point to point mapping of the external space into Retina and at a second level, a *function space representation* level as it occurs with Hubel and Wizel feature detectors or the so called *bug detectors* of Lettvin, Maturana, Pitts and McCulloch (1959).

It can be shown that an abstract not semantically significant representation of what later on will be identified as objects can be built using as components of the representation, the so called *geons* (Simões da Fonseca, 1999). If we ask ourselves a question about the meaning of those representations in terms of linguistic declarative statements, the answer, in most cases, will be that at this level all those identifiable components of the representation have no meaning- and nevertheless these meaningless representations give birth to meaningful representations of objects.

In other words, at the first level we have the representation of individual or singular components of stimuli, at a second level we have identified partial invariants we call features and at a third level we have the assembling of these elementary components as complex relational ensembles.

At this point begins the semantic interpretation of the symbolic representations which were constructed in this visual domain. We can speak at this level for instance of visual meaning or auditory meaning or else somato sensory meaning that remains at a procedural nondeclarative latent unconscious knowledge.

A task of *building a visual meaning* is comparable to the construction of a dictionary from a foreign language into the native language of the user. The elements of the *relevant symbolic components* of the representations are then *fixed in memory through learning processes* and is rendered accessible to further uses and generalized to a manifold of distinct adaptive subsystems. Finally a *second dictionary* integrates the *psychological representations and computations* at the service of *social relationship schemata*.

Our claim is that in most of the brain processes, if we ask during a *meaningful task* what is being represented and processed, the answer will be *nonverbally representable symbols of distinct systems of invariants*- much as Picasso representations of a person produce a portrait of peculiarly assembled and placed particular distortions of components of faces and bodies which nevertheless do not loose their *linguistic declarative meaning*.

So, besides of the problem of how knowledge architecture can be represented as a semantic network or else by a representation with components in the frequency domain, or else by geometric representations and still *Treelets* representations (Marcus, 2001), the truth is that *no verification will be feasible* unless we find the *abstract concepts* that are the *correlates of the symbols and relationships* implemented by the CNS.

So we can divide our research in four distinct domains: 1) the domain of nonverbal but nevertheless operative components of knowledge architectures that in some measure will give rise to semantic systems; 2) systems of symbols; 3) intermodality coordinate transformation systems; 4) individual representations and concept representations as classes of equivalence.

2.2 Cognitive Competence with or without Conscious Awareness

It is known since the pioneering work of Helmholtz (1954) that percuting an object, for instance a table, with a metallic object or else with a book and so producing two different sounds, if we ask a listener which of the two sounds was more acute, the subject will immediately answer that it was the first one produced by the metallic object. He does so in a *completely automatic manner* that does not require any *conscious judgement*.

Helmholtz called this phenomenon *unconscious inference*.

Pavlov himself referred to *conditioned responses* as being the result of *unconscious inference*.

In the forties of the XX century, Leeper made a very competent revision of a class of phenomena that express what he denoted as *unconscious inference*. This kind of problem appears immediately and in a variety of forms if we make the introspection of our voluntary acts. Namely, if we want to grasp a book on a table that is at a distance that requires the execution of a certain number of steps, we may clearly formulate our intention and proceed performing the complex behaviour that leads to the attainment of the goal.

At first sight it seems trivial, but it isn't.

First we must formulate a conscious plan of action with the orientation of the eyes, we must leave the chair on which we are seated, stand up, make a detour around the table, walk a small distance, approach with the right hand, in general, the book, under the control of the eyes, grasp it the right way adapting the tension of the muscles of our fingers to the weight we predicted and finally produce the utterance "I got it!".

Let we consider this situation in some more detail: - what is initially in our conscious experience is the *perception of the book*, the place on which it lies, the awareness of the position on which we are and of our intention of grasping the book.

In no moment the *clear consciousness* of our *intention* is accompanied by any *conscious awareness of any subjective knowledge* of any one of the details of the *plans* and of any one of the movements we must execute.

Nevertheless this simple behaviour requires in a summarised form that we orient our eyes adequately all along this behaviour; to move from our initial place we must coordinate in an exact manner the muscles of our inferior member as well as muscles around spinal column, of our abdomen and still we must maintain an adequate position of our neck and head. To make these acts in a true conscious form and under a true voluntary control, we should take all these *decisions in a conscious way*, but that is not the case.

As a matter of fact an *unconscious*, *procedural translation* is capable of attaining a *complete syntactic translation* of our *cognition* into a 'machine language' understanding of our *intention*. Paradoxically what gives sense to our *conscious experience* is an *unconscious experience* that *is not accessible for us*.

It does it so in an extraordinarily competent form – that is, it solves what philosophers call the *mind-body* problem.

The picture is still not complete because first our position in space must be evaluated by Parietal Area 5, these data must be coordinated with visual information in Parietal Area 7 which coordinates also the information it receives from Visual Frontal Area.

All this information is sent to the Pre-Motor Area and the Supplementary Motor Area that plan the coordinations involved in the action.

This information is sent to the Putamen Nucleus whose neurons are desinhibited. Thereafter they inhibit the neurons of the Globus Pallidus.

In consequence, the cells in the Dorsal Lateral Nucleus of the Thalamus are deshinibited and give permission to start action and the Motor Area enters in action and sends the commands for movement that still will be coordinated by Cerebellum and finally by Brainstem and the Spinal Cord.

Let us stop here and reflect on the Neurocomputational requirements of the translation of our *Intention* into *Behaviour* that leads to the achievement of the *Goal* we did choose.

Again we must putatively admit that an *Unconscious System* made all the choices about possible *Trajectories* of our *Body* which implies a complex ensemble of transformations from the visual and somatic mappings together with motor mappings to render available the necessary information to the Frontal Cortex and in particular to its *Executive Agencies*.

This is the true syntactic and semantic meaning of producing a plan and executing it. Our acceptance of the *adequacy of the plan during its performance* that expresses itself in not making any correction is accessible at the level of *Conscious Awareness*, but what justifies this *Conscious Experience* is not any other set of *Conscious Experiences*.

Rather an *Unconscious Experience* gives rise, content and meaning to the *Conscious Experience*. The final satisfaction of grasping the book will result from the reduction of drive that is consequence of fulfilling the initial motivation.

It is accompanied by some form of *subjective evaluation* that again is produced by some components of the Limbic and Neocortex systems.

The moral of this story is that our problem in Neurocomputation is not that we don't know the mechanisms and processes but rather we *don't have the ideas* necessary to understand what is being computed.

The real problem is that we are lacking the understanding of the *architecture of knowledge* that is submitted to *Neurocomputational* processes.

We know nevertheless that a first rank tennis player like Agassiz has no possibility of deciding his hits using verbal declarative processes and perhaps he even does not see the ball when he beets it.

If we think about extraordinary animal performances like the possibility that pigeons possess for forming a class of equivalence with all landscape that have trees and distinguishing neatly such stimuli from landscapes without trees it is reasonable to conclude they behave according with nondeclarative but nevertheless cognitive operations.

This makes us to return to our argument. We may putatively admit that the *two cognitive systems*, the *verbal declarative* and the *procedural implicit* coexist in humans, the verbal declarative being philogeneticaly more advanced and differentiated.

The bulk of the question about *consciousness* would need, to be answered, some knowledge about what *implicit processes* do that *declarative verbal processes* cannot do.

We know furthermore that they work together and relevant interactions occur during concept formation, planning, commands of execution and control processes along the execution of the successive stages of an intentional plan.

A promising hypothesis would be that for most of human actions, thoughts and feelings, *work memory* would take care of most of the interactions and would be responsible for most *conscious experiences*.

If we look into philosophical, mathematical and artistical cognitive processes, a strong argument in favour of this hypothesis is found in the *processes of mental training* of athletes of high competition who train hours a day thinking and imagining the complete sequence of movements they would perform if they were competing. They don't make any move and solely imagine it, and results prove that *pure imagination is translated into better command, control and strength* (Jeannerod, 1997, Hall and Schmidt, 1995, Orlick, 1986).

Psychiatric patients may be trained the same way to enhance the quality of their performances either professional, interpersonal or motivational. They *succeed better* in controlling distractors, environmental variations, social interferences and many sorts of interference that *would reduce, without this mental training, their efficiency*.

2.3 What to do?

Next we summarise some of the results that Marcus (2001) has obtained using a multilayer perceptron approach that are in some way a first step towards a new approach to Neuroscience and Cognitive Science.

3 Symbolic Processing versus Connectionism

3.1 Localist and Distributed Representations

Some input and output representations are *localist* and others are *distributed*. In localist representations each input node corresponds to a specific word or concept. For example in Elman's (1990, 1991, 1993) syntax model, each input unit corresponds to a particular word and also each output unit corresponds to a particular word.

Other localist representational schemes include those in which a given node corresponds to a particular location in a retina like visual array (Munakata, McClelland, Johnson and Siegler, 1997).

In *distributed* representations any particular input is encoded by means of a set of simultaneously activated nodes, each of which can participate in the encoding of more than one distinct input. For example in a model of the inflection of the English past tense proposed by Hare, Elman and Daugherty (1995), input features correspond to speech segments in particular positions: 14 input nodes correspond to 14 possible onsets (beginnings of syllables), six input nodes correspond to six possible instantiations of the nucleus (middles of syllables), and 18 input nodes correspond to 18 possible codas (ends of syllables). The word *bid* would be represented by the simultaneous activation of 3 nodes corresponding to *b* in the initial position, *i* in the nucleus position, and *d* in the coda position. *Each of those nodes would also participate in the encoding of other inputs.*

In some models *input nodes* do not correspond to anything *obviously meaningful*, for example consider each node as the digit of a numerical input in some base.

3.2 Multilayer Perceptrons and Symbol-Manipulation

The vast majority of the connectionist models that have been used in discussions of cognitive science are *multilayer perceptrons*, either *feedforward* or *simple recurrent networks* that have the advantage of the *learning capability about sequences of elements presented over time*.

Among the many domains in which such models have been used are the acquisition of linguistic inflection (e.g., Rumelhart and McClelland, 1986a), the acquisition of grammatical knowledge (Elman, 1990), the development of object permanence (Mareschal, Plunkett and Harris, 1995; Munakata, McClelland, Johnson and Siegler, 1997), categorization (Gluck and Bower, 1988; Plunkett, Sinha, MØller and Strandsby, 1992; Quinn and Johnson, 1996), etc.

4 How *Multilayer Perceptrons* Appeared in Discussions of Cognitive Architecture

The idea that *connectionist networks* might offer an *alternative to symbol-manipulation* started to become prominent with the work of J. A. Anderson and Hinton (1981). This idea became even more prominent in 1986 with the publication of an influential paper by Rumelhart and McClelland (1986a). Rumelhart and McClelland presented a *two-layer perceptron* that captures certain aspects of children's acquisition of the English past tense. Similarly, Bates and Elman (1993) suggest that their particular connectionist approach "runs directly *counter to the tendency in traditional cognitive and linguistic research to seek 'the rule' or 'the grammar' that underlies a set of behavioural regularities*". And Seidenberg(1997) writes that the kind of network he advocates "incorporates a novel form of knowledge representation that provides an *alternative to equating knowledge of a language with a grammar. Such networks do not directly incorporate or implement traditional grammars*".

Still, although such claims have received a great deal of attention, *not everyone who advocates multilayer perceptrons denies that symbol manipulation plays a role in cognition.* A weaker but commonly adopted view holds that *symbol manipulation exists but plays a relatively small role in cognition* (Touretzky and Hinton, 1988).

5 Theoretical Considerations for the Elimination of Symbol Manipulation

One reason that *multilayer perceptrons* seem especially attractive is that they are "*more compatible* than symbolic models with what we *know of the nervous system*" (Bechtel and Abrahamsen, 1991). Nodes are loosely modeled on neurons, and the connections between nodes are loosely modeled on synapses. Conversely, *symbol manipulation* models do not, on their surface, *look much like brains*, and so it is natural to think of *multilayer perceptrons* as perhaps being *more fruitful ways of understanding the connection between brain and cognition.*

A different reason for favoring *multilayer perceptrons* is that they have been shown to be able to represent a *very broad range of functions*. For virtually any given function there exists some multilayer perceptron with some configuration of nodes and weights that approximate it (Hadley, 2000).

Still others favor multilayer perceptrons because they appear to require relatively *little in the way of innate structure*. For researchers drawn to views in which a *child enters the world with relatively little initial structure*, *multilayer perceptrons offer a way of making their view computationally explicit*.

Multilayer perceptrons are also appealing because of their intrinsic ability to learn (Bates and Elman, 1993) and because of their ability to *gracefully degrade:* they can tolerate limited amounts of noise or damage without dramatic breakdowns (Rumelhart and McClelland, 1986b).

Since multilayer perceptrons have *context-independent representations of categories* we may count them as having symbols. In this sense we may say that a multilayer perceptron is a *symbol manipulation system*. But the important question that we will address next is whether the mind is a system that represents variables, operations over variables, structured representations, and a distinction between *kinds* and *individuals*.

6 Relations Between Variables

Although it seems clear enough that *we can manipulate algebraic rules* in *serial, deliberate reasoning,* not everybody agrees that abstract relationships between variables play an important role in other aspects of language and cognition. For example, Rumelhart and McClelland's (1986a) two layer perceptron was an attempt to explain how children might acquire the past tense of English *without using anything like an explicit rule*.

What we want to do here is to *clarify the relationship between multilayer perceptrons and devices that perform operations over variables*. A better understanding of that relationship will help *clarify whether the mind does in fact make use of operations over variables* and also clarify how such operations can be implemented in a neural substrate.

The distinction between encoding a variable with a single bucket and encoding a variable with a set of buckets is helpful because the relationship between multilayer perceptrons and operations over variables can be understood in similar terms. *The key question is whether a given input variable in a particular network is encoded using one node or a set of nodes.*

This difference – in whether a particular variable is encoded by one node or by many nodes – is not the same as the difference between localist and distributed representations. While all models that use distributed representations allocate more than one variable per node, *it is not the case that all localist models allocate a single node per variable.* In Elman's sentence prediction model the input to the model is a single variable that we might think of as *current word.* Although any given instantiation of that variable (say, *cat*) will activate only a single node, every input node can potentially indicate an instantiation of the variable current word. For example, the node for *dog* might not be active at this moment, but it might be active during the presentation of another sentence. Elman's sentence prediction model is thus an example of a localist model that allocates multiple nodes to a single input variable.

We are now ready to consider the relation between multilayer perceptrons and systems that represent and generalize operations over variables. We argue neither that *multilayer perceptrons cannot represent abstract relationships between variables nor that they must represent abstract relationships between variables.* Simple claims like "Multilayer perceptrons cannot represent rules" or "Multilayer perceptrons always represent 'concealed' rules" simply are not correct. The real situation is more complex.

7 Models that Allocate one Node to Each Variable

Models that allocate a *single node* to each *input variable* behave very differently from models that allocate *more than one node to each input variable*. Models that allocate a single node to each input variable are (with some caveats) simpler than models that allocate multiple nodes to each variable. *One node per variable model can only represent universally quantified one-to-one mappings (UQOTOM).* It follows that all that a learning algorithm can do is choose between one UQOTOM and another. *Such models cannot learn arbitrary mappings.* For example they *cannot learn to map an input number that specifies the alphabetical order of a person in a phonebook to an output that specifies that person's telephone number.* In this way they provide a *candidate hypothesis* for how *operations over variables* can be *implemented in a neural substrate* and not for a *mental architecture that eliminates the representation of abstract relationships between variables.*

8 Models that Allocate More than One Node per Variable

Models that allocate *more than one node per variable too can represent UQOTOM, but they do not have to.* When such a *network represents identity or some other UQOTOM,* it represents an abstract relationship between variables, which is to say that such a network implements an *algebraic rule*. Models that allocate a *single node to each variable only can represent abstract relationships between variables*, whereas models that allocate multiple nodes to each variable sometimes represent abstract relationships between variables and sometimes do not: *what they represent is a function of what their connection weights are*. In these latter perceptrons *some connection weights represent UQOTOM*, others *represent many to one mappings*, and still others can *represent purely arbitrary mappings. In this way multilayer perceptrons that allocate more than one node to each variable are quite flexible.*

One might ask whether this flexibility suggests that *multiple-nodes-per-variable multilayer perceptrons* are the best way of *implementing abstract relationships between variables* in a *neural-like substrate*. What we suggest next is that their flexibility is both an asset and a liability and this liability justify a search for alternative ways in which abstract relationships between variables can be implemented in a neural substrate.

The flexibility in what multiple nodes per variable models can represent leads to a flexibility in what they can learn: *they can learn UQOTOMs and they can learn arbitrary mappings. But what they learn depends on the nature of the learning algorithm*. The learning algorithm most commonly used, backpropagation, does not allocate special status to UQOTOMs. Instead a many nodes per variable multilayer perceptron that is trained by backpropagation can learn a UQOTOM, such as identity, multiplication, or concatenation, *only if it sees that UQOTOM illustrated with respect to each possible input and output node*.

Many nodes per variable multilayer perceptrons that are trained by backpropagation can generalize one-to-one mappings within the training space, but assuming that the inputs are binary, they cannot generalize one to one mappings outside the training space. For example (Marcus, 1998c) found that if his MLP *is trained only on inputs with a rightmost digit of 0, it will not generalize identity to inputs with a rightmost digit of 1*. Instead, whether the rightmost digit is a 1 or a 0, the model always returns an output in which the rightmost digit is 0. For example, given the input 1111, the model generally returns 1110, an inference that is mathematically justifiable but totally different from what humans typically do.

It can be shown (Marcus, 2001) that any *many-nodes-per-variable multilayer perceptron that are trained by backpropagation cannot generalize one-to-one mappings between nodes*. This is because the *learning that results from backpropagation is, in an important sense, local*. This *localism* has the consequence that if a model is exposed to a simple UQOTOM relationship for *some subset of the inputs* that *leaves some nodes untrained,* as in Elman's (1990,1991,1993) syntax model, in which each *input unit* corresponds to a *particular word* and also each *output unit* corresponds to a *particular word.*

A localist algorithm is a liability only *if it is used to capture phenomena in which an organism can freely generalize*. In cases where organisms cannot freely generalize, it is possible that localist algorithms may be appropriate. But in some cases it appears that humans can freely generalize from restricted data, and in these cases *many nodes per variable multilayer perceptrons* that are trained by *backpropagation are inappropriate*. We point out this fact because the literature on connectionist models of cognitive science is filled with distributed multilayer perceptron models that are trained by backpropagation, and many of those models are aimed at accounting for aspects of mental life in which humans do appear to be able to freely generalize from incomplete input data. *Humans can freely generalize one to one mappings but distributed multilayer perceptrons that are trained with localist learning algorithms cannot.* For these cases we must seek alternative models.

9 Alternative Models that *Generalize* from Restricted Data

In general, what is required is a system that has five properties. First the system must have a way to distinguish *variables* from *instances*. Second the system must have a way to *represent abstract relationships between variables like an equation*. Third the system must have a way to bind a *particular instance* to a *given variable*. Fourth the system must have a way to *apply operations to arbitrary instances of variables*. Finally the system must have a way to *extract relationships between variables* on the *basis of training examples*.

9.1 Conjunctive Coding

In MLPs the *current instantiation of a given variable is indicated by a pattern of activity*. There are a number of other possible ways to indicate the *binding between a variable and its current instance*. One possibility is to *devote specific nodes to particular combinations of a variable and an instance*.

It seems likely that conjunctive coding plays some role in our mental life. For example, experiments with single-cell recordings by Funashi et al(1993) have indicated that certain neurons are most strongly activated when a particular object appears in a particular position. We may assume that these neurons conjunctively encode combinations of objects in particular positions.

But the brain must rely on other techniques for variable binding as well. Conjunctive codes do not *naturally* allow for the representation of binding between a variable and a novel instance. Moreover, conjunctive encoding schemes may require an unrealistically large number of nodes, proportional to the number of variables times the number of possible instances.

9.2 Tensor Products

A more general, more powerful way of doing conjunctive binding is the *tensor product* (Smolensky, 1990). A tensor product is a way of representing a binding between a variable and an instance. A tensor product is not by itself a way of representing a relationship between variables or a way of applying operations to variables. Further machinery would be required to represent or extend relationships between variables. We do not discuss such machinery here but instead focus only on how tensor products represent bindings between variables and instances.

In the tensor product approach, each possible instance and each possible variable is represented by a vector. A particular binding between a particular variable and a particular instance is represented by applying a process analogous to multiplication. The resulting combination, a *tensor product*, is a vector of greater dimensionality.

One way in which tensor products differ from the simple conjunctive scheme is in the role of a given node. In the conjunctive scheme each node is dedicated to the representation of a single particular binding. In contrast, in the tensor product scheme, every node participates in every binding.

The tensor product scheme has two fundamental advantages. First, it is potentially more efficient. The simple conjunctive scheme requires $i*v$ nodes, where i is the number of instances and v is the number of variables. The tensor product scheme requires $a*b$ nodes, where a is the length of the vector encoding the instance and b the length of the vector encoding the variable. For example, if there are 128 instances and 4 variables, the tensor product scheme is considerably more efficient, requiring 7 + 2 + 14=23nodes, 7 nodes to represent the instance (in binary code we have 2^7=128 possible combinations), 2 to represent the variable and 14=7*2 to represent the possible combinations of the two. The simple conjunctive scheme requires 128*4=512 nodes. Second, the tensor product scheme can more readily cope with the addition of new instances. Assuming that the new instance can simply be assigned a new vector, representing a binding containing that instance is simply a matter of plugging a new vector into the preexisting tensor product machinery. Nonetheless, despite these advantages, we think that tensor products are not plausible as an account of how we represent recursively structured representations.

9.3 Registers

A limitation of the binding schemes discussed so far is that none provides a way of *storing* a binding, they are all transitory and last while a certain input is constant. One obvious way to implement this *memory* is to use devices that have two or more stable states; digital computers use memories with two stable states (0,1). If registers are used in the Brain, they might be bistable or have more than two stable states; we are not aware of any evidence that directly bears on this question.

We agree with Marcus(2001) that registers may be central to human cognition. Trehub(1991) proposed that autaptic cells- cells that feed back into themselves-could serve as rapidly updatable bistable devices. A related proposal comes from Calvin(1996), who proposed a set of hexagonal self-excitatory cell assemblies that could serve as registers.

Although MLPs do not provide for registers, it is an easy matter to construct bistable registers out of nodes and connections. All that is needed is a single node that feeds back into itself. If the input is 0, the output tends to go to 0; if the input is 1, the output tends to go to 1. If the input is 0.5, which we can think of as the absence of a write-to-memory operation, the output tends to remain unchanged. Once the input is taken away, the model tends to remain stable at one or another *attractor point* (0 or 1). In this way we may use the self-feeding node as a memory component within a more structured network.

Although it is often assumed that knowledge is stored in terms of changes in synaptic connection weights, it is logically possible that knowledge is stored within cells. For example the reciprocal modulation of ion channels could provide an intracellular basis for registers.

Registers, however they are implemented, can provide a basis not only for variable binding but also, more generally, for the kinds of memory in which we learn things on a single trial. Such rapidly updatable memory clearly plays an important role throughout our mental life. Whatever rapidly updatable neural circuitry supports these kinds of everyday experiences could also be used to support registers that store instances of variables.

9.4 Temporal Synchrony

Although that at least some registers will be defined in terms of physically isolable parts of the brain (cells, circuits, or subcell assemblies), several other possibilities have been proposed in the literature. Most prominent among these alternative possibilities is *temporal synchrony*, also known as *dynamic binding* (Shastri and Ajjanagadde, 1993), which we can think as a framework for representing registers in time rather than in space.

In the temporal synchrony framework, both instances and variables are represented by nodes. Each of these nodes oscillates on and off over time. A variable is considered to be bound to its instance if both fire in the same rhythmic phase.

This proposal is motivated by the suggestions of neuroscientists such as von der Malsburg (1981) and Singer et al (1997) that the synchronization of the activity of neurons may play an important role in neural computation. One potential limitation of temporal synchrony is that it is likely to be able to keep distinct only a small finite set of phases, typically estimated as less than 10. Hence such a system can simultaneously represent only a small set of bindings. Nevertheless Shastri and Ajjanagadde (1993) have suggested that the limitation on the number of phases can capture limits in rapid reasoning, while Hummel and Holyoak (1997) have suggested that the limitation on phases can help to account for some phenomena in our computation of analogy. But it is clear that temporal synchrony is inadequate for representing *long-term* bindings between variables and their instantiations. We may probably memorize millions of bindings in long-term memory, yet on nobody's account can the brain keep distinct millions of phases.

One possible alternative is to combine frequency coding and phase coding. The problem is that the distance between two neighbor frequencies must be very small to represent millions of bindings, and that would increase the 'decodification' or 'reading' errors.

10 Temporal Synchrony and Our Results

In the same vein our group did find evidence in favor of the relevance of phase synchrony for the representation of cognitive-affective states in distinct areas of the brain (Isabel Barahona da Fonseca et al, 2001). On the other hand we proposed a further more comprehensive model based on Quantum Chemical Valence Theory interpreting representations of events by periodic oscillatory phenomena in dendritic trees (José Barahona da Fonseca, 2003).

Beyond the simple enunciation of the possible meaning of frequency and phase synchrony it was proposed that the mathematical apparatus of Quantum Chemistry, involving eigen values and functions may be used to represent classes of equivalence which define concepts in Brain processing.

The model that was used took as a referent the theory of isolobal organo-metallic components of Hoffmann. This proposal relates phase synchrony with a complete field of mathematical relationships which are useful to match psychological events, as for instance the generation of qualitative attributes by lower level valence attributes.

11 Discussion

Connectionist models can tell us a great deal about cognitive architecture but only if we carefully examine the differences between models. It is not enough to say that some connectionist model will be able to handle the task. Instead, we must ask what architectural properties are required. Marcus(2001) showed that models that include machinery for operations over variables succeed and that models without such machinery do not.

12 Conclusions

Sensory information processing, procedural implicit as well as verbal declarative cognitive processes are discussed extensively to show the relevance of concepts and hypotheses for Neuroscience research. An example of an alternative to current empirical discussions is given using Marcus's results. It is mentioned our contribution to this type of alternative approach to Neuroscience.

References

Anderson J. A. and Hinton G. E. (1981). Models of Information Processing in the Brain. Parallel Models of Associative Memory. Edited by G. Hinton and J. A. Anderson, Erlbaum, Hillsdale, NJ.

Barahona da Fonseca J., Barahona da Fonseca I., Suárez Araujo C. P. and Simões da Fonseca J. (2003). New Paradigms for Information Processing: A New Approach to Dendritic Computation. Quantum Chemistry Models for Cognition. paper submitted to Casys 2003.

Barahona da Fonseca I., Barahona da Fonseca J., Suárez Araujo C. P. and Simões da Fonseca J. (2001). Cognition: Characteristic Waveform Correlates. Journal of Computing Anticipatory Systems, Vol 10, pp. 265-282.

Bates E. A. and Elman J. L. (1993). Connectionism and the Study of Change. Brain Development and Cognition. Edited by M. H. Johnson, Basil Blackwell, Cambridge, MA.

Bechtel W. and Abrahamsen A. (1991). Connectionism and Mind: An Introduction to Parallel Processing in Networks, Basil Blackwell.

Elman J. L. (1990). Finding Structure in Time. Cognitive Science, 14, pp. 179-211.

Elman J. L. (1991). Distributed Representations, Simple Recurrent Networks, and Grammatical Structure. Machine Learning, 7, pp. 195-224.

Elman J. L. (1993). Learning and Development in Neural Networks: The Importance of Starting Small. Cognition, 48, pp. 71-99.

Funashi S. M., Chafee M. V. and Goldman-Rakic P. S. (1993). Prefrontal Neuronal Activity in Rhesus Monkeys Performing a Delayed Anti-Saccade Task. Nature, 365, pp. 753-756.

Gluck M. A. and Bower G. H. (1988). Evaluating an Adaptive Network of Human Learning. Journal of Memory and Language, 27, pp. 166-195.

Hadley R. L. (2000). Cognition and the Computational Power of Connectionist Networks. Connection Science, 12.

Hall C., Bernoties L. and Schmidt D. (1995). Interference of mental imagery on a mental motor task. British Journal of Psychology, 86, pp. 181-190.

Hare M., Elman J. and Daugherty K. (1995). Default Generalization in Connectionist Networks. Language and Cognitive Processes, 10, 601-630.

Helmholtz H. L. F. [1877], (1954). On the Sensations of Tone as a Physiological Basis for the Theory of Music. Dover, New York.

Hubel D. H. and Wiesel T. N. (1963). Shape and Arrangement of Columns in Cat's Striate Cortex. J. Physiol., 165.

Hummel J. E. and Holyoak K. J. (1997). Distributed Representations of Structure: A Theory of Analogical Access and Mapping. Psychological Review, 104, pp. 427-466.

Jeannerod M (1997). The Cognitive Neuroscience of Action. Oxford, Blackwells.

Koch C. and Crick F (2000). Some Thoughts on Consciousness and Neuroscience. The New Cognitive Neurosciences. Edited by M. S. Gazzaniga. MIT Press. Mass. pp.1285-1294.

Lettvin J. Y., Maturana H. R., Pitts W. H. and McCulloch W. S. (1959). What the Frog's Eye Tells the Frog's Brain. Proc. Inst. Radio. Eng., 47.

Marcus G. F. (2001). The Algebraic Mind: Integrating Connectionism and Cognitive Science. The MIT Press, Cambridge, Massachusetts.

Mareschal D., Plunkett K. and Harris P. (1995). Developing Object Permanence: A Connectionist Model. Proceeedings of the Seventeenth Annual Conference of the Cognitive Science Society. Edited by J. D. Moore and J. F. Lehman, Erlbaum, Mahwah, NJ, pp. 170-175,.

Munakata Y., McClelland J. L., Johnson M. H. and Siegler R. S. (1997). Rethinking Infant Knowledge: Toward an Adaptive Process Account of Successes and Failures in Object Permanence Tasks. Psychological Review, 10(4), pp. 686-713.

Orlick T. (1986). Psyching for Sport. Mental Training for Athletes. Champaign, Illinois

Plunkett K, Sinha C, MØller M F and Strandsby O (1992). Symbol Grounding or the Emergence of Symbols? Vocabulary Growth in Children and a Connectionist Net. Connection Science, 4, 293-312.

Quinn P C and Johnson M H (1996). The Emergence of Perceptual Category Representations During Early Development: A Connectionist Analysis. in G W Cottrel (Ed), *Proceedings of the Eighteen Annual Conference of the Cognitive Science Society*, Erlbaum, Hillsdale, NJ.

Rumelhart D E and McClelland J L (1986a). On Learning the Past Tenses of English Verbs. in J L McClelland, D E Rumelhart and the PDP Research Group (Eds.), Parallel Distributed Processing: Explorations in the Microstructures of Cognition, Vol. 2, Psychological and Biological Models, The MIT Press, Cambridge, MA.

Rumelhart D E and McClelland J L (1986b). PDP Models and General Issues in Cognitive Science. in J L McClelland, D E Rumelhart and the PDP Research Group (Eds.), Parallel Distributed Processing: Explorations in the Microstructures of Cognition, Vol. 1, Foundations, pp. 110-146, The MIT Press.

Seidenberg M S (1997). Language Acquisition and Use: Learning and Applying Probabilistic Constraints. Science, 275, 1599-1603.

Shastri L and Ajjanagadde V (1993). From Simple Associations to Systematic Reasoning: A Connectionist Representation of Rules, Variables, and Dynamic Bindings using Temporal Synchrony. Behavioral and Brain Sciences, 16, 417-494.

Simões da Fonseca, J (1999). Almada, Camões e o Número. author's edition, Lisbon.

Singer W, Engel A K, Kreiter A K, Munk M H J, Neuenschwander S and Roelfsema P R (1997). Neuronal Assemblies: Necessity, Signature and Detectability. Trends in Cognitive Sciences, 1(7), 252-261.

Smolensky P (1990). Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems. *Artificial Intelligence*, 46, 159-216.

Touretzky D S and Hinton G E (1988). A Distributed Connectionist Production System. *Cognitive Science*, 12, 423-466.

Trehub A (1991). The Cognitive Brain. The MIT Press, Cambridge, MA.

von der Malsburg C (1981). The Correlation Theory of Brain Function. Technical Report 81-2, Department of Neurobiology, Max-Planck-Institut for Biophysical Chemistry.