

Indispensability of Computational Modeling in Cognitive Science*

Igor Farkaš

Comenius University in Bratislava

farkas@fmph.uniba.sk

The concept of computation remains a frequently discussed topic in cognitive science, but there is no consensus about its meaning and the role in this field. I discuss this concept in wider sense, also including nonclassical computation, in the light of Marr's three levels of analysis and their relevance for main modeling frameworks pursued in cognitive science – symbolic, connectionist, dynamic and probabilistic. I point to differences between these approaches and argue, providing empirical and theoretical arguments, that connectionism, out of the existing approaches, holds the promise of providing the most plausible and detailed accounts of human cognition. Connectionism also benefits from the emerging field of cognitive developmental robotics that aims at designing autonomous cognitive robots using the synthetic bottom-up approach. I conclude with emphasizing the key role of computational modeling that will help advance the field of computational cognitive science as an indispensable core component.

Key words: *computation, cognition, modeling frameworks, levels of analysis, connectionism, learning, representation, developmental cognitive robotics*

1. Introduction

The birth of cognitive science, triggered by the Cognitive revolution in the 1950s, was an outcome of parallel endeavors to study the human mind and mental processes (Gardner, 1987). Since then, each of the relevant research

*This work was supported by Slovak Grant Agency for Science, grant no. 1/0439/11.

areas known as mother disciplines has proposed answers and has formulated new questions related to various aspects of the mind and cognition. Ancient philosophers were the first ones to put these topics on their agenda, and (much) later the natural sciences came to help using experimental methods. Cognitive psychologists formulate theories and research hypotheses and test them using the behavioral experiments. Language faculty, as a gate to the mind of a human subject who can report about his/her internal mental states or processes, is undoubtedly crucial for studying cognition. Psycholinguists focus on the language-related behavioral experiments, whereas neurolinguists are interested in relations between the brain lesions and the linguistic behavior, attempting to reveal the neural correlates of language. The invention of the modern digital computer brought a new dimension to the study of the mind and cognition in terms of the computational modeling. Cognitive anthropology also has its place in studying the human behavior in the social context. Last but not least, since 1990s, the experiments can be efficiently enriched by the modern brain imaging methods (such as fMRI) that enable to shed light on the brain correlates in 3D at a higher spatial resolution.

Since various mother disciplines use different methods and tools in research, the question pops out whether cognitive science is a single discipline with a single object of the study (cognition, mind–brain), or whether there are multiple disciplines. The Google search for “cognitive science” reveals around 6 million hits, whereas “cognitive sciences” return roughly 1,7 million hits (in November 2012). So it seems that cognitive science is perceived, at least when judged from the perspective of the usage of both terms, more often as a single research discipline. I consider this evidence as additional support for my view, assuming that the statistics covers the opinions of people involved in cognitive science(s). Therefore, I think this statistics has some explanatory value.

I think it is good that cognitive science is more often perceived as a single discipline because the study of such a complex multifaceted entity as the mind does require heterogeneous, interdisciplinary research strands, and it should be a continuing challenge for involved research approaches to inform and inspire each other. Spreading the view of cognitive science as a single discipline could facilitate the researchers’ intentionality not

only to make progress in their particular area of expertise but also to try to go beyond, at least by absorbing what's new and relevant for their own research in cognitive science and by identifying the links, analogies or similarities.

In this paper, rather than trying to provide a unifying view at cognitive science from various angles, I will focus on one of these angles. I will discuss the (still often debated) role of the computational approaches in cognitive science and will emphasize their inevitability for advancing the field.¹ Specifically, I will advocate the role of neural computation, or connectionism, that I see the most plausible candidate for approaching the mind–brain. I will conclude with highlighting the emerging field of cognitive developmental robotics that allows testing theoretical ideas in real settings.

2. Historical Context of Computation

The birth of a concept of digital computation and of the modern digital computer in 1950s was clearly an important step that triggered the progress of several disciplines, including the artificial intelligence (AI) and cognitive science. However, the world of computation came into being long before the Turing machine (Turing, 1936). Sloman (2002) introduced the historical context by offering his view, that the digital computer was a result of a convergence of two strands of development with a long history: development of machines for automating various physical processes (e.g. clocks, weaving machines, sorting machines, etc.) and machines for performing abstract operations on abstract entities (e.g. doing numerical calculations, various operations on symbols). The universality of a digital computer makes it a powerful programmable device, that has been since its invention widely used as a tool in numerous applications.

On a theoretical side, Sloman also discusses the role of the Turing machine and its usefulness in the context of abstract mathematical operations. The Turing machine, currently being revived at the centenary of its

¹ I basically agree with McClelland's (2009) view related to surveyed computational approaches but here I make a special emphasis on connectionism by listing several arguments in its favor, along with remaining challenges.

inventor (e.g. Cooper, 2012; French, 2012), is argued to be irrelevant for AI and cognitive science, because it does not help in solving practical AI tasks, nor in understanding how the brains could work (Sloman, 2002). The distinction between a powerful concept of the Turing machine (operating in unlimited time using an unbounded tape) and a computational device with limited resources attempting to account for human cognition is reminiscent, I think, of the distinction between competence and performance in the theory of human language (Chomsky, 1965). Of course, abstractions are not only useful but often inevitable in research and it is indeed fascinating that our brains (with limited resources) are able to work with and to reason about highly abstract concepts in infinite spaces and structures.² But any computational account of human cognition should consider an entity with limited resources, so in this regard it is reasonable to consider the (finite) physical digital computer as a potential candidate for explaining human cognition (if one is a proponent of the symbolism as discussed further).

It is interesting that although the two abovementioned strands were very different in their objectives, they had some features in common (Sloman, 2002). For instance, each strand involved both discrete and continuous machines. Machines could work either with continuous variables (e.g. speed governors) or discrete variables (e.g. sorting machines), and a calculator could either rely on continuous slider-ules or discrete devices (e.g. ratchets). The other commonality was related to the degree and nature of human involvement in the interaction with the machine (e.g. where the human is involved in taking decisions and feeding control information, or merely provides the energy, once the machine is set up for the task).

So the technology provided us with a variety of machines before the universal digital computer with its vast application potential took over. While nobody doubts that the modern digital computer is also a machine, there are still disagreements about what distinguishes the digital computer from earlier machines, whether classical machines compute in the same way as the digital computer, or what is the class of systems that compute. Finite

² Understanding this ability (in skilled individuals) also belongs to the ultimate goals of cognitive science, but before achieving that we must first understand more “rudimentary” cognitive abilities, like perception, motor skills, or language.

state automata, Turing machines, calculators and digital computers are commonly assumed to compute. But how about other entities? The answer depends on assumptions. These range from very strict (e.g. Pylyshyn, 1984), that disqualify even finite state automata, to very soft assumptions (e.g. Scheutz, 1999), according to which also hurricanes or digestive systems compute. Piccinini (2007) holds the intermediate view that “only the right things compute, the others don’t.” I think that in principle, all physical systems whose variables can be measured, can be *viewed as* computational systems. The distinguishing feature for *truly* computational systems may be that they must be artifacts (physical or virtual), constructed in order to perform computations whose results are interpretable by humans.

The related issue concerns the type of computation. The classical (narrow) view of computation is associated with discrete computations in discrete time. In the paper, I adhere to the view that computation should be considered in wider sense to also embrace nonclassical computation, such as analog computation, quantum computation or probabilistic computation.

3. The brain as a Computational Device

The modern digital computer seems to be the only entity whose computational nature is not questioned by anyone. In the context of cognition that aims to explain how the mind–brain and its underlying cognitive processes work, the following question has been formulated: Is the brain, as the assumed physical host, a computational device? If it is, then the level at which these computations take place, should be identified. But unlike the digital computer whose level of computations is well-defined (logic gates), in the case of the brains no such level is evident. As argued by Churchland and Sejnowski (1992), the brains were not constructed to compute but rather evolved to allow the organisms to survive in a dynamic environment. If we want to look at the brain as a computing entity, we must ask: What is the most plausible level at which the brain computations could be defined and formalized? Thank to the known brain anatomy, the first hypothesis that comes to mind is the level of neurons. But it could be argued that it is only a rough approximation because subcellular effects are known to influence neuron firing, and neuron functioning is also affected by various

(slower) hormonal and neuromodulating systems (Kaczmarek & Levitan, 1987). In addition, due to the functional redundancy and permanent loss of neural cells, the contribution of each cell probably does not matter, so maybe a higher level of brain organization (e.g. the neural assemblies) would be more appropriate. For instance, the large-scale brain networks are a frequent approach to computationally study the brain functions at the system level (e.g. Sporns, 2011). From the perspective of information flow, the brain functions within a plethora of feedback loops operating at various spatiotemporal scales in continuous time (Bell, 1999), which makes the effort to capture these processes computationally much harder. Similarly to many other natural complex systems, the brain and its function(s) have a highly nonlinear dynamic character.

Fortunately, it is possible to simulate continuous dynamics with very high precision using discrete systems such as digital computers. Since the real world is noisy and so are the biological neurons, infinite precision is not realistic but useful, as will be argued later in the context of the recurrent neural networks. However, the principles of neural coding used by the brain have not yet been reliably identified (Rieke et al., 1999), because different rate coding and spike coding theories both provide plausible accounts of various investigated phenomena. Despite that uncertainty, in order to be able to proceed in cognitive science, we must view the brain as a computational device, without worrying much about philosophical questions associated with this assumption. Indeed, the devoted research field of computational neuroscience flourishes, as witnessed by numerous conferences and the growing number of published journal papers.

4. Importance of Computational Modeling

Computational modeling is viewed by many as the central pillar in modern cognitive science (McClelland, 2009; Chalmers, 2011). However, its role still seems controversial since there is no consensus in the literature about the meaning of key concepts such as computation or information processing (Piccicini & Scarantino, 2011).

For instance, Chalmers (2011) defends his view of the so-called minimal computationalism as based on the proposed two theses: computational suf-

ficiency, stating that the right kind of computational structure suffices for the possession of a mind and for the possession of a wide variety of mental properties; and computational explanation, stating that computation provides a general framework for the explanation of cognitive processes and of behavior. The argument for computational sufficiency can arguably be questioned (Ritchie, 2011) since it is unclear what mentality entails. Are computational correlates sufficient for explaining consciousness? Is mentality an inherent feature of a complex biological matter? The answers depend on definitions of concepts. I am not going to deal with highly debatable issues related to the nature of mentality or conscious states. Instead, my goal is to look at the computational approach as the description, discuss its features in some more detail, and emphasize its inevitability for building cognitive systems.

Whatever modeling approach is to be taken, it cannot go without formalization. Computational models are crucial mainly for three reasons: First, they allow to formally describe cognitive processes that gives them rigor and explicitness that are in principle not achievable by verbal theorizing (mainly due to the ambiguity of language). This can help understand mechanisms underlying cognition. Second, the computational models provide testable hypotheses and inspire further empirical research. Third, they are used in building autonomous intelligent robots, that will interact with the humans, which is a progressive research strand. I see the last reason as having not only the added, but crucial value, being extremely desirable in the computational cognitive science. I will not speculate about the ontological and social status of robots in the future, but from the cognitive science perspective, the inclusion of these intelligent agents in our environment may trigger new research questions, and may lead to redefinition of some key concepts and/or to adding new ones.

In his recent overview of computational modeling in cognitive science, McClelland (2009) asserts that “the essential purpose of cognitive modeling is to allow investigation of the implications of ideas, beyond the limits of human thinking.” This is true I think, because the human cognition is very complex, the human ability to study it (i.e. ourselves) rather limited, so we need external tools that will help us understand this complexity. On the other hand, a word of caution is due here, regarding the explanatory power

of the models. McClelland admits that “even if a modeler can show that a model fits all available data perfectly, the work still cannot tell us that it correctly captures the processes underlying human performance in the tasks that it addresses.” Put in other words, the computational models provide only candidate explanations. The well-known quote of Box and Norman (1987, p. 424) suggests that “all models are wrong, but some are useful” may reflect reality, but again, I argue that good models are not only useful but crucial. The role of models touches upon the philosophical questions of ontology and our epistemological tools to learn about the nature of cognition. In terms of the rigor of qualitative and quantitative description, there is no better alternative to the mathematical description of the models and their application.

5. Brief Review of Modeling Frameworks in Cognitive Science

As already mentioned in Section 2, the term computation became biased since the birth of cognitive science, to refer to digital computations, implemented in a digital computer. When people talk about computing, they usually mean discrete mental operations, needed for tasks like adding of multiplying numbers. At this point, I will make a brief survey of the main modeling frameworks as they were formed and have been used in cognitive science, and I will briefly compare them.

Prior to doing so, I find it useful to introduce here the levels of analysis, as proposed by the influential theory of Marr (1982) who postulated three levels of analysis – computational, algorithmic and implementational, by forging the analogy with the digital computer. Computational level defines computations that should be performed (a mathematical function, or a task specification), without saying how to do them. Algorithmic level specifies the used representations and algorithms for performing the desired computation. Implementational level specifies the processes of information processing that are tied to the specific hardware used for running the algorithm. In Marr’s view, the levels are basically independent of one another, and a lower level typically provides multiple realizability of the specification at a higher level.

5.1. Symbolism

It is no surprise that the invention of the concept of digital computation inspired the symbolic paradigm in cognitive science. Symbolic paradigm conceptualizes the mind as a (digital) computational device, separable from the environment, that manipulates internal symbols according to logical rules, in the same way as a computer.³ Being inspired by earlier works (Tur-ing, 1950; Newell & Simon, 1976; Fodor & Pylyshyn, 1988), Harnad (1990) reconstructed an accurate definition of a symbol system, where the crucial role is played by the syntax, and where the semantics can be added to it from outside.

The symbolic view is still considered by many (e.g. Fodor, 2000; Pinker 2009) as the right level of abstraction, because it provides elegant and powerful formalisms for representing knowledge, it captures important human intuitions (biases) about the symbolic character of cognition, and it could be implemented in AI. Symbolism views cognition as manipulation of amodal (i.e. independent of modalities) symbols in an algorithmic way which renders the cognition as the classical computation that can be implemented also in standard computers. The prominent Computational Theory of Mind, introduced by Putnam (1961) and developed mainly by philosophers (Fodor, 2000), evokes the view of the mind as a universal computing device (information processing system). The crucial feature of symbolism is that the mind is driven by programs, which are realized by algorithms. How these algorithms are realized in the brain, is in the symbolic view merely a matter of implementation. The irrelevance of the implementation has also been enforced by Marr's theory. The computer is truly a dual entity with an independent hardware and software, where computer algorithms can easily be turned into implementations by the completely automatic process of compilation (translation to a hardware-dependent machine code). In contrast, in the brain (see also Section 3) the (neural) implementation is certainly not derived automatically from some higher-level description

³ This paradigm is in the literature also referred to as cognitivism, but I will avoid that term because it evokes (I think undesirable) implication that this paradigm is closest to the nature of cognitive processes.

(O'Reilly & Munakata, 2000, p. 5). It turns out that implementational level plays an important role and should be appropriately linked to the higher level description, in order to facilitate the interpretation of the function.

5.2. Connectionism

Connectionism challenges a basic assumption of much of AI that mental processes are best viewed as algorithmic symbol manipulations. It represents a spectrum of methods that arose within AI and were biologically inspired. The key concepts that we do not need to explain here include: a high number of simple processing elements (neurons), with trainable connections (synapses) between them, parallel processing of information and distributiveness of knowledge in the system. The communication is numeric, rather than symbolic, which leads to subsymbolic representations.⁴ A neural network can process continuous variables in the form of vectors, subjected to nonlinear transformations (between layers of units). This leads to the lack of transparency of its function ("black box"), which requires some effort to debug information, using the techniques of clustering and visualization (O'Brien & Opie, 2006).

The connectionist approach, as opposed to symbolism, includes also the implementational level of analysis. Some philosophers suggest that the explained phenomena be separated from mechanisms of their origin (Abrahamsen & Bechtel, 2006), but in connectionist models these two aspects are connected, so there exists dependence between implemented mechanisms (given by neural interactions) and cognitive phenomena (O'Reilly & Munakata, 2000). In other words, the specification of the neural network can be provided at the computational level (in terms of a function) but the elementary functions (related to neurons) already provide room for direct implementation (of neuron's input-to-output functions) in neural hardware, that bears some similarity to biological networks.

The crucial property of neural networks is that the designer's approach

⁴ It should be noted that localist models of neural networks (as opposed to models with distributed representations) resemble more symbolic systems, in the sense that what each neuron represents does not matter (the form and content are separated), they could be swapped without affecting the function (Page, 2000).

is restricted because it is possible to take advantage of learning paradigms and the different levels of detail in feedback from the environment (Haykin, 2008). Neural networks are also commonly assumed to compute. How they compute, will be described in Section 6.

5.3. Dynamical Systems

The operation in a continuous space and time is typical for dynamical systems framework, which, unlike symbolism and connectionism, puts emphasis on the situated and embodied nature of human behavior (Schoener, 2008). Cognitive processes are viewed as mental activity unfolding in real time, which can be described, at the computational level, by the system of coupled nonlinear differential equations. The real time processing within the brain, and the coupling between the brain, body and environment, with continuous interactions and mutual reciprocal causation, renders the view on the mind as a permanent dynamic process taking place in a high-dimensional state space, being complementary in all aspects to the discrete sequential symbolic process. The dynamical systems framework forms a pillar of enactive approach to cognition (Maturana & Varela, 1992) that is built on concepts like experience, autonomy, emergence and sense-making.

The radical thesis of the dynamical systems framework disregards representations as unnecessary (Thelen & Smith, 1994), but softer versions of this framework (e.g. Kelso, 1995) are compatible with the representational-computational view of the mind-brain. The dynamical systems framework has found empirical support in numerous experiments (Spivey, 2007) and has been used to account for certain types of human behavior (Tschacher & Dauwalder, 2003). It uses two types of mutually connected variables: the *collective* variables that express the relations between the interacting components of the system and explain the behavior, and the *control* variables, whose quantitative changes can lead to qualitative changes in system behavior (phase transitions). Whereas dynamical systems definitely brought new ideas into the understanding of cognition (mainly in terms of agent's interaction with the physical world), it is these control parameters, whose change, with the goal to account for certain aspects of cognitive development, often comes from outside the model (the problem of learning). Finding the ways to allow the changes in these variables to arise as a result of

experience and behavior would clearly enhance the dynamical systems framework (McClelland, 2009).

5.4. Probabilistic Framework

The probabilistic approach has become popular in cognitive science especially in the last decade (Perfors, Tenenbaum, Griffiths & Xu, 2011). This principled (Bayesian) approach that exploits a broad spectrum of representations (trees, vectors, logical rules, etc.), combines them with statistical learning and making inferences under uncertainty, provides explanations at the computational level (in Marr's sense). Implementational level is considered less important, leaving the room open for links to potential neural mechanisms. Rather than going bottom-up, starting with mechanisms (the case of connectionism), probabilistic approach goes top-down, starting from the function that we want to explain, looking for an optimal representation for explaining the data. It also incorporates the nature–nurture aspect in terms of inductive biases (prior distribution over the set of hypotheses) that enter into computations of the posterior distribution.

The probabilistic approach has proven successful in explaining the wide range of human behaviors (see Griffiths et al., 2010, and references therein). The explanation at the computational level is not seen merely as a computational abstraction of underlying mechanisms, but is claimed to have an independent validity as an account.

This approach stems from the belief that it is possible to understand the human cognition as an optimal response (leading to rational behavior) to the constraints placed on the cognizing agent by a situation or a set of situations. McClelland (2009) raises two arguments. Namely, that rationality depends on agent's cost function and is context-dependent, each of which renders the behavior suboptimal from observer's perspective (but not the agent itself). Another difficulty with Bayesian approaches is computational intractability for larger problems, and the lack of online or incremental learning versions. On the other hand, I think, the increased interest in this approach predicts its importance in the future, and it will be interesting to find the potential links between the probabilistic computational level and the neural implementational level in understanding human knowledge.

5.5. Comparison of Approaches

Probabilistic approach shares several features both with symbolism and with connectionism. Like symbolism, it is a top-down approach, using a very powerful and expressive framework in terms of symbolic data representations. With connectionism, it shares continuous representations, the strong emphasis on learning, and some other commonalities. For instance, some forms of connectionist learning have natural Bayesian interpretations (e.g. Mackay, 1996), and the subclass of generative connectionist models uses the probabilistic framework, while preserving its connectionist style (e.g. deep architectures; Bengio, 2009). Recurrent neural networks can basically be viewed as dynamical systems. They have a continuous state space (typically on the hidden layer), where the unit activations unfold in time (activation dynamics), combined with parameter change (learning dynamics), that can give rise to interesting emergent phenomena in the organization of the state space.

Connectionism is claimed to make strong assumptions about the nature of human mental representations and inductive biases based on a certain view of neural mechanisms and development, such as graded distributed representations, lacking explicit structure, being shaped almost exclusively by experience (Griffiths et al., 2010). That is basically true, although connectionist models also use localist (symbolic) representations (Page, 2000). Nevertheless, connectionist commitments to certain types of representation and learning mechanisms should be viewed as principled, intentional and biologically inspired. Of course, this assumption imposes potential difficulty for connectionist models to account also for higher aspects of cognition, that are typically dealt with by symbolic or probabilistic models.

As argued by the proponents of probabilistic approach, the representational variety is useful because it allows to identify the types of representation that best account for certain cognitive behavior. This seems justified at the computational and algorithmic levels, where the type of ideal representation could depend on the task as suggested by numerous examples cited in Griffiths et al. (2010). From the evolutionary perspective, however, one could ask, whether it is plausible to look for qualitatively different, task-dependent accounts, or whether evolutionarily older principles were used

in the service of higher functions. In this context, Clark (2001) introduces the concept of *cognitive incrementalism* and considers it a remaining big issue in cognitive science. Considerable empirical evidence, covered by the umbrella of grounded (=embodied+embedded) cognition, suggests that higher cognition is embodied in the lower-level sensorimotor processes (Barsalou, 2008). Connectionism seeks to provide a uniform account, constrained by the implementational level, whereas probabilistic approach speaks in favor of plurality, predicting that ultimately, for all computational accounts, the corresponding neural mechanisms may be found. Symbolic account also takes a uniform approach, focusing on higher functions of the mind, and not caring much about lower aspects of cognition (i.e. those related to sensorimotor behavior). Dynamical systems framework also operates at the computational level, but its concepts are useful also for connectionism (in particular, recurrent neural nets) as argued below.

6. Computation and Neural Networks

The above mentioned frameworks are all meant to perform some kind of computations in the common sense, but a closer look reveals that they do not involve the same types of computation. Actually, all four frameworks refer to the computational level of analysis, but only connectionism makes commitments to the implementation. Algorithmic level is relevant mostly for symbolism and connectionism. Yet, these two frameworks, and I agree, are qualitatively different accounts (the view of eliminative connectionism), and neural networks provide a better, more accurate account of the human cognition than symbolic models (see also Feldman & Ballard, 1982; Greco, 1998; Churchland & Sejnowski, 1992; Bechtel & Abrahamsen, 2002).⁵ Before continuing with argumentation, it is useful to realize what kind of computation is performed by neural networks.

6.1. Computation in Neural Networks

As recently reviewed by Piccinini (2008), connectionism does not embrace

⁵ Connectionist models are also frequently listed among the greatest modeling hits as judged by (around 70, personal comm.) cognitive modelers (Cottrell, 2012).

a homogeneous group of methods, because it covers both the models with discrete and continuous variables that operate in discrete or continuous time. The first neural networks composed of logical neurons (McCulloch & Pitts, 1943) resembling the logic gates operated in discrete time, used discrete activation values, were not able to learn, so they can be viewed as equivalent to the symbolic system. The function of such a network can be described by a program that realizes it.

Perceptrons (Rosenblatt, 1962) already introduced continuous activation functions, but remained functioning in discrete time. Importantly, these models can be trained, so there exist learning rules implementing arbitrarily small changes of model parameters. The learning rule is a key element in an adaptive system, because it allows for *elementary changes in knowledge representation*, and this can only be well dealt with in the domain of (real valued) numeric representations. In this case, we deal with analog, *non-classical computation*, which implies that the model realizes the required function (i.e. it assigns desired output to a given input) without having to execute a program.⁶

The third category of neural networks are those that operate in continuous space and time. These are typically represented by spiking neural network models (Maass & Bishop, 1999), that have become more popular during the last two decades and formed the new field of computational neuroscience. In this case we cannot talk about computations, because these require discrete time. A simulation of these models with infinite precision can only be achieved using analog computers, but fortunately the digital computers will often do the job sufficiently well.

Computational neuroscience as a research discipline differs from a more psychologically oriented connectionism by focusing on a more detailed computational description (being also constrained by implementational level). It emphasizes the features of functional and biologically realistic

⁶ Let us consider the often cited example of multiplying two numbers. A neural network can be trained to learn to multiply two (also larger) integers, without having to run a program (as opposed to symbolic systems, and children that learned the procedure at school). In a neural network, the result comes out in a single step (nonlinear transformation).

neurons (and neural systems) and their physiology and dynamics. However, these models cover multiple spatial-temporal scales, ranging from membrane currents, protein and chemical coupling to network oscillations, columnar and topographic architecture and learning and memory. Computational neuroscience attempts to model not only purely neural phenomena but also tries to relate these to mental processes. Therefore, computational neuroscience should be considered as one of the perspectives of computational cognitive science, even though it has not been commonly identified as one of the modeling frameworks.

6.2. Computational and Explanatory Power of Neural Networks

Some theoretically interesting ideas arose from earlier papers focusing on the computational properties of neural networks. The universal approximation theorem provides a proof that a feedforward neural network with one hidden layer can approximate an arbitrary continuous function with arbitrary precision (Hornik, Stinchcombe & White, 1989). Similarly, the often cited equivalence proof states that a recurrent neural network with sigmoidal units (of which the Elman network is the best known example; Elman, 1990) is computationally equivalent to the universal Turing machine (Siegelmann & Sontag, 1991). These rigorous results are of high theoretical importance, but not so much of practical importance in computational modeling, because they say nothing about how to construct a network (i.e. to set its weights) that would perform desired computations. This is important because the explanation of phenomena in cognitive science includes not only the processing and representation of information but also the acquisition of knowledge, and this requires learning.

Nevertheless, there is another important theoretical argument that is relevant in the context of recurrent neural networks, as opposed to (symbolic) Turing machines. Using sequences of symbols (with one-hot encoding) for training recurrent nets (dynamic systems) allows one to investigate the underlying dynamics in the parameter space. In line with earlier results by Siegelmann (1999), Tabor (2009) provides the analysis of the so-called affine dynamical automata (a simplified linearized version of an artificial recurrent neural net) showing that they can exhibit a range of symbol processing behaviors some of which are *not achievable* by a symbolic sys-

tem (Turing machine). This makes the connectionist account more general, incompatible with the symbolic account.⁷ As Tabor (2009) explains, the super-Turing capability of neural nets allows a more complex dynamics that provides explanations of various cognitive phenomena, e.g. in the field of language modeling (Christiansen & Chater, 2001). The real-valued metric relations in network parameter space leads to computation with infinite precision, which, as Tabor argues, “allows us to discover principles of linguistic and logical organization like compositionality which would be out of reach if only finite state computation were employed.” This contradicts the claim of some symbolists (e.g. Fodor, 2000) that connectionist and symbolic accounts are equivalent because they can be transformed to each other. As shown, this transformation is possible only in some cases (when the generating partition of the state space exists).

6.3. Challenges for Neural Networks

Neural networks have undoubtedly proven successful in modeling various lower-level cognitive tasks, but they have been criticized for being fundamentally inappropriate for modeling higher cognition (Fodor & Pylyshyn, 1988). The area of “combinatorial symbol combination” has been claimed to be the core computational feature of symbolic theories, and plausibly a foundational element in empirically defensible properties of the human mind: productivity (allowing to generate unlimited number messages with limited vocabulary), systematicity (understanding “X loves Y” implies understanding “Y loves X”), and compositionality (parts are combined to larger chunks, contributing to the overall meaning). However, even in this combinatorial domain, connectionist models have demonstrated considerable empirical success. Productivity requires recursiveness and this has been studied extensively with recurrent neural network using the classes of artificial languages in Chomsky’s hierarchy (Christiansen & Chater, 1999). Compositionality has been investigated as well. For instance, Greco and Caneva

⁷ Mathematically speaking, the explanation is based on structural correspondence between continuous dynamics (the state space of a neural network) and symbolic dynamics (Turing machine), given by non/existence of the so-called generating partition of the state space (beim Graben, 2004).

(2010) point to advantages of compositional grounded representations for motor patterns using behavioral experiments, supported by connectionist simulations. An interesting phenomenon revealed about neural networks is the so-called functional compositionality, as a qualitatively different alternative to concatenative compositionality, typical for symbolic models (van Gelder, 1990). Systematic behavior in neural networks has been studied in relation with both the syntax (e.g. Farkaš & Crocker, 2008) and the semantics (Frank, Haselager & van Rooij, 2009). All these models provided qualitatively different explanations for the studied phenomena, without being mere implementations of symbolic systems.

Another important issue that remains a challenge for neural networks is the binding problem.⁸ In his very recent review, Feldman (2012) points out that the well-known (neural) binding problem comprises at least four distinct problems with different computational and neural requirements. These are: general (sensorimotor) coordination, visual feature-binding (in perception), variable binding (in language), and the unity of perception (subjective experience). As Feldman puts it, there has been significant continuing progress, partially masked by confusing the different versions of the binding problem. More concretely, coordination and visual feature-binding demonstrated the most significant progress, variable binding remains a challenge and the question of subjective unity of perception remains intractable. Interestingly, some of the proposed neural accounts borrow the ideas from the computational neuroscience (temporal synchrony).

Catastrophic interference⁹ is another common argument against connectionism. How to acquire new knowledge without erasing old one? Interference is assumed to be a negative consequence of using distributed representations that support generalization and robustness in neural networks. Several solutions to this problem have been proposed (French, 1999), of which two seem to be the most convincing. One is based on interleaving the newly acquired knowledge with old knowledge during learning. The other

⁸ Thanks to one of the reviewers for pointing to this issue.

⁹ It is also related to the stability-plasticity dilemma, i.e. how could the system learn fast without becoming unstable (children are known to be able to learn certain things based on very few exemplars).

enforces the formation of less distributed representations towards sparse coding, hence alleviating interference among different representations.

In the context of modeling language acquisition and processing, symbolic and connectionist accounts are usually viewed as qualitatively different alternatives. Interestingly, some connectionist models attempt to reconcile them, for instance, a recent model of sentence generation, that draws on sensorimotor information processing, incorporates certain Chomskyan ideas about innate syntactic knowledge and parameter setting (Takáč, Beňušková & Knott, 2012).

6.4. More and Shared and Distinguishing Features among Paradigms

Symbolism shares with connectionism the concept of representations, albeit with different properties (e.g. symbolic versus subsymbolic). On the contrary, the radical thesis of the dynamic paradigm denies representations, because there is no room for them in the continuously changing coupling system. This view can be acceptable in the context of agent's sensorimotor interactions with the environment but hardly in the wider context of cognitive processes, most of which do require internal states (mental representations), either cued (triggered by an external stimulus) or detached (Gärdenfors, 1996). Hence, the representational-computational approach is compatible with a softer version of the dynamical paradigm.

Connectionist, dynamical and probabilistic approaches may become in the future more integrated within the computational cognitive science (McClelland, 2009). Symbolic paradigm stands in a striking contrast with these paradigms and will probably preserve its role of a "coarse-grained," but understandable description of cognitive phenomena. For instance, the explanation of the past tense formation of English verbs ("ed" rule and exceptions) is a clear symbolic explanation that can be enriched by a process of acquisition as observed in developing children (Pinker & Ullman, 2002).

Regardless of the future of cognitive science, the continuing effort to use the symbolic approach will remain very useful, especially when it comes to designing software systems that will serve humans, such as knowledge-based (expert) systems for various domains, or the semantic web (Davies, 2006). Their potential is far from having been fully exploited.

To summarize, we should understand the concept of computation in

wider sense, to accommodate also nonclassical computation. Practically, all computational models can be well approximated by discrete computations, implemented in standard computers, but this is only a technical issue, not crucial for cognitive science as such. In other words, the fact that a digital computer can implement a neural network does not imply anything about the nature of human cognition. Various modeling paradigms appear to be useful for cognitive science in general. With respect to the three levels of analysis and all reviewed frameworks, the algorithmic level seems important only for connectionism and symbolism due to specified representations. In connectionism, the representations are also constrained by the implementational level. In symbolism, an important role is played also by involved algorithms (programs).

7. Importance of Learning

In trying to model the complex behavior, it is crucial to find the proper balance along the nature–nurture dimension. Actually, in an effort to use the models in autonomous embodied agents, there must be a strong focus on learning. The traditional symbolic approaches are powerful in using the representational framework but the incorporation of learning mechanisms is in principle difficult. Machine learning as a research field that has been developing independently from cognitive science, has much to offer. It covers a variety of techniques some of which are biologically inspired, the others are not. Machine learning covers both subsymbolic methods (e.g. neural networks) as well as symbolic methods (e.g. graphical or probabilistic models). It appears that all three categories of learning algorithms (supervised, unsupervised and reinforcement learning) are important, are biologically relevant, and together form a powerful framework (Doya, 1999; O'Reilly & Munakata, 2000).

The importance of learning is also one of the hallmarks of the related research field, labeled *computational intelligence* (CI; e.g. Engelbrecht, 2007), that focuses more on the so-called soft computing (i.e. various types of heuristics used in bio-inspired models), rather than hard computing that is traditionally anchored in logic-based core of AI.¹⁰ Computational intelli-

¹⁰ It should be noted that there is no complete agreement regarding the differences

gence covers primarily three core areas: neural networks, fuzzy systems and evolutionary methods. It is a bottom-up approach that focuses on numerical (rather than symbolic) data, it emphasizes pattern recognition, learning (adaptation), autonomy of agents, and several other features (Craenen & Eiben, 2003). On the other hand, despite common foci, computational intelligence is currently quite heterogeneous. Duch (2007) proposes his vision of computational intelligence, as it should proceed towards exploring a variety of learning methods, efficiently combining them into committees of models, with incorporation of meta-learning, towards an integrative theory of adaptability.

8. Design of Artificial Cognitive Systems

The necessity to apply learning mechanisms is reflected in the design of artificial cognitive systems, embodied in physical hardware. The field of robotics has progressed nicely during the last half century, delivering robots of all kinds, ranging from artificial insects up to complex humanoids. Given the knowledge of inverse kinematics and well designed control, these robots can execute motions with very high speed and precision, making the designer's approach very successful.

To meet the requirements posed by the bottom-up approach, the cognitive developmental robotics (Asada et al., 2009) emerged as a natural way of applying the constructivist principles to building autonomous artificial agents (robots) with higher degrees of autonomy (Pfeifer & Scheier, 2000). This research area benefits from increased attention, driven by the vision that the modeling of cognitive processes cannot be achieved without embodied agents, embedded in the environment the agent is interacting with (Pfeifer, Lungarella & Iida, 2007). The design of learning mechanisms is often inspired by cognitive neuroscience and developmental cognitive psychology, that abound with empirical data. Recent achievements are relatively modest, focusing on small scale problems with fewer degrees of free-

in the scope between AI, CI and the role of machine learning in them (Craenen & Eiben, 2003). Some proponents of the modern AI include CI tools under the umbrella of AI (Russell & Norvig, 2009).

dom, but I believe this trend will continue towards more complex scenarios. Cognitive robotics imposes constraints on learning mechanisms and representations, and it is quite difficult to achieve higher degrees of autonomy. The role of the designer is typically reduced to setting up the architecture, i.e. the modules and their connectivity, and the overall functional description. The suitable parameters are found by learning.

A number of cognitive architectures have been proposed in the last two decades, and of these, three major categories can be identified: symbolic architectures, subsymbolic (emergent) architectures, and hybrid approaches as the mixture of the two (Vernon, Metta & Sandini, 2007, and Weng, 2012, provide recent reviews). The popularity of the hybrid approach results from the difficulty of the purely bottom up approach and will probably remain the main constructivist approach in the nearest future. Within emergent approaches, artificial neural networks play a crucial role.¹¹

Empirical literature will remain an important source of constraints that could/should be considered in the design of artificial systems. For instance, the traditional sense–think–act cycle may be replaced with the view consistent with the common coding theory of perception and action, according to which these two modules are intertwined, deeply interacting, rather than being separate (Prinz, 1984). Similarly, the traditional distinction between cognitive processes (including the abstract ones) and sensorimotor processes is being replaced by the grounded cognition perspective, according to which cognitive processes are inseparable from sensorimotor processes, with respect to the underlying neural substrate (Barsalou, 2008). The discovery of the mirror neuron system in some higher species and in humans can be considered a big step forward, offering various interesting implications for modeling social cognition (Gallese et al., 2011).

9. Computational Modeling: Memento and Manifesto

The importance of models for cognitive science has been recognized not

¹¹ Using neural nets in robotics provides an added value for the connectionist paradigm towards the enactive approach (Varela, Thompson & Rosch, 1991), that lends itself to acquiring autonomous emergent behavior.

only by modelers but also by experts from “outside.” As recently mentioned by the cognitive psychologist D. Gentner (2010), since its birth in 1950s, cognitive science has gone through the development that reveals certain trend. The proportion of papers starts to be imbalanced because cognitive psychology is becoming a dominant discipline. Gentner (2010) comes with a memento that nowadays cognitive psychology accounts for more than 50% of scientific production in cognitive science (with its growth having started roughly in 1978, with a proportion of 26%), whereas the other two disciplines – artificial intelligence and linguistics – each account only for 20%. As Gentner puts it, “if present trends continue, then by 2038, psychology will have completed its conquest in cognitive science” (p. 330). However, this would be considered a Pyrrhic victory by most researchers (including cognitive psychologists) who are aware of the importance of interdisciplinarity of this field. Therefore, this memento should stimulate modelers in cognitive science. We have a lot of empirical data, thank to behavioral experiments and data imaging, but we need more computational models.

It is good to have computational models, but another issue is to ask how useful they are and whom they serve. Addyman and French (2012) recently came with a manifesto for change in computational modeling in cognitive science. As they argue, unlike the computer and software technologies that have advanced significantly during the last decades, modeling methodology in cognitive science has remained rather old-fashioned, nontransparent, preventing wider exchange of information within the community, not only to other modelers, but also outside the core modeling community. The programmers typically write their code using one of the preferred available programming languages (of various kinds), but rarely provide a well-documented source code on their website hence prompting other researchers to use it. This reluctant approach can probably be explained, I think, by the absence of motivation to do so, because this would require an extra energy (the academia forces the modelers to publish papers which does not imply the need to care about the visibility and understandability of their code). This is then only a matter of person’s own initiative and motivation to try to become more visible. Of course, nice exceptions do exist, and Myung and Pitt (2010) recently founded cognitive modeling repository on

the web which provides room for sharing computational models using various frameworks.

Consistently with the above manifesto, the important fact is that most people involved in cognitive science are non modelers. Addyman and French (2012) propose that a presentation of any model should consider three categories of users: (1) casual users – who want to observe the model running essentially as a demo, (2) motivated users – who want to run their own data on the model and/or explore the parameters of the model; (3) modelers – i.e. skilled programmers who want access to the code, in order to potentially modify the structure of the model itself and test it.

Maybe the way to enforce the availability of the modeling software could be to require the well documented code as a part of the paper submission, is the same way as asking the experimenters for supplementary material in experimental papers, which helps the readers to grasp the details of the described method and to potentially replicate the experiment. It is a common practice to repeat experiments, it could also be the case with models. Unlike the experiments with human subjects, the replication of the (deterministic) simulations should be exactly possible, being an advantage, because there should be no hidden factors involved. Time will tell whether this manifesto will evoke sufficient response.

10. Conclusion

I have advocated the view that computational approaches in cognitive science are not only important but crucial. More modelers are welcome to the field to balance the participation of various approaches to study the mind and cognition. I have dealt with the concept of computation that should be understood in wider sense, relaxed of discrete computations, operating on symbols. I reviewed four computational frameworks in cognitive science, and have presented the view on connectionism as the most promising approach in the field. Cognitive developmental robotics will presumably represent a suitable platform for designing, implementing and testing intelligent physical artifacts, with increasingly higher levels of autonomy.

References

- Addyman, C., & French, R. 2012. Computational Modeling in cognitive Science: A Manifesto for Change. *Topics in Cognitive Science* 4, 332-341.
- Asada, M. et al. 2009. Cognitive Developmental Robotics: A Survey. *IEEE Transactions on Autonomous Mental Development* 1(1), 12-34.
- Barsalou, L. 2008. Grounded Cognition. *Annual Reviews of Psychology* 59, 617-645.
- beim Graben, P. 2004. Incompatible Implementations of Physical Symbol Systems. *Mind and Matter* 2(2), 29-51.
- Bell, A. 1999. Levels and Loops: The Future of Artificial Intelligence and Neuroscience. *Philosophical Transactions of the Royal Society London B* 354, 2013-2020.
- Bengio, Y. 2009. Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning* 2(1), 11-27.
- Box, G.E.P., Norman R.D. 1987. *Empirical Model Building and Response Surfaces*. John Wiley & Sons.
- Chalmers, D.J. 2011. A Computational Foundation for the Study of Cognition. *Journal of Cognitive Science* 12(4), 1-21.
- Clark, A. 2001. *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford University Press.
- Christiansen, M., Chater, N. 1999. Toward a Connectionist Model of Recursion in Human Linguistic Performance. *Cognitive Science* 23, 157-205.
- Christiansen, M., Chater, N. (Eds.) 2001. *Connectionist Psycholinguistics*. Ablex Publishing.
- Cooper, S.B. 2012. Turing's Titanic Machine? *Communications of the ACM* 55(3), 74-83.
- Craenen, B., Eiben, A. 2003. Computational intelligence. In: *Encyclopedia of Life Support Sciences*. EOLSS Publishers Co.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Churchland, P., Sejnowski, T. 1992. *The Computational Brain*. Cambridge, MA: MIT Press.
- Cottrell, G. 2012. Cognitive modeling greatest hits. <http://cseweb.ucsd.edu/users/gary/CogSciLiterature.html>
- Davies, J. 2006. *Semantic Web Technologies: Trends and Research in Ontology Based Systems*. John Wiley & Sons.
- Doya, K. 1999. What are the Computations of the Cerebellum, the Basal Ganglia, and the Cerebral Cortex? *Neural Networks* 12, 961-974.

- Duch, W. 2007. Towards comprehensive foundations of computational intelligence. *Lecture Notes in Computer Science*, Springer 63, 261-316.
- Elman, J. 1990. Finding Structure in Time. *Cognitive Science* 14, 179-211.
- Engelbrecht, A. 2007. *Computational Intelligence: An Introduction* (2nd ed.), John Wiley & Sons.
- Farkaš, I., Crocker, M. 2008. Syntactic Systematicity in Sentence Processing with a Recurrent Selforganizing Network. *Neurocomputing* 71, 1172-1179.
- Feldman, J. 2012. The Neural Binding Problem(s). *Cognitive Neurodynamics*. doi: 10.1007/s1157101292198.
- Fodor, J., Pylyshyn, Z. 1988. Connectionism and cognitive Architecture: A Critical Analysis. *Cognition* 28, 3-71.
- Fodor, J. 2000. *The Mind Doesn't Work That Way*. MIT Press.
- Frank, S., Haselager, W., Van Rooij, I. 2009. Connectionist Semantic Systematicity. *Cognition* 110, 358-379.
- French, R. 1999. Catastrophic Forgetting in Connectionist Networks. *Trends in Cognitive Science* 3(4), 128-135.
- French, R. 2012. Dusting off the Turing Test. *Science* 336(6078), 164-165.
- Gardner, H. 1987. *The Minds New Science: A History of the Cognitive Revolution*. New York: Basic Books.
- Gärdenfors, P. 1996. Conceptual Spaces as a Basis for Cognitive Semantics. In: *Philosophy and Cognitive Science*, A. Clark et al. (Eds.), Kluwer, Dordrecht, 159-180.
- Gallese, V., Gernsbacher, M., Heyes, C., Hickok, G., Iacoboni, M. 2011. Mirror Neuron Forum. *Perspectives on Psychological Science* 6(4), 369-407.
- Gentner, D. 2010. Psychology in Cognitive Science: 1978-2038. *Topics in Cognitive Science* 2, 328-344.
- Greco, A. 1998. Connectionist Models can Reveal Good Analogies. *Psychology* 9(47) Connectionist Explanation (31).
- Greco, A., Caneva, C. 2010. Compositional Symbol Grounding for Motor Patterns. *Frontiers in Neurobotics* 4:111, doi: 10.3389/fnbot.2010.00111.
- Griffiths, T., Chater, N., Kemp, C., Perfors, A., Tenenbaum, J. 2010. Probabilistic Models of Cognition: Exploring Representations and Inductive Biases. *Topics in Cognitive Science* 14, 357-364.
- Harnad, S. 1990. The Symbol Grounding Problem. *Physica D* 42, 335-346.
- Hornik, K., Stinchcombe, M., and White, H. 1989. Multilayer Feedforward Networks are Universal Approximators. *Neural Networks* 2(5), 359-366.
- Kaczmarek, L., Levitan, I. (Eds.) 1987. *Neuromodulation: The Biochemical Control of Neuronal Excitability*. NY: Oxford University Press.
- Ma, W.J., Beck J., Latham P. and Pouget A. 2006. Bayesian Inference with

- Probabilistic Population Codes. *Nature Neuroscience* 9(11), 1432-1438.
- Maass, W. & Bishop, C. (Eds.). 1999. *Pulsed Neural Networks*. MIT Press.
- MacKay, D.J.C. 1995. Probable Networks and Plausible Predictions – A Review of Practical Bayesian Methods for Supervised Neural Networks. *Network: Computation in Neural Systems* 6, 469-505.
- Marr, D. 1982. *Vision*. San Francisco, CA: W.H. Freeman.
- Maturana, H. & Varela, F. 1992. *The Tree of Knowledge: The Biological Roots of Human Understanding* (Rev. ed.). Boston: Shambhala.
- McClelland, J. 2009. The Place of Modeling in Cognitive Science. *Topics in Cognitive Science* 1(1), 11-38.
- Myung, J.I., & Pitt, M.A. 2010. *Cognitive Modeling Repository*. Columbus, OH: Department of Psychology, Ohio State University (Distributor). Retrieved from <http://cmr.osu.edu/>
- Newell, A., Simon, H. 1976. Computer Science as Empirical Inquiry: Symbols and Search. *Communications of the ACM* 19(3), 113-126.
- O'Brien, G. & Opie, J. 2006. How do Connectionist Networks Compute? *Cognitive Processing* 7, 30-41.
- O'Reilly, R. & Munakata, Y. 2000. *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge: MIT Press.
- Page, M. 2000. Connectionist Modelling in Psychology: A Localist Manifesto. *Behavioral and Brain Sciences* 23(4), 443-467.
- Perfors, A., Tenenbaum, J.B., Griffiths, T.L., Xu, F. 2011. A Tutorial Introduction to Bayesian Models of Cognitive Development. *Cognition* 120, 302-321.
- Pfeifer, R., Scheier, C. 1999. *Understanding Intelligence*. MIT Press, Cambridge, MA.
- Piccinini, G. 2007. Computing Mechanisms. *Philosophy of Science* 74, 501-526.
- Piccinini, G. 2008. Some Neural Networks Compute, Others don't. *Neural Networks* 21(23), 311-321.
- Piccinini, G. & Scarantino, A. 2011. Information Processing, Computation, and Cognition. *Journal of Biological Physics* 37, 1-38.
- Pinker, S. 2009. *How the Mind Works*. W.W. Norton and Co.
- Pinker, S., Ullman M. 2002. The Past Tense Debate. *Trends in Cognitive Sciences* 6(11), 456-474.
- Port, R., van Gelder, T. (Eds.) 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.
- Prinz, W. 1984. Modes of Linkage between Perception and Action. In W. Prinz & A.F. Sanders (Eds.), *Cognition and Motor Processes* (pp. 185-193). Berlin: Springer.
- Pylyshyn, Z. 1984. *Computation and Cognition: Toward a Foundation for*

- Cognitive Science*. Cambridge: MIT Press.
- Rieke, F. et al. 1999. *Spikes: Exploring the Neural Code*. MIT Press.
- Ritchie, J. 2011. Chalmers on Implementation and Computational Sufficiency. *Journal of Cognitive Science* 12, 401-417.
- Rosenblatt F. 1962. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Washington: Spartan Books.
- Russell, S. and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach* (3rd ed.) Prentice Hall.
- Scheutz, M. 1999. When Physical Systems Realize Functions *Minds and Machines* 9, 161-196.
- Schoener, G. 2008. Dynamical Systems Approaches to Cognition. In R. Sun (Ed.), *Cambridge Handbook of Computational Psychology* (pp. 101-126). New York: Cambridge University Press.
- Siegelmann, H., Sontag, E. 1991. Turing Computability with Neural Nets. *Applied Mathematics Letters* 4(6), 77-80.
- Siegelmann, H. 1999. *Neural Networks and Analog Computation: Beyond the Turing Limit*. Birkhauser, Boston.
- Slovan, A. 2002. The Irrelevance of Turing Machines to Artificial Intelligence. In M. Scheutz (Ed.), *Computationalism: New Directions* (pp. 87-128). Cambridge: MIT Press.
- Spivey, M. 2007. *The Continuity of Mind*. Oxford: Oxford University Press.
- Sporns, O. 2011. *Networks of the Brain*. MIT Press.
- Sprevak, M. (Ed.) 2010. Computation and Cognitive Science. *Studies in History and Philosophy of Science* 41, 223-334.
- Tabor, W. 2009. A dynamical Systems Perspective on the Relationship between Symbolic and Nonsymbolic Computation. *Cognitive Neurodynamics* 3(4), 415-427.
- Takáč, M., Beňušková, L., Knott, A. 2012. Mapping Sensorimotor Sequences to Word Sequences: A Connectionist Model of Language Acquisition and Sentence Generation. *Cognition* 125(2), 288-308.
- Thorisson, K.R., Helgasson, H.P. 2012. Cognitive Architectures and Autonomy: A Comparative Review. *Journal of Artificial General Intelligence* 3(2), 1-30.
- Tschacher, W., Dauwalder, J.P. (Eds.) 2003. *The Dynamical Systems Approach To Cognition: Concepts and Empirical Paradigms Based on Self-Organization, Embodiment, and Coordination Dynamics*. World Scientific.
- Turing, A. 1936. On computable Numbers, with an Application to the Entscheidungs Problem. *Proceedings of the London Mathematical Society* 2(42), 230-65.
- Turing, A. 1950. Computing Machinery and Intelligence. *Mind* 59, 433-460.
- Van Gelder, T. 1990. Compositionality: A Connectionist Variation on a Classical

Theme. *Cognitive Science* 14, 355-384.

Varela, F., Thompson, E., & Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Vernon, D., Metta, G., Sandini, G. 2007. A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents. *IEEE Transactions on Evolutionary Computation* 11(2), 151-180.

Weng, J. 2012. Symbolic Models and Emergent Models. *IEEE Transactions on Mental Development* 4(1), 29-54.

Manuscript received: Sep 28, 2012, in revised form: Oct 29, 2012.

