Searching for Emergent Representations in Evolved Dynamical Systems

Thomas Hope, Ivilin Stoianov, and Marco Zorzi

Computational Cognitive Neuroscience Lab, University of Padova,
Via Venezia 12/2, Padova 35131, Italy
{thomas.hope, ivilin.stoianov, marco.zorzi}@unipd.it

Abstract. This paper reports an experiment in which artificial foraging agents with dynamic, recurrent neural network architectures, are "evolved" within a simulated ecosystem. The resultant agents can compare different food values to "go for more," and display similar comparison performance to that found in biological subjects. We propose and apply a novel methodology for analysing these networks, seeking to recover their quantity representations within an Approximationist framework. We focus on Localist representation, seeking to interpret single units as conveying representative information through their average activities. One unit is identified that passes our "representation test", representing quantity by inverse accumulation.

1 Introduction

In 1963, Feyerabend ([1] and [2]) claimed that improvements in our scientific understanding of the mind will eventually undermine our basic concepts of mental states. His position – Eliminative Materialism – stemmed from the intuition that “folk psychology” [3] is merely our current best “theory” of mind. Like any other theory, folk psychology may eventually be falsified, perhaps in favour of a neuro-biological account of cognition. Significantly, there is no requirement that this replacement must “explain” the theory that it replaces – like phlogiston and alchemy, mental states may simply disappear [4] in the face of scientific progress. In cognitive science, a debate has recently emerged that adds a practical dimension to Feyerabend’s position.

Since the emergence of the digital computer during the 1940’s, the Computer Metaphor (CM – [5]) has dominated the way in which scientists study intelligent behaviour. One of the principal methodological commitments of the CM is Functional Decomposition (FD), which implies that, like computers, cognitive systems can be understood as networks of functional “modules” [6]. The acceptance of this intuition is nearly ubiquitous in contemporary cognitive science; most experimental paradigms are designed explicitly to isolate and manipulate these putative modules. Seeking to account for experimental data, computational cognitive models have tended to be directed along similar lines. This division naturally emphasises the concept of “representation” in contemporary cognitive theories, since the specification of a module’s interfaces (input and output representations) has a critical impact on its empirical behaviour (e.g. [7]).

S. Nolfi et al. (Eds.): SAB 2006, LNAI 4095, pp. 522 – 533, 2006.

Yet for all its evident utility, FD carries a heavy burden of explanation; functional modules must be *integrated* before they can reasonably be said to account for cognitive behaviour. Recognising the critical role that module interfaces play in designing these modules, some researchers (e.g. [8], [9]) have questioned whether this integration could ever be successful.

Dynamicism [10] offers an alternative. The Dynamicist programme construes cognition as intrinsically embodied, emerging from the interaction of adaptive behaviours. These behaviours are the Dynamicist equivalent of functional modules – the atomic components of intelligence [8]. To the extent that this framework addresses the "integration problem" mentioned previously, it must surely be welcomed even by researchers entrenched within the CM. The problem is that, at least as commonly construed [11], the Dynamicist programme is resolutely Eliminativist.

Dynamicist models are thought to be best understood in terms of the temporally situated causal processes that manage sensor-motor integration [11]. Traditionally critical concepts, like representation, simply do not appear to "fit" with how these systems work. Recognising this, many Dynamicist researchers have been moved to claim that the CM-inspired distinction between data and process – the very concept of representation itself – must go the way of phlogiston [8]-[11].

This proposition is antithetical to many neuroscientists because it simply does not appear to correspond with the observed structure of cognition in the brain. Selective neural disorders and brain imaging experiments [12] provide convincing evidence of functional specialisation in human cognition; though consistent with Dynamicism, a functionally specialised cognitive system naturally encourages analysis by FD. Concepts of representation are a powerful and intuitive tool in accounting for well-confirmed experimental data (e.g. [13], [14]); in some cases, such as [15], these accounts can even draw on "observations" of the representations themselves. In some respects at least, brains just do appear to represent and to compute. If Dynamicism must be accepted at the expense of the concept of representation, few cognitive neuroscientists will accept it.

To manage this tension, we propose an Approximationist response. Approximationism articulates the intuition that computational accounts of cognitive processes may be useful and approximately correct, without necessarily capturing every detail of the underlying causal processes [16]. One implication is that we might usefully search for – and discover – representations in neural systems, while at the same time accepting that the implied "computational story" will be at best a good approximation to the underlying "causal story".

Following the logic of [17], we evolve artificial agents to perform a "representation-heavy" task – a task for which some kind of representational structure appears to be required. Section 2 describes the artificial ecosystem and agents, as well as the representation-heavy task that they evolve to perform. The goal is then to recover the agents' evolved representations. Section 3 describes and applies a methodology designed to achieve this goal.

2 Through Foraging to Quantity Comparison

The focus for the current project is "quantity comparison", a common theme of study within the cognitive neuroscience of numeracy. Representation plays a critical role in contemporary accounts of the way in which subjects (humans and animals) manipulate numbers and numerosities [13] – indeed, there is great debate in this field concerning the precise format of that representation (e.g. [13], [18], [19]). Quantity comparison therefore meets our requirement for a representation-heavy process. Further, a growing body of evidence indicates that certain facets of the "number sense" may be inherited [20]. The implication is that evolution might engender a preparedness in humans and animals to represent quantity in a particular way, raising an independent question about what kinds of quantity representations can emerge "spontaneously" during evolution. The current work approaches quantity comparison as an evolutionary by-product of selection for quantity-sensitive foragers.

2.1 The Artificial Ecosystem

The environment is a 2-dimensional, toroidal grid, composed of 100x100 square cells. Agents navigate the grid by moving between neighbouring cells. Each cell can contain "food", construed as appearing in "bundles" of some specified numerosity (1-9). Any number of agents can co-exist in the same cell: the only upper limit is the size of the population itself (200), which remains constant throughout the run. Food can also "grow", in the sense that its numerosity can increase. A record is kept of the total depletion of food during the run, and this food is periodically reinserted by sharing it among randomly selected cells.

The ecosystem proceeds by iterative update. During each iteration, every individual is updated, with sensor activity propagated through the neural network and effector units interpreted to identify if any action has been made. The update order for agents is randomly specified at the beginning of each iteration.

The agents are recurrent, asymmetrically connected neural networks. In a network of N units, the activity u of the i-th unit (u_i) at time step t is calculated by

$$u_i(t) = S(\sum_{j=1}^{N} w_{ij} u_j(t-1)).(1-m) + u_i(t-1).m \tag{1}$$

where w_{ij} is the weight of the connection from unit j to unit i, $S()$ is the sigmoid function and m is a fixed momentum term with a value of 0.5.

A subset of units act as sensors, which are clamped according to the salient features of the environment around the agent. Agents have a 3 cell field of view, and are also sensitive to food in the cell that they currently occupy (see Fig. 1), for a total of four sensor "fields". Each field represents its corresponding food quantity using a "Random Position Code"; this was used in [18], among others, to capture quantity information without biasing models in favour of specific representational strategies. To represent food numerosity N, the code requires that N (randomly chosen) sensor units (positions) should be active. The scheme is illustrated in Fig. 2, where N = 5.

Another subset of units are effectors, whose activity determines how the position / orientation of the agent's body is updated, as well as defining when agents try to eat. The remaining units are hidden and do not interact directly with the environment. Sensor units receive no input from the rest of the network, and have no direct connection to effector units, but the hidden layer is universally connected – every unit is connected to every other unit, and to itself.

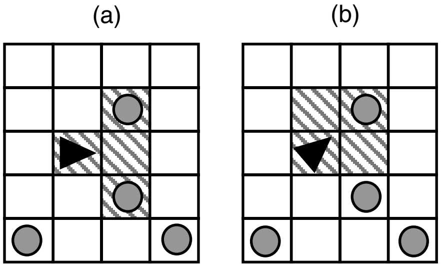

Fig. 1. An agent in its environment. (A) The agent – a black triangle – is facing right and can sense food (grey circles) in its right and left-most sensor fields. (B) The same agent, after making a single turn to the left. It can now sense only one cell containing food.

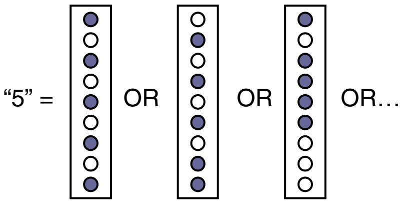

Fig. 2. The Random Position Code. For food = N, exactly N units are active (activity clamped to ‘1’), while the remainder are inactive (activity clamped to ‘0’). Each sensor field has a total of 9 units.

2.2 Evolution in the Ecosystem

All of the agents in the initial population are specified with random numbers of hidden units, and random weights. The "fitness" (F) of the i-th agent is just the rate at which it has consumed food,

$$F(i) = f_i / a_i \tag{2}$$

where f is the total food consumed and a is the agent’s age (expressed as the number of iterations since the agent's creation).

At the end of each iteration, two “parents” are drawn at random from the population and their fitnesses compared. The structure of the “child” is defined by randomly mixing the parents’ weight vectors (cross-over), followed by mutation. The mutation operator will usually increment or decrement a randomly selected weight value by a small constant (0.01), but may also add or remove a hidden unit. The resultant child replaces the least fit of the two parents.

The best signal that agents are discriminating quantity is high food collection efficiency: food collected per moves made in the environment. The agents’ food collection efficiency rises above 5 after about 10 million iterations (~50,000 generations), indicating that genuinely “discriminating” foraging behaviour has evolved.

2.3 Quantity Comparison Performance

The evolved agents are not merely models of quantity comparison; to capture that facet of their behaviour, we need a methodology that can effectively isolate it. Fortunately, examples of the required kind of methodology already exist. In [21], the “subjects” (salamanders) were placed at the base of a clear Perspex T-maze. Two clear jars of *drosophilia* fruit flies, which the salamanders eat, were placed at the end of each branch of the maze – two flies in one jar and three flies in the other. The authors reported that twice as many salamanders “chose” the jar with more flies (signified by walking toward and touching that jar). Our methodology emulates this experiment.

The experimental environment is a 3x3 “mini-world”. Two cells, the top left and top right of the world, contain food of varying quantity. In its initial position at the centre of the world, the agent can “see” both of these food quantities, though it may turn without constraint once the trial begins. Food “selection” occurs when the agent moves onto one or other of the filled cells – the only cells onto which the agent is allowed to move. A correct choice is defined as the selection of the larger of the two food groups. Every agent in the population was tested using this methodology, with 50 repetitions of every combination of food quantities (1-9, 72 combinations in all), for a total of 3,600 trials per agent. The results are displayed in fig. 3.

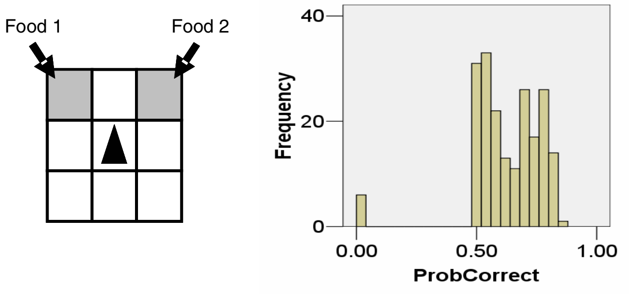

Fig. 3. (Left) The schematic structure of the comparison experiment. The agent is placed in the centre of the mini-world, facing “up”. (Right) A histogram of the population performance in the quantity comparison experiment.

A few of the agents perform extremely badly, indicating that the evolved foraging solutions are brittle in the face of “evolutionary” change – perhaps emphasised as a consequence of a mutation bias against specialised structures [22]. The main bulk of the population distribution is also apparently bimodal; agents in the left-most cluster perform at roughly chance levels, whereas agents in the right-most cluster perform significantly above chance – only this latter group appear to discriminate quantity. The persistence of non-discriminating agents is unsurprising, since high rates of food collection can be achieved by sacrificing decision quality for decision speed. A visual inspection of the performance scores for agents in this cluster indicates strong asymmetry in their behaviour; many simply “choose” the right-hand square regardless of the food quantities presented.

2.4 Single-Subject Comparison Behaviour

Using the results of the previous section, we selected the best “discriminator” from the population and recorded its empirical performance in more detail. The results are displayed in fig. 4.

The agent’s empirical behaviour displays certain characteristic phenomena that are also reliably found when both humans and animals compare quantities. As the minimum of the two quantities-to-be-compared increases (fig. 4a), there is an increase in discrimination error, this is an instance of the “Size Effect” [13]. As the numerical

Fig. 4. Accuracy scores are rates of correct choices per 1,000 trials. (a) Mean accuracy vs. minimum quantity of food (Min) in a given trial. (b) Mean accuracy vs. numerical distance (Split) between food quantities. (c) Mean "reaction time" vs. numerical distance between quantities.

distance between the quantities increases (fig. 4b), there is a corresponding improvement in discrimination accuracy; this is an example of the "Distance Effect" [13].

Surprisingly, this agent also displays a Distance Effect for *reaction times*, defined as the number of time steps between the start of a comparison trial and the agent's selection of one of the two food quantities. Non-discriminating foragers can persist by sacrificing decision accuracy for decision speed, but this agent is capable of reversing that tradeoff, sacrificing decision speed in order to more reliably "go for more".

Though the agents are too simple to support a meaningful comparison with biological organisms, this behavioural correspondence is nevertheless encouraging. As mentioned prevously, most contemporary theories of quantity manipulation account for empirical phenomena (such as the Size and Distance effects) as a consequence of the way in which subjects represent quantities. A representational account of this agent's behaviour could therefore add a new dimension to this traditional debate, expanding the space of representational strategies that can account for the empirical phenomena.

3 Approximationist Representation

The key property of a representation is that it tracks some property of the environment. During the comparison experiment, the most salient properties are the values of the two food "options", which remain static throughout each trial. But a visual inspection of the agent's network dynamics reveals nothing remotely static – nothing that seems appropriate to represent these food values.

Our thesis is that although the behaviour of each unit is subject to chaotic variation, its *average* activity may still be interpreted as conveying representative information.

If a unit's average activity (relative to some food value) is a *functionally significant* representation, it should be possible to "fool" the agent by interfering with that activity by fixing the unit's activity to its average for a *different* food value. In deference to previous work (e.g. [23] and [24]) on lesion types in network analysis, we refer to this kind of interference as "Partial Informational Lesion" (PIL). Our Approximationist method relies on statistical analyses of the impact of PIL's on the agent's behaviour.

In the material that follows, we restrict the analysis to Localist "theories" of representation – to theories that *single* units can be interpreted as conveying functionally significant, representative information. This is a simplifying assumption, but not without biological justification; recent single-cell recording experiments [25] suggest that single neurons in the primate parietal cortex may be selectively tuned to quantity.

3.1 Defining "Representation-Like" Deviation from Normal Behaviour

Consider a simple example in which a hypothetical agent manages a quantity comparison using just two hidden units – unit A and unit B. Each unit is subject to extensive noise, but represents a corresponding quantity (Food A and Food B) by its average activity. As in the current work, our hypothetical system uses food values 1-9, signified by average unit activities of 0.1, 0.2…., 0.8, and 0.9. Suppose that we interfere with unit A so that its activity is *always* 0.5 – now, the agent will always perceive food A as taking the value '5'.

In some circumstances, this discrepancy between "actual" and "perceived" value of food A should reduce the agent's comparison performance. If the actual value of food A is '2', and the value of food B is '4', our agent will "think" (wrongly) that food A is larger and could make the wrong decision. There are also circumstances in which this intervention should improve the agent's performance. Suppose now that the true value of food A is in fact '2' and food B is '1'; the perceived and actual comparisons between the two quantities both have the same "answer" (i.e. food A is larger), but the perceived comparison is arguably easier because the numerical distance between '5' and '1' is greater than that between '2' and '1'.

This logic leads us to define two groups of comparison trial, relative to particular PIL's; *Consistent* trials are those for which PIL's should improve comparison performance, whereas *Inconsistent* trials are those for which PIL's should reduce comparison performance. The hypothesis that some unit's average activity does in fact "represent" can then be judged by reference to two "Representation Scores"; one for Consistent trials, and one for Inconsistent trials, calculated as in equation 3.

$$
\begin{aligned}
R_j^{+} &= \sum_{i=1}^{n} P_j(C,L) - \sum_{i=1}^{n} P_j(C,N) \\
R_j^{-} &= \sum_{i=1}^{n} P_j(I,L) - \sum_{i=1}^{n} P_j(I,N)
\end{aligned}
\tag{3}
$$

where R_j is the representation score for the j-th unit, and $P_j()$ is a function that counts the number of correct comparison choices that the agent makes under the four possible conditions. The four conditions are:

1) (C, L): Consistent comparisons made while the unit was subject to a PIL (i.e. is Lesioned).
2) (C, N): Consitent comparisons, but where the unit is allowed to change freely (i.e. is Normal).
3) (I, L): Inconsistent comparisons made while the unit was subject to a PIL (i.e. is Lesioned).
4) (I, N): Inconsistent comparisons, but where the unit is allowed to change freely (i.e. is Normal).

A "positive" result is observed when R^+ (the "Consistent Score") is significantly positive, and R^- (the "Inconsistent Score") is significantly negative.

3.2 Identifying "Representation-Like" Deviation from Normal Behaviour

The current results are derived from the hypothesis that *single* units may represent food quantities. The average activities of these units – putatively localist representations – were collected during 100 repetitions of the comparison experiment described in section 2. This process associates each unit with one average activity for each food value (1-9) in each position (Food 1 or Food 2) – 18 values in total.

Fig. 5 displays scatter plots for the 2-dimensional representation scores of each of the agent's 25 hidden units. Each unit has a data point in both graphs. Each component of a unit's representation score is a mean average value. Scores that pass our "Filter test" will lie in the top-left quadrant of each graph; when applied to the units that correspond to these data points, PIL's improve the agent's performance during Consistent comparison trials, and reduce that performance during Inconsistent trials.

T-tests for paired samples (Lesioned vs. Normal in both Consistent and Inconsistent comparison conditions, N = 90 in both cases) confirm that the marked data points represent significant deviation from normal performance after the lesion ($p < 0.05$) in both Consistent and Inconsistent conditions.

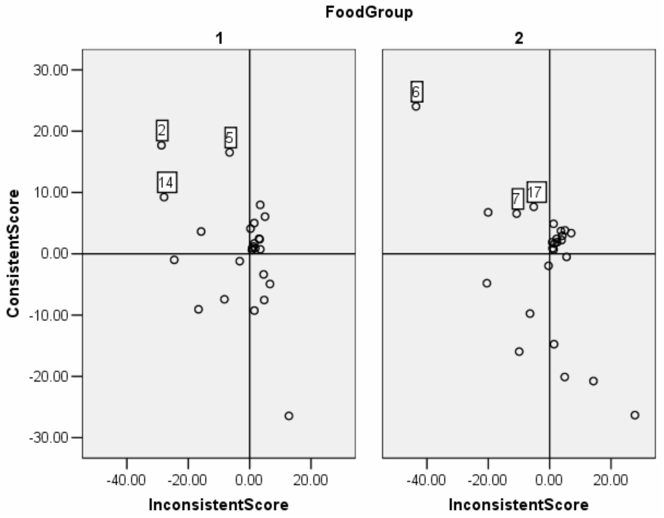

Fig. 5. Representation scores for the agent's 25 hidden units, relative to the hypotheses that each unit represents either food group 1 (left) or food group 2 (right). Scores in the top-left quadrant of the graph indicate a positive result. Data labels denote the unit numbers of associated data points.

Both of the food groups support multiple theories of unit-centric representation; relative to each food group there are three units that, when lesioned with PIL's, engender the behavioural deviations that we hoped to discover. A visual inspection of the associated representation scores is not sufficient to adjudicate between them. To make that step, we need to extend the method.

3.3 Comparing Candidate Theories

Returning once more to our hypothetical agent, consider what happends when we fix the value of unit A to '0.1'. In this case, the agent will "think" that Food A takes the value '1', regardless of the true value of Food A, and its comparison performance should reflect that perception. This logic is the foundation for our extension.

We can collect a pair of performance scores based on this hypothesis; in the "Lesioned" condition, the PIL is applied and Food A takes some value other than '1', whereas in the "Unlesioned" condition, no PIL is applied and Food A is always equal to '1'. After collecting analogous pairs for every other food value, relative to each of the two food groups, we have two sets of paired series of performance scores. To the extent that the PIL's have captured the agent's representational strategy, there should be a significant relationship between these paired series; we can capture that relationship using linear regression. If the relationship is significant, its "variance explained" (R^2) provides the metric that we need to compare competing theories.

Table 1 displays the results when this Comparison test is applied to the data-points highlighted in Fig. 5. The values are derived from series generated by 10 repetitions of each experiment; regressed series are 90 elements long.

Table 1. Linear regression results for each of the unit-centric theories identified by the Filter test. The shaded column corresponds to a theory that passes the Comparison test.

	Food 1			Food 2		
Unit	2	5	14	6	7	17
Significance	0.127	< 0.001	0.176	0.943	0.544	0.062
R^2	0.026	0.219	0.021	< 0.001	0.004	0.039
Beta	0.162	0.476	0.144	0.008	-0.065	-0.198

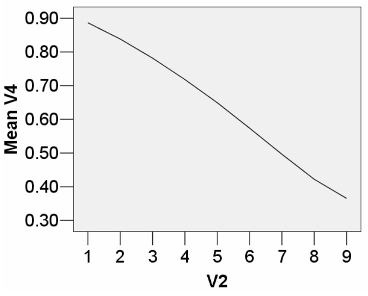

Fig. 6. The average activity of the agent's hidden unit '5', relative to the value of food group 1

Only one of the units passes this Comparison test – unit 5 does justify an interpretation of representing food group 1 by its average activity. The quality of that justification depends on the R^2 and Beta values associated with the regression; the quality increases as these variables approach the value ‘1’. This unit's activity (Fig. 6) is inversely proportional to the value of “Food 1”, a pattern consistent with the “Accumulator” theory [26] of neural quantity representation.

4 Discussion

Conventional approaches to cognitive modeling tend to focus on the isolated, functional components of cognitive behaviour; methods that could facilitate more behaviourally integrated models have failed to attract great support among cognitive neuroscientists. One reason for this is the sense that the gulf between human / animal and artificial agent behaviour is too wide to permit useful comparisons. Another reason is that these techniques can challenge our fundamental notions of representation, which remains an important conceptual tool for understanding cognitive systems.

The current work lays the foundations of a methodological framework designed to address these apparent inconsistencies. We have evolved agents whose behavioural performance is reminiscent of that found in biological organisms, and offered an analytical framework that permits the recovery of classically “recognisable” representations from those agents. Critically, our method quantifies the extent to which a representational account of the agent’s behaviour can be justified – the extent to which the theory captures the underlying causal process.

We chose to base the analysis on the thesis the average unit activities can be interpreted as conveying functionally significant representative information. Though potentially controversial, this thesis is at least consistent with the practice of single-cell recording experiments (e.g. [25]), which emphasise average neural behaviour at the expense of the apparently random [27] variation in specific spike trains. The real justification for the choice flows from the results of the analysis itself; despite its restricted (Localist) scope, we have identified a unit that appears to represent quantity by its average activity. We expect that the strength of this result can be improved by relaxing the Localist restriction, and work to implement that extension is currently underway.

Despite its limitations, the current version of this system yields an interesting implication concerning the symmetry assumption in conventional cognitive modeling. Contemporary models of quantity comparison are invariably “functionally symmetrical” in that both of the quantities-to-be-compared are treated in the same way; the current result exposes that assumption to unfavourable scrutiny. As mentioned in section 2.3, the evolved population reliably contains “non-discriminating” foragers. These agents display strongly asymmetrical behaviour, selecting food group 2 (initially on the agent’s right) regardless of the food values. This strategy emerges rather earlier in evolutionary runs than does the more “discriminating” variant; the implication is that quantity comparison processes emerge within a behaviourally asymmetrical context.

Given the context of behavioural asymmetry, it seems natural to predict that an agent's representational strategy will also be asymmetrical – though preliminary, our results do support this prediction because there is no equivalent of unit ‘5’ for food group 2. In other words, though we cannot say with confidence how the agent “represents”, we can predict that its distributed representations will not be symmetrical.

References

1. Feyerabend, P.: Materialism and the mind-body problem. Review of Mataphysics, Vol. 17 (1963) 49-66
2. Feyerabend, P.: Mental events and the brain. Journal of Philosophy, Vol. 60 (1963) 295-296
3. Stich, S., Ravenscroft, I.: What is Folk Psychology? Cognition, Vol. 50 (1994) 447-68
4. Rorty, R.: Mind-Body Identity, Privacy, and Categories. Review of Metaphysics, Vol. 19 (1965) 24-54
5. Crowther-Heyk, H.: George Miller, language, and the computer metaphor. History of Psychology, Vol. 2 (1999) 37-64
6. Bechtel, W., Richardson, R. C. Discovering complexity: Decomposition and localization as strategies in scientific research. Chapter 2, Princeton: Princeton University Press (1993)
7. Stoianov, I., Zorzi, M., Umiltà, C.: The role of semantic and symbolic representations in arithmetic processing: Insights from simulated discalculia in a connectionist model. Cortex, Vol. 40 (2004) 194-196
8. Brooks, R. A.: Intelligence without representation. Artificial Intelligence, Vol. 47 (1991) 139-159
9. Beer, R. D.: Dynamical approaches to cognitive science. Trends in Cognitive Sciences, Vol. 4(3) (2000) 91-99
10. van Gelder, T. J.: Dynamic approaches to cognition. In Wilson, R., Keil, F. (editors), The MIT Encyclopedia of Cognitive Sciences. Cambridge MA: MIT Press (1999) 244-6
11. Harvey, I.: Untimed and misrepresented: Connectionism and the computer metaphor. Newsletter of the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB Quarterly), Vol.96 (1996) 20-27
12. Cabeza, R., Nyberg, L.: Imaging cognition II: an empirical review of 275 PET and fMRI Studies. Journal of Cognitive Neuroscience, Vol. 12 (2000) 1-47
13. Zorzi, M., Stoianov, I., Umiltà, C.: Computational modeling of numerical cognition. In Campbell, J. editor, Handbook of mathematical cognition, chapter 5 London: Psychology Press (2005) 67-83
14. Dahaene, S., Cohen, L., Sigman, M., Vinckier, F.: The neural code for written words: a proposal. Trends in Cognitive Science, Vol. 9(7) (2005) 335-341
15. Naccache, L., Dahaene, S.: The priming method : imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. Cerebral Cortex, Vol. 11 (2001) 966-974
16. Smolensky, P.: On Variable Binding and the Representation of Symbolic Structures in Connectionist Systems. Technical report CU-CS-355-87, Department of Computer Science, University of Colorado at Boulder, 1987
17. Beer, R. D.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J., Pollack, J., Wilson, S., editors, Proceedings of the 4th International Conference on Simulation of Adaptive Behavior, MIT Press (1996) 421-429
18. Verguts, T., Fias, W.: Representation of number in animals and humans: a neural model. Journal of Cognitive Neuroscience, Vol. 16(9) (2004)
19. Gelman, R., Gallistel, C. R.: Language and the origin of numerical concepts. Science, Vol. 306(5695) (2004) 441-443
20. Butterworth, B.: What Counts, chapter 3, The Free Press: New York (1999) 144-147
21. Uller, C., Jaeger, R., Guidry, G., Martin, C.: Salamanders (Plethodon cinereus) go for more: rudiments of number in an amphibian. Animal Cognition, Vol. 6 (2003) 105-112

22. Watson, R.A., Pollack, J.B.: Coevolutionary Dynamics in a Minimal Substrate. Proceedings of the 2001 Genetic and Evolutionary Computation Conference, Spector, L., et al., editors, Morgan Kaufmann (2001)
23. Keinan, A., Sandbank, B, Hilgetag, C. C., Meilijson, I., Ruppin, E.: Fair attribution of functional contribution in artificial and biological networks. Neural Computation, Vol. 16(9) (2004) 1887-1915
24. Aharanov, R., Segev, L., Meilijson, I., Ruppin, E.: Localization of function via lesion analysis. Neural Computation, Vol. 14(4) (2003) 885-913
25. Nieder, A., Freedman, D. J., Miller, E. K.: Representation of the quantity of visual items in the primate prefrontal corext. Science 6 (2002) 1708-1711
26. Meck, W. H., Church, R. M.: A mode control model of counting and timing processes. Journal of Experimental Psychology: Animal Behaviour Processes 9 (1983) 320-324
27. Tomko G., Crapper, D.: Neuronal variability: non-stationary responses to identical visual stimuli. Brain Res, Vol. 79 (1974) 405-418