

Contextual Learning in the Neurosolver

Andrzej Bieszczad¹, Kasia Bieszczad²

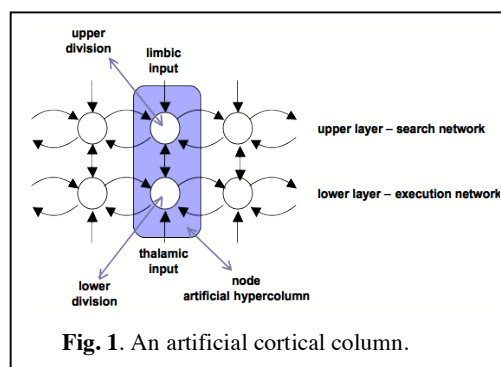
¹Computer Science, California State University Channel Islands
One University Drive, Camarillo CA 93012
(aj.bieszczad@csuci.edu)

²Center for the Neurobiology of Learning and Memory, U.C. Irvine
320 Qureshey Research Laboratory, University of California, Irvine, CA 92697
(kbies@uci.edu)

Abstract. In this paper, we introduce an enhancement to the Neurosolver, a neuromorphic planner and a problem solving system. The enhanced architecture enables contextual learning. The Neurosolver was designed and tested on several problem solving and planning tasks such as re-arranging blocks and controlling a software-simulated artificial rat running in a maze. In these tasks, the Neurosolver learned temporal patterns independent of the context. However in the real world no skill is acquired in vacuum; Contextual cues are a part of every situation, and the brain can incorporate such stimuli as evidenced through experiments with live rats. Rats use cues from the environment to navigate inside mazes. The enhanced architecture of the Neurosolver accommodates similar learning.

1 Introduction

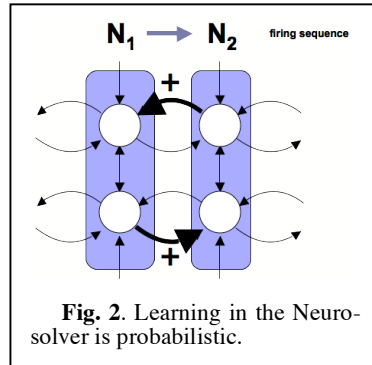
The goal of the research that led to the original introduction of Neurosolver, as reported in [1], was to design a neuromorphic device that would be able to tackle problems in the framework of the state space paradigm [2]. The research was inspired by Burnod's monograph on the workings of the human brain [3]. The class of systems that employ state spaces to present and solve problems has its roots in the early stages of AI research that derived many ideas from the studies of human information processing; e.g., on General Problem Solver [2]. This pioneering work led to very interesting problem solving (e.g. SOAR [4]) and planning systems (e.g. STRIPS [5]).



1.1 Architecture of the Neurosolver

The Neurosolver is a network of interconnected nodes. Each node is associated with a state in a problem space. A problem is presented to the Neurosolver by two signals: a goal associated with the desired state and a sensory signal associated with the current state. The goal is delivered through the limbic input to the upper division of the node, and the percept, through the thalamic input to the lower division of the node. A sequence of firing nodes represents a trajectory in the state space. A solution to a problem is a succession of firing nodes starting with the current node and ending with the goal node.

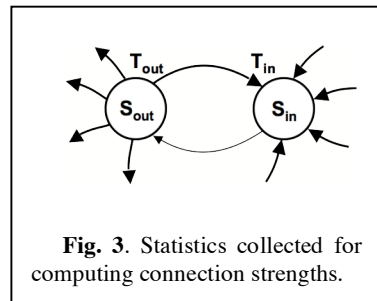
The node used in the Neurosolver is based on a biological cortical column (references to the relevant neurobiological literature can be found in [1]). It consists of two divisions: the upper and the lower, as illustrated in Fig. 1. The upper division is a unit integrating internal signals from other upper divisions and from the control center providing the limbic input (i.e., a goal or - using more psychological terms - a drive or desire). The activity of the upper division is transmitted to the lower division where it is subsequently integrated with signals from other lower divisions and the thalamic input. The upper divisions constitute a network of units that propagate search activity from the goal, while the lower divisions constitute a network of threshold units that integrate search and sensory signals, and generate sequences of firing nodes. The output of the lower division is the output of the whole node. An inhibition mechanism prevents cycles and similar chaotic behavior. Simply, a node stays desensitized for a certain time after firing.



1.2 Learning in the Neurosolver

The Neurosolver learns by receiving instructional input. Teaching samples representing state transitions are translated into sequences of firing nodes corresponding to subsequent states in the samples. For each state transition, two connections are strengthened:

one, in the direction of the transition, between the lower divisions of the two nodes, and another, in the opposite direction, between the upper divisions (Fig. 2). The strength of all inter-nodal connections is computed as a function of two probabilities: the probability that a firing source node will generate an action potential in this particular connection and the probability that the target node will fire upon receiving an action potential from the connection.



To compute the probabilities, each division and each connection collects statistics (Fig. 3). The number of transmissions of an action potential T_{out} is recorded for each connection. The total number of cases when a division positively

influenced other nodes S_{out} is collected for each division. A positive influence means that an action potential sent from a division of a firing node to another node caused that node to fire in the next cycle. In addition, we also collect statistical data that relate to incoming signals. T_{in} is the number of times when an action potential transmitted over the connection contributed to the firing of the target node and is collected for each connection. S_{in} , collected for each division, is the total number of times when any node positively influenced the node. With such statistical data, we can calculate the probability that an incoming action potential will indeed cause the target node to fire. The final formula that is used for computing the strength of a connection (shown in Eq. 1) is the likelihood that a firing source node will induce an action potential in the outgoing connection, multiplied by the likelihood that the target node will fire due to an incoming signal from the connection:

$$P = P_{out} \cdot P_{in} = (T_{out}/S_{out}) \cdot (T_{in}/S_{in}) \quad (1)$$

1.3 Learning sequences

As we already mentioned, the function of the network of upper divisions is to spread the search activity along upper-to-upper connections (Fig. 4) starting at the original source of activity, the node associated with the goal state that receives the limbic input. This is a search network, because the activity spreads in hope that at

some node it will be integrated within the activity of the same lower division that receives a thalamic input. If the activity exceeds the output threshold, then the node will fire triggering a resolution. The thalamic input is applied to the node corresponding to the current state. The process of spreading activity in a search tree is called goal regression [5].

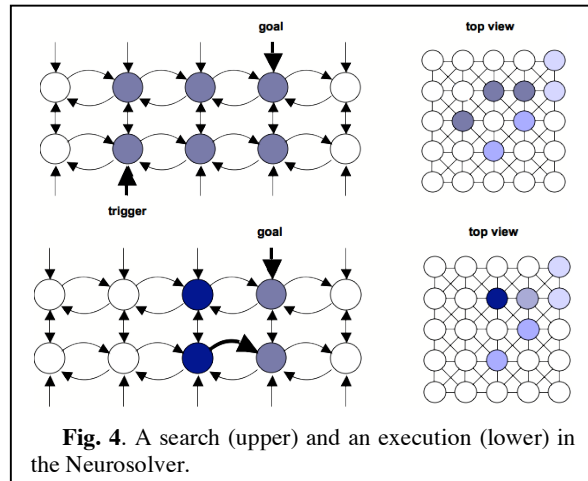


Fig. 4. A search (upper) and an execution (lower) in the Neurosolver.

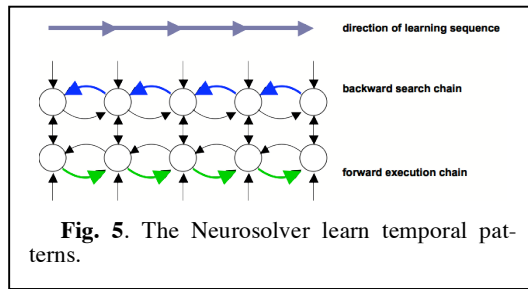


Fig. 5. The Neurosolver learn temporal patterns.

at some node it will be integrated within the activity of the same lower division that receives a thalamic input. If the activity exceeds the output threshold, then the node will fire triggering a resolution. The thalamic input is applied to the node corresponding to the current state. The process of spreading activity in a search tree is called goal regression [5].

The purpose of the network composed of lower divisions and their connections is to generate a sequence of output signals from firing nodes (along the connections shown in Fig. 5). Such a sequence corresponds to a path between the current state and the goal state, so it can be considered a solution to the problem. As we said, a firing

of the node representing the current state triggers a solution. Each firing node sends actions potentials through the outgoing connections of its lower division. These signals may cause another node to fire if its attention (i.e., the activity in the upper division) is sufficiently high. In a way, the process of selecting the successor in the resolution path is a form of selecting the node that is activated the most. The purpose of inhibiting a node for some time after firing is to avoid oscillations. The length of the inhibition determines the length of cycles that can be prevented. The Neurosolver exhibits goal-oriented behavior similar to that introduced in [6].

2 Simulated Rat Maze

We tested the Neurosolver capabilities on several abstract data sets of various sizes and have successfully applied the Neurosolver as a planner for rearranging blocks in a block world. Our analysis of the Neurosolver capabilities in the context of symbolic planners proves that it is a promising adaptive mechanism for general planning and problem solving. We collected these and many other ideas from several earlier publications in a summary article published as [1].

Currently, to explore Neurosolver's cognitive capabilities, we apply it to control a rat running in a maze. The Neuroscience and Cognitive Science community commonly use rats to gain insight into brain processes. One area of research is concerned with rats running in several types of mazes (e.g., an excellent Web site on rat behavior [7]). Evidently, rats are good learners in this domain (e.g., [8]). In one type of maze, a rat is allowed to explore the maze with food placed in constant locations. In subsequent runs, the rat improves its ability to locate the food faster. One of the currently supported hypotheses says that the rat builds a topology of the maze in its brain (cognitive map) encoded by place cells in hippocampus ([9]).

Rats evidently use topological maps of the maze to construct best possible paths to food locations. Neurosolver has been constructed on the premise that temporary associations are made in cortex. Some researchers hypothesize that the entorhinal cortex is a link between labile episodic storage in the hippocampus to more permanent and semantic storage in the cortex ([11]).

To conduct experiments, we built a simulated rat maze. The user interface is illustrated in Fig. 6. The simulator consists of the maze area (left side of the frame¹) and

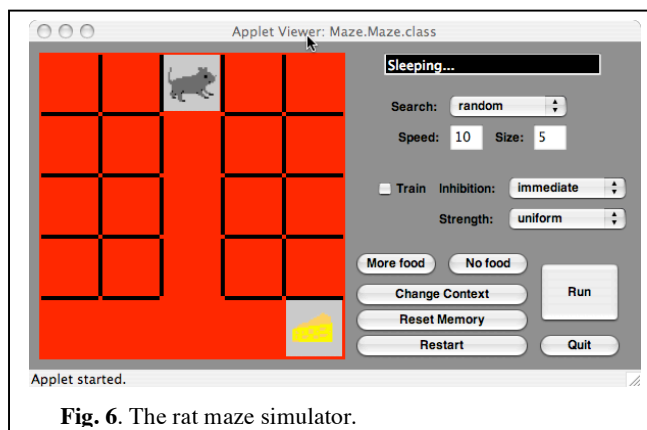


Fig. 6. The rat maze simulator.

¹ We will use just this area in the subsequent illustrations.

the control area. The maze area consists of a number of locations separated by walls. Mazes of various shapes can be built by modifying the walls. The rat can be positioned by dragging its image to a desired location. Food (illustrated by a piece of cheese) can be placed in multiple locations also by dragging the image from the location in which it is posted when the More Food button is pressed. Clicking on the Run button activates the metabolic system that gradually increases the sensation of hunger. If the motivation to find food exceeds a certain threshold, the rat's limbic system is activated. The rat starts to run. The Run button changes to Pause, so the simulator can be stopped and restarted at will.

The Change Context button changes the color of the floor that is used in experiments on contextual learning. We will discuss the details later in this paper.

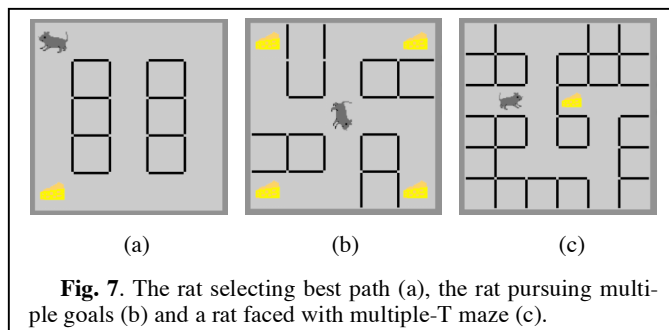
The rat is capable of moving in four directions: up, right, down and left. It can move only if there is no wall between two neighboring locations. The simulated rat uses an instance of Neurosolver to determine the next move in the maze. Other search modes such as random, constrained, and depth-first search are implemented to control the rat if the Neurosolver cannot find a solution .

When the rat finds food its motivation level decreases to below threshold, another parameter controlled and the rat may stop running. The threshold for this is one of the parameters that can be controlled during experiments.

Rats use cues from the environment to activate a specific place cell in the brain ([12]). A similar associative mechanism allows a rat to learn and remember where the food was found ([11]). Similarly, in the simulator, the rat remembers the locations of successful feedings, so it can use them as goals in the future runs. If the rat finds food again in the location of a goal, the strength of that goal increases, so the rat is more likely to pursue this goal in the future. Conversely, if the rat gets to a location that had food in the past, but the food is not there, the strength of the goal is lowered. If such negative experience is consistent, the goal will be completely inhibited, a behavioral process called extinction ([13]). We are planning to examine the goal management theories to a greater extent in the future.

Using our simulator, we performed a number of tests with the artificial rat running in various mazes. A T-maze is the simplest type of maze used in experiments with rats. A passage in a shape of the letter T forces the rat to choose the direction, left or right, at the crossing as shown in Fig. 6. If food is placed consistently in one arm of the T, then this is the arm that will be selected by the rats in the subsequent runs. If the rat obtained food from both arms then it will choose the one that has a better trace in memory.

In another experiment shown in Fig. 7 (a), there are three paths leading from the original position of the rat to the food. The rat selects the shortest path if the uniform learning is selected. If probabilistic learning



is chosen, then the path most often followed in the past and therefore the most probable, will be taken. This behavior comes from the fundamental characteristics of the Neurosolver. If a wall is created along the shortest path, then the rat reconsiders the plan and selects an alternate path backtracking as necessary. The rate of degradation can be controlled from the simulator. Live rats may exhibit aberrant behaviors during such stressful situations that may lower the motivation induced by hunger. As long as the threshold to search for food is met, the rat will try to get to the goal by any means putting more effort in recollections of past passages and food locations. Fig. 7 (b) shows a rat in a maze with four branches (“star-shaped”). The simulated rat trying to get all food again performs similarly to live rats. If food is removed from certain locations, then the rat will tend to move to the branches that provided consistent food-reward. There is a more complex T-maze used in tests with rats as shown in Fig. 7 (c). The rat is faced with multiple choices (T’s) on their path to the food. This is a more challenging task to live rats. It also takes a longer training session for the artificial rat to build a map, and higher motivation to find a path to the food.

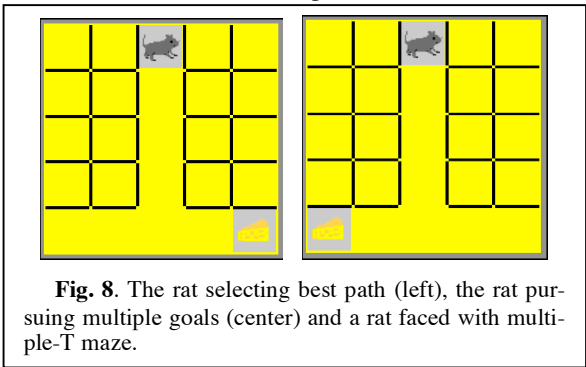


Fig. 8. The rat selecting best path (left), the rat pursuing multiple goals (center) and a rat faced with multiple-T maze.

3 Applying Context

In spite of these successful experiments, a problem arises if the probability of finding food is the same in both arms of T-maze as shown in Fig. 8. For example, if the number of successful feeding in both arms of T-maze is the same, then the rat is confused. Where should the rat go left or right?

Figure 9 illustrates a top view of the Neurosolver trying to solve the rat’s problem. The activity spreads from the cells corresponding to both goals (food locations) along

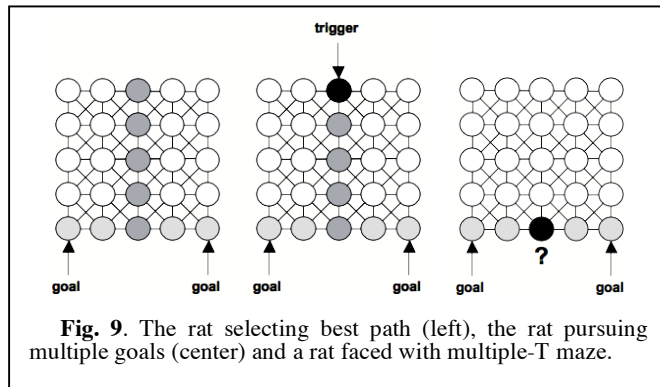


Fig. 9. The rat selecting best path (left), the rat pursuing multiple goals (center) and a rat faced with multiple-T maze.

the learned paths. The common segment of the search path is activated a bit more, because the search activity from the left goal is integrated with the activity coming from the right goal. That does not have much of an impact on the trigger in the cell corresponding to the cur-

rent position of the rat in the maze. The columns fire along the common segment, but then there is no clear-cut choice whether the next firing column should be to the left or to the right from the fork location. The Neurosolver tries to increase the activity from the goals, but it is increased by the same amount on both sides. The current version of the Neurosolver has a mechanism that arbitrarily selects one of such alternative columns and forces it to fire. The mechanism is not biologically plausible. Animals – as well as humans – make choices, and the choices are made using context.

Neuroscientists working with rats routinely incorporate context in the learning process (e.g., [7]). The Morris water maze is used as evidence for spatial contextual learning. In such a maze, the rat cannot see a hidden platform submerged in white, opaque water, but it nevertheless learns its position by using visual cues from the environment. The results (e.g., [16], [17]) indicate that rats navigate using direct cues to the goal (taxon navigation) or memories based on places activated by the environment clues (praxic navigation), although the true nature of rats' maze running skills is still under discussion (e.g., [18]). When the environment changes, or obstacles like walls are used to obstruct the view, the appropriate place cells do not fire, implying that the rat is unaware of its true location and needs to re-discover the platform through more explorations. As we mentioned earlier, our simulator already includes such a discovery mechanism.

At this stage, we are interested in praxic navigation, because at the moment we do

not experiment with real-time visual information that could be used to guide movements (like in hand-eye coordination).

To introduce a context, the color of the maze floor or the color of the light illuminating the maze might be changed (let's say yellow and red)

depending on the food location. This is of course inspired by experiments with live rats, as they tend to follow the paths leading to goals associated with the color of the floor or the color of the light.

We have extended the maze simulator, so a colored floor can be used in the rat maze simulator as illustrated in Fig. 10. Through the addition of the color cue, we can conduct experiments involving contextual learning.

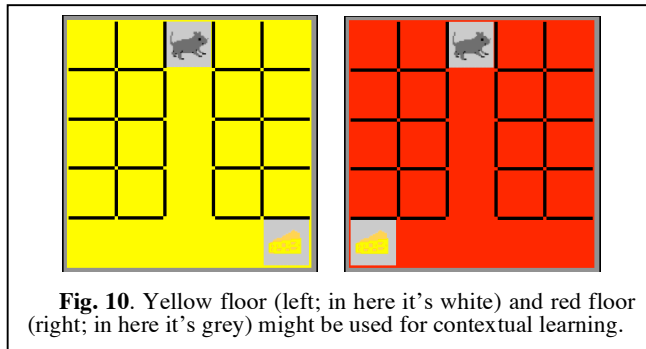


Fig. 10. Yellow floor (left; in here it's white) and red floor (right; in here it's grey) might be used for contextual learning.

4 Architecture for Contextual Reasoning

The new capabilities of the simulated maze amount to just one element of the required modifications. More fundamentally, the Neurosolver needs new functionality to accommodate context. One way to do this would be to incorporate the contextual information as one of the dimensions that determine the place cell – in our jargon, a hy-

percolumn that corresponds to the location of the rat. In that way, the identical path learned with red floor would be represented in the Neurosolver separately from the same path that was learned with yellow light, because the place cells would be different². That solution would be a waste of the Neurosolver resources, because each new cue would lead to yet another space. Therefore, we propose a different architecture that uses auxiliary contextual cells instead. A context cell is activated by contextual cues³. It is debatable which solution is more biologically plausible, but the use of contextual (or snapshot as some researchers call them) cells makes sense from a pragmatic perspective, because the storage requirements are drastically lower in the approach that we have taken.




Fig. 11. Contextual association with red floor (left) and yellow floor (right).

As illustrated in Fig. 11, the color of the floor results in activation of the context cell corresponding to that color. At times, the contextual cells will be co-activated with the firing nodes of the Neurosolver corresponding to the movements of the rat.

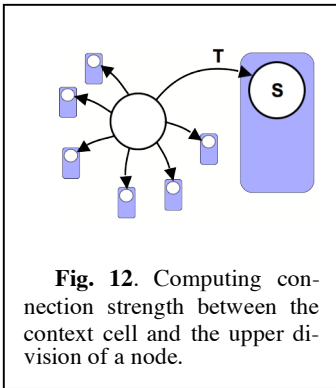


Fig. 12. Computing connection strength between the context cell and the upper division of a node.

As described earlier, the Neurosolver columns fire in response to the thalamic input reflecting changes in the rat's location. Through the use of Hebbian learning rules, each node in the temporal pattern becomes gradually associated with the co-activated context cell as shown in Fig. 12. Similarly to the connection strengths in the Neurosolver, we apply statistical learning here. The strength of a connection is based on the co-activation statistics. If any given hypercolumn fired S times, and a specific context cell was co-activated T times, then the strength of the connection is the co-activation probability P_c :


$$P_c = T/S \quad (2)$$

After the learning, during Neurosolver's operation, any activity in the context cell is modulated by that probability and projected into the upper division of the hypercolumn.

With such associations in place, certain nodes of the Neurosolver are activated even in absence of any search activity (see Fig. 13, upper left). The nodes in the search

² They would be two different points in the state space that includes a color axis, and the projections on this axis ("red" or "yellow") would be different for both paths.

³ For example, it can be selected through a competitive, winner-takes-all network. Creating such a network is not the objective of the work reported in this paper. Instead, we assume that a recognizing mechanism is in place, and red color activates one cell, while yellow activates the other one.



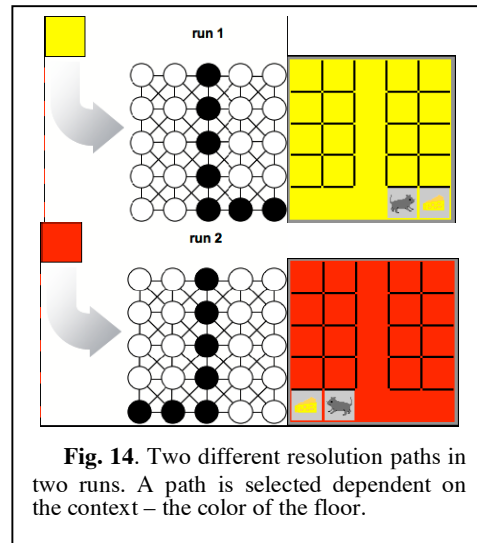
turns right and is shown just before consuming the reward after making the correct choice. In the second run, with the red floor, the rat turns to the left, and again is shown just before feeding.

5 Conclusions and future work

In spite of the obvious simplifications, the Neurosolver used as a brain of the artificial rat performs analogously to live rats when tested on similar tasks. We use the analysis of any discrepancies between the behavior of the Neurosolver and live rats to evolve the architecture of the Neurosolver. One difficulty with that arises from the variety of behavioral patterns identified in lab animals. The problem of studying learning and memory in animals is that we can only infer memories from their behavior. Careful analysis of the enhanced architecture of the Neurosolver leads us to the conclusion that it indeed provides the capabilities for contextual learning. The pragmatic aspect of the new mechanism is evident, so from the engineering perspective it is a significant improvement. Nevertheless, an interesting generic question to ask is whether the behavioral patterns observed on animals can be replicated by simulated

path corresponding to a particular learned sequence will now be activated not only by the activity coming from the source of the search, but also by the activity coming from the context cell. When two competing goals are applied, the search progresses as described earlier, but the nodes in the path associated with the current context are activated at a higher level, because the search activity integrates with the contextual activity. When the trigger activates a resolution path, the firing nodes will follow the highest activity in the neighborhood, so when the node corresponding to the fork in the T-maze fires, the next node to fire will more likely be the one on the path whose nodes are associated with the currently active context.

Figure 14 illustrates the corresponding behavior of the rat in the maze with two different contexts. With the yellow floor, the rat



whether the behavioral patterns observed on animals can be replicated by simulated

creatures controlled by Neurosolver to the same degree as the more basic alignment of behaviors reported in our previous papers.

There are plenty of other interesting venues worth pursuing in research with Neurosolver, and in more general terms in computational architectures inspired by progress in neuroscience. For example, we plan to further analyze the fundamental simplification that we consciously made in the Neurosolver that is the integration of the functionalities of the prefrontal cortex (where logical sequencing and executive function takes place) and hippocampus (where place cells are located). A related decision to map every possible location as a place cell affects the scalability of the Neurosolver, so we are exploring ideas to reduce the state space. For example, only the important places could be mapped (e.g., forks, turns, etc.), but how to modulate such a mapping is an open question.

Another idea that is very interesting for a computer scientist is “programmability” in the hippocampus. When an animal learns a new environment, he creates a map of that space that is encoded by the hippocampus. When he is out of that environment, the map, the “code”, is transferred elsewhere for “storage”, and the hippocampus is available for the “encoding” of a new environment. When he returns to the first environment, the “code” is “loaded” back into the hippocampus and the place cells are once again engaged. The idea that there are “grandmother cells” in the hippocampus is generally not applied in neuroscience circles.

At this moment, the searches that the Neurosolver performs are conducted off-line with locked knowledge. Allowing for real-time dynamic modifications that may alter the search or execution process is both appealing and biologically plausible. The capability to accomplish taxon navigation could be one result of such work. That in turn, could be the basis for Neurosolver-based controllers (for example, for hand-eye navigation). Another alluring idea is to explore learning by analogy that would require an association and generalization mechanism. That would allow knowledge re-use in context that engages concepts similar to those with which the current knowledge was acquired.

The path optimization capability of the Neurosolver can be viewed both positively and negatively. At its current incarnation, the Neurosolver cannot store higher order sequences. That is, for example, a sequence 1-2-3-2-4 is optimized to 1-2-4. Some applications may need such capability, so we are looking into designing an appropriate mechanism that can provide that.

Numerous technical aspects of the implementation are also worth pursuing. For example an alternative to software implementation could accommodate much larger state spaces. To provide required granularity in finding best paths and simulate analog workings of the brain, the propagation of activity is very inefficient as implemented in software. Furthermore, the Neurosolver uses statistical learning, so each of the connected elements is a little processor churning numbers. We plan to experiment with other types of weight management, perhaps more aligned with the mainstream research on Neural Networks.

Goal management is another interesting research area. How do rats choose where to go? Do they ever forget the associations between the pleasures of feeding and the location? Neuroscientists have been feeding us with new data in this area (somewhere between the amygdala, hippocampus and the cortex), and we are eager to explore the new findings.

We would like to express our gratitude to the reviewers, who raised many important and interesting issues that we tried to address in this section.

6 References

1. Bieszczad, A. and Pagurek, B. (1998). Neurosolver: Neuromorphic General Problem Solver, *Information Sciences: An International Journal* 105: 239--277, Elsevier North-Holland, New York, NY.
2. Newell, A. and Simon, H. A. (1963). GPS: A program that simulates human thought, in Feigenbaum, E. A. and Feldman, J. (Eds.), *Computer and Thought*. New York, NJ: McGrawHill.
3. Burnod, Y. (1988). *An Adaptive Neural Network: The Cerebral Cortex*. Paris, France: Masson.
4. Laird, J. E., Newell, A. and Rosenbloom, P. S. (1987). SOAR: An architecture for General Intelligence, *Artificial Intelligence*, 33: 1--64.
5. Nillson, N. J. (1980). *Principles of Artificial Intelligence*. Palo Alto, CA: Tioga Publishing Company.
6. Deutsch, M. (1960). The Effect Of Motivational Orientation Upon Trust And Suspicion, *Human Relations*, 13: 123--139
7. <http://ratbehavior.org/>
8. Hodges, H. (1996). Maze Procedures: The Radial-Arm And Water Maze Compared, *Cognitive Brain Research* 3 (1996) 167 – 181, Elsevier North-Holland, New York, NY.
9. Fenton, A. A. and Muller, R.U. (1998). Place Cell Discharge Is Extremely Variable During Individual Passes of The Rat Through The Firing Field, *Proc. Natl. Acad. Sci. USA*, Vol. 95, pp. 3182--3187.
10. Poucet, B. and Save, E. (2005). Attractors in Memory, *Science*, 308: 799--800, AAAS Press.
11. Fyhn, M., Molden, S. and Witter, M. P. (2004). Spatial Representation in the Entorhinal Cortex Marianne, *Science*, 305: 1258--1264, AAAS Press.
12. O'Keefe, J. and Dostrovsky, J. (1971). The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-Moving Rat, *Brain Research*, 34: 171-175.
13. Pavlov, I. P. (1927). *Conditioned Reflexes*. Routledge and Kegan Paul, London.
14. Bitterman, M. E., Lolordo, V. M., Overmier, J. B. and Rashotte, M. E. (Eds.) (1979). *Animal Learning: Survey And Analysis*, New York, NJ: Plenum Press.
15. Charles C. Kemp (2001). Think Like a Rat. Paper for MIT EECS Area Exam (http://people.csail.mit.edu/cckemp/cckemp_place_cells_area_exam_2001.pdf).
16. Hartley T., Burgess N. (2002) *Models Of Spatial Cognition*, Encyclopaedia of Cognitive Science, MacMillan.
17. Burgess, N. and O'Keefe, J. (2002). Spatial Models of the Hippocampus in: *The Handbook of Brain Theory and Neural Networks*, 2nd Edition Ed: Arbib M A, MIT press, Cambridge MA.
18. Chavarriaga, R., Strösslin, T., Sheynikhovich, D. and Gerstner, W. (2005). Competition Between Cue Response And Place Response: A Model Of Rat Navigation Behavior, *Connection Science*, Vol. 17, Nos. 1--2, March--June 2005, 167--183.