



Internal simulation of perception: a minimal neuro-robotic model[☆]

Tom Ziemke^{a,*}, Dan-Anders Jirenhed^b, Germund Hesslow^b

^a*School of Humanities and Informatics, University of Skövde, P.O. Box 408, 54128 Skövde, Sweden*

^b*Department of Experimental Medical Science, Lund University, BMC F10, 22184 Lund, Sweden*

Received 20 December 2002

Abstract

This paper explores the possibility of providing robots with an ‘inner world’ based on internal simulation of perception rather than an explicit representational world model. First a series of initial experiments is discussed, in which recurrent neural networks were evolved to control collision-free corridor following behavior in a simulated Khepera robot and predict the next time step’s sensory input as accurately as possible. Attempts to let the robot act blindly, i.e. repeatedly using its own prediction instead of the real sensory input, were not particularly successful. This motivated the second series of experiments, on which this paper focuses. A feed-forward network was used which, as above, controlled behavior and predicted sensory input. However, weight evolution was now guided by the sole fitness criterion of successful, ‘blindfolded’ corridor following behavior, including timely turns, as above using as input only own sensory predictions rather than actual sensory input. The trained robot is in some cases actually able to move blindly in a simple environment for hundreds of time steps, successfully handling several multi-step turns. Somewhat surprisingly, however, it does so based on self-generated input that is not particularly similar to the actual sensory values.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Inner world; Simulation of perception; Sensory anticipation; Cognitive robotics; Representation

[☆] An earlier version of this paper appeared as a technical report [39].

* Corresponding author. Tel.: +46 500 448330; fax: +46 500 448399.

E-mail addresses: tom.ziemke@his.se (T. Ziemke), dan-anders.jirenhed@mphy.lu.se (D.-A. Jirenhed), germund.hesslow@mphy.lu.se (G. Hesslow).

1. Introduction

“The world is its own best model” was one of the central tenets of Brooks’ formulation of the behavior-based approach to robotics and artificial intelligence (AI) [3]. Accordingly, his notion of situatedness was that robots should or need “not deal with abstract descriptions, but with the here and now of the world directly influencing the behavior of the system”. However, although Brooks was certainly right in his critique of traditional AI in general, and the traditional view of representation in particular, many researchers have argued that the notion of representation cannot be abandoned completely, or that it would have to be replaced or complemented by alternative theories of how memory and internal states affect behavior (e.g. [6–8,12,29,38]). Obviously, relying on the “here and now of the world” for behavioral guidance at each and every point in time is problematic in cases where autonomous agents need to remember the past, plan the future, or deal with the present in the absence of sufficient, reliable sensory input. In such cases, it has been argued, some kind of ‘*inner world*’ would have to stand in for the real world.

This paper is organized as follows: the following background section discusses what is meant by an ‘inner world’, some of the evidence that humans and other animals actually have such an inner world, and possible accounts for how it might be realized. Sections 3 and 4 then report on our experiments exploring the possibility of an inner world based on simulation of perception, using a simulated Khepera robot guided by artificial neural networks (ANNs). Section 5, finally, summarizes and discusses the work presented in this paper.

2. Background

2.1. Motivation: the inner world

Introspection seems to tell us that we are able to have sensory (or quasi-sensory) experiences in the absence of external stimuli. Examples of this phenomenon could be the visual experience of ‘seeing an apple’ upon reading the word ‘apple’ in a text or the auditory experience of ‘hearing the Moonlight Sonata’ upon a request to imagine a piece of music by Beethoven. This kind of ability has further been implicated by experimental psychological results of, e.g., Lee and Thompson [18]. In a series of experiments they demonstrated the accuracy with which humans can guide their behavior based solely on internally generated sensory experiences. A group of subjects were first instructed to look at their surroundings (a room) and direct specific attention to certain objects, such as marks on the floor and different obstacles. They were then blindfolded and asked to perform different tasks such as walking to the marked locations, avoiding the obstacles and throwing objects at different targets in the room. The subjects performed these tasks almost as accurately with eyes closed as when they were free to look. One explanation of these results could be that the subjects actually did ‘see’, ‘walk’ and ‘throw’, but did so internally—in their inner world where sensory experiences and consequences of different behaviors may be anticipated.

Evidence from experiments with rats seems to imply a similar interpretation. When rats are allowed to move freely in a maze environment on several trials without receiving food reward, presumably no navigational behaviors are reinforced and thus not learnt [35]. Later, when food is placed at one location for a number of trials, these rats will soon find that location and avoid dead ends just as reliably as a group of continually rewarded rats (a phenomenon known as *latent learning*). The evidence implies that maze topography is actually learnt while moving around although no rewards are given. It has been argued that such observations should be understood as the result of an internal *cognitive map* being built up in the rat nervous system from sensory experiences. The map is then supposedly consulted when navigating to the food location.

The traditional cognitive-scientific explanation for how such an inner world might be organized is based on symbolic world models, internal maps and planning mechanisms. This view has been questioned since the mid-1980s by many researchers (e.g. [1,2,5,6,30,31,36,38]) who de-emphasize the role of internal world models in the traditional sense, and instead emphasize the situated and embodied nature of intelligence as well as the reactive nature of many of the mechanisms underlying the interaction between agents and their environments. Accordingly, robots or other autonomous agents, controlled by subsumption architectures, ANNs, classifier systems, or similar mechanisms, have in many cases exhibited good performance in tasks that could be solved in a more or less reactive fashion. Purely reactive mechanisms, however, seem ill suited to explain the above capacity to simulate/anticipate sensory experience and to behave appropriately in the absence of external stimuli. Hence, it is unclear how such more or less reactive systems could handle tasks that might require planning and action in the absence of sufficient external sensor input.

An alternative to representational world models in the traditional sense is Hesslow's *simulation hypothesis* which accounts for the inner world in terms of *internal simulation* of perception and behavior [13,14]. Similarly, Clark and Grush in [7] argue that "the capacity to think about the distal and absent is grounded in the use of systemic stand-ins and emulation-based strategies". They discuss the notion of internal *emulation* based on *forward models* in their account of a "minimal robust representationalism". Whether or not an internal simulation or emulation, or the results thereof, should be referred to as an internal *representation* is, of course, a controversial issue (cf. [7,12]). This question is not addressed in further detail in this paper, which in the following uses Hesslow's non-representational account and terminology.

Neurophysiological findings indicate that the neural structures involved in *perception* and initiation of *overt behavior* are also responsible for *covert behavior* and *mental imagery*. As illustrated in Fig. 1, an agent's overt behavior when interacting with the environment could be described as a sequence of stimulus-response pairs. In the initial situation a stimulus (S) triggers a response (R), which changes the environment from the agent's perspective into a new situation (S') which functions as a stimulus triggering another response (R'), which in turn causes S'' and R'' and so on. All these steps go via internal processes where sensory states (s) cause motor response preparations (r).




Fig. 1. Stimulus-response sequence during overt behavior. S refers to the stimulus or situation, s to the neural sensory activity, r to the neural motor activity, and R to the behavioral response.

The simulation hypothesis [13,14] assumes three mechanisms in order to explain the ‘inner world’ phenomenon. The first assumption is *covert behavior*, i.e. an ability to generate neural motor responses that do not become observable bodily actions but only neural activation patterns, which remain purely internal. Secondly, the existence of a *sensor reactivation* or *imagery* mechanism is assumed. This allows for internally generated activation of sensory areas in the brain, so as to produce the simulated experience of a stimulus, but without the presence of that external stimulus. Finally, the existence of an *anticipation* mechanism is assumed, i.e. an ability to predict or simulate the sensory consequences of a motor response (a bodily movement). Support for each of these assumptions can be found in the neuroscience literature. Covert behavior and imagery (sensor reactivation) are discussed by, e.g. Jeannerod in [15], and imagery further by Kosslyn et al. in [17]. Finally, several neuroscientists consider the cerebellum a prime candidate module involved in prediction and anticipation (e.g. [22,23,34,37]).

With these three mechanisms in place (covert behavior, sensor reactivation and anticipation), it would be possible, to *internally simulate* the above overt behavior sequence (cf. Fig. 1) as illustrated in Fig. 2. In the initial situation, a stimulus (S) causes a sensory activation (s) which triggers a motor response preparation (r). But instead of causing the overt response (R), the motor response preparation could cause a new (anticipated) sensory activation (s') via a mechanism for anticipation of the sensory consequences and, in turn, a sensor reactivation mechanism. This new (internal) sensor activation triggers another motor response preparation (r'), which in turn causes a new sensory activation (s'') which triggers a motor response preparation (r'') and so on. Instead of overtly interacting with the environment, the agent internally simulates this interaction.

2.2. Related robotics work

Sensory prediction or anticipation has been addressed in a number of robotic and neurocomputational studies. Tani and Nolfi [32,33], for example, investigated how a recurrent ANN robot controller could acquire an internal model of the world




Fig. 2. Sequence of internally simulated sensorimotor interaction.

through training on sensor prediction while moving around in a two-room environment. After learning, the controller had developed internal dynamics corresponding to environmental features (or concepts) such as ‘corner’ and ‘corridor’ and specific sequences of such low-level concepts were also used to form higher-level concepts such as ‘room A’ and ‘room B’.

Drawing on psychological, physiological and neuroanatomical data on the properties of the cerebral cortex, Gross et al. in [11] presented a neurocomputational model of perception (called MASIM) in which sensor anticipation/prediction played a central role. In a series of robot experiments the MASIM architecture was used to guide the selection of motor responses based on anticipated sensor flow. Their results showed that anticipating robots displayed better navigation abilities, with higher speeds and fewer collisions when compared to reactive robots. Perhaps most significantly, evasive maneuvers in order to avoid obstacles could be initiated much sooner.

However, the work that influenced our initial experiments most are Meeden, McGraw and Blank’s experiments [20,21] with a toy-car-like robot controlled by a so-called simple recurrent network (SRN; [9]). The robot’s task was to alternate between approach and avoidance of a light source in one corner of its rectangular environment while avoiding obstacles. It was shown that through the use of internal feedback (short-term memory) the robot could exhibit behavior that Meeden et al. described as “plan-like” for the following reasons [21]: (a) the robot associated abstract behavioral goals with sequences of primitive actions; (b) its behavior could be described in hierarchical terms (e.g., *light seeking* was comprised of the sub-behaviors *orient towards light* and *go to light* which in turn were realized as primitive actions of moving and turning); (c) the robot maintained its overall strategy even when reacting flexibly to the environmental conditions. On the other hand, the behavior was clearly *not* plan-like at all in the traditional sense that the robot would actually ever explicitly anticipate or plan future situations.

Meeden et al. also conducted experiments in which the robot was additionally trained on predicting the next time step’s sensory input [21]. The results showed that this additional training did have a positive effect on the behavior learning. However,

no analysis was presented regarding the quality of the robot's sensor predictions or the robot's capacity to actually make use of its own prediction instead of external input. Tani and Nolfi in [33], on the other hand, illustrated that the output of their sensory prediction networks in the above experiments was indeed similar to the actual sensory input. Using phase space plots they also showed that the networks' dynamics during interaction with the environment were to some degree similar to the networks' intrinsic dynamics, i.e. the dynamics when the networks used their own prediction as input in the next time step. Like Meeden et al., however, they did not analyze if the robot's sensory prediction was accurate enough to allow 'blindfolded' interaction with the world. These issues were therefore addressed in our initial experiments which will be discussed in the following section.

3. Initial experiments

As mentioned above, our initial experiments (see [16] for details) are built on the work of Meeden et al. [21]. The basic idea behind our extension of their work was this: If, using the ANN control architecture shown in Fig. 3, the robot's sensory predictions were sufficiently accurate, then it should be able to use them instead of actual sensory input. If this were done repeatedly in consecutive time steps, as illustrated in Fig. 4 (cf. Fig. 2), then it should enable the robot to behave appropriately, at least for some time, in the absence of external stimuli, i.e. based on self-generated or simulated stimuli alone, just as Lee and Thompson's [18] human subjects presumably did.

Our initial experiments were carried out using a Khepera robot [25] simulator (cf. Fig. 5) based on sensor and motor measurements obtained from a real Khepera robot [24]. The experiments have been documented in detail in [16], but here are the main points: Recurrent ANN controllers were trained using a fairly standard evolutionary algorithm (EA), similar to those used in [10,26,27]. The controller networks received as input the sensory activation values of the robot's eight infrared proximity sensors, and the motor outputs controlled its two motors directly




Fig. 3. Robot control architecture used in the initial experiments. Solid arrows indicate fully connected layers of weights between layers of units (indicated by dotted lines). The dashed arrow represents a 1:1 copy connection.

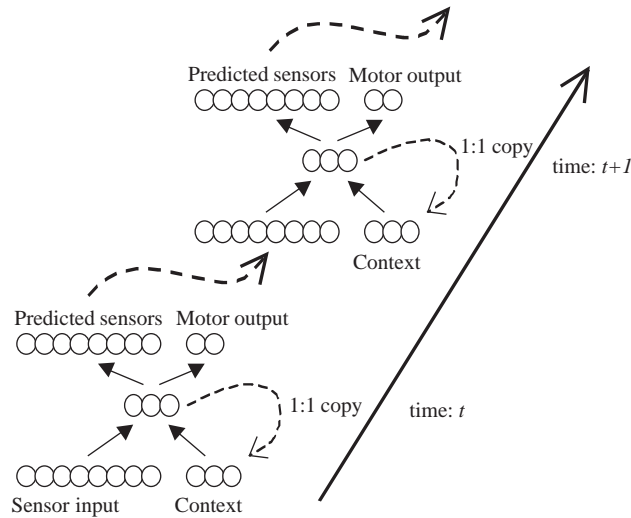


Fig. 4. Internal simulation of perception. Instead of real sensor readings the previously predicted sensor activation vector was used as input in each time step. Overt behavior during a simulation phase was thus solely based on the robot's predictions about sensor states. Solid arrows indicate fully connected layers of weights, dashed arrows indicate 1:1 copy connections.




Fig. 5. (a) Khepera robot built by K-Team SA (www.k-team.com). (b) Schematic drawing of the robot with infrared proximity sensors (1–8), left and right wheel (controlled by independent motors). The robot's diameter and sensor range are about 55 mm. The direction of forward motion in the reported experiments is indicated by the arrow.

(cf. Fig. 5). Only weight values were evolved during training, while the network topology remained static. Three experiments were carried out in two environments (one of which is shown in Fig. 6). First, in order to establish that the architecture was suitable for the task, controller networks without sensory prediction were evolved (successfully) to simply control collision-free corridor following. Then networks of the type illustrated in Fig. 3 were evolved (from scratch) to control the above behavior and to additionally predict the next time step's sensory input as accurately

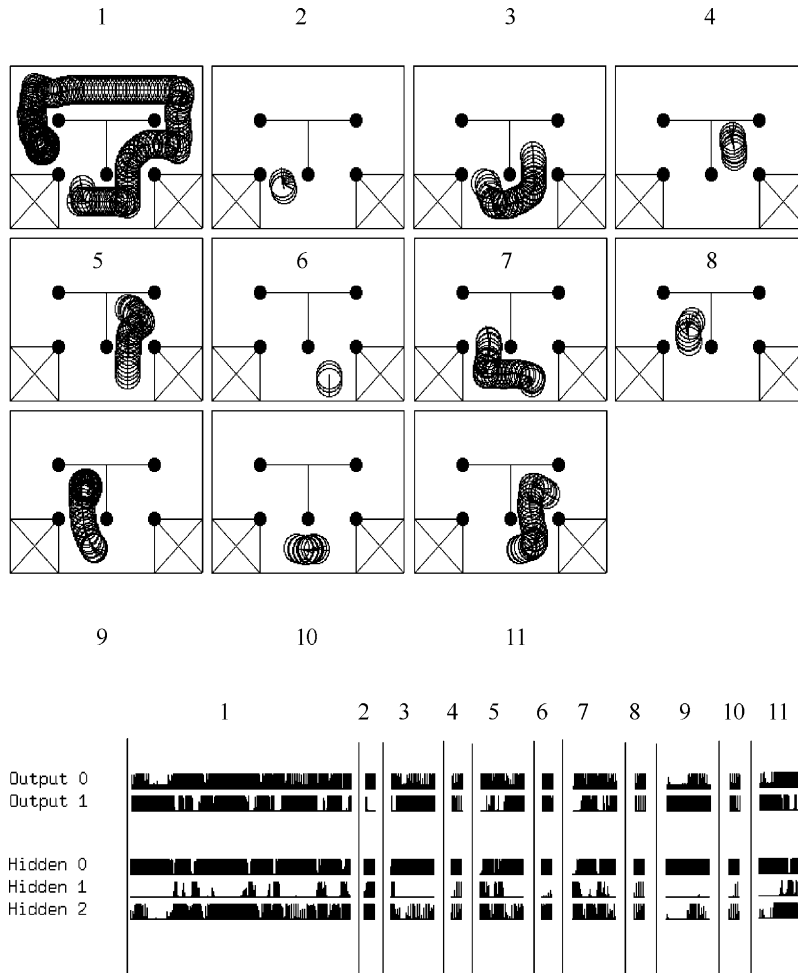


Fig. 6. Robot during context building phase (frame 1; 200 time steps), internal simulation phases (even-numbered frames; 10 steps each) and context-rebuilding phases (odd-numbered frames from 3; 40 steps each). The top part illustrates the robot's behavior in the environment (each time step represented by a circle, with the line indicating the heading). The bottom part illustrates the controller network output and hidden activation values during the above phases.

as possible (a two step selection/fitness function was used; for details see [16]). Finally, we let the robots trained in the above second experiment act blindly, repeatedly using their own prediction instead of the real sensory input, as shown in Fig. 4. Fig. 6 illustrates the behavior generated by a representative controller network. The odd-numbered frames show its behavior during context (re-) building phases during which it receives the real sensory input, whereas even-numbered frames depict its behavior in the internal simulation phases during

which it receives only self-generated sensor input but still acts (blindly) in the environment.

As the even-numbered frames in Fig. 6 show, the third experiment was not particularly successful. The robot seems to correctly continue its turn in frame 10, but in both frames 4 and 6 it fails to initiate the turn in time, and in frame 8 it turns in the wrong direction. A more detailed analysis (cf. [16]) of the evolved sensory prediction revealed that the networks managed to capture the most relevant sensors to the left and to the right, but they typically failed to anticipate when the frontal sensors became active, i.e. when it was about to collide with an obstacle and a turn was required. However, as the hidden and output activation values show (cf. Fig. 6), the controllers do not lack internal dynamics. That means, it is not the case that the robot during internal simulation phases just blindly keeps repeating the same action. Instead, the problem seems to be that, due to the short range of the robot's proximity sensors (approximately the same as its diameter), the frontal sensors are active so seldom that for the EA it does not pay off sufficiently to predict them accurately. Hence, the controllers seem to learn just not to predict any changes in front sensor activation, but instead to make the fairly good 'guess' that the activation remains the same as in the current time step (similar problems have been reported in [28] and [19]).

4. New experiments

4.1. Experimental setup

The set of experiments documented in this section was directly motivated by the problems encountered in the above initial experiments. These new experiments were carried out using our YAKS Khepera robot and ANN simulator [4]. However, instead of the previously used short-range proximity sensor, the long-range 'rod sensor' illustrated in Fig. 7 was used. This sensor is to some degree similar to the Khepera robot's K213 linear vision turret, but some of the properties are significantly different. The rod sensor covers an angle of about 30 degrees and has a range of about 30 cm. It distinguishes between black and white objects, but detects no gray levels. It also provides information about the proximity to objects, ranging from 0 (maximum distance) to 1 (touching). Hence, the most significant property for the experiments documented in the following is that, unlike the short-range sensors changing rather abruptly in the previous experiments, now the rod sensor activations change relatively continuously during a longer period of time while approaching rods. In the following experiments activation values of a simulated retina with a 10-bit resolution were mapped to 10 input units.

In addition to different sensors, a simpler world was used (cf. Fig. 8), consisting of four equally long corridors. It should be noted that the walls are below the rod sensor level, such that the rod sensor can only see the four rods which are placed outside the corridors, one at the end of each.




Fig. 7. Simulated 'rod sensor' (top) and an input pattern (bottom). See text for detailed explanation.

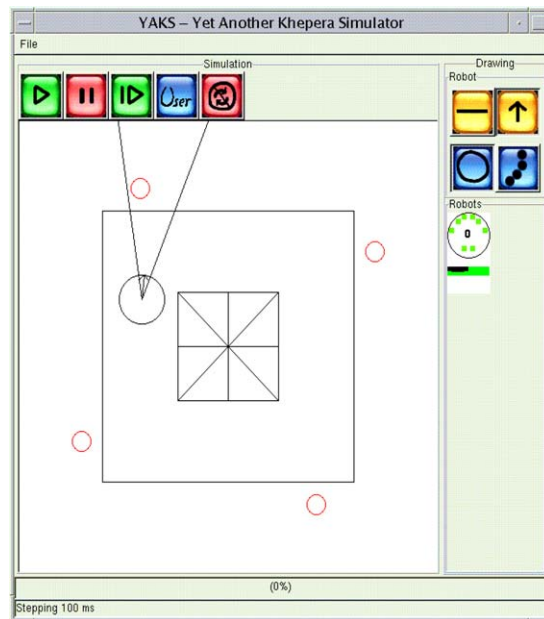


Fig. 8. Robot's environment, containing walls ('invisible' to the robot controller) and four circular rods.

4.2. Training

Experiments were carried out with different recurrent and feed-forward ANN architectures, and best results were achieved using the purely reactive feed-forward

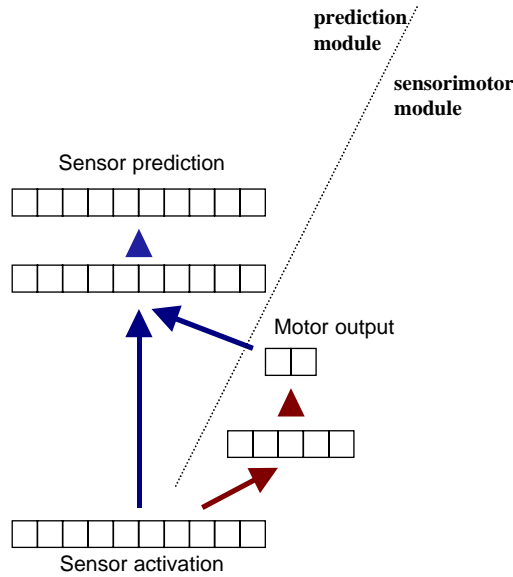


Fig. 9. Robot control architecture, consisting of two separately trained feed-forward modules: one sensorimotor module (to the right) mapping sensor input to motor output (trained first), and one prediction module (to the left) mapping sensor input and motor output to a prediction of the next time step's sensory input (trained second). When the robot receives real sensory input, it is controlled by the sensorimotor module alone. When the robot is 'blindfolded' and no external sensory input is available, the prediction module is used to generate/predict the next time step's sensory input.

controllers illustrated in Fig. 9. Again, an EA¹ was used to evolve network weights only. Different training regimes were used, and best results were achieved (somewhat counterintuitively; cf. final section) when training the sensorimotor module first and the prediction module afterwards.²

The fitness criterion for the sensorimotor module's training was based on obstacle avoidance and fast straightforward motion, resulting in collision-free corridor following. More specifically, the fitness was calculated in each time step as follows:

$$f = v_l + v_r - |v_l - v_r|,$$

where v_l and v_r are the velocity of left and right wheel, respectively (a velocity of 1.0 corresponds to full speed forward motion, 0.0 to full speed backward motion, and

¹Population size: 200. Generations: 10,000 or more. Elitist selection was used. Gaussian distributed weight mutations (variance = 0.5; 68% of mutations between -1.5 and $+1.5$). Connection weights were limited to the range -10 to $+10$.

²For the behavior module experiments were carried out with different rod sensor resolutions (4, 6 or 10 inputs) and different numbers of hidden units (0, 3, 5). The 10-5 combination illustrated in Fig. 9 achieved best results. For the prediction module different corresponding numbers of inputs (6, 8, 12) and outputs (4, 6, 10) were tested, as well as different numbers of hidden units (3, 8, 10). The 10-10-10 combination illustrated in Fig. 9 gave best results.

0.5 to standing still). The overall fitness of each individual sensorimotor module was calculated by summing up the fitness over a number of time steps according to the above formula during ten test epochs starting from 10 different random positions.

It should be noted that the robot and the fitness function do not make use of the infrared proximity sensors, i.e. the robot has to avoid the walls although it cannot actually see them. This makes the task more difficult than it might seem at first. The fitness criterion for the training of the prediction module was actually exactly the same as above. That means, accuracy in sensor prediction was no longer rewarded as such,³ but instead the sensor prediction capacity was evaluated indirectly by evaluating the quality of the behavior it allowed the robot to produce when real sensors were replaced by predicted ones during internal simulation. Since behavior was evolved with the robot using real sensor input, the predictions necessarily have to capture some properties of the real perceptions, otherwise the correct responses cannot be produced. Successful simulated perception therefore cannot be realized as just any randomly constructed sequence of activation vectors if the behavior is to be viable. Real sensor experiences and the corresponding responses thus restrict the set of viable sensor sequences which may be internally simulated.

4.3. Results

Fig. 10 shows a representative example of collision-free corridor-following behavior as guided by the trained sensorimotor module receiving real sensor input. It might be worth noting once more that the robot controller only sees the rods, not the walls. The turns in Fig. 10 are smooth enough to allow internal simulation, as shall be documented in the following. Another behavioral strategy was produced by other sensorimotor modules, a straightforward approach of the rod followed by a sharp turn, but no prediction module was ever found that could successfully predict/simulate the corresponding sensory sequence.

Fig. 11 is an extension of Fig. 10. In addition to the previously shown first round controlled by a trained sensorimotor module receiving real input, the ‘blindfolded’ robot’s trajectory during the second round is shown. It should be noted that during that second round the robot acted blindly, i.e. with no sensory input whatsoever from the environment, but solely based on its own internal sensory simulation. Nevertheless, the second round’s trajectory is collision-free and very similar to the first round (as indicated by the fact that they are very difficult to keep apart in Fig. 11).

4.4. Analysis

Fig. 12 illustrates the activation values in the sensorimotor module over time during the two-round trajectory illustrated in Fig. 11. To the left, the values during

³Evolving prediction accuracy had turned out to be a problem in the initial experiments documented in Section 3. The problems with the frontal sensors might have been overcome by weighting them higher, or by generally weighting changing sensors higher, but all these solutions would have forced us to set these parameters manually.

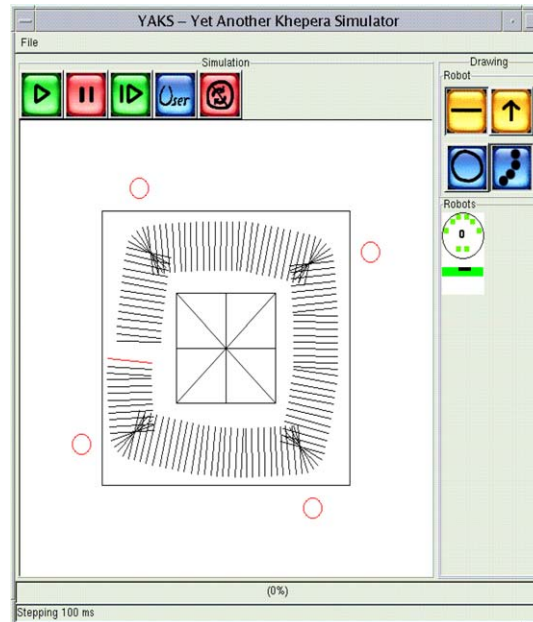


Fig. 10. One-round trajectory (moving clockwise) produced by a trained sensorimotor module receiving real input from the environment (no prediction/simulation is used).

the first round (based on real input) are shown, and in the center the values during the second round (based on internal simulation). It should be noted that although the robot's behavior is almost identical during the two rounds, the self-generated input does not appear to be as similar to the real input as one might have expected. In fact, in our experiments we have found no prediction module that produced self-generated input patterns with surface similarities to the real input.

To analyze further why self-generated and real sensory input are dissimilar, but nevertheless lead to almost identical behavior, we had a closer look at the example illustrated in Fig. 12 to find out exactly how the behavior in both these cases is controlled by the sensorimotor module. We found that while the robot is using real sensory input, it is mostly hidden unit h2 that controls the switch between forward motion and right turning. Hidden unit h2 has a very large positive connection weight to the right motor output (RM), such that the robot starts to turn right as soon as h2 becomes less active. That happens at the end of the corridor when the robot has already turned so much that it 'loses sight' of the rod which leads to 'deactivation' of h2, which is otherwise activated by sensory input. Consequently, the robot turns right for five time steps. When the robot uses only simulated, self-generated input, on the other hand, hidden unit h2 plays no role at all in controlling turns but is constantly fully activated (cf. Fig. 12). Instead it is now hidden units h1 and h3, which both have strong inhibitory connections to the right motor output, that control right turning behavior. In this case, the turning behavior actually only lasts

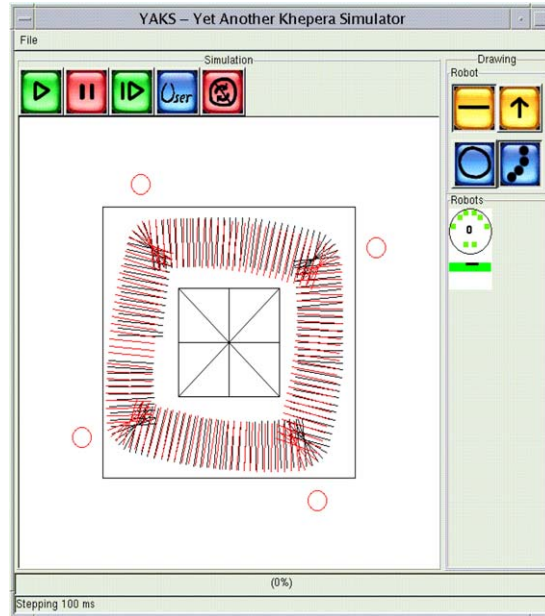


Fig. 11. Two-round trajectory. During the first round (black lines) the robot was controlled by a trained sensorimotor module receiving real input (using no prediction/simulation). During the second round (red lines) the robot was ‘blindfolded’, i.e. it received no sensory input at all from the environment, but acted solely based on self-generated sensory simulations.

four time steps instead of the five steps it takes when using real sensory input. That means, the turning behaviors are actually not the same, although they are functionally equivalent in the above observed sense that they lead to very similar motion trajectories. Another, more crucial difference that explains why self-generated and real sensory input are not more similar, and to some degree cannot be, is that while using real sensory input the robot’s behavior is purely *reactive*, controlled by the feed-forward sensorimotor module alone (cf. Fig. 9). When using self-generated input, on the other hand, the robot cannot just react to the environment, but has to time its turns very accurately, and to do that it has to ‘count’ the number of steps between turns. To some degree, the sensory inputs can play the role of a counter due to the fact that they increase monotonically as long as they are active. However, for several time steps during the ‘seeing’ robot’s turn, when none of the rods is in sight, all sensory activations remain at the same (low) level until suddenly the sensors become active again, one by one, when they detect the next rod. Obviously for the ‘blindfolded’ robot this transition (from not sensing any rod for several time steps to sensing a rod again) would not be possible to simulate with a purely reactive input–output mapping. The self-generated input therefore has to act as some kind of *short-term memory* that keeps track of how many of the required steps the robot has already carried out. This is reflected in self-generated inputs i_2 , i_3 ,




Fig. 12. Activation values over time in the controller network's sensorimotor module during the two-round trajectory illustrated in Fig. 12. The left third of the figure show the values (in black) during the first round (based on real input). The center shows the values (in red/gray) during the second round (based on internal simulation). Somewhat surprisingly, the real input is, in fact, very different from the self-generated input, although they lead to almost identical behavior.

i5 and i9 which exhibit continuous changes in activation, rather than the abrupt changes that can be observed in the real sensory inputs (cf. Fig. 12).

As a final piece of analysis/illustration of the results, it should be added that even for the best networks of course not all attempts of internal simulation are 100% successful. Many of them lead to collisions sooner or later. However, in many of the cases we have analyzed the mistakes are relatively minor. Fig. 13 shows such an example. The first round is the same as in Figs. 10 and 11 and was carried out in the same environment (containing walls and rods) based on real sensory input. The second round, based on internal simulation, was unsuccessful in the sense that the robot would have collided relatively quickly through a mis-timed, too early turn to the right. But, as Fig. 13 shows, when controlling the second round blindly in an environment without walls or rods, the controller network is still able to simulate the shape of the environment (or its own trajectory in it) pretty exactly; it just happens to get the timing and angle of the first turn wrong.

5. Discussion and conclusions

This paper has presented a number of experiments on internal simulation of perception using ANN robot controllers, as well as some of the theoretical background (cf. [13,14,7]). Obviously, the positive results presented in the previous




Fig. 13. Another two-round trajectory. The first round (based on real sensory input) is the same as in Figs. 11 and 12 and was carried out in the same environment (containing walls and rods). The second round was based on internal simulation and carried out without walls or rods; it would have led to a collision otherwise.

section have been achieved in a very simple environment, and our analysis showed that the successful behavior was achieved in a different way than one might have expected. Nevertheless, we think that the results presented so far are interesting, and the approach taken here is worth further investigation.

Some of our future work will be concerned with more realistic training and test regimes. The example of the ‘unsuccessful’ run illustrated in Fig. 13 points to one of the problems: our current fitness evaluation scheme could easily have abandoned this (partial) solution, due to the fact that in the original environment it would have crashed and ‘died’ early on. Instead, individuals might be given the chance to recover from mistakes, similar to the way Lee and Thompson’s [18] human subjects might have corrected their direction after brief contact with a wall. Furthermore, it may seem somewhat counterintuitive that no positive results were achieved using an integrated evolutionary training process for sensorimotor and prediction modules (cf. Section 4.2). We would have expected such an integration to facilitate the evolution of behavior that can actually be predicted. Further experiments will be carried out to investigate this question in more detail. Other questions are raised by the differences between the initial experiments (cf. [16]) and the new experiments documented here. In the first case (cf. Section 3), the fitness criterion networks were selected for the accuracy of their sensory predictions (only one time step ahead), and

this turned out not to lead to the evolution of a simulation capacity reliable enough to guide blind behavior. In the second case (cf. Section 4.2), prediction accuracy was not used as a fitness criterion, but only the quality of the resulting blind behavior. The resulting ‘simulation’ capacity turned out to generate input without obvious surface similarities to the actual sensory input, but nevertheless it was in many cases able to reliably guide the ‘blindfolded’ robot by producing functionally equivalent behavior patterns. Again, this will have to be investigated further with more elaborate training and test regimes. Finally, the issue of *abstraction* will have to be addressed. Obviously, prediction at the level of individual time steps is useful for tasks such as motor control, but anticipation and planning should also be possible at a higher level of abstraction.

The result illustrated in Figs. 11 and 12, i.e. the fact that, at least to our surprise, the robot manages to act blindly for more than 100 time steps, covering a whole round including four well-timed turn sequences (in an admittedly simple environment), *without* obvious surface similarities between self-generated input and real sensory values, leads us to ask the following question: is this a case of successful internal simulation of perception, in the sense discussed in Section 2, or is it not? It certainly is successful in the sense that it solves the given behavioral task (and that was the only criterion for the evolutionary training process). But, is it a *simulation of perception*? Although the self-generated input successfully stands in for the real sensory values, it can perhaps not be considered a simulation in a narrow sense because the input patterns are simply too dissimilar (cf. Fig. 12). But then again, let us go back to Lee and Thompson’s [18] human subjects (cf. Section 2): when they saw the world before their ‘inner eye’ they did not see it with their eyes, but perceived it ‘in’ their brains. Obviously, the comparison makes only limited sense, but it seems that the robot’s equivalent to a perception might have to be found in the hidden unit patterns rather than the input patterns. A closer comparison of the hidden unit activations does indeed show that the patterns there could be considered more similar than at the input level (they are likely to be since they have to produce relatively similar outputs). The self-generated hidden unit patterns for h1 and h3 could be considered rough approximations of the respective ‘real’ patterns. But, does it make sense to consider them perceptions?

To sum, despite many open questions, we think that the results presented here are encouraging, and we believe that our experimental approach will turn out to be a useful complement to neuroscientific and philosophical investigations of the inner world.

Acknowledgements

The experiments documented in this paper were carried out by Dan-Anders Jirenhed during the autumn of 2001, while employed at the Department of Computer Science, University of Skövde. This work was supported by grants from the Knowledge Foundation (1507/97) and the Swedish Medical Research Council (09899).

References

- [1] R.A. Brooks, A robust layered control system for a mobile robot, *IEEE J. Robotics Automat.* 2 (1986) 14–23.
- [2] R.A. Brooks, Intelligence without representation, *Artif. Intell.* 47 (1991) 139–159.
- [3] R.A. Brooks, Intelligence without reason, in: *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA, 1991, pp. 569–595.
- [4] J. Carlsson, T. Ziemke, YAKS—yet another Khepera simulator, in: S. Rückert, Witkowski (Eds.), *Autonomous Minirobots for Research and Entertainment—Proceedings of the Fifth International Heinz Nixdorf Symposium*, HNI-Verlagsschriftenreihe, Paderborn, Germany, 2001.
- [5] W.J. Clancey, *Situated Cognition: on Human Knowledge and Computer Representations*, Cambridge University Press, New York, 1997.
- [6] A. Clark, *Being There*, MIT Press, Cambridge, MA, 1997.
- [7] A. Clark, R. Grush, Towards a cognitive robotics, *Adaptive Behavior* 7 (1) (1999) 5–16.
- [8] G. Dorffner, Radical connectionism—a neural bottom-up approach to AI, in: G. Dorffner (Ed.), *Neural Networks and a New Artificial Intelligence*, International Thomson Computer Press, London, UK, 1997, pp. 93–132.
- [9] J. Elman, Finding structure in time, *Cognitive Science* 14 (1990) 179–214.
- [10] D. Floreano, F. Mondada, Evolution of Homing navigation in a real mobile robot, *IEEE Trans. Systems Man Cybernet.—Part B: Cybernetics* 26 (3) (1996) 396–407.
- [11] H.-M. Gross, A. Heinze, T. Seiler, V. Stephan, Generative character of perception: a neural architecture for sensorimotor anticipation, *Neural Networks* 12 (1999) 1101–1129.
- [12] R. Grush, The architecture of representation, *Philos. Psychol.* 10 (1) (1997) 5–23.
- [13] G. Hesslow, Will neuroscience explain consciousness?, *J. Theoret. Biol.* 171 (1994) 29–39.
- [14] G. Hesslow, Conscious thought as simulation of behaviour and perception, *Trends Cognitive Sci.* 6 (6) (2002) 242–247.
- [15] M. Jeannerod, The representing brain: neural correlates of motor intention and imagery, *Behav. Brain Sci.* 17 (1994) 187–245.
- [16] D.-A. Jirnehed, G. Hesslow, T. Ziemke, Exploring internal simulation of perception in mobile robots, in: K. Arras, A.-J. Baerfeldt, C. Balkenius, W. Burgard, R. Siegwart (Eds.), *2001 Fourth European Workshop on Advanced Mobile Robotics—Proceedings*, Lund University Cognitive Studies, vol. 86, Lund, Sweden, 2001, pp. 107–113.
- [17] S.M. Kosslyn, M. Behrmann, M. Jeannerod, The cognitive neuroscience of mental imagery, *Neurophysiologia* 33 (11) (1995) 1335–1344.
- [18] D.N. Lee, J.A.I. Thompson, Vision in action: the control of locomotion, in: D. Ingle, M.A. Goodale, R.J.W. Mansfield (Eds.), *Analysis of Visual Behavior*, MIT Press, Cambridge, MA, 1982.
- [19] F. Linåker, L. Niklasson, Extraction and inversion of abstract sensory flow representations, in: *Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, vol. 6, MIT Press, Cambridge, MA, 2000, pp. 199–208.
- [20] L. Meeden, An incremental approach to developing intelligent neural network controllers for robots, *IEEE Trans. Systems Man Cybernet. Part B: Cybernetics* 26 (3) (1996) 474–485.
- [21] L. Meeden, G. McGraw, D. Blank, Emergent control and planning in an autonomous vehicle, in: *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, NJ, 1993, pp. 735–740.
- [22] R.C. Miall, Sequences of sensory predictions, *Behav. Brain Sci.* 20 (1997) 258–259.
- [23] R.C. Miall, D.M. Wolpert, Forward models for physiological motor control, *Neural Networks* 9 (8) (1996) 1265–1279.
- [24] O. Miglino, H.H. Lund, S. Nolfi, Evolving mobile robots in simulated and real environments, *Artificial Life* 2 (4) (1995) 417.
- [25] F. Mondada, E. Franzi, P. Ienne, Mobile robot miniaturization: a tool for investigating in control algorithms, in: *Third International Symposium on Experimental Robotics*, Kyoto, Japan, 1993.
- [26] S. Nolfi, Using emergent modularity to develop control systems for mobile robots, *Adaptive Behav.* 5 (3–4) (1997) 343.

- [27] S. Nolfi, D. Floreano, *Evolutionary robotics—The Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press, Cambridge, MA, 2000.
- [28] S. Nolfi, J. Tani, Extracting regularities in space and time through a cascade of prediction networks: the case of a mobile robot navigating in a structured environment, *Connection Science* 11 (2) (1999) 125–148.
- [29] M. Peschl, The representational relation between environmental structures and neural systems: autonomy and environmental dependency in neural knowledge representation, *Nonlinear Dyn. Psychol. Life Sci.* 1 (2) (1997) 99–121.
- [30] R. Pfeifer, C. Scheier, *Understanding Intelligence*, MIT Press, Cambridge, MA, 1999.
- [31] L.A. Suchman, *Plans and Situated Action: the Problem of Human-Machine Communication*, Cambridge University Press, New York, 1987.
- [32] J. Tani, S. Nolfi, Learning to perceive the world as articulated: an approach for hierarchical learning in sensory–motor systems, in: R. Pfeifer, B. Blumberg, J.-A. Meyer, S. Wilson (Eds.), *From Animals to Animats*, vol. 5, MIT Press, Cambridge, MA, 1998.
- [33] J. Tani, S. Nolfi, Learning to perceive the world as articulated: an approach for hierarchical learning in sensory–motor systems, *Neural Networks* 12 (1999) 1131–1141.
- [34] W.T. Thach, What is the role of the cerebellum in motor learning and cognition?, *Trends Cognitive Sci.* 2 (9) (1998) 331–337.
- [35] E.C. Tolman, Cognitive maps in rats and men, *The Psychol. Rev.* 55 (4) (1948) 189–208.
- [36] F.J. Varela, E. Thompson, E. Rosch, *The Embodied Mind Cognitive Science and Human Experience*, MIT Press, Cambridge, MA, 1991.
- [37] D.M. Wolpert, R.C. Miall, M. Kawato, Internal models in the cerebellum, *Trends Cognitive Sci.* 2 (9) (1998) 338–347.
- [38] T. Ziemke, The construction of ‘reality’ in the robot, *Found. Sci.* 6 (1) (2001) 163–233.
- [39] T. Ziemke, D.-A. Jirnehed, G. Hesslow, Blind adaptive behavior based on internal simulation of perception, Technical Report HS-IDA-TR-02-001, Department of Computer Science (School of Humanities and Informatics), University of Skövde, Sweden, 2002, available online at: http://www.ida.his.se/ida/research/tech_reports/tech_reports2002.html.



Tom Ziemke is Professor of Cognitive Science in the School of Humanities and Informatics at the University of Skövde, Sweden. He received his doctorate from the University of Sheffield, UK, with a thesis on “Situated Neuro-Robotics and Interactive Cognition”. His main research interests concern theories and robotic models of embodied/situated/distributed cognition and agent-environment interaction. He is Editor of the journal *Connection Science* and associate editor of the journal *New Ideas in Psychology* (from 2005).



Dan-Anders Jirnehed is a Ph.D. student in the Department of Biomedical Sciences, Lund University, Sweden. He received a masters degree in cognitive science from Linköping University and after that worked as a research assistant at the University of Skövde. His current research focuses on cerebellar microcomplex dynamics in classical conditioning.



Germund Hesslow is Professor of Neuroscience in the Department of Physiological Sciences, Lund University, Sweden. He holds two doctoral degrees in philosophy and physiology, both from Lund University. His main research interest is behavior and the physiological mechanisms behind it. Much of his research is focused on the physiology of the cerebellum and the mechanisms underlying classical conditioning, but he has also worked in the philosophy of science, mainly causality, and in the philosophy of mind.