UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Representation as Internal Simulation: A Minimalistic Robotic Model

Permalink
https://escholarship.org/uc/item/2363n7fw

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN
1069-7977

Authors
Morse, Anthony
Svensson, Henrik
Ziemke, Tom

Publication Date
2009

Peer reviewed

Representation as Internal Simulation: A Minimalistic Robotic Model

Henrik Svensson (henrik.svensson@his.se)
Anthony Morse (anthony.morse@his.se)
Tom Ziemke (tom.ziemke@his.se)
Informatics Research Center, Cognition & Interaction Lab, University of Skövde
PO. Box 408, SE-541 28 Skövde, Sweden

Abstract

Embodied cognition theorists have in recent years proposed that a cognitive agent's "representations" or "inner world" can at least partly be constituted by internal emulations or simulations of its sensorimotor interaction with the world, i.e. covert perception and action. This paper recapitulates some of the empirical evidence, distinguishes between implicit, internal and external anticipation, and discusses possible neural correlates. Furthermore a robotic neurocomputational model of external anticipation is presented and analyzed.

Keywords: Anticipation; Emulation theory; Inner world; Simulation theory.

Introduction

Cognitive scientists have for a long time considered some kind of internal model a conditio sine qua non for (higher-level) cognition. The possible advantage of having internal models was described early on by Kenneth Craik (1943)

> If the organism carries a "small scale model" of external reality and of its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise… (Craik, 1943, p. 61)

Traditionally much artificial intelligence (AI) research took this quite literally and equipped robots with internal maps as analogs of the "small scale models" in our heads, which explicitly labeled the topology and content of the world with conceptual/semantic descriptions. Simulation (or emulation) theories have approached the same problem from a biological and embodied perspective taking much inspiration from neuroscience (Cotterill, 2001; Grush, 2004; Hesslow, 2002; Jeannerod, 2001), which have inspired new robot models in AI (e.g. Ziemke, Jirenhed, & Hesslow, 2005). These theories view thinking as mental simulations rooted in perception and action. Simulation theories claim that *thinking is the (predictive) coupling of covert actions and perceptions, i.e., reactivations of perceptions and actions* (cf. e.g. Hesslow, 2002, see also Figure 1). That means, covert actions can internally generate sensory feedback (a covert perception) without actually having to perform the corresponding action in the environment and a covert perception may generate a covert or overt action. Coupling covert actions and perceptions may lead to long chains of simulated interactions.

The existence of covert actions and perceptions in cognition is supported by a wide range of empirical research (Svensson, 2007; Svensson, Lindblom, & Ziemke, 2007). For example, motor imagery experiments have shown that mentally (covertly) simulating an action is similar to overt action in the following aspects: execution time including the reproduction of Fitt's law and isochrony (e.g. Grush, 2004; Guillot & Collet, 2005), physiological effects (e.g. Yue & Cole, 1992), PET, fMRI, and TMS responses (e.g. Jeannerod, 2001). Although the reactivation of actions is most pronounced in motor imagery, similar effects have been found in many other cognitive abilities (for a review see Svensson et al., 2007). Simulation theories address many cognitive functions, but this paper focuses on the general ability to construct an inner world that can be used instead of the real world, i.e., the ability to construct mental simulations of previous and future agent-environment interactions. We argue that mental simulations consist of the neural processes used for *implicit*, *internal*, and *external* anticipation. The view of simulation presented here incorporates ideas from a number of other models of simulation theories (Baldassarre, 2002; Cotterill, 2001; Hesslow, 2002; Shanahan, 2006). Furthermore, we propose a novel architecture, based on Echo State Networks (ESNs), for computationally modeling the external anticipation aspect of simulation processes in a simple mobile robot and present some initial results.

Neural pathways of simulation

As simulation processes are likely to be part of many different cognitive functions their implementation in the brain requires the recruitment of several areas and neural mechanisms. In this section, we distinguish three different functional parts of the simulation process, partly based on their neural substrate: implicit anticipation, internal anticipation, and external anticipation. Implicit anticipation refers to sensorimotor anticipations that are only behaviorally instantiated, internal anticipations involve predictions of future bodily states, and external anticipation refers to predictions of exteroceptive inputs.

Implicit Anticipation

An implicitly anticipatory animal is capable of anticipatory behavior without having a model of its own body, cognitive system, or the external environment (Butz, Sigaud, & Gérard, 2003). In the mammalian brain, the cerebellum learns implicit predictions in the form of associations between a stimulus and action (e.g. Downing, in press). Several models suggest that also the basal ganglia learn to match actions with sensory situations on the basis of the expected future rewards of performing that action. Even though the animal learns to behave as if it has access to

what the reward is, the learning process does not establish representations of what the reward is but only behavioral programs (Downing, in press).

An animal capable of simulation processes is likely to reuse the mechanisms for implicit anticipation. The need to select actions is just as important for a simulation process as it is when interacting with the environment (cf. Humphries & Gurney, 2002). It is not viable to simulate every possible action and its consequence. The same neural circuitry might also control whether or not an action should be executed or merely simulated (Cotterill, 2001). The idea that simulation processes are dependent on implicit anticipations produced by cerebellar and basal ganglia circuits is consistent with neuroimaging studies which show activation of both the cerebellum and basal ganglia in motor imagery (Jeannerod, 2001). Single-neuron studies have also provided evidence for predictively activated motor representations in the brain (Cisek & Kalaska, 2004). They could show by recording neurons in the dorsal premotor cortex of monkeys that these neurons fired both during the performance of the task and in anticipation of the task. Furthermore, the predictive and performance related activity was strikingly similar. Also, canonical neurons could be seen as implementing a prediction of an action (cf. Miall, 2003).

Internal Anticipation

Internal anticipation means that the animal has a model of its own body. Thus, the animal can use predictions of the future state of its body to influence its behavior. It is widely believed that, in mammalian brains, the cerebellum implements an internal model of the bodily consequences of performing a certain action in a certain situation (i.e., a forward model) (e.g. Doya, 1999; Wolpert, Miall, & Kawato, 1998). The model can be used to produce predictions in the form of a state representation or in the form of sensory information (Wolpert et al., 1998). The motor system needs to act on this predictive knowledge of future states to, for example, compensate for feedback delays (Wolpert et al., 1998) Furthermore, it has been suggested that the sensory predictions can be feed back to the cerebellum to generate further state or sensory predictions, perhaps in cooperation with the basal ganglia (e.g. Doya, 1999; Grush, 2004; Wolpert et al., 1998). Thus, mental simulations that are about, or need information about, the details of performing an action is likely to be mediated by simulation processes in a cerebellar-cortico-basal ganglia circuit. For example, this would be the case in athletes engaging motor imagery of their sports.

External Anticipation

An animal capable of external anticipation is an animal with the ability to generate predictions that describe the state of the world. Suggestive evidence has been found for external anticipation in rats. In one type of experiment, a rat is first subjected to a learning scheme where one action (e.g. lever pressing) results in one type of reward and the other action (e.g. chain pulling) results in another type of reward.

Subsequently, outside the previous learning context, one of the rewards is made less desirable for the rat. The result is that the rat chooses the action associated with the now more desirable effect, which suggests that the rat is able to predict the resulting reward stimulus of performing lever pressing and chain pulling (cf. e.g. Balleine & Dickinson, 1998; J. Hoffmann, 2003). The existence of response-effect predictions has been found in several experiments with humans as well (reviewed in detail in Kunde, Elsner, & Kiesel, 2007). Furthermore, some of these experiments suggest that the effects are in the form of covert perceptions as suggested by simulation theories (Kiesel & Hoffmann, 2004; Kunde, 2003 in Kunde et al., 2007).

External anticipations are likely to involve the neocortex and the hippocampus (e.g. Cotterill, 2001; Downing, in press; Hawkins & Blakeslee, 2004). The hierarchical structure of the motor and sensory cortices and the reciprocal connections between them at various levels (Fuster, 1997, 2004) suggest the possibility of the cortex implementing both predictions from motor to sensory activity and the reverse. Cotterill (2001) argued that premotor areas send information back to the sensory cortex by way of axon collaterals. He further noted that "there are three such efference copy routes…One goes directly, another passes through the anterior cingulate, and the third goes via the thalamic ILN" (p. 22). Efference copy routes might indeed be a ubiquitous property throughout the sensorimotor hierarchy (Hesslow, personal communication cf. Fuster, 2004). Gomez et al. (2004) have, based on their own experiments with the contingent negative variation and other corroborating studies, suggested that there exists an attentional-anticipatory system that "include[s] not only the frequently described prefrontal, SMA, and primary motor cortices, but posterior parietal cortex, cingular cortex, and pulvinar thalamic nuclei too. The neural substrate of the perceptual domain is not so well-described, but, of course, the participation of primary sensory areas has been hypothesized" (p.67). Gomez et al.'s studies do not, however, show decisively how the preparatory activity of the sensory cortex is elicited, i.e., directly via the sensory cue or indirectly by preparatory activity of the motor related cortices. The study by Kastner et al. (1999) shows influence from frontal and parietal areas on extrastriate cortex during covert attention shifts, suggesting the possibility of motor areas modulating the activity of sensory areas in an anticipatory manner.

The mirror neuron system might also be implicated in external anticipations. Projections from area F5 of the ventral premotor cortex, through area PF, and to STS, essentially "convert[s] the motor plan back into a predicted visual representation (a sensory outcome of the action)" (Miall, 2003). However, it should also be pointed out that Miall (2003) argued similar transformations might be implemented by pathways incorporating the cerebellum.

The likely existence of explicit predictions and the observation of reactivations of covert perceptions

throughout the sensory cortices suggest that they have an important function in mental simulations.

Experiment

The main purpose of the experiment in this paper is to test whether a model of external anticipations based on an echo state network (ESN), could be implemented in a mobile robot and be made to generate "mental" simulations. Although, the robot model at this point does not map entirely to the possible simulation pathways of the neocortex described above, the ESN has been shown to exhibit some of the computational properties of the neocortex. Thus, the focus is here to come up with an existence proof rather than a full implementation and account of the view of simulation outlined above.

Related Work

Our robot experiment takes inspiration from the previous robot experiments performed in our lab (Jirenhed, Hesslow, & Ziemke, 2001; Ziemke et al., 2005) as well as some other neurobiologically inspired models (Baldassarre, 2002; Gross, Heinze, Seiler, & Stephan, 1999; Shanahan, 2006; Tani & Nolfi, 1999). These models have shown that it is possible to produce accurate predictions, which can be used to establish simulations, and they outperform reactive models. Functionally, our and the previous models tries to learn a forward model, i.e., sensory and motor states (or context) are used as input and the architecture outputs future sensory states (and motor states). Previously, both supervised learning and global reinforcement learning algorithms (evolutionary learning) have been used to establish the predictions.

Even though there are successful models of simulations in various tasks, producing simulations is not trivial. For example, Jirenhed et al. (2001) were not able to generate extended simulation behaviors in h- and T-shaped mazes and Ziemke et al. (2005) were not able to generate successful simulations based only on proximity sensors of a simple robot in a simple environment. Tani and Nolfi (1999), on the other hand, achieved almost perfect prediction performance using a recurrent version of the adaptive mixture of local experts approach. Baldassarre's (2002) model with a similar type of architecture produced higher prediction errors, but was still able to perform well. On the other hand, prediction errors can be problematic. Jirenhed et al. (2001) found that their robot failed to predict sensors that were seldom active, which could explain the robot's failure to develop extended simulations.

Experimental Setup

The task of the robot in our setup is to navigate blindly through a square shaped environment consisting of four equally long corridors (see Figure 2, left), similar to the world used by (Ziemke et al., 2005). In the learning phase, the robot is controlled by a simple behavior program that allows the robot to move along the corridors counterclockwise. During this time the robot learns to predict the sensory inputs and the motor outputs at the next time step. In the test phase, the robot, instead of receiving inputs from the environment uses its learnt predictions.
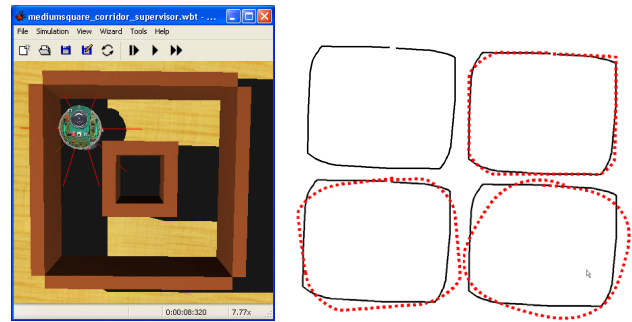

Figure 2. (Left) Webots simulator. (Right) Motion trajectories of the robot. See text for details

The robot experiment was performed using a simulated E-puck robot (www.e-puck.org) running on a robot simulation platform (Webots 5.10.0). The e-puck has a circular body with a diameter of 70 mm, and is equipped with 8 infrared proximity sensors around its body with a range of approximately 6 cm. It has two motors, which independently control the two wheels, one to each side. The wheels can rotate forward and backward independently, such that the robot can turn on the spot if they rotate in opposite directions.

Architecture

The architecture of the robot controller is based on an ESN. ESNs and liquid state machines (LSM) introduce new computational features that may allow us to model even more complex tasks. The ESN provides a somewhat biologically plausible model for external anticipations in mental simulations, which is described in this section.

Although its exact make up and function is debated in neuroscience, the neocortex consists of units called microcolumns (Hawkins & Blakeslee, 2004; Mountcastle, 1997). Our aim here is not to model the specific circuitry and make up of cortical micro-columns, but rather to model the following four properties of real cortical micro-columns (Markram, Wang, & Tsodyks, 1998). (1) Cortical micro-columns are observed to be non-chaotic. (2) Cortical micro-columns do not display stable attractor dynamics (their activity quickly decays on cessation of input). (3) Input size to cortical micro-columns is sparse relative to the size of the micro-column. (4) The state space achieved by an active 'firing' micro-column is large and sensitive to its input.

The reason we focus on these specific properties of biological cortical micro-columns is that they turn out to have very useful computational implications by making highly non-linear features linearly separable much as a kernel warping function does (cf. Cristianini & Shawe-Taylor, 1999), and also by acting as a fading memory (Maass, Natschläger, & Markram, 2002).

As a starting point then we use the ESN to model the role of the cortex in simulation processes. The ESN that was

used to perform the robot experiment reported here is derived from a random weights matrix populated with 20% connectivity and adjusted so as to have a spectral radius < 1, i.e. $|\lambda_{\max}| < 1$, where $\lambda_{\max}$ is the eigenvalue of w which has the largest absolute value, thus the ESN is uniquely controlled by the input and the effect of initial states disappears. By observation this would also seem to be true of cortical micro-columns. The ESN reservoir is cycled according to standard DTRNN equations: $a_i = \Sigma y_j\, w_{ij} + i_i$ where neuron output is computed by: $y_i = \tanh(a_i)$. and input to the reservoir is provided via weights generated by the same method as the ESN weights, except a higher connectivity (80%).

Controller The architecture controlling the robot consists of an input and output layer of perceptrons and a hidden layer consisting of the ESN just described. The input layer receives eight inputs from the robot's proximity sensors, as well as two motor inputs (normalized to values between 0 and 1).

The output layer consist of 10 units that use the sigmoid activation function, i.e., the outputs are between 0 and 1, which for the two motor units means that 0 corresponds to full speed backward rotation, 0.5 corresponds to no motion, and 1 corresponds to full speed forward rotation. The output layer is trained online by a simple supervised learning scheme using a standard delta rule: $\Delta w_i = \alpha(a^P(1 - a^P))(t^P - a^P)x_i^P$. At every time step the motor units of the output layer is feed back to the ESN and used as additional inputs.

In the testing phase, there are two modes of behavior in which activations of the output layer units influence the behavior of the robot in different ways. In *blind sensory* mode, activations of the 8 output units that correspond to sensor input are copied to the input layer and used in place of the robot's sensory input. In this mode, the robots motor output is then generated by the predefined motor program. In *blind all* mode, both the sensor and motor output units are copied back to the input layer and used to control the behavior of the robot.

Results and Analysis

Figure 2 (right) shows a plot of the path traversed by the robot during four different conditions, normal (top left), blind sensory with walls (top right) and with the surrounding walls removed,(bottom left) and blind all with no walls (bottom right). The behavior of the robot in the blind modes (red dotted lines) closely corresponds to the behavior of the normal mode (black full line). The slightly skewed path seen in the blind behaviors without the restricting walls (lower half of Fig. 2, right) is partly due to the predefined motor program which allows the robot to slide against the walls during the turns, which means that while the motor program dictates a certain path it becomes restricted by the surrounding walls.

Figure 2 (right) shows only one lap for each blind mode, but the simulations are able to carry on for more than one lap, although it is worth noting that a single lap already

consists of simulations that are 200 time steps long. The graph in Figure 3 shows the difference between the internally generated input and the actual input summed over all ten inputs (8 sensory and 2 motor) during 1300 time steps. The robot is able to generate relatively accurate simulations for roughly 800 time steps. However, as seen in the Figure 3 the difference between the real and simulated increases rapidly after that.
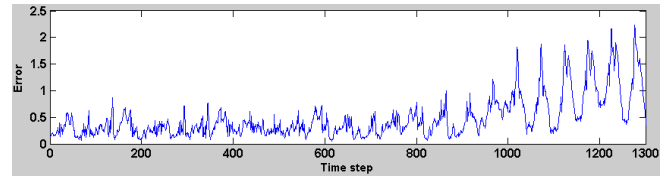

Figure 3. Difference between normal and blind all mode for all output units at each time step

Hoffmann & Möller (2004) also reported an accumulation of error as the chains of predictions increase in length, but that the errors still are very small. Even though the increase in difference might suggest that the robot stops simulating the inputs, it is more accurately described as the simulation going out of phase. More specifically, it is due to the timing of the corners getting out of sync with the real world. This is illustrated in Figure 4, which plots the two front sensors. The peaks indicates that the robot approaches a corner and we can see that around the 800[th] time step the peaks of the simulated and real front sensors starts to go out of sync. However, it is noticeable that the simulated sensors continue predicting a wall even though they do not get the timing right. The average one-step predictions errors were quite low (0.02/output unit).
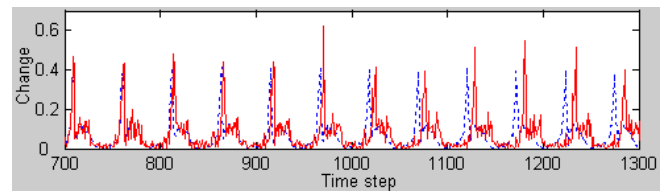

Figure 4. Front sensors: normal (red) and blind all mode (blue, dashed).

Discussion and Conclusion

There are at least two criteria which can be used to asses whether our model generates simulation processes. Internally generated states should be similar to the states observed in normal interactions. A weaker criterion consistent with the reactivation hypothesis is the reuse of the same mechanisms or the same resources used for online interactions in simulations. The predicted sensory and motor states as well as the internal states of the ESN (not illustrated, due to space restrictions) have been shown to closely correspond to the "real" ones. Thus, according to the first criterion the model models simulation processes. The model also uses the same resources since the same neural networks and weight matrices are used in covert and overt operation modes, except for the blind sensory mode where

the external behavior rules are used to generate the motor output.

The model implements external anticipation since (1) ESNs model aspects of neocortical processing, which, as discussed above, is particularly involved in such predictions and (2) it is used to predict the sensory input generated by the environment. The generation of overt and covert actions was achieved by a hardwired behavior program or the ESN (in blind sensory mode). Thus, the model can be said to in this mode to minimally model implicit anticipation in the form of simple stimulus-response associations mediated by cerebellar circuits. However, it lacks the learning mechanisms associated with cerebellar and basal ganglia mediated implicit anticipation. To more fully incorporate implicit anticipation, the current architecture could be extended to also include a separate neural network for learning and generating overt and covert actions based on reinforcement learning principles (e.g. Baldassarre, 2002). The ability to model the distinction between external and bodily prediction is, however, limited when using a robot with a very simple morphology such as the E-puck robot.

Ziemke et al. (2005), who used a similar kind of simulated robot, world, and task, but used a computationally simpler architecture were not able to generate successful simulations based only on proximity sensors of a simple robot in a simple environment. By using another sensor they were able to establish blind navigation, but curiously, this time the generated sensory predictions did not correspond to the real sensory input, i.e., it did not actually simulate the sensory input. Still, it was able to travel the world blindly. A possible explanation of why the robot in the current robot experiment managed to internally generate sensory input that closely matched the input from the proximity sensors is the fading memory of the echo state network. The particular 150 neuron ESN reported here has a "memory trace" of roughly 20 time steps, which exceeds the number of time steps without sensory input from a corner wall.

A major issue for robotic models of simulation theories has been the ability and usefulness of abstractions, i.e., changing the level of granularity of the simulations. This has earlier been done by a number of different methods, such as chunking similar sensory input to some form of concept (e.g. Holland & Goodman, 2003; Stening, Jakobsson, & Ziemke, 2005) or simply train to predict a situation with a sensory situations x time steps ahead. The current model did not change the level of granularity, but the fading memory of the ESN allows the robot to have a more holistic view of the situation it is currently in, that spans backwards in time. Thus, although it does not allow the simulations to occur on a faster time scale, by predicting several time steps ahead in each simulation step, it might embody some of the benefits of a hierarchical system in which higher levels of abstraction influences the predictions at lower levels. This, and the issue of abstraction in mental simulations, remains to be investigated in more detail.

Although the *task* is relatively simple it is comparable to mental imagery, in particular motor imagery of cyclic movements such as walking and paddling, which have been shown to closely match the time the take to actually perform them (cf. Guillot & Collet, 2005). However, it has also been shown that a number of different variables affect the speed at which an action is imagined. For example, a higher level of expertise as well as simpler tasks leads imagery processes with the same speed as performing the action (Guillot & Collet, 2005). Thus, given that we investigated a simple task involving cyclic actions and extensive training the robot should and were able to "imagine" the "walk" at a similar speed as the actual "walk". Prior robot studies of simulation theory have put little emphasis on matching the speed of mental imagery processes except for speeding up the simulations by decreasing the level of detail. While it is likely that the level of detail decreases when a task is imagined faster than it is performed, the actual speed of the imagery process might also be increased. Although not analyzed here in much detail, we observed in many of our trials that the simulation process speed increased without decreasing the level of detail of the simulations. Future studies should investigate how the speed of the simulations is affected by different parameters, such as task complexity.

As indicated above, the robot experiment described here uses only a very simple square-shaped environment and only serves as a first proof-of-concept of the architecture. It should be noted though that the robot experiment has also been replicated successfully with more complex environments, e.g. with corridors of different lengths and L-shaped environment, where not all turns are in the same direction. Due to space restrictions, these results will need to be documented in detail elsewhere (e.g. Svensson, under preparation).

References

Baldassarre, G. (2002). A biologically plausible model of human planning based on neural networks and Dyna-PI models. *Proceedings of the Workshop on Adaptive Behaviour in Anticipatory Learning Systems (ABiALS-2002)*, 40-60.

Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action. *Neuropharmacology, 37*(4-5), 407-419.

Butz, M. V., Sigaud, O., & Gérard, P. (2003). Internal models and anticipations in adaptive learning systems. In *Anticipatory Behavior in Adaptive Learning Systems* (pp. 253-273).

Cisek, P., & Kalaska, J., F. (2004). Neural correlates of mental rehearsal in dorsal premotor cortex. *Nature, 431*, 993-996.

Cotterill, R. M. J. (2001). Cooperation of the basal ganglia, cerebellum, sensory cerebrum and hippocampus. *Progress in Neurobiology, 64*(1), 1-33.

Craik, K. (1943). *The nature of explanation*. Cambridge: Cambridge University Press.

Cristianini, N., & Shawe-Taylor, J. (1999). *An introduction to support Vector Machines*. New York: Cambridge University Press.

Downing, K. L. (in press). Predictive models in the brain. *Connection Science*.

Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks, 12*, 961-974.

Fuster, J. M. (1997). Network memory. *Trends in Neurosciences, 20*(10), 451-459.

Fuster, J. M. (2004). Upper processing stages of the perception-action cycle. *Trends in Cognitive Sciences, 8*(4), 143-145.

Gomez, C. M., Fernandez, A., Maestu, F., Amo, C., Gonzalez-Rosa, J. J., Vaquero, E., et al. (2004). Task-specific sensory and motor preparatory activation revealed by contingent magnetic variation. *Cognitive Brain Research, 21*(1), 59-68.

Gross, H.-M., Heinze, A., Seiler, T., & Stephan, V. (1999). Generative character of perception. *Neural Networks, 12*, 1101–1129.

Grush, R. (2004). The emulation theory of representation. *Behavioral and Brain Sciences, 27*(03), 377-396.

Guillot, A., & Collet, C. (2005). Duration of mentally simulated movement. *Journal of Motor Behavior, 37*(1), 10-20.

Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. New York: Henry Holt.

Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences, 6*(6), 242-247.

Hoffmann, H., & Möller, R. (2004). Action selection and mental transformation based on a chain of forward models. In S. Schaal, A. Ijspeert, S. Vijayakumar, J. C. T. Hallam & J. A. Meyer (Eds.), *Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior* (pp. 213–222). Cambridge, MA.: MIT Press.

Hoffmann, J. (2003). Anticipatory Behavioral Control. *Lecture Notes in Computer Science, 2684*, 44-65.

Holland, O., & Goodman, R. (2003). Robots with internal models - A route to machine consciousness? *Journal of Consciousness Studies, 10*(4-5), 77-109.

Humphries, M. D., & Gurney, K. N. (2002). The role of intra-thalamic and thalamocortical circuits in action selection. *Network: Computation in Neural Systems, 13*(1), 131-156.

Jeannerod, M. (2001). Neural Simulation of Action. *NeuroImage, 14*(1), 103-109.

Jirenhed, D. A., Hesslow, G., & Ziemke, T. (2001). *Exploring internal simulation of perception in mobile robots*. Paper presented at the 2001 Fourth European Workshop on Advanced Mobile Robotics Lund.

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron, 22*(4), 751-761.

Kunde, W., Elsner, K., & Kiesel, A. (2007). No anticipation–no action. *Cognitive Processing, 8*(2), 71-78.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states. *Neural Computation, 14*(11), 2531-2560.

Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences, 95*(9), 5323-5328.

Miall, R. C. (2003). Connecting mirror neurons and forward models. *NeuroReport, 14*(17), 2135.

Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain, 120*(4), 701-722.

Shanahan, M. (2006). A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition, 15*(2), 433-449.

Stening, J., Jakobsson, H., & Ziemke, T. (2005). Imagination and abstraction of sensorimotor flow. Paper presented at the *Symposium on Next Generation Approaches to Machine Consciousness - Imagination, Development, Intersubjectivity and Embodiment*.

Svensson, H. (2007). *Embodied simulation as off-line representation*. Unpublished Licentiate Thesis, Linköping.

Svensson, H. (under preparation). *Embodied simulation: theory, evidence, models. (working title)*. Doctoral dissertation to be defended in 2009

Svensson, H., Lindblom, J., & Ziemke, T. (2007). Making sense of embodiment. In T. Ziemke, J. Zlatev & R. Frank, M. (Eds.), *Body, language and mind. Volume 1: Embodiment* (pp. 241-270). Berlin: Mouton de Gruyter.

Tani, J., & Nolfi, S. (1999). Learning to perceive the world as articulated. *Neural Networks, 12*(7-8), 1131-1141.

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences, 2*(9), 338-347.

Yue, G., & Cole, K. J. (1992). Strength increases from the motor program. *Journal of Neurophysiology, 67*(5), 1114-1123.

Ziemke, T., Jirenhed, D. A., & Hesslow, G. (2005). Internal simulation of perception. *Neurocomputing, 68*, 85-104.