Towards a Neural Model of Mental Simulation

Matthew Hartley and John Taylor

King's College London, The Strand, London, WC2R 2LS
mhartley@mth.kcl.ac.uk, John.G.Taylor@kcl.ac.uk

Abstract. We develop a general neural network-based architecture for the process of mental simulation, initially treated at a somewhat abstract level. To develop the theory further it is shown how the theory can handle observational learning as a specific form of mental simulation: simulations are presented of simple paradigms and results obtained on children undergoing tests on observational learning. Questions of learning and other aspects are treated in a discussion section.

1 Introduction

Mental simulation has become an area of interest to a number of branches of research: philosophy, psychology, engineering, brain imaging, military studies and many more. It is of particular relevance in trying to understand how human subjects can empathise with those they meet, and build a .theory of mind. about these others. There are two distinct branches to these studies of mental simulation. There are general analyses of how such mental simulation can help people empathise with others, and develop understanding of what decision-making had occurred, and what range of beliefs it could have been based on [4]. This is to be regarded as .putting yourself in someone else's place.. Such an approach has led to numerous studies of the manner in which beliefs and desires can be involved in such internal simulations. In particular mental simulation has been recognised as central in planning, decision-making, hypothesis generation and testing and in belief revision. On the other hand there has been a flurry of interest associated with mirror neurons and the associated brain processes discovered when people (or monkeys) watch an actor perform some salient action [7]. Initially it was considered that such mirror neurons were in very specialized areas of the monkey (and human) cortex. However more recently it has been realised that a considerable amount (although not all) of those sites active during execution are also active during observation of the same actions of another [6]. Because of this more extended view of mirror neurons we will term neurons active in the paradigms associated with observation of others executing actions as .simulation neurons.

These two approaches are especially different in terms of the cognitive level of processing occurring in the brain of the subject: the first (related to an approach through the "theory of mind") is at a much higher cognitive level than paradigms used to observe simulation neurons. However we suggest that the higher level can be regarded as a more sophisticated version of the lower one. For cognitive

V. Kůrková et al. (Eds.): ICANN 2008, Part II, LNCS 5164, pp. 969–979, 2008.

simulations of the results of actions or of how to achieve certain goals must depend on internal simulations of these actions based on those goals. The further mental simulations associated with decisions between various goals or actions and on beliefs as the bases of such decisions will depend on various long-term memories and decision processes (such as choosing between various courses of action or goals) which are beyond the scope of this paper. However they can be seen only as biases to the mental simulation process itself, so that simulation can be regarded as the place to start analysis of possible appropriate brain architectures.

In this paper we therefore begin an attack on possible brain neural networks involved in mental simulation by starting at the lower level. Thus we consider how goals, as objects or actions, can be mentally simulated by observation of an actor achieving the goals; later the actors actions and goals achievements can be imitated (although the actions may not be exactly as those carried out by the actor). This analysis thus covers more fully the process of observational learning and imitation. We then go on to consider how mental simulation as internally driven could occur with this architecture. This is a further step beyond the externally driven process of observational learning, involving as it does internally created goals and their manipulation leading to reasoning (which should be included as a part of mental simulation). However we will not consider reasoning per se, but only note how it fits into the architecture we are considering.

In the next section a general architecture is presented which we propose as being at the basis of the low-level mental simulation powers of humans. In the following section we present a simulation of this architecture for simple tasks being performed by infants, and relate this to results obtained by colleagues [2]. The paper finishes with a set of conclusions and further work.

2 A General Neural Architecture

We start by extending the architecture used in a simulation of data on observational learning on infants [3] to the more general case of internal mental simulation. The former architecture is shown in figure 1; containing the set of modules:

The extended architecture, shown in figure 2, uses much of this except for the addition of two specific features:

1) The mental simulation loop of figure 1 is expanded to allow for looping through a sequence of actions and states as part of the process of "imagining" the action needed to cause a state to change to another, and the result of the action on the state to generate the next state ahead in time to function as a goal so as to generate a further action to achieve it. Such goal generation and action creation require the use of well-trained forward and inverse models, the training of which we will not discuss in detail here (but see the discussion).

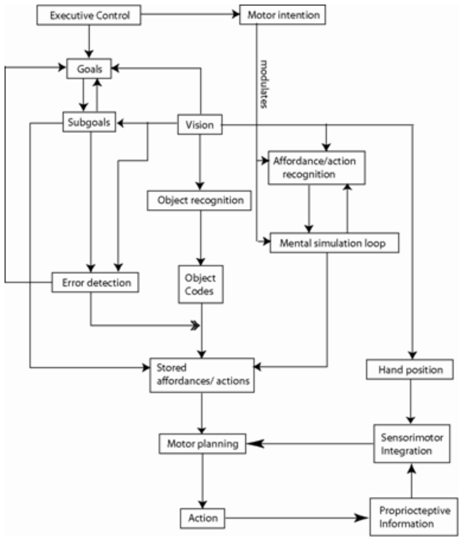

Fig. 1. Neural architecture for observational learning and production of actions based on the learning of affordances

2) The addition of various working memory buffers (possibly with some episodic or longer term powers) so as to hold the results of these computations of sequences of states and actions. Such computations will very likely require the bringing to bear of attention (of both motor and sensory form) so as to enable them to be employed at various levels. Thus if only imitation of the goal of the observed actions is required then only the final state of the generated sequence will be needed, whilst if both final state and sequence of actions is required then either or both of the sequences of internally generated actions and states will need to be held in suitable working memory sites.

We note that the architecture of figure 2 can perform internal simulations totally on its own, provided it has built up a suitable set of memories. It can also be driven by outside inputs to simulate observed actions of another, so perform in an observational learning paradigm. Naturally it can also execute a series of actions with goals set up by the system itself. Thus the architecture can handle all three of the important processes involved in internal simulation processes: internal simulation as self-driven "imagining", internal simulation through observing another in action and internal simulation as part of action planning before and during execution.

It was noted in the introduction that mental simulation is basic to a number of mental activities: planning, decision-making, hypothesis generation and testing and belief revision. Let us consider in general how the architecture of figure 2 can provide a basis for such mental activities. We begin with planning.

Planning is based on the attempt to find a route through a suitable space (of concepts or as physical space itself in the case of planning a trip); a final

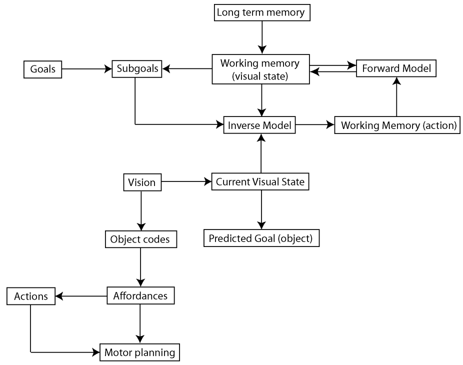

Fig. 2. Figure 2 - Extended mental simulation architecture containing internal models for generating a sequence of actions

goal is given as to where the voyage will end. This can be solved by use of the IMC alone in figure 2, if such a one has been created. For given the goal and the present state of the subject, the IMC can generate a (possibly sequence of) action to achieve the goal. Then the action (sequence), stored in a suitable working memory, therefore provides the plan.

Decision making requires further modules to be added to the architecture of figure 2. One of these is a reward map of goals, so that several of them can be differentiated between by their value. Then any decision module (modelled as a competitive net, for example) would function by having as input a set of goals biased by their reward values, and as output the most highly valued of the goals. A similar mechanism could function to provide a decision between different actions. In this case mental simulation, using suitable FMs of figure 2, would allow assessment of the final goals reached by the various possible actions; choice between these goals by the previous decision mechanism would thus lead to a choice between the actions.

Hypothesis generation and testing can also be handled by the architecture of figure 2 with suitable further modules. We denote here a hypothesis as an assumption conjectured in order to test its empirical or logical implications. The testing we consider under the heading of mental simulation is purely at a mental level, but this is important since it could lead to results already contradicted by experience or which are later discovered to be contradicted by experience.

Consider for example the hypothesis “water is lighter than air”. This would lead us to predict that water floats in the air or even above it. This is clearly contrary to experience, so the original hypothesis must be false. But we can also consider this hypothesis as a counter-factual, and explore its consequences. One of these is that we would expect the seas to be floating in or above the sky above us, a situation which we can visualise (although knowing it is not true). This counterfactual situation can be simulated in the architecture of figure 2 by using the general (learnt) causal law that “if A is lighter than B then A moves above B”. To move A above B from its present position of A being below B (when

A = water, B = air) a suitable IMC would generate the action of moving A from below to above B. This action would then change the state on a working memory site as an FM holding A below B, into A being above B on the site. If we identify imagining the scene of A and B as arising from activity on this working memory module, then we have the process of testing a hypothesis and arriving at an imaginary world in which the seas literally do float in or above the skies.

Thus the process of hypothesis testing may be handled by the architecture of figure 2. However the process of hypothesis generation is outside the scope of this paper (involving a number of more complex processes using long-term memory, salience and possible outside inputs, such as being in a group playing an .as if. game). The same should be said about belief revision, although it can partake of the same processes as hypothesis testing in some situations.

Consider an observer of an actor performing some action towards a goal. Is the observer also undergoing mental simulation as well as observation learning? The answer is that they are not in an autonomous sense since the observer is being mentally stimulated from outside, they obviously are simulating in an externally-driven sense however. Later they can then perform a mental simulation of the situation they had observed as if they were doing it themselves (thus performing mental planning). This aspect emphasises the important part played by memory in the process of autonomous mental simulation, be it short or long-term. Provided the final goal and a suitable IMC is available to the subject then they can call on this memory of the goal to generate the required action (or action sequence) in their minds. Thus the process of observational learning will be very important to expand the repertoire that can be called on for autonomous mental simulation provided that the suitable FM/IMC pairs are created (trained) as part of the observational learning process, or are already available to be used in the new observational learning context.

Before concluding this general discussion, it is important to point out that sensory attention will also have a role to play: it is unlikely that one can perform mental simulation without attending to the ongoing processes in one.s mind. Thus the visual states would be those very likely on a working memory buffer so be available for report. These visual states will thus have been attended to as stimulus inputs to be able to attain the working memory sites for use in mental simulation.

At the same we note that the mental simulation loop itself is very likely at the heart of the motor attention (or intention) control system. Motor attention has been studied over a number of years by brain imaging, such as by Rushworth and colleagues [8], as well as by others. A neural model was proposed for this [9] but suffered the defect that there was no clear link between the motor attention and the visual attention control systems except for the feeding of attended visual input to bias the motor control system. In the architecture of figure 1 we see that there is much better fusion now (as compared to that in the Taylor-Fragopanagos model), in that the motor IMC generates what can be termed the motor attention control signal; that can be used or stored internally in the case of mental simulation but also be sent to lower level motor planning systems if execution is to be performed; that was at the basis of the motor attention model of [9]. Now however we have, in the mental simulation loop of figure 2 a more natural fusion, since the FM allows for the internal action of this motor attention-based action signal to modify the visual state of the system. Thus we can regard the component of the output of the IMC sent to the FM as a corollary discharge of the main signal (to be sent to the lower level motor system to bias a motor plan) under execution.

We also need to turn back to the visual attention system as a further site for mental simulation. If we consider purely spatial rearrangements in one.s mind of various structures in space, such as moving a ball from the floor up to the ceiling, this may be done purely by spatial attention. The visual attention goal to achieve that is clear (the corresponding trajectory in a frontal site such as the frontal eye fields), and the resulting bias of the visual attention IMC would thus produce a movement of the focus of attention vertically upwards. This would have a corollary discharge to achieve this on a suitable buffer working memory (the visuospatial sketchpad of [1]). The corresponding movement would then results, using visual attention throughout. This covert attention movement (with the eyes fixed) breaks the similarity with the motor imagination system above, since the action sequence corresponding to the imagined action could in actuality be taken if inhibition to the execution system was cancelled.

3 Model Details

Unless specified otherwise, all dedicated nodes consist of graded neurons, the membrane potential of which obeys the equation:

$$C\frac{dV}{dt} = g_{leak}(V_{leak} - V) + I_{input} \tag{1}$$

Where C is the capacitance of the neuron, g_{leak} its leak conductance, V_{leak} its equilibrium potential and I its input current. The output of these graded neurons follows the form:

$$I_{out} = \frac{I_{base}}{1 + e^{\frac{V}{V_{scale}}}} \tag{2}$$

Here I_{base} and V_{scale} are constants controlling the maximum neuron output and its scaling. Connections between modules are subject to a time delay of 250ms.

The visual state working memory module consists of dedicated nodes coding for the possible stages of box opening (Closed box, Part open box (latch closed), Part open box (latch open), Open box). Initially the Closed box node is primed by the visual system, later these nodes are activated by the forward model.

The inverse model (IM) takes input from the current visual state and the goal and uses these to produce an action. Actions are again coded as dedicated nodes (Pull cover, Unfasten latch, Open box, Remove reward), and weights are chosen so that the correct action is activated by the combination of appropriate goal and visual state (we discuss how the IM might be trained later).

The action working memory buffer holds representations of the actions generated by the IM so that they can be passed to the FM. These representations are dedicated nodes with recurrent connections to maintain their activity.

The forward model (FM) takes an action provided by the IM and a description of the current state and uses these to calculate the next state that would result from performing the action. In this simple simulation, weights are chosen so that the correct state is generated by connectivity from the action/state inputs.

The visual working memory holds the next state calculated by the forward model and represents these states as dedicated nodes (with the same coding as the current visual state module). If some information about the next state must be filled in from memory (such as the contents of the box), this is done by the bidirectional connection to the memory module.

4 Specific Simulations

We apply the architecture of figure 2 to the paradigm mentioned briefly in (ICANN2008). In this paradigm children open a box which requires several stages of manipulation. In the simplest example, these stages are:

1) Remove a cover by grasping and pulling.
2) Unfasten a latch.
3) Open the box.
4) Remove a reward from a transparent tube inside the box.

To operate the system in full mental simulation mode, we need to activate the goal of opening the box and provide the system with the initial visual state of the closed box. Our inputs to the system are therefore to the goals module where we prime the goal node corresponding to the desire to open the box and extract the reward, and the current visual state of the closed box (it would be possible to perform mental simulation with no external stimulus but then some other method of providing the desire to simulate would be needed, and the initial visual state would have to be provided by memory). The goal node activates a suitable subgoal based on memory (the knowledge that to obtain a reward, the box must be opened), and together, these provide the necessary initial conditions to activate the IMC. The mental simulation "loop". Current state→IMC→Buffer action WM→FM→Buffer state WM→IMC) then supplies the rest of the information with assistance from other modules.

5 Simulation Results

We can examine the output from the nodes representing goals, visual states and actions to look at the time progression of activations. In the first figure we can see the initial stages of simulation . the goal of opening the box and obtaining the reward combined with the visual state of the closed box generate the action of pulling the top of the box, and simulation continues from there:

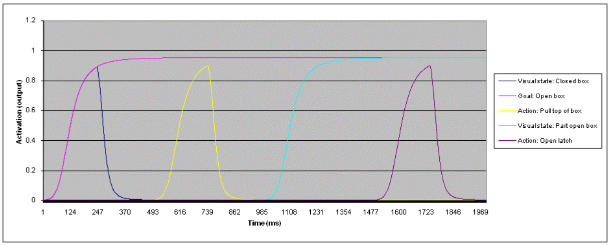

Fig. 3. First 2000ms of simulation

In the second figure we see the final states of the simulation:

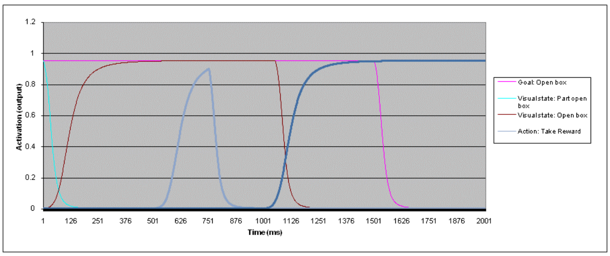

Fig. 4. Final 2000ms of simulation

Another way to represent the system.s operation is to look at the flow of activations of components of the mental simulation. In this we can see how the IMC and FM work through the progression of states needed to simulate the stages of opening the box.

In the diagram, we can also see that the long term memory fills in information about the projected visual states to assist the forward model. After the initial visual state, these later visual states are imagined and held in a buffer visual working memory so they can be acted on.

6 Discussion

One of the important questions about the model's operation is how we could train the inverse and forward models (since in our simulation these are pre-wired). One possible system for training the IM by observation of another.s actions is shown here:

The visual input and goals module prime the inputs to the IM. A buffer working memory holds the visual description of the movements taken by whoever is performing the demonstration. These are then passed through a classifier which extracts an action code based on the actions known to the observer. This

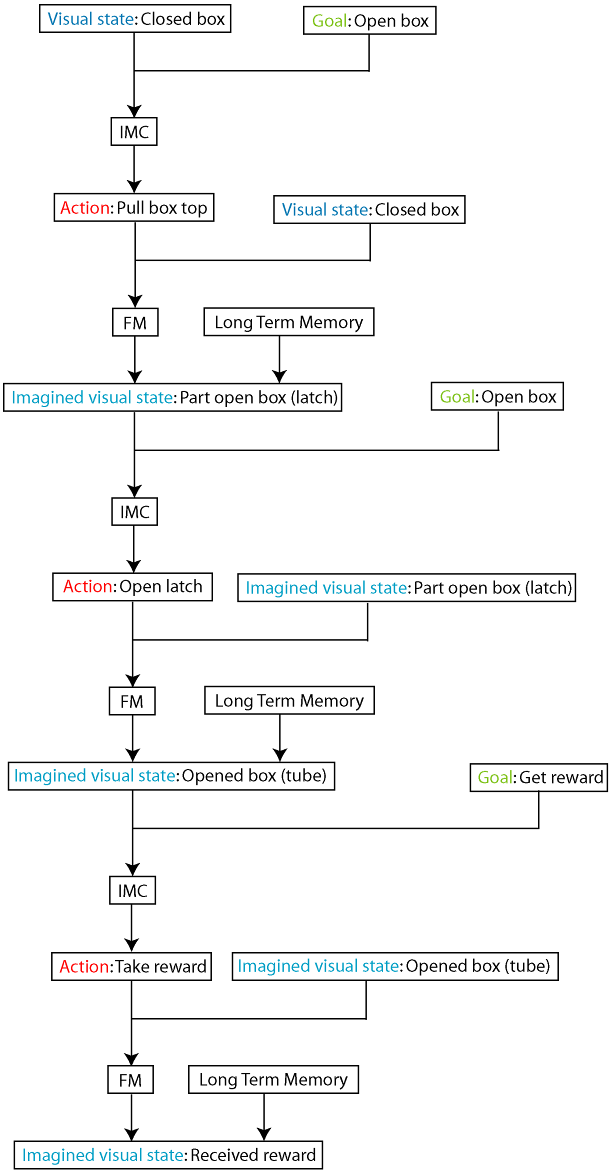

Fig. 5. Flow of activations during model operation showing the process by which the IM and FMC generate the next state/action from the previous state/action

action code primes one of the actions available to the IM, and associative learning between the inputs and output form a suitable connection, such that when presented with the inputs at a later time, the correct action results.

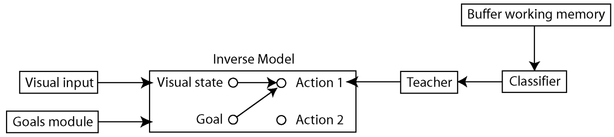

Fig. 6. Proposed architecture for learning the connectivity of the internal model, by priming of the correct action output by a teacher module

This question of training the IM is related to the idea of observational learning, during which performance at motor tasks can be improved by observing others performing those tasks [5]. The mechanism described above, of allowing the observed action to prime part of the IM for associative learning provides a possible mechanism for some parts of observational learning.

We assume that the forward model is based on an existing internal prediction model based on the physics of the world. It operates based on spatial transformations to determine what will result from performing a given action. The long term memory can then fill in needed information to complete the description of the next visual state.

We can also use the model, particularly the action of the mental simulation loop, to make predictions for verification. Since we suggest that each stage of mental simulation involves use of the whole simulation loop, it may be possible to use event related FMRI to detect activations occurring during these different stages (for example, by examining the difference in activations between mentally simulating a two stage task and a three stage task).

7 Conclusions

We have described an architecture for mental simulation based on internal models and extending our existing neural architecture for observational learning. A version of the system using graded neurons was used to simulate a simple mental simulation task based on an infant learning paradigm. We also suggested a method of associative learnin for the inverse model as well as presented some predictions for experimental verification based on the timings of activations to be studied using event related FMRI.

Acknowledgements

Both authors would like to thank the EC for financial support under the MATHESIS Project on Observational Learning, and our colleagues J Nadel et al for communicating their experimental results to us.

References

1. Baddeley, A.D.: Working memory. Oxford University Press, Oxford (1986)
2. Nadel, J., et al.: Studies of children with ASD. Unpublished data (in preparation, 2008)

3. Hartley, M., et al.: Observational versus trial and error effects in a model of an infant learning paradigm. In: ICANN 2008 (submitted, 2008)
4. Goldman, A.I.: Simulating minds. Oxford University Press, Oxford (2006)
5. Hodges, N.J., Williams, A.M.: Current status of observational learning research and the role of demonstrations in sport. J. Sports Sci. 25, 495–496 (2007)
6. Raos, V., Evangeliou, M.N., Savaki, H.E.: Observation of action: grasping with the mind's hand. Neuroimage 23, 193–201 (2004)
7. Rizzolatti, G., Fadiga, L., Gallese, V., Fogassi, L.: Premotor cortex and the recognition of motor actions. Brain Res. Cognit. 3, 131–142 (1996)
8. Rushworth, M.F.S., Ellison, A., Walsh, V.: Complementary localization and lateralization of orienting and motor attention. Nat. Neurosci. 4, 656–661 (2001)
9. Taylor, J.G., Fragopanagos, N.: Simulations of attention control models in sensory and motor paradigms. In: Proceedings of the International Joint Conference on Neural Networks. IEEE, Los Alamitos (2003)