FP6-511931

Mind RACES

from Reactive to Anticipatory Cognitive Embodied Systems

DELIVERABLE D4 (D3.1)

Critical Comparison of Attention, Monitoring and Control Architectures

Due date of deliverable:
31 / 05 / 2005

Actual submission date:
13 / 07 / 2005

Start date of project:
01 / 10 / 2004

Duration:
36 month

Organization name of lead contractor for this deliverable
LUCS

Revision:
3

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	Public	**X**
PP	Restricted to other programmes participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Document identifier:	**DELIVERABLE_WP3_N_1**
Date:	**13/07/2005**
Work package:	**WP3**
Partner(s):	*IDSIA, IST, ISTC-CNR, LUCS, OFAI, NBU, NOZE, UW-COGSCI*
Lead Partner:	*LUCS*
Document status:	*Approved*
Deliverable identifier:	**WP3_D3.1**

Delivery Slip

	Name	Partner	Date	Signature
From	CHRISTIAN BALKENIUS	LUCS	10/06/2005	
Verified	RINO FALCONE	ISTC-CNR	30/06/2005	
Approved by	RINO FALCONE	ISTC-CNR	08/07/2005	

Files

Software Products	User files
MS Word™	D3_1.DOC

Project information

Project acronym:	Mind Races
Project full title:	MIND RACES: from Reactive to Anticipatory Cognitive Embodied Systems
Proposal/Contract no.:	IST-511931
Project Manager:	ISTC_CNR
Name:	Rino Falcone
Address:	CNR-ISTC via S. Martino della Battaglia,44 00185 Rome ITALY
Phone:	+39 06 44 595 253
Fax:	Fax: +39 06 44 595 243
Mobile:	
E-mail	rino.falcone@istc.cnr.it

TABLE OF CONTENTS

Summary

This deliverable provides a critical comparison of a number of current attention, monitoring and control architectures and gives a theoretically analysis and comparison of different architectures to identify their weaknesses and strengths. Models concerned with bottom-up attention, top-down attention, reinforcement control of attention, emulation, schemas, context processing and priming are considered. It is shown that no current model is able to handle all these mechanisms but that there exist many subsystems that could potentially be combined into a more capable system.

The goal of this text is to provide a critical comparison of architectures for attention, monitoring and control. This comparison is put within the overal goal of the MindRACES project which is to (1) incorporate anticipatory functionalities into existing cognitive models; (2) improve anticipatory functionalities of existing cognitive models; (3) integrate different anticipatory functionalities of cognitive models. There will thus be an emphasis on the role of prediction and anticipation in the different models that are compared.

As few current models aim at including all aspects of attention, monitoring and control it is not possible to compare such architectures with each other in a straight forward manner. Instead, the path taken here is to list and compare models that address some of the mechanisms related to attention, monitoring and control. Since most attention systems to date relate to visual attention rather than attention in general, many of the systems will have vision as the target modality.

The current overview will form a base for improving existing models and architectures and will provide a theoretical comparison by identifying the currently best models for attention, epistemic actions, constructive perception with context effects and priming effects. Possible improvements will also be stated for future integration into a more complete architecture.

1 Introduction

The goal of this text is to provide a critical comparison of architectures for attention, monitoring and control. This comparison will be put within the overal goal of the MindRACES project which is to (1) incorporate anticipatory functionalities into existing cognitive models; (2) improve anticipatory functionalities of existing cognitive models; (3) integrate different anticipatory functionalities of cognitive models. There will thus be an emphasis on the role of prediction and anticipation in the different models that are compared.

As few current models aim at including all aspects of attention, monitoring and control it is not possible to compare such architectures with each other in a straight forward manner. Instead, the path taken here will be to list and compare models that address some of the mechanisms related to attention, monitoring and control. Since most attention systems to date relate to visual attention rather than attention in general, many of the systems will have vision as the target modality.

The current overview will form a base for improving existing models and architectures and will provide a theoretical comparison by identifying the currently best models for attention, epistemic actions, constructive perception with context effects and priming effects. Possible improvements will also be stated for future integration into a more complete architecture.

2 Bottom-Up Models of Attention

In bottom-up models of attention, the input which typically consists of an image or an image sequence is processed by one or several operators which are combined to give regions of interest or key-points in the image to which attention is assigned. The different models differ in what operators are used, in how they are combined and in how the regions of interest are selected.

Itti and Koch One of the most influential bottom-up models of visual attention was described by Itti and Koch (2000) (Fig. 1). It has its basis in the feature-integration theory of Treisman and Gelade (1980). The model consists of a number of parallel channels that process the visual input to find regions of high contrast in intensity, color or orientation. The activity in each of the parallel processing stages are added together to form a saliency map which codes positions in the image to which attention will be directed using a winner-take-all mechanism. The model used 7 different features and 6 maps per feature and works on a number of different spatial scales.

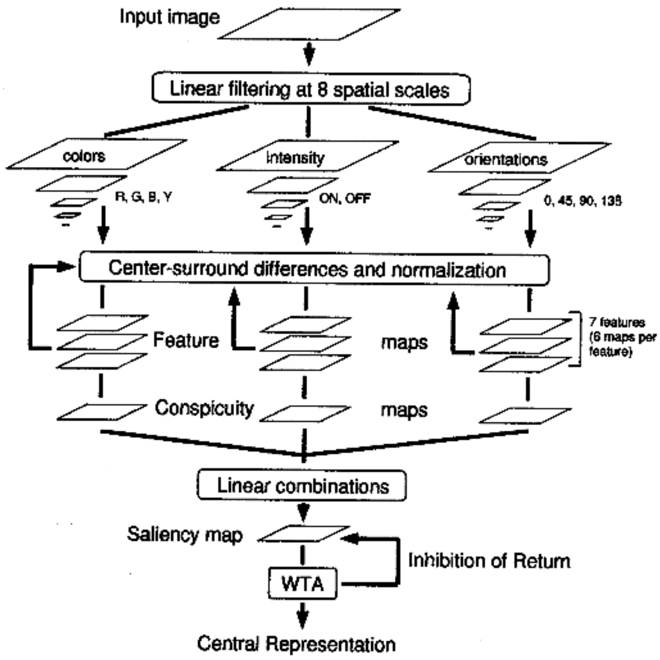

Fig. 1. The bottom-up model of attention described by Itti and Koch (2000).

The main strength of the model is that it is able to produce results similar to those of human subjects in visual search tasks. A less satisfactory aspect of the model is its large demand on computational resources. Although the different filters used claim to parallel those in the human visual system, it is hard to get anywhere near real-time performance on general purpose computing hardware. This currently limits the use of the model for technical applications.

Key-Point Extraction Although not intended as a models of attention, different methods for extraction of interest points (key points) in images play a similar role to the filters in the Itti and Koch model. The Harris detector (Harris and Stephens, 1988) is often used in image processing to quickly find regions of interest in an image. Subsequent processing is performed in these regions. An alternative method which is invariant to scale and rotation was described by Lowe (2004). It uses difference of Gaussians at different scales to extract extrema points which are then tested for stability (Fig. 2).

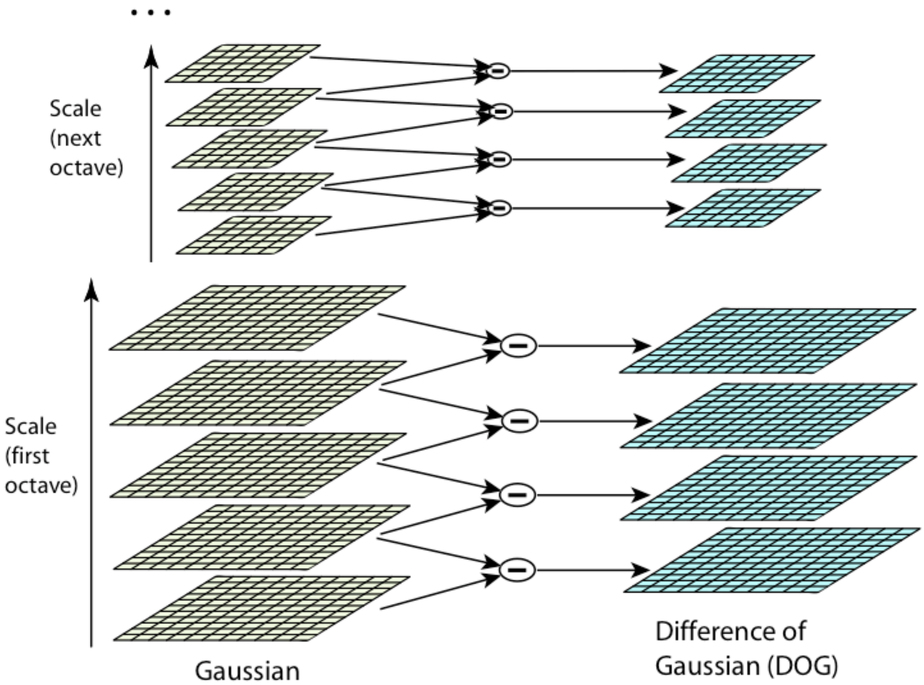

Fig. 2. The calculation of scale-space extrema plays a central role in the key-point extraction scheme described by Lowe (2004).

These type of methods are good candidates for technical applications where real-time performance is required. It is also possible to use even simpler filters to extract oriented contrast at a single scale in many cases (e.g. Balkenius, 1998, Kopp 2004).

Generalized Symmetry Operator An early system for directing attention by bottom-up means used a generalized symmetry operator (Reisfeld, Wolfson and Yehezkel, 1995). It was demonstrated that the operator would find regions in images that naturally attracts attention such as faces, eyes and other symmetrical objects. The model demonstrates that more complex properties such as symmetry can be useful in directing bottom-up attention. It is possible that a bottom-up attention system would be more useful if this type of features could also be detected at a lower level.

Discussion Bottom-up processing is a necessary component of any model of attention. The main question to address in future work is what types of low-level operators are necessary to achieve appropriate performance. Another question is how different preattentive channels should be combined for best performance.

Top-Down Models of Attention An obvious shortcoming of bottom-up attention models is that they do not adapt to the current situation, since they always direct attention to the same visual properties. This limitation has been addressed in various ways by different models by including also top-down influences on attention. Note that most top-down models also includes a bottom-up component which is often the larger part of the system.

Westin et al. Westin et al. (1996) describes an early system using both bottom-up and top-down processes for the control of attention. The bottom-up processing uses normalized convolution to filter the input image. This method has the advantage that it makes it possible to include the certainty of an estimate in the calculations thus avoiding, for example, border effects and allow sparse or heterogeneous sampling of the image. It also allows top-down influences to modulate the bottom-up processing in a simple way by changing the certainty of each measurement or signal.

Balkenius and Hulth Another form of simple top-down influence was described by Balkenius and Hulth (1999). A bottom-up attention system was enhanced with a simple learning ability that would make the system direct attention to the type of stimulus that had previously been rewarded. The system used a simple form of reinforcement tuning to learn the desired target stimulus. The system was limited in that only a single stimulus type could be learned and the system would thus not adapt to different tasks or situations, but the idea of using reinforcement learning to tune a bottom-up attention system can be useful also in more complex models.

Choi, Ban and Lee A minimal form of top-down influence in a basically bottom up model was included in the model of Choi, Ban and Lee (2004). In the model, a human operator can train the system to ignore certain features based on top-down inhibition. The top-down system consists of a modified ART network that receives input from the saliency map and may categorize it as uninteresting and subsequently inhibit it.. The top-down part of the model is very simple as essenitally categorizes the independent components used for the saliency maps. A more useful form of top-down inhibition requires contextual control (cf. Balkenius, 2000).

Corchs and Deco Many models of top-down attention have been inspired by the human neurophsyiology. A typical model of this type is the one proposed by Corchs and Deco (2002). It includes a model of the primary visual cortex as well as visual area 2-4 and an object recognition system corresponding to inferiotemporal cortex (Fig. 3). It also includes spatial coding system that works in parallel to the object recognition pathway resulting in one "what" system and one "where" system. Both systems receives top-down biases from a module corresponding to prefrontal cortex.

The model aims at reproducing data at a neuronal level, and as is usual with this type of model, it is not clear how it could be extended to more realistic inputs and outputs. Nevertheless, it suggests many ideas for the architecture of a biologically inspired attention system. In particular, the disassociation into identification and spatial processing is probably a required component of any attentional system as is the local competition within each module of the system.

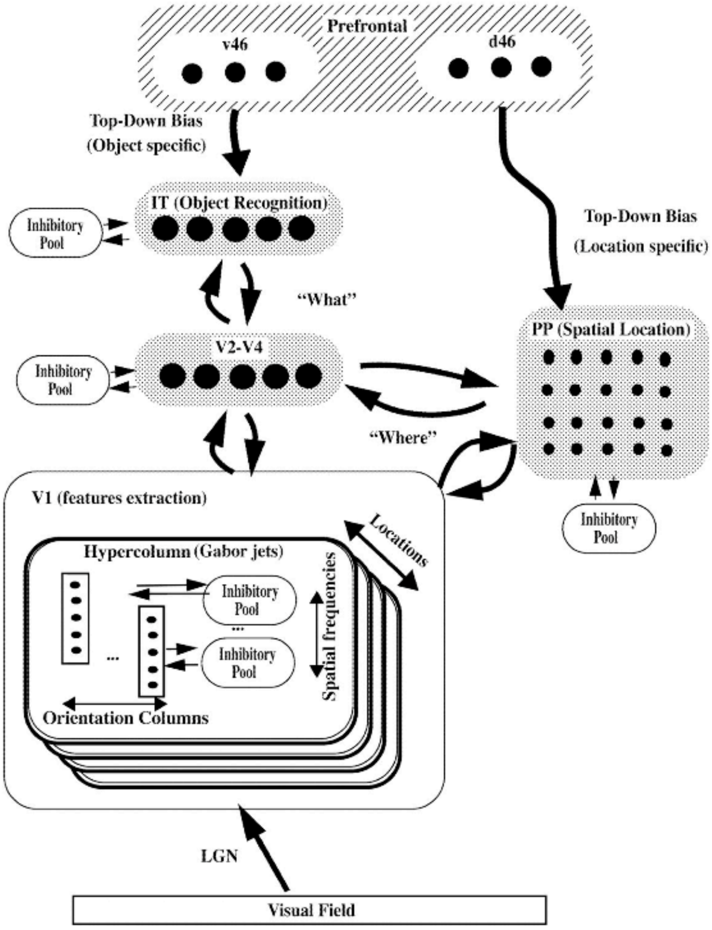

Fig. 3. The model of Corchs and Deco (2002) is a typical biologically motivated model of visual attention.

Decon and Zihl Deco and Zhil (2001) describe a model that combines bottom up and top-down processing that addresses the problem of select both the scale and area of the attended location in an image. It consists of a number of modules that process the input at different spatial levels and can be viewed as hierarchical predictors at the different spatial levels (Fig. 4). This results in a sequential coarse to fine analysis of the image. The main contribution of the model is that it shows that it is necessary to process visual input at several simultaneous scales to be able to direct attention either to global patterns or to local features.

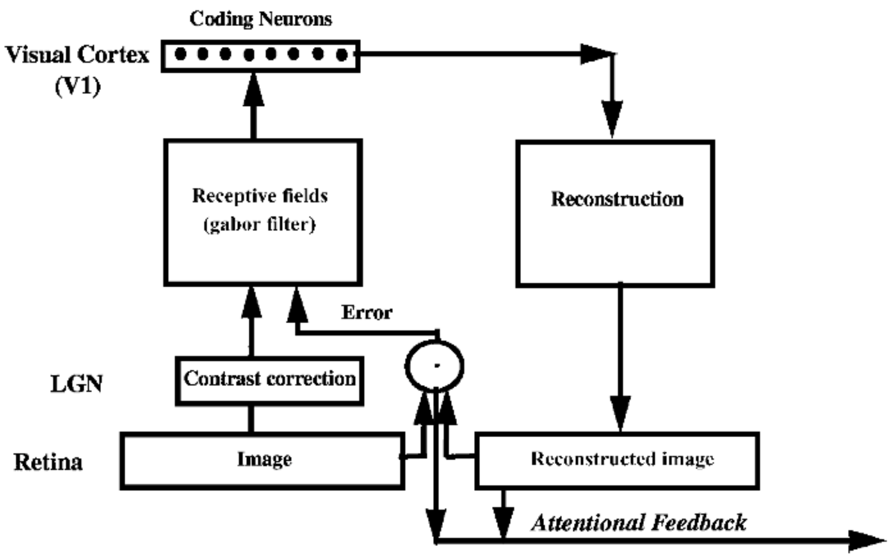

Fig. 4. The model of Deco and Zihl (2001) integrates bottom-up and top-down processing of visual stimuli.

Rao and Ballard Rao and Ballard (1996, 1999) introduced the idea of the visual cortex as a hierarchical predictor (Fig. 5). A number of processing levels make up a hierarchy of Kalman filters where the feedback from the higher levels act as prediction for the lower levels (Fig. 6). The model attempts to explain the behavior of classical and non-classical receptive fields in the visual cortex.

In their method, uncertainty about anticipations by components of other components is estimated, affecting the amount of influence levels have on each other (when the higher level is very certain that an event is going to happen in a lower layer, it has more control over the lower layer than when it is not certain at all).

Additionally, we imagine the inclusion of attention in the framework in that attention may increase the top-down influence dependent on current motivations, emotions, or intentions. More advanced variants of such systems can be implemented, using techniques such as the gating mechanisms in LSTMs and Kalman filters.

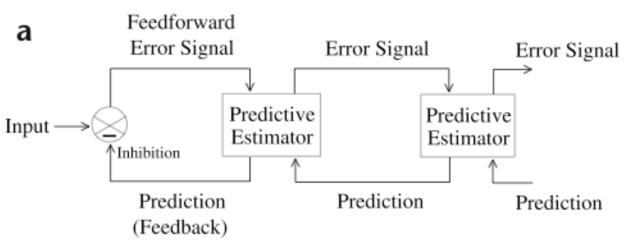

Fig. 5. The basic operation of the model proposed by Rao and Ballard

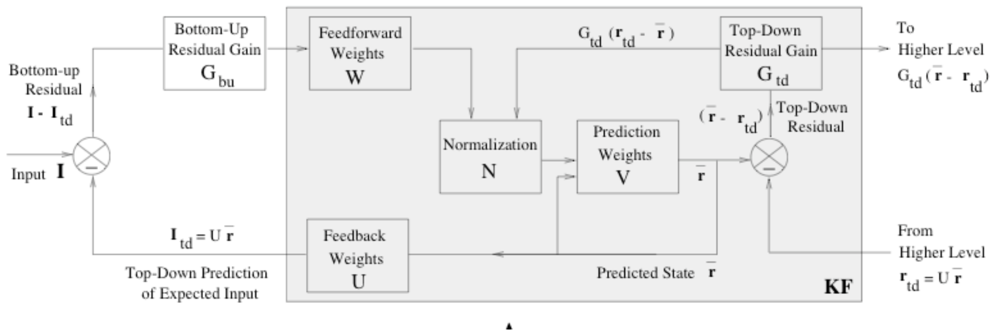

Fig. 6. Hierachical prediction seen as a Kalman filter

Prediction is thus a fundamental aspect of these models. However, the top-down influences are very direct and it is not clear how the model can handle complex task-requirements since there are no temporal aspect in the basic model. Nevertheless, the framework of hierarchical predictive systems appears to be very fruitful for future attention systems. Especially if anticipation over time could be included.

Spratling and Johnson Spratling and Johnson (2004) developed a neurobiological model of attention in cortex which is based on levels of interacting cortical modules (Fig. 7). There is a biased competition within each module which is influenced by top-down and bottom up inputs. The model can reproduce some instances of foreground/background segmentation and contextual cueing.

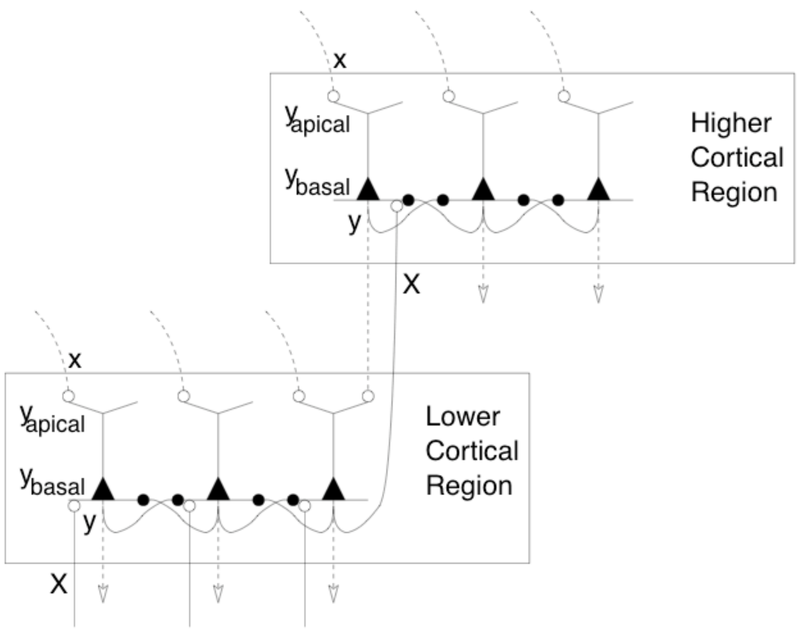

Fig. 7. The interaction between two levels of cortex in the model of Spratling and Johnson, 2004.

Navalpakkam and Itti Navalpakkam and Itti (2003, 2005) extended the previous bottom-up model of Itti and Koch with two types of top-down influences (Fig. 8). One is a task specification that is used to decide on which features are important in the scene and also on what spatial locations should be attended to. The other is a system that attempts to find the "gist" of the scene that can further guide the selection of a particular image region. The model also includes a symbolic component that models long-term and working memory.

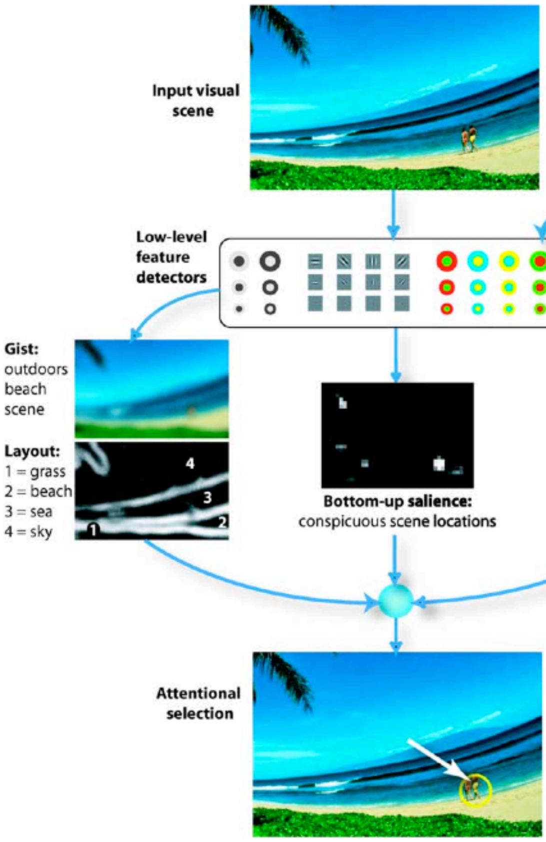

Fig. 8. The model of Navalpakkam and Itti (2005) allows the task to influence the attentional system.

These additions are necessary extensions of the original model, but it appears that some of the additions follow a very different design philosophy than the original model. However, the symbolic top-down systems suggests how more complex task related control systems could interact with a top-down attention system.

Balkenius, Aström and Eriksson In the model by Balkenius, Åström and Eriksson (2004), a number of pre-attentive processing stages were selected based on their computational efficiency and utility in finding targets for the attention system. These bottom-up systems consisted of horizontal and vertical contrast at a single scale, curvature detection using the Harris operator (Harris and Stephens, 1988), and a foreground estimator (Stauffer and Grimson, 1999) (Fig. 9).

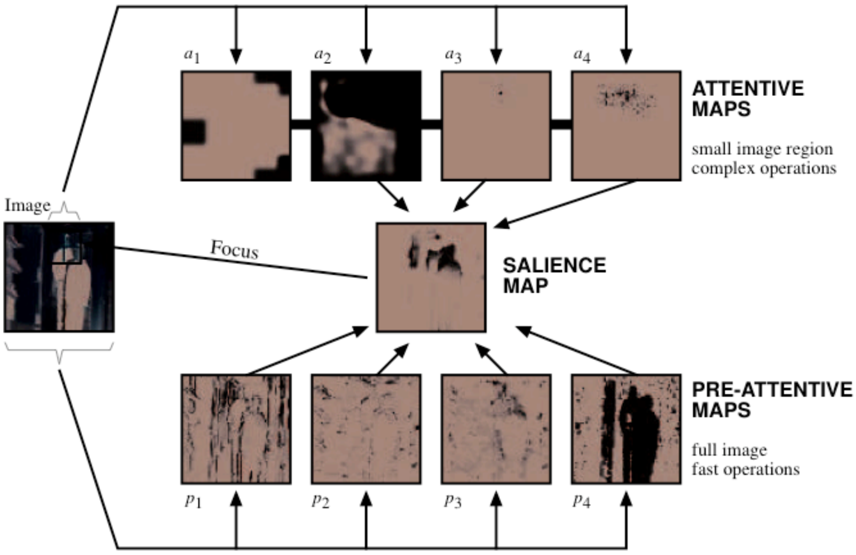

Fig. 9. Combining botto-up and top-down influences on attention in the model described by Balkenius, Åström and Eriksson, 2004.

In addition, the model included a number of attentive, or top-down, systems that learn to predict image regions where targets are likely to occur or relations between different features in the image. For example, on attentive system learns associations between target stimuli and visual cues that predict the location of target stimuli. A number of different relations between cues and target were studied including absolute and relative spatial and temporal relations between a single or multiple cues and a target.

Discussion

Adding hierarchies to anticipatory systems could enhance the system's ability to abstract and generalize, while reducing the search space enormously. Hierarchically combined recurrent neural network systems where lower layers feed into higher layers have been shown to effectively enhance certain abstraction and generalization properties of such systems as a whole (Riesenhuber, Poggio, 1999).

However, systems where higher layers do also influence the lower layers are even more interesting. Imagine a cascaded RNN where every layer is predicting its own state at the next time step. Higher level layers try to model the behavior of the lower level layer, correcting lower layers when the lower levels do not have a high-enough level of abstraction to reliably predict their own state. In other words, we envision systems where high-level components, where possible, anticipate and correct the state of lower level components.

Additionally, several hierarchies should be implemented, such as one for action and motor control and one for visual processing, as proposed recently (Poggio, Bizzi, 2004). Interactions between the two can be manipulated with Rao and Ballard's original proposion (1997), or more recent approaches, such as the mechanisms in Grimes and Rao (2005).

Rao and Ballard (1997) present a probabilistic system. Hierarchical structures should be able not only to predict one possible next state, but a probability distribution over next possible states. Using RNNs, we can use extensions of standard adaptations of the error function in order to achieve this.

Although this system features continuous predictions, the hierarchical structure enables some form of discretization. For example, Boden (2004) discretizes the internal state of the lower-level network according to prototype vectors before feed-forwarding it to the next level. A more sophisticated approach, where discretization is achieved by, for example, a gates-based structure like LSTM has, seems more general than his approach.

Another interesting aspect is that hierarchies can achieve a form of abstraction. Although this is less general than purely using bigger networks, the probability of finding appropriate representations might be higher (might be – this needs to be investigated). Essentially it is expected – as often claimed before – that hierarchically, decomposed structures are ubiquitous in our environment (Simons, 1969; Gibson, 1979). Since our project is concerned with embodied, cognitive systems, as long as sensory representations are chosen wisely, the typical structural constraints in the environment should carry over to the internal hierarchical structural predispositions.

The different models of top-down attention has often concentrated on either spatial or feature/object selection of attention and it is clear that any complete model must include both these types of top-down control.

The different models show that there are a number of issues that must be addressed by a complete top-down attention model:

- There must exits control both of spatial (where) and feature/object based attention (what).
- The model may be able to reconstruct its input from its predictions
- It is necessary to handle both coarse and fine analysis of the input, that is, attention must work also in scale
- Hierarchies appear to be necessary for more complex control of attention
- For prediction and anticipation, a number of possible relations between cues and targets must be considered

3 Reinforcement Learning in Attention Control

Schmidhuber and Huber Schmidhuber and Huber (1990, 1991) built an artificial fovea controlled by an adaptive neural controller. The fovea had high resolution in the center and low resolution in the periphery. Without a teacher, it learned trajectories causing the fovea to find targets in simple visual scenes, and to track moving targets. The controller used an adaptive input predictor (a limited kind of world model) to optimize its action sequences.

The only goal information was the shape of the target - the desired final input. Since this reinforcement learning task is of the `reward-only- at-goal' type, it involves a complex spatio-temporal credit assignment problem. The latter was solved using a recurrent network training algorithm.

Grupen et al. Q-learning was used in the model of Goncalves et al. (1999) to control attention based on multimodal input and reinforcement signals. The model includes subsystems for long term memory, what and where processing and attention control. An interesting feature of the model is that it integrates attention and action in a straight forward way.

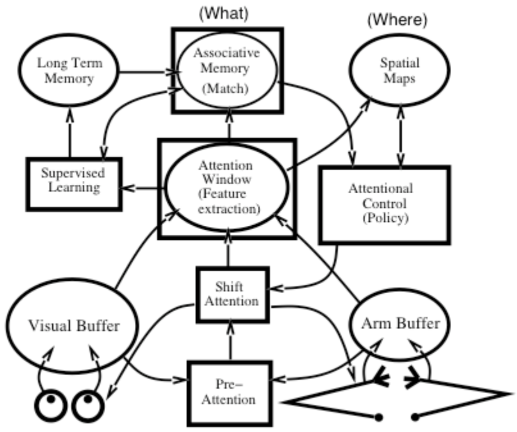

Fig. 10. The model of Goncalves et al. 1999.

Balkenius Balkenius (2000) proposes four principles for attention (Fig. 11). Attention to object B is inhibited which disengages the focus of attention (attention-as-inhibition). Shifting the attention and gaze g from object B to A is an action s (attention-as-action). Object B is selected for the action a by directing the focus of attention, and gaze, toward it (selection-for-action). The focus of attention is used as an implicit argument for the action. The focus of attention refers to the object B without explicitly representing all of its properties (deictic reference).

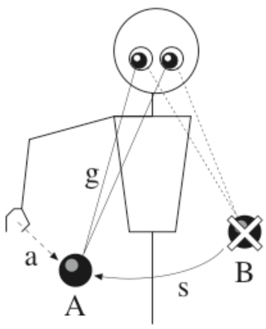

Fig. 11. Principles behind the control of attention and action

The principle of attention as action implies that attention is controlled in a way similar to actions and it thus becomes natural to use reinforcement learning also to attention. This was tested in a number of simple visual search tasks.

Minut and Mahadevan One model that uses reinforcement learning to control visual attention was described by Minut and Mahadevan (2001). The Q-learning algorithm is used to select the target locations in the image.

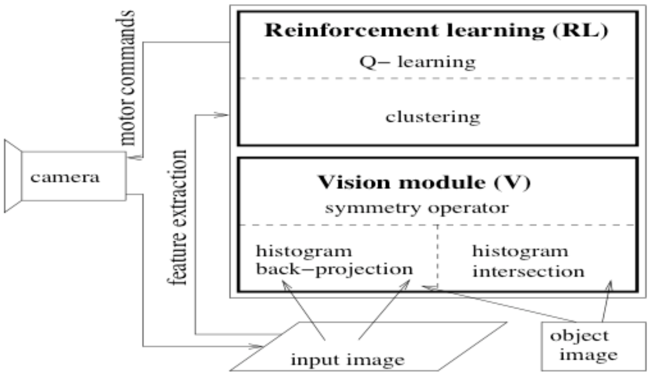

Fig. 12. The model of Minut and Mahadevan.

K. Shibata Reinforcement learning was also used by K. Shibata et al. (1995) to control the movement of a visual sensor over an image. The goal of the system was to find the optimal location of the sensor for object recognition. The same neural network was used both for object recognition and to preoduce the sensory motion output.

Context Sensitive Reinforcement Learning In many situations, it makes sense to divide the input to a reinforcement learning system into two parts: one that codes for the stimulus and one that codes for the context. Current thinking in animal learning theory suggests that the stimulus and the context do not play symmetric roles in learning. Initial learning appears to be insensitive to the context, while relearning makes behavior increasingly context sensitive. By using an asymmetric learning rule of this kind, a reinforcement learning system can be designed that initially generalizes maximally between contexts and later restricts the selection of actions to contexts where they are successful. If the stimulus and context are selected and represented in an appropriate way for the task, this scheme can lead to very fast learning. It also avoids catastrophic forgetting if the context inputs are capable of coding the different learning situations.

These ideas have been applied to Q-learning to develop ContextQ which learns using a context sensitive linear approximator (Fig. 13). Initial learning operates as an ordinary linear approximator, but relearning invokes the context to cancel out inappropriate associations. Apart from making the learned behavior context sensitive, the method also limits catastrophic forgetting in cases when the context can accurately predict which set of behaviors is appropriate in the different learning instances.

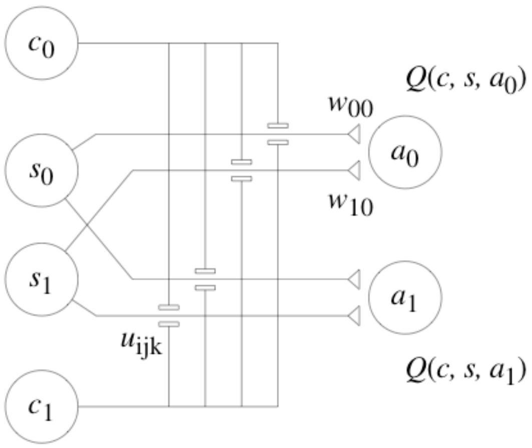

Fig. 13. The context sensitive linear associator forms the basis of the ContextQ algorithm.

Discussion

The use of reinforcement learning in the control of attention offers many interesting possibilities. This is especially the case if the control of attention can be integrated with motor control in a natural way. A reinforcement learning framework offers the possibility to use a common currency to describe the success of both attention and action.

An important future development will be to allow reinforcement learning systems to learn attention control that is valid in many different situations and under different environmental demands. This will make it necessary to use learning systems that can take the current task, context and motivations into account.

4 Dynamic Models of Attention

Although most models of attention primarily address situations in which the viewed image is stationary, some models instead look at the pursuit of moving target. A number of such models have been developed within active vision and humanoid robotics and here we will only describe two recent models that illustrates some important aspects of such models. These models deal primarily with the "where" process rather than the "what" process and do thus complement the other attention models.

T. Shibata et al. The model of smooth pursuit eye-movements of Shibata et al. (2001) uses a predictive model based on linear or non-linear regression networks to predict the behavior of the target stimulus. A cascaded control scheme is used in which the predictor sets the desired path for the eye controller (Fig. 14).

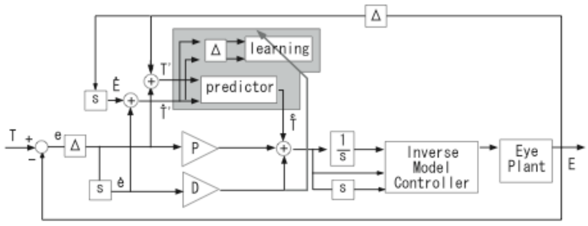

Fig 14. The model of smooth pursuit eye-movements of Shibata et al. (2001).

The model clearly illustrates that anticipation is absolutely necessary for any system using active vision in a realistic environment. Since the visual processing is delayed, it is not possible to control the movements of the camera using traditional feedback control. Instead, the direction of the gaze must be controlled by a model that predicts where the target is now.

One limitation of this model is that it only allows a single model of target motion although it would not be too complicated to allow several such models that could be chosen depending on the circumstances. An open question is whether there should be many models for different targets or if a single simple model is sufficient for all cases. For example, a model which simply extrapolates the motion of the target, possibly taking acceleration into account, may be sufficiently good. It is also possible that much more complex models should be learned when the system has to observe complicated but regular motion.

Balkenius and Johansson An alternative model of smooth pursuit eye-movements was proposed by Balkenius and Johansson (2005). The main strength of the mode is that it includes both continuous and discrete target predictions (Fig. 15). The model assumes that a linear predictor is sufficient for normal pursuit of the target, but when certain events occurs, the model learns specific expectations of the behavior of the target.

The model is thus a compromise between a continuous and a discrete model of target motion. The main strength of the model is that the hybrid approach allows it to learn complex scenes where the target may disappear and reappear or change motion abruptly at certain locations. It also acknowledges that the modeling of target behavior must occur in global and not egocentric coordinates.

The main limitation is that only a single scene can be learned but this will be addressed in the next version of the model that will include a context sensitive predictor for both the continuous and discrete models.

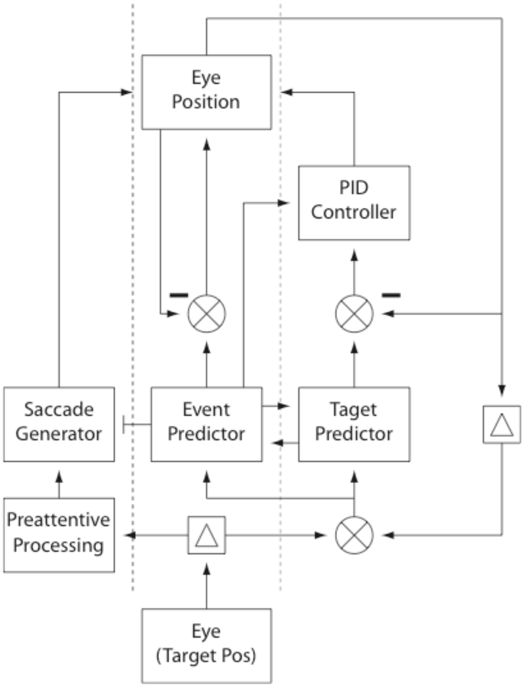

Fig. 15. The model of target tracking described by Balkenius and Johansson, 2005.

Discussion

Dynamic models of visual attention clearly show the need for anticipation in the control of attention and are thus a good starting point for models of attention and control that wish to implement anticipation. They also connect the control of attention to emulator models, since an emulator is a necessary component for such control systems.

5 Emulator Models

Although emulator models have been mainly described in relation to motor control they are also of interest to models of attention and control.

Grush Grush (2004) describes a mechanism based on Kalman filters. An emulator exploits efferent copies of the control signal for producing a feedback to the system. The emulator, running simultaneously with the perception-action cycle of the system, provides expectations in the sense of expected results of the possible actions of the system; both the input and the output of the emulator are in the same format of the representation used by the system to perceive and act (and thus they can be directly compared), as shown in figure 16.

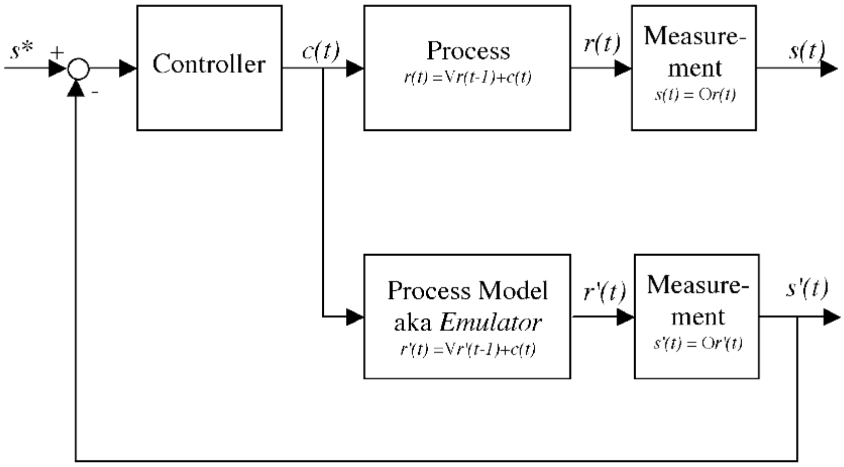

Fig. 16. Using a model to emulate the behavior of a process. From Grush, 2004.

The same mechanism, the emulator, can be used even for generating pseudo-proprioceptive information that is at the basis of motor imagery; and, assuming that the data produced by the emulator can run even the sensory areas, even visual imagery can be produced. This is an important aspect of an anticipatory system that must be taken into account in real applications since this will generate predications of the outcome of an action that can be matched against the actual outcome. Of particular interest is the use of several emulators working in different coordinate systems resulting in modal and amodal emulation (Fig. 17).

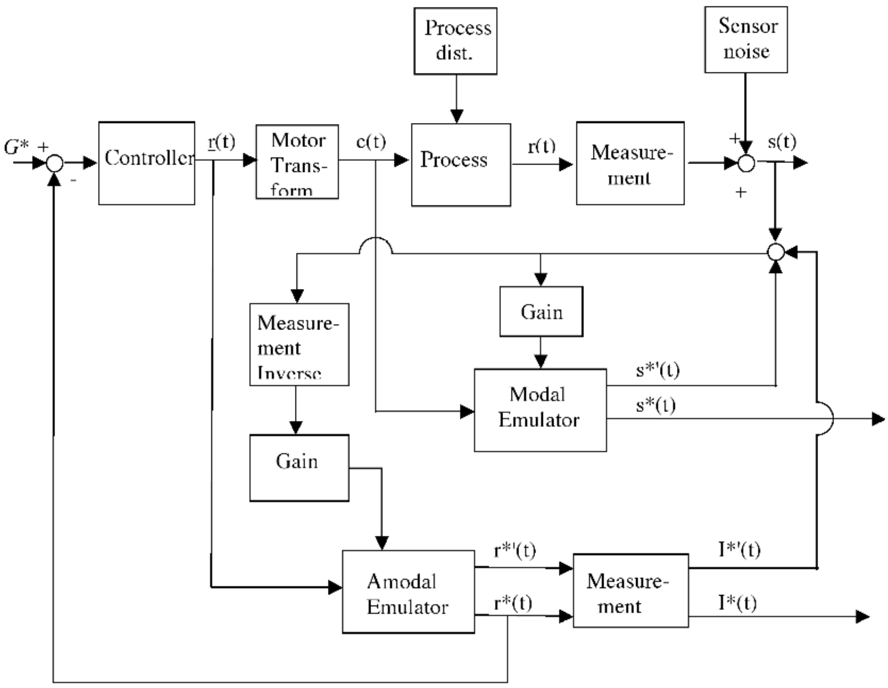

Fig. 17. Using both modal and amodal emulators. From Grush, to appear.

Barsalou Barsalou (1999) describes (although he does not implement) the perceptual symbol system that is heavily based on the idea of producing modal representations by using Simulators, that re-enact sensory patterns. Differently from emulators, simulators can also produce expectations by running offline: "perceptual symbol systems cannot be implemented adequately until it is possible to run simulations either in isolation or simultaneously with perception".

Recently the discovery in neurobiology of mirror neurons has strongly revitalized and influenced all the research about "productive" systems such as simulators and emulators, focusing on the strong coupling between motor activity and perceptual stimuli.

Doya et al. A main limitation of the current emulator models is that they assume there is a single process that should be emulated. A major future development will be to allow many simultaneous emulators that can be selected as needed . The first steps toward such an architecture has been taken by Doya et al. (2000) (Fig. 18).

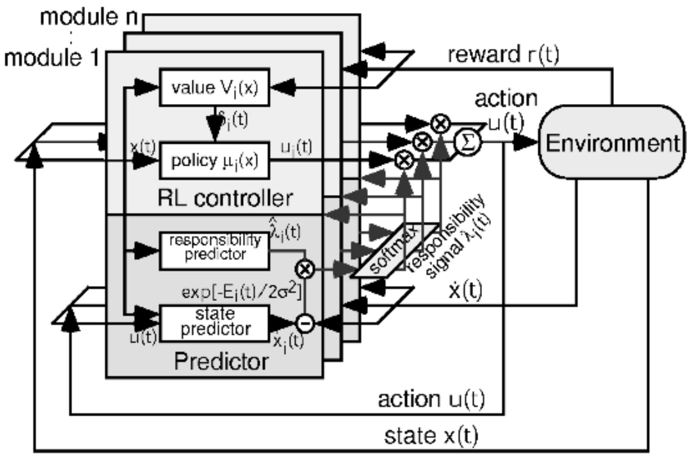

Fig. 18. Using multiple paralell models to decompose a reinforcement learning problem.

Discussion

Emulators are central for systems with anticipatory abilities. To be able to react to the future state of the environment, it is necessary to anticipate that future state and this requires some form of model whether it is implicit or explicit. The emulator framework offers a nice way to think about such anticipation.

The major future task will be to investigate how emulators can be designed that can handle more complex situations and many situations in parallel.

6 Schema Models

Schematic theories (Arbib 1995, Arkin 1998) have their roots in the work of Kant (1781) and Piaget (1954). In artificial intelligence they share resemblances with the concepts of frames in Minsky (1975) and scripts in Schank and Abelson (1977). Schematic theories are strongly motivated by biological and ethological models; some of the first works are in fact the attempt to replicate into robots the behavior of cockroach, mantis and toad. Schemas are conceived as coarse grained modeling unities and opposed e.g. to neural networks that are more fine grained (in some implementation they are however decomposed and ultimately realized as neural networks; however, it has to be stressed that the descriptive level is coarser). Schemas include perceptual and motor elements to form coordinate control programs: a schema can be activated by sensory input or internally generated goal states, and many arbitration mechanisms are described in literature for selecting the more appropriate one among the active ones (they share some resemblances with some action selection mechanisms (e.g. the behavior networks, Maes 1990) as well as with the subsumption architecture (Brooks 1991).

Schematic theories stress the procedural knowledge (versus declarative): "a schema constitutes the long term memory of a perceptual and/or motor skill, or the structure coordinating such skills". They are especially well suited for parallel and distributed systems, since they can be seen as concurrent computing unities. Lyons and Arbib (1989) describe RL, a language for implementing schemas in robotic applications. Schematic theories were not originally conceived for anticipation; in fact, Arkin describes their use mainly in the perspective of reactive and behavior based robotics. However, schemas embed a predictive component that is used for action selection (it could be in principle even used for offline planning and for online control of action, although in literature these aspects are not stressed).

Roy The recent work of Roy (2005) extends the schematic approach by coupling sensorimotor engagement, expectations, and active perception (Fig. 19). Schemas are also more sophisticated: they are actually networks of beliefs (both analog or discrete), linked by six possible actions, called projections (sensors, actions, transformers, categorizers, intentional projections, and generators). The anticipatory nature of schemas is stressed: "Beliefs are both memories of what has transpired, and also predictions of what will transpire (contingent on action)". Schemas are parametric and can be concatenated or abstracted. The whole apparatus is constructive and follows the 'active perception' paradigm: perception depends on understanding the effects of movements on sensory stimulation, thus it is intrinsically based on anticipations. Ripley (as well as other robots) are equipped with (hand-made) schemas that provide paths from actions to anticipated sensations; for example, the following picture shows a schema for active sensing of compliance through grasping.

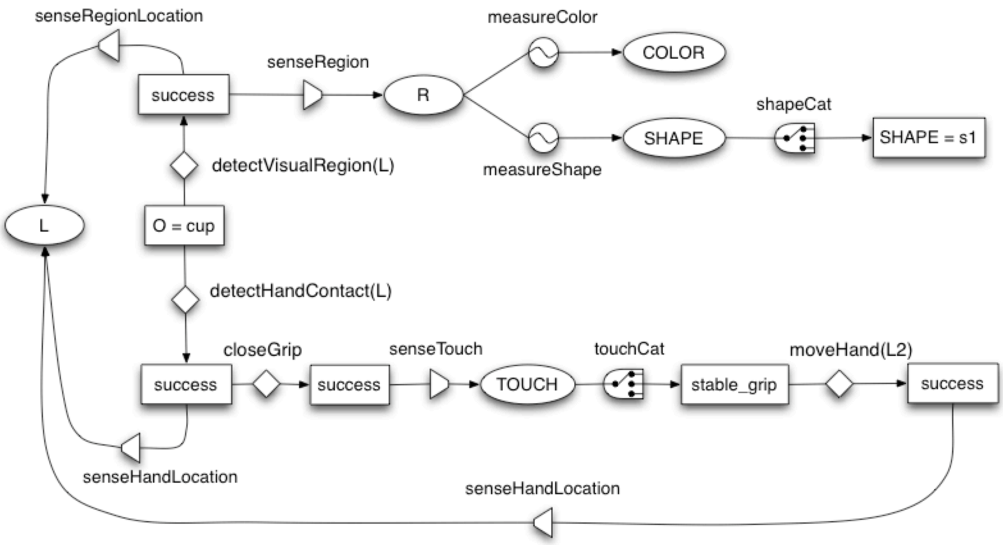

Fig. 19. A example schema for a tangible graspable object such as a cup.

Roy introduces schemas for objects, properties, events and situation; all these representational primitives expressed in the same formalism: in this way "beliefs about concrete objects (e.g., cups) [can] be efficiently translated into expectations with respect to actions". For example, in a way that is similar to Drescher's SM, objects are represented as the set of possible interactions with it (also capturing in a natural way the notion of affordances, Gibson 1979).

Roy aims to provide a grounding for language, "a computational path from sensing and motor action to words and speech acts". Grounding itself is defined as a causal-predictive loop, stressing the anticipatory nature of action and even sensing.

Kopp An attention system based on schema theory was designed by Kopp (2003). It uses schemas for learned objects or scenes to direct attention to locations where it expects to find relevant visual stimuli. In the earlier implementations, this system was purely top-down controlled, but has now been extended with a bottom-up stages that recodes the input image using a visual alphabet.

In object recognition, the schemas are implemented as visual agents that attempts to direct processing to different locations in the image. By reading the letters from the visual alphabet at each location where object relevant information is predicted to be located, each agent gradually tries to confirm its expectation of an object. These ideas could potentially be very powerful in any system that has to interact with the real world. According to this view, perception is the act of confirming predictions of the world.

Drescher Drescher's Schema Mechanism (Drescher 1991) is a constructivist mechanism inspired by the work of Piaget (1954); it represents an object as a set of expected interactions with the environment. The SM learns to anticipate the consequence of its actions by creating three parts schemas "context/action/result"; contexts and results are items (propositions) that can assume the values "On" or "Off".

The Schema Mechanism performs mainly two operations: Empirical learning and Concept Invention. Empirical learning consists in finding a result item that is relevant to an action, as well as context items that make the result reliable (it uses a statistical method called marginal attribution). Concept Invention consists in postulating/reifying synthetic items (objects, not directly sensed) as the "cause" of the perceptions; this property permits to improve the ontology of the system.

The original SM accommodates only binary-valued sensors; but recently it has been extended to continuous values. It has also no parameters: each schema is unique (this poses many scalability problems). It can in learn POMDP environments, as shown in Holmes and Isbell (2005). It focuses on sensory rather than reward prediction.

Rybak et al. Another type of anticipatory mechanism is included in the model of attention-guided visual perception proposed by Rybak et al. (1998). This model is able to recognize visual stimuli using behavioral recognition programs that suggests attentional movements and the expected image fragments that will be perceives at each image location.

DUAL DUAL (Kokinov, 1994) is a general-purpose cognitive architecture which comprises a unified description of mental representation, memory structures, and processing mechanisms. All these aspects of the architecture are organized around a small set of principles:

(1) Emergent computation. All processing and knowledge representation in the architecture is carried out by a cohort of small entities called DUAL micro-agents. There is no centralized executive that controls the whole system, or makes large-scale decisions, allocates resources, resolves conflicts and so on. Instead, small-scale DUAL agents interact with one another, locally, and the global behavior of the system emerges from the self-organizing pattern of these interactions.

(2) Micro-Level Hybridism. DUAL integrates the symbolic and connectionist approaches in a horizontal manner at the level of micro-agents, this means that each micro-agent has two parts: the symbolic part of the agent represents its knowledge while the connectionist part of the agent represents the current relevance of this piece of knowledge. The micro-agents are dualistic in another respect as well: they represent some declarative knowledge, but at the same time they are also processing units that act by sending messages, activation, etc. and in this way embody some procedural knowledge.

(3) Dynamics and context-sensitivity. An important feature of DUAL's operation is that it is constantly changing in response to influences from the environment. This is possible due to the emergent nature of the processing that underlies DUAL's operation and to the lack of rigid and pre-programmed specification of the computation. In particular, there is no sharp boundary between the 'task' or the 'problem' given to the sys¬tem to solve, and the 'context' that accompanies this problem. One and the same problem can be solved in different ways during two successive runs of a DUAL-based model, in spite of the strictly deterministic character of the architecture.

AMBR The AMBR model (an acronym for Associative Memory-Based Reasoning) is a cognitive model of analogy-making based on the DUAL cognitive architecture (Kokinov, 1988). It has been implemented in LISP and simulation experiments have been performed with it and compared with human data. A number of predictions have been generated with this model and later on confirmed by psychological experiments. The AMBR model:

(1)Integrates memory and reasoning by implementing both retrieval of episodes and mapping between the base episode and the target episode within a single model and using the same basic cognitive structures and operations (from DUAL).

(2) Has internal dynamics – implementing a mechanism of state evolution that allows the previous states of the system to influence the specific computation that will emerge with given input. In this way the model explains and predicts priming effects.

(3) Is context-sensitive – changes in the environment that are reflected in the input of the system can influence the emerging computations even if they are not directly relevant to the task. The model has predicted certain context effects that have been confirmed in psychological experiments.

(4)Is parallel and interactive – this means that various subprocesses (like retrieval and mapping) run in parallel and influence each other thus allowing complex interactions to be explained.

Butz et al. Butz et. al. have developed a schema-related mechanism in their anticipatory classifier system (ACS) architecture. The system develops three-part situation-action-effect schemata. The learning algorithm a combination of an information theoretic specialization and a genetic generalization mechanism showed to be able to solve challenging maze, blocks-world, as well as classification problems (Butz, Stolzmann, Goldberg, 2002). Additionally, the mechanism is able to ignore irrelevant inputs. It identifies those features that are necessary to predict feature change, caused by the specified action, accurately. Attentional mechanisms may thus be incorporated naturally by using the representation (1) for enhanced feature processing as well as (2) faster relevant feature detection. Besides the learning and generalization capabilities, it should also be mentioned that Butz and Hoffmann (2003) showed that the ACS can simulate the behavior of rats. The simulation showed that in order to account for the observed behavioral patterns, anticipatory adaptive processes are mandatory, since similar behavior could only be generated if the predictive model in ACS was used to adjust behavior.

Discussion

Schema models are the most expressive type of system that can be used for anticipation. The main question for schema based systems is how schemas can be automatically created, learned and generalized to new situations. It is possible that hierarchical reinforcement learning methods together with anticipatory components may be able to learn complex schemas in the required way, but this is an area for future research.

Schema based models offers many interesting possibilities for attention and control but most of these have not yet been investigated. An interesting aspect of schemas models is that they have certain similarities to hierarchical models and it would be very useful to investigate to what extent hierarchical models can be seen as instances of schema architectures.

7 Context

Although `context' is a term that originates in literary studies, it has recently gained an increased interest also in other disciplines. The role of context in different forms of cognitive processing has been extensively described in relation to speech perception (Altmann and Steedman, 1988, Lee, 1990), conecpt formation (e. g. Medin and Schaffer, 1978), semantic processing (e. g. McDonald and Brew, 2002), learning theory (Bouton and King, 1983, Bouton, 1993) and psychiatric disorders (Cohen and Servant-Schreiber, 1992, Balkenius and Björne, 2004). Recently the concept has been related to scene perception and the control of visual attention (Balkenius, 2003).

Balkenius & Morén Although context plays a role in many models of attention and control, the source of the context has seldom been investigated. The goal of the context processing model first described by Balkenius & Morén (2000) is to build codes for the current context based on sequences of sensory stimuli. One way to build the context codes is to use an attention mechanism. The current focus of attention acts as the stimulus while a sequence of attention states make up the context. This implies that the context can be controlled by choosing how attention is allocated.

The model consisted of four interacting modules (Fig. 20). The binding stage (BIND) used distinct events to build bindings coding for stimuli and place relations, which where then combined into context codes (CONT). The context in turn was used to form expectations of the stimuli that would be present in each environment. These predictions which were stored in a memory (MEM) and were used to match the current input to the expected inputs (MATCH). When a mismatch occurred new bindings and possibly also a new context was created.

The main target for the mode was to explain how negative priming and inhibition operates in classical conditioning, but the mode has also been used to explain contextual cueing and spatial habituation. In the later development of the model the binding is done using tensor product coding which removes the creation of context codes from the behavioral situation. In this model, contexts can still be created when expectations are not met but they can also be constructed passively from sequences of input stimuli.

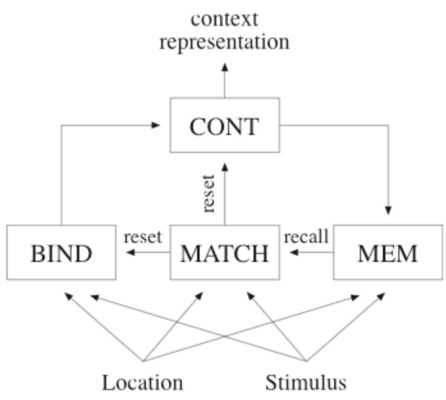

Fig. 20. The context processing model of Balkenius and Morén, 2000.

Torralba et al. Torralba and coworkers (Torralba, 2002, 2003) have developed a way to holistically code a visual context. The method first extracts the main spectral components at different locations in the image suing complex Gabor filters tuned to different spatial frequencies which produces a high dimensional representation of the image. In a second step, this representation is mapped onto a low dimensional representation using principal component analysis. The resulting representation is used to prime spatial attention and object recognition. It is used to produce probability density functions for object likelihood, spatial attention and local target appearance.

The use of top-down context priming in object detection offers an alternative to the bottom-up saliency based approach.

LSTM Another aspect of context is that it can reach over longer times and it is necessary to include in the context events that took place some time ago. LSTM (Long Short-Term Memory) is a recurrent neural network architecture with a gates-based structure (Fig. 21): input gates and output gates guard input/output access to the internal states of neurons, enabling the algorithm to maintain memory over longer periods of time (Hochreiter and Schmidhuber, 1997). This method has been shown to effectively deal with the problem of vanishing gradients. LSTM can remember and relate events far apart in history, and could thus be used as a prediction tool for anticipatory systems that require long-term memory.

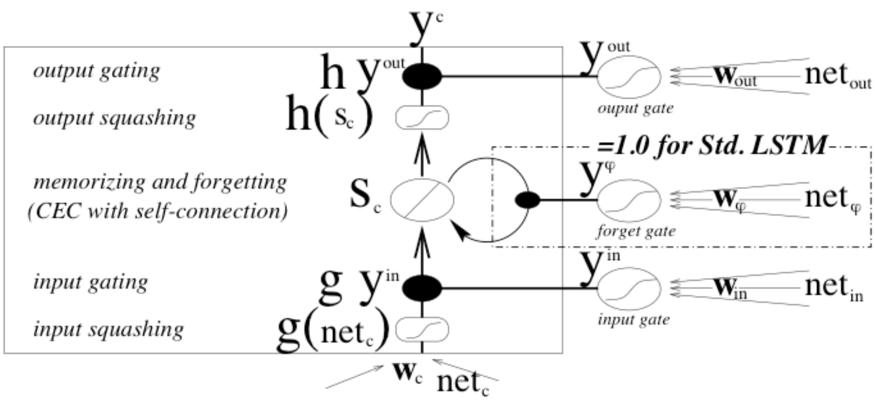

Fig. 21. The LSTM cell.

The method has the capability to capture long-term time dependencies, but in order to predict events far in the future or in order to predict future abstract states, the method probably needs to be combined with hierarchical approaches. It is a multi-purpose predictive system, with the ability to handle continuous inputs and outputs. It could be used for sensory prediction, but also for utility prediction (as has been done for reinforcement learning purposes), among other things. Since it is an RNN, standard methods from RNN literature can be used, such as error functions that allow us to output probability distributions instead of pure predictions per se, or methods to control regularization. LSTM is thus an ideal candidate for context that reaches over longer times.

Discussion

The different models offers three types of context that are of important in models of attention and anticipation:

- Context as the "gist" of a scene, that is, a spatial simultaneous pattern
- Context as sequences of stimuli
- Context as point events that should influence later behavior.

So far, no model has attempted to include all three types of context and such a system would constitute a clear step forward. It is possible that context could be included in a similar way to other sensory inputs for predictions and anticipation, but is more likely that context must be handled in a separate way.

8 Priming

Priming occurs when one stimulus has a positive or negative effect on the processing of another stimulus at a later stage. If the two stimuli are the same and the processing is enhances, it is a case of repetition priming. If the processing is reduced, it is instead a case of habituation. Other types of priming can be categorized as perceptual or semantic depending on whether the enhancement is a direct consequence of the perceptual properties of the stimulus or of its meaning.

Interactive activation models The earliest connectionist models to investigate priming were the different instances of the interactive activation models of McClelland and Rumelhart (1981, 1982). In these models priming was seen as the result of associations between nodes in a connectionist network that coded for different concepts. Positive and negative priming would result depending on whether the associations were excitatory or inhibitory.

In this interpretation of priming it is natural consequence of any associative network model and there is thus no need to explicitly mention all models that produces priming in this way.

Plaut Plaut (1995) presented a distributed attractor network that was able to model semantic and associative priming (Fig. 22). The model produces stronger associative priming than semantic priming and associative priming that last even when an intervening item is presented.

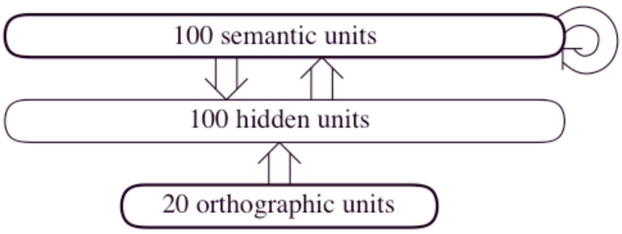

Fig. 22. The model of Plaut (1995).

Balkenius, et al. Balkenius and coworkers have developed a model of attention control based on reinforcement learning (Fig. 23). The goal of the model is to explain how different learning processes contribute to the control of visual attention. There are five main components in the model (Balkenius, 2000). A sensory buffer codes the visual input in different ways. This coding ranges from the detection of oriented contrasts to object identity. The sensory buffer also codes the visual input using a spatial code that allows attention to be directed toward the location of a stimulus. A fixed response system R reacts to the location coding in the sensory buffer to produce overt orienting reactions.

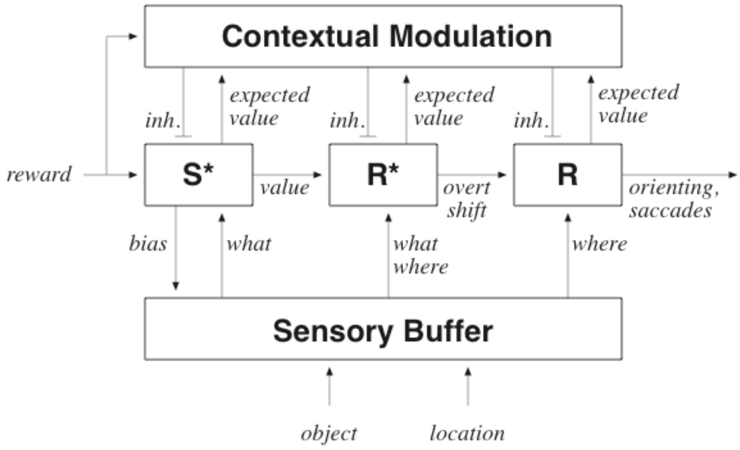

Fig 23. Overview of the Model.

The shift of attention is controlled by two basic learning systems. The stimulus evaluation system S* assigns a value to each stimulus based on its reward history and the response evaluation system R* assigns values to stimulus-response pairs. Together, the S* and R* systems implement an actor--critic (or two-process) architecture for learning (Sutton & Barto, 1998, Mowrer, 1960/1973). The learning process in S* is classical conditioning and in R* it is instrumental conditioning.

In addition, the three modules, S*, R* and R, are influenced by the context system that codes the current visual context or situation (Balkenius & Morén, 2000). The inhibitory role of the context (Balkenius, 2000, Balkenius & Morén, 2000, Morén, 2002) as been investigated as weell as excitatory influences of the context on the response evaluation system R* (Balkenius, 2003, Chun & Jiang, 1989, Chun, 2000). The model can produce many attention phenomena and can operate both in top-down or bottom-up mode.

The model produces priming as a consequence of learning (Balkenius, 2004). When an observed stimulus is not rewarded, an inhibitory connection will form from the current context to that stimulus representation which will results in negative priming by the context. The model is also able to produce contextual cueing when an attended stimulus has been reinforced in a specific context (Fig. 24). The context will form excitatory connections to the response system that produces saccades to expected target location in the learned visual context. This can be seen as a form of positive priming.

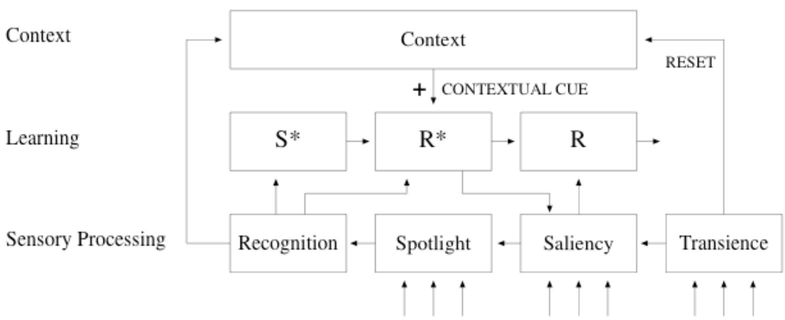

Fig. 24. The extended context model applied to sequential contextual cueing. A transience detector is added to reset the context when the visual scene changes.

Discussion

Although priming in some form is present in many models, there have been few attempts to include all types of priming that can be seen in psychological experiments. A minimal requirement would be to include positive and negative priming from at different hierarchical levels (perceptual/semantic) and between sensory representations ranging from single stimuli to contexts. In relation to the goals of the MindRACES project, it would also be interesting to study to what priming can be seen as a form of anticipation.

9 Deliberative Agents and Epistemic Actions

BDI models are models of deliberative agents. They have been introduced in philosophy (Bratman 1987), inspiring the development of computational models of cognitive agents (Rao & Georgeff 1991b) and logical theories of pragmatic reasoning (Cohen & Levesque 1990, Singh & Asher 1993, Meyer et al. 1999). The main properties of this approach is the interest to characterize the functional properties of mental states such as beliefs, desires, goals and intentions and to analyse their interrelationships, to define the way an agent selects actions to achieve his goals and gets committed to execute the action. In BDI models is assumed some kind of decision procedures between actions (as means for achieving the goals of the agent) based on probability associated with the beliefs and expectations and utility values associated to the goal states (Rao & Georgeff 1991a).

Recent theoretical and formal developments in cognitive robotics and especially in the logical theory of belief revision and information acquisition have had a considerable impact on modelling the way an intentional agents can acquire information from the external world and revise his pre-existent knowledge on the basis of new acquired information (De Giacomo et al. 2000; Herzig et al. 2000, Del Val et al. 1997, Reiter 2001; Bacchus et al. 1995). These kinds of theoretical developments implie the necessity to characterize the different sources of information that an agent can question (by means of Epistemic Actions) in order to verify his beliefs and predictions (Castelfranchi 1995; Van Linder et al. 1997), in order to update his pre-existent knowledge or in order to gather further data, and consequently the necessity to provide models of bounded reasoning integrating the concept of attentional focus (Grant et al. 2000) with respect to stored data, temporally bounded inferential process (Elgot-Drapkin et al. 1999) and the distinction between explicit beliefs and implicit (potentially derivable) beliefs.

The integration of the notion of Epistemic Action in the BDI framework raise important questions with respect to the issue of deliberation. Deliberative agents should be capable of choosing between pragmatic alternatives and/or epistemic alternatives. At each round they should be able to decide to do some pragmatic action (movements) or to get new information by doing appropriate tests to the external environment and by focusing the attention on stored data.
The problem of concurrence between epistemic actions and pragmatic actions in the selection of plans and policies has already been considered in the design of Partially Observable Markov Decision Processes (Kaelbling et al. 1998), and give insights with respect to the way a deliberative agents evaluates a piece of information as a relevant (and valuable) one.

Finally, the development of theoretical and practical models for deliberative agents using confidence and implicit trust is a very new and challenging field of study in the AI (Castelfranchi and Falcone, 1998a; Castelfranchi and Falcone, 1998b; Falcone and Castelfranchi, 2001). In particular, a very relevant problem is the analysis of an agent action as a sort of (implicit or explicit) reliance on the world with a given quantitative value about the expected results (and governing both the decision and the epistemic control).

Discussion

Integrating a model of Epistemic Actions with the general BDI models represents a relevant achievement of the research in computational modelling of deliberative agents. Theoretical improvements are needed in order to model agents with the capacity to actively test the external environment (in order to gather external data), to actively focus their attention to internal data and to make inferences on the internal focused data. Moreover, an extension of BDI models to cover Epistemic Actions is relevant for designing agents that can plan to execute standard pragmatic actions in order to verify (by means of epistemic actions) both their ability to do those actions (this is the essential structure of every *intentional trial* in the sense of Anzai & Simon 1979) and at the same time the reliability of the environment and of mediating processes (with an explicit/implicit feedback on their trustworthiness).

10 Conclusion

Several promising research directions come to mind considering anticipatory mechanisms in attention, monitoring and control.

One interesting development would to be the enhancements of the rule-based schemas in ACS to (1) more NN-based structures and (2) to hierarchical structures. It is intended to combine the expertise of ACS on discrete processing mechanisms with processing mechanisms in continuous domains. Rules in ACS provide confidence values, so that another challenge is to combine rule bases hierarchically, where the higher levels predict activitiy in the lower layers. Research is in progress towards these directions.

The most interesting aspect is the coupling of action control mechanisms with sensory processing control. Schema mechanisms are the most straight forward, rather symbolic-based approach. The coupling of schema-based ACS structures with lower level forward model, emulator-like structures poses the biggest challenge to the MindRACES project and may yield most promising results in the mean time.

Another important development will be in the area of contextual anticipatory control of real-time attention in dynamical scenes. This involves both an individual robot looking at complex events and multiple robots interacting in a dynamical environment. The use of context sensitive hierarchical anticipatory systems controlled by reinforcement appears to be an especially promising area for future research.

In summary, the aim of the future research will be to merge the abilities of systems for (1) top-down and bottom-up control, (2) reinforcement learning, (3) dynamical perception, (4) schemas, and (5) context. It is clear that anticipation plays an important role in this development.

References

Abbot, L. F. & Blum, K.I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. Cerebral Cortex, Vol. 6, pp. 406-416.

Altmann G. & Steedman M. (1988). Interaction with context during human sentence processing. Cognition, 30, 3, 191-238.

Anzai, Y., & Simon, H. A. (1979). *The theory of learning by doing*. Psychological Review, 86, 124-140.

Arbib, M. A. (1995). Schema Theory, in M. A. Arbib (Ed.) The Handbook of Brain Theory and Neural Networks. Cambridge, MA: MIT Press.

Arkin, R. C. (1998). Behavior-Based Robotics, Cambridge, MA: MIT Press.

Bacchus, F., Halpern, J. Y., Levesque, H. J. (1995). Reasoning about Noisy Sensors in the Situation Calculus. In *Proceedings of. IJCAI 1995*, pp. 1933-1940.

Balkenius, C. (1998). Spatial learning with perceptually grounded representations. Robotics and Autonomous Systems, 25, 165-175

Balkenius, C. (2000). Attention, habituation and conditioning: toward a computational model. Cognitive Science Quarterly, 1, 2, 171-214.

Balkenius, C. (2003). Cognitive processes in contextual cueing. In Schmalhofer, F., Young, R. M., & Katz, G. (Eds.), Proceedings of the European Cognitive Science Conference 2003 (pp. 43-47). Mahwah, NJ: Lawrence Erlbaum Associates.

Balkenius, C. (1998). Spatial learning with perceptually grounded representations. Robotics and Autonomous Systems, 25, 165-175.

Balkenius, C. & Hulth, N. (1999). Attention as selection-for-action: a scheme for active perception. In Schweitzer, G., Burgard, W., Nehmzow, U., & Vestli, S. J. (Eds.), Proceedings of EUROBOT '99 (pp. 113-119). IEEE.

Balkenius, C. & Johansson, B. (2004). The Coding of Context in Visual Scenes. unpublished manuscript.

Balkenius, C. & Johansson, B. (2005). Event Prediction and Object Motion Estimation in the Development of Visual Attention. In Proceedings of EpiRob '05.

Balkenius, C., & Björne, P. (2001). Toward a robot model of attention-deficit hyperactivity disorder (ADHD). In Balkenius, C., Zlatev, J., Kozima, H., Dautenhahn, K., & Breazeal, C.

(Eds.), Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems. Lund University Cognitive Studies, 85.

Balkenius, C., & Morén, J. (2000). A computational model of context processing. In Meyer, J-A., Berthoz, A., Floreano, D., Roitblat, H. L., Wilson, S. W. (Eds.), From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour. Cambridge, MA: MIT Press.

Balkenius, C., & Morén, J. (2000). A computational model of context processing. In Meyer, J-A., Berthoz, A., Floreano, D., Roitblat, H. L., Wilson, S. W. (Eds.), From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour. Cambridge, MA: MIT Press.

Balkenius, C., & Morén, J. (2000). Emotional learning: a computational model of the amygdala. Cybernetics and Systems, 32, 6, 611-636.

Balkenius, C., & Winberg, S. (2004). Cognitive modeling with context sensitive reinforcement learning. In Proceedings of AILS '04. Lund: Dept. of Computer Science.

Bandura, A. (1999). A social cognitive theory of personality. In L. Pervin & O. John (Ed.), Handbook of personality (2nd ed., pp. 154-196). New York: Guilford Publications. (Reprinted in D. Cervone & Shoda [Eds.], The coherence of personality. New York: Guilford Press.)

Barsalou, L. W. (1999). Perceptual symbol systems. Behavioral & Brain Sciences, 22, 577-600.

Bjöne, P., & Balkenius, C. (2005). The role of context and extinction in ADHD. Behavioral & Brain Sciences, in press.

Boden, M. (2004). Generalization by symbolic abstraction in cascaded recurrent networks. Neurocomputing.

Bouton, M. E. (1993), Context, time, and memory retrieval in the interference paradigms of Pavlovian learning, Psychol Bull. 1993 Jul;114(1):80-99.

Bouton, M. E. & King, D. A. (1983). Contextual control of the extinction of conditioned fear: tests for the associative value of the context. J Exp Psychol Anim Behav Process. 1983 Jul;9(3):248-65

Bratman, M. E. (1987). *Intentions, plans, and practical reason*, Cambridge, MA: Harvard University Press.

Brooks, R.A., (1991). How to build complete creatures rather than isolated cognitive simulators," in K. VanLehn (ed.), Architectures for Intelligence, pp. 225-239, Lawrence Erlbaum Assosiates, Hillsdale, NJ.

Butz, M. V. (2002). Anticipatory learning classifier systems. Boston, MA:
Kluwer Academic Publishers.

Butz, M. V., & Goldberg, D. E. (2003). Generalized State Values in an Anticipatory Learning Classifier System In Butz, M. V., Sigaud, O., & Gérard, P. (Eds.), Anticipatory Behavior in

Adaptive Learning Systems: Foundations, Theories, & Systems (LNCS 2684), 282-301, Springer Verlag, Berlin.

Butz, M. V., Goldberg, D. E., & Stolzmann, W. (2002). The anticipatory classifier system and genetic generalization. Natural Computing, 1, 427-467.

Butz, M. V., & Hoffmann, J. (2002). Anticipations control behavior: Animal behavior in an anticipatory learning classifier system. Adaptive Behavior, 10, 75-96.

Castelfranchi, C. (1995). Representation and integration of multiple knowledge sources: issues and questions. In Cantoni, Di Gesu', Setti e Tegolo (Eds.), *Human & Machine Perception: Information Fusion*, Plenum Press.

Castelfranchi, C., Falcone, R., (1998a) Towards a Theory of Delegation for Agent-based Systems, *Robotics and Autonomous Systems*, Special issue on Multi-Agent Rationality, Elsevier Editor, Vol 24, Nos 3-4, pp.141-157.

Castelfranchi C., Falcone R., (1998b) Principles of trust for MAS: cognitive anatomy, social importance, and quantification, *Proceedings of the International Conference on Multi-Agent Systems (ICMAS'98)*, Paris, July, pp.72-79.

Choi, S-B., Ban, S-W. and Lee, M. (2004). Biologically Motivated Visual Attention System Using Bottom-up Saliency Map and Top-down Inhibition. Neural Information Processing-Letters and Reviews. 2, 1, 19-25.

Chun, M. M. (2000). Contextual cueing of visual attention. Trends in Cognitive Science, 4 (5), 170-178.

Chun, M.M. & Jiang, Y. (1998) Contextual cueing: implicit learning and memory of visual context guides spatial attention. Cognit. Psychol. 36, 28–71.

Cohen, P. R. , Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, pp. 213-261.

Cohen, J. D. & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. Psychological Review. 1992 Jan Vol 99(1) 45-77.

Doya K., Samejima K., Katagiri K. & Kawato M. (2000). Multiple model-based reinforcement learning. Japan Science and Technology Corporation.

De Giacomo, G., Lespérance, Y., Levesque, H. J. (2000). ConGolog, a concurrent programming language based on the situation calculus. *Artificial Intelligence,* 121, pp. 109-169.

Deco, G. & Zhil, J. (2001). A Neurodynamical Model of Visual Attention: Feedback Enhancement of Spatial Resolution in a Hierarchical System. Journal of Computational Neuroscience 10, 231–253.

Del Val, A., Shoham, Y., Maynard-Reid II, P. (1997). Qualitative Reasoning about

Perception and Belief. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence.* San Mateo, CA: Morgan Kaufmann.

Drescher. G. Made-up minds: constructivist approach to artificial intelligence. MIT Press, 1991.

Elgot-Drapkin, J., Kraus, S., Miller, M. Nirkhe, M., Perlis, D. (1999). Active logics: A unified formal approach to episodic reasoning. *Technical report*, Department of Computer Science, University of Maryland.

Falcone R., Castelfranchi C., (2001). Social Trust: A Cognitive Approach, in *Trust and Deception in Virtual Societies* by Castelfranchi C. and Yao-Hua Tan (eds), Kluwer Academic Publishers, pp. 55-90.

Farmer, J.D, Packard H.N & Perelson A.H. (1986). The Immune system, adaption and machine learning. Physica D, 22:187-204, 1986.

Gers, F. A. & Schmidhuber, J. & Cummins, F. (2000). Learning to Forget: Continual Prediction with LSTM. Neural Computation, 12(10):2451--2471.

Gibson, J. J. (1979). The ecological approach to visual perception.Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Gibson, J.J. (1979). The Ecological Approach to Perception, Houghton Mifflin.
Grant, J., Kraus, S., Perlis, D. (2000). A logic for characterizing multiple bounded agents. *Autonomous Agents and Multi-Agent Systems journal*, 3(4):351-387.
Graves, A., Eck, E., Beringer, N., & Schmidhuber, J.(2004). Biologically Plausible Speech Recognition with LSTM Neural Nets. In J. Ijspeert (Ed.), First Intl. Workshop on Biologically Inspired Approaches to Advanced Information Technology, Bio-ADIT 2004, Lausanne, Switzerland, p. 175-18.

Grimes, D. B., & Rao, R. P. N. (2005). Bilinear sparse coding for invariant vision. Neural Computation, 17, 47-73.

Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. Behavioral & Brain Sciences, 27 (3): 377-396.

Goncalves Luiz MG, Gilson A. Giraldi, Oliveira Antonio AF 2004 Grupen RA. (1999), Learning Policies for Attentional Control , Proceedings of the 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA '99), Monterey, CA, USA

Harris, C. and Stephens, M. (1988). A combined corner and edge detector. In Fourth Alvey Vision Conference, Manchester, UK, pp. 147-151.

Hart E., Ross P., Webb A. & Lawson A. (2003). A Role for Immunology in Next Generation Robot Controllers. In Proceedings of ICARIS 2003. Springer Verlag Heidelberg.

Herzig, A., Lang, J., Polacsek, T. (2000). A modal logic for epistemic tests. In *Proceedings of European Conference on Artificial Intelligence (ECAI'2000)*, Berlin, August.

Hochreiter, S. & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8):1735-1780.

Holmes, M. & Isbell, C. (2005). Schema Learning: Experience-Based Construction of Predictive Action Models NIPS 17.

Jerne N.K. (1974). Towards a network theory of the immune system. Ann. Immunol (Inst. Pasteur) 125C. (p.373-389).

Kaelbling, L. S., Littman, M. L., Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101, pp. 99-134.

Kant, I. (1781 / 1998). Critique of Pure Reason. Cambridge University Press.

Kirsh, D., Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, pp. 513-549.

Kokinov, B. (1988). Associative Memory-Based Reasoning: How to Represent and Retrieve Cases. In: T.O'Shea & V.Sgurev (eds.) Artificial Intelligence III: Methodology, Systems, Applications, Elsevier Science Publ.

Kokinov, B. (1994). A Hybrid Model of Reasoning by Analogy. Chapter 5. in: K. Holyoak & J. Barnden (eds.) Analogical Connections, Advances in Connectionist and Neural Computation Theory, vol.2, Ablex Publ. Corp.

Kokinov, B. (1994). The DUAL Cognitive Architecture: A Hybrid Multi-Agent Approach. In: A. Cohn (ed.) Proceedings of ECAI'94. John Wiley & Sons, Ltd., London.

Kokinov, B. (1994). The Context-Sensitive Cognitive Architecture DUAL. In: Proceedings of the 16th Annual Conference of the Cognitive Science Society. Erlbaum, Hillsdale, NJ.

Kokinov, B. (1997). Micro-Level Hybridization in the Cognitive Architecture DUAL. In: R. Sun & F. Alexander (eds.) Connectionist-Symbolic Integration: From Unified to Hybrid Architectures, Lawrence Erlbaum Associates, Hilsdale, NJ, (First published in: Sun, R. & Alexander, F. (eds.) Proceedings of the IJCAI-95 Workshop on Connectionist-Symbolic Integration, 1995.)

Kokinov, B. (1998). Analogy is like Cognition: Dynamic, Emergent, and Context-Sensitive. In: Holyoak, K., Gentner, D., Kokinov, B. (eds.) – Advances in Analogy Research: Integration of Theory and Data from the Cognitive, Computational, and Neural Sciences. Sofia: NBU Press.

Kokinov, B., Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: Gentner, D., Holyoak, K., Kokinov, B. (eds.) The Analogical Mind: Perspectives from Cognitive Science, Cambridge, MA: MIT Press.

Kokinov, B., Grinberg, M. (2001). Simulating Context Effects in Problem Solving with AMBR. In: Akman, V., Thomason, R., Bouquet, P. (eds.) Modeling and Using Context. Lecture Notes in Computer Science (Lecture Notes in Artificial Intelligence), vol. 1775, Springer Verlag.

Kopp, K. (2003). Natural Vision for Artificial Systems: Active Vision and Thought. Lund University Cognitive Studies, 100.

Lee, K.-F. (1990). Context-dependent phonetic hidden Markov models forspeaker-independent continuous speech recognition. Acoustics, Speech, and Signal Processing, 38, 4, 599-609.

Littman, M. L., Sutton, R. S. & Singh, S. . Predictive representations of state. In Advances in Neural Information Processing Systems, pages 1555–1561. MIT Press, 2002.

Lowe, D. G. (2004). Distinctive Image Features from Scale-Invariant Keypoints. International Journal of Computer Vision.

Maes, P. (1990), Situated Agents Can Have Goals, Robotics and Autonomous Systems, 6:49-70.

Meyer, J.-J. Ch., van der Hoek, W., van Linder, B. (1999). A Logical Approach to the Dynamics of Commitments. *Artificial Intelligence*, 113(1-2), pp. 1-40.

McCallum, A. K. (1996). Reinforcement learning with selective perception and hidden state. Ph.D. Thesis.

McCarthy, J. & Buvac, S. (1997). Formalizing Context (Expanded Notes). available at cogprints.

McClelland, James L. & Rumelhart, David E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. Psychological Review. 1981 Sep Vol 88(5) 375-407.

McClelland, James L. & Rumelhart, David E. (1982). An interactive activation model of context effects in letter perception: II. The contextual enhancement effect and some tests and extensions of the model. Psychological Review. Vol 89(1) 60-94

McDonald, Scott & Brew, Chris (2002) A distributional model of semantic context effects in lexical processing. avilable at cogprints.

McFarland D. & Bosser, T. (1994). Intelligent Behaviour in Animals and Robots. MIT Press, Cambridge Ma.

Medin, D. L. & Schaffer, M. M. (1978). Context theory of classification learning. Psychological Review. 1978 May Vol 85(3) 207-238.

Michelan, R. & von Zuben, F. (2002). J. Decentralized control system for autonomous navigation based on an evolved artificial immune system. In C.L. Giles et al, editor, Proceedings of CEC-02 (p.1021-1026). IEEE Press.

Minsky, M..(1975). A framework for representing knowledge. In P. Winston, editor, The Psychology of Computer Vision, pages 211–277. McGraw-Hill, New York.

Morén, J. (2003). Emotion and Learning. PhD Thesis. Lund University Cognitive Studies 93.

Mowrer, O. H. (1960/1973). Learning theory and behavior. New York: Wiley

Navalpakkam, V. 2002 Itti, L. (2005). Modeling the influence of task on attention. Vision Research 45 (2005) 205-231.

Peirce, C. S. (1897 / 1940). Logic as semiotic: The theory of signs. In Philosophical writings of Peirce. Dover.

Petrov, A. & Kokinov, B. (1999). Processing Symbols at Variable Speed in DUAL: Connectionist Activation as Power Supply. In: Dean, T. (ed.) Proceedings of the 16th International Joint Conference on Artificial Intelligence. San Francisco, CA: Morgan Kaufman, p. 846-851.

Piaget J. (1970). La Formation du Symbole Chez L'Enfant. Delachaut et Nestlie.

Piaget, J.. The Construction of Reality in the Child. Ballentine, 1954.

Plaut, D. C. (1995). Semantic and associative priming in a distributed attractor network. Proceedings of the 17th annual conference of the cognitive science society (pp. 37-42). LEA.

Poggio, T., & Bizzi, E. (2004). Generalization in vision and motor control. Nature, 431, 768-774.

Pomerleau, D. (1989). "ALVINN: An Autonomous Land Vehicle In a Neural Network," Advances in Neural Information Processing Systems 1, Morgan Kaufmann.

Rao, R- P- N. & Ballard, D. H. (1996). The Visual Cortex as a Hierarchical Predictor. Technical Report 96.4. National Resource Laboratory for the Study of Brain and Behavior. Department of Computer Science, University of Rochester.

Rao, R. P. N. & Ballard, D. H. .(1997). Dynamic Model of Visual Recognition Predicts Neural Response Properties in the Visual Cortex. Neural Computation. 9(2).

Rao, A. S., Georgeff, M. P.(1991a). Deliberation and Intentions. In *Proceedings of 7th Conference on Uncertainty in Artificial Intelligence*, Los Angeles.

Rao, A. S., Georgeff M. P. (1991b). Modelling rational agents within a BDI-architecture. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann Publishers, San Mateo, CA.

Rattenberger J., Poelz P.M. & Prem E. (2004). Artificial Immune Networks for Robot Control. In Proceedings of Epigenetic Robotics 2004.

Reisfeld, D., Wolfson, H. and Yehezkel, Y. (1995). Context free attentional operators: the generalized symmetry transform. International Journal of Computer Vision, 14, 119-130.

Reiter, R. (2001). *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. Cambridge, MA: MIT Press.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. Nature Neuroscience, 2, 1019-1025.

Roy, D. (2005). Grounding Language in the World: Schema Theory Meets Semiotics. Submitted to Artificial Intelligence.

Rybak I. A., Gusakova V. I., Golovan A. V., Podladchikova, L. N. & Shevtsova N. A. (1998). A model of attention-guided visual perception and recognition. Vision Res. 38, 15-16, 2387-2400

Schank, R. & Abelson, R. (1977). Scripts, Plans, Goals and Understanding. Erlbaum.

Schmidhuber, J. and Huber, R. (1991). Learning to generate artificial fovea trajectories for target detection. International Journal of Neural Systems, 2 (1 & 2): 135-141.

Schmidhuber, J. and R. Huber, R. (1991) Using sequential adaptive neuro-control for efficient learning of rotation and translation invariance. In T. Kohonen, K. Mäkisara, O. Simula, and J. Kangas, editors, Artificial Neural Networks, pages 315-320. Elsevier Science Publishers B.V., North-Holland.

Schmidhuber, J. and Huber, R. (1990). Learning to generate focus trajectories for attentive vision. Technical Report FKI-128-90, Institut für Informatik, Technische Universität München.

Schmidhuber J. (1992). Learning unambiguous reduced sequence descriptions. In Moody J.E., Hanson S.J., Lippman R.P. (Eds.), Advances in Neural Information Processing Systems 4 (NIPS-1992), San Mateo, CA: Morgan Kaufmann, pp. 291-298.

Schmidhuber, J. (1991). A possibility for implementing curiosity and boredom in model-building neural controllers. In J. A. Meyer and S. W. Wilson, editors, Proc. of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats, pages 222-227. MIT Press/Bradford Books.

Schmidhuber, J. (1991). Curious model-building control systems. In Proc. International Joint Conference on Neural Networks, Singapore, volume 2, pages 1458-1463. IEEE.

Schmidhuber, J. (1992). Learning complex, extended sequences using the principle of history compression. Neural Computation, 4(2):234-242.

Shibata, K., Nishino, T. and Okabe, Y. (1995). Active Perception Based on Reinforcement Learning. Proc. of WCNN'95 Washington D. C., 2, pp.170-173.

Shibata, T., Vijayakumar, S., Conradt, J, & Stefan Schaal, S. (2001). Humanoid Oculomotor Control Based on Concepts of Computational Neuroscience. Proceedings of Humanoids 2001, Second IEEE-RAS Intl. Conf. on Humanoid Robots.

Simon, H. A. (1969). Sciences of the Artificial. Cambridge, MA: MIT Press.

Singh, M. P., Asher, N. M. (1993). A Logic of Intentions and Beliefs. *Journal of Philosophical Logic*, 22, pp. 513-544.

Singh, S., Barto, A. G. & Chentanez, N. (2005). Intrinsically Motivated Reinforcement Learning. Proceedings of Advances in Neural Information Processing Systems 17 (NIPS).

Stauffer, C. & Grimson, (1999). W.E.L. Adaptive Background Mixture Models for Real-Time Tracking, CVPR, 2, 2, 2246.

Sutton, R. S. & Barto, A. G. (1998). Reinforcement Learning: An Introduction. Cambridge, MA: MIT Press.

Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. Proceeding of the Seventh International Conference on Machine Learning, pp. 216-224.

Sutton, R.S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. In Working Notes of the 1991 AAAI Symposium, pp. 151- 155. AAAI Press.

Sutton, R.S., Precup, D., Singh, S. (1999). Between MDPs and semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. Artificial Intelligence, 112, pp. 181-211.Trappl, R., Petta, P., Payr, S. (2003). Emotions in Humans and Artifacts. Cambridge, MA: MIT Press.

Torralba A. (2002). Contextual modulation of target saliency. in T. G. Dietterich, S. Becker, & Z. Ghahramani (eds.) Advances in Neural In]ormation Processing Systems 14. MIT Press, Cambridge, MA.

Torralba, A. (2003). Contextual priming for object detection. International Journal of Computer Vision, 52, 2, 169-191.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. Cognitive Psychology, 12, 97-136.

Watanabe, Y, Ishiguro, A., Shiari, Y., & Uchikawa, Y. (1998). Emergent construction of behavior-arbitration mechanism based on the immune system. In Proceedings of the 1998 IEEE World Congress on Computational Intelligence, (pp. 481-486). Los Alamitos, CA: IEEE Computer Society Press.

Van Linder, B., van der Hoek, W., Meyer, J.-J. Ch. (1997). Seeing is believing (and so are Hearing and Jumping). *Journal of Logic, Language and Information*, 6(1), pp. 33-61.

Webb A., Hart E., Ross P., & Lawson A. (2003). Controlling a simulated Khepera with an XCS classifier system. In Proceedings of 7th European Conference on A-Life. Springer-Verlag.

Westin, C.-F. Westelius, C.-J. Knutsson, H. Granlund, G. (1996). Attention control for robot vision. Computer Vision and Pattern Recognition, 1996. Proceedings of CVPR '96 (pp. 726-733).