FP6-511931

Mind RACES

from Reactive to Anticipatory Cognitive Embodied Systems

DELIVERABLE 6.3
"Final report on integration"

Dissemination and Use Plan

Due date of deliverable:
30 / 07 / 2007

Actual submission date:
15 / 09 / 2007

Start date of project:
01 / 10 / 2004

Duration:
36 month

Organization name of lead contractor for this deliverable
ISTC-CNR, Rino Falcone

Revision:
1

Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)		
Dissemination Level		
PU	**Public**	V
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_WP6_N_6.3**
Date:	**15/11/20077**
Work package:	**WP6**
Partner(s):	*All*
Lead Partner:	*ISTC-CNR*
Document status:	*VERIFIED*
Deliverable identifier:	**DELIVERABLE_WP6_N_6.3**

Delivery Slip

	Name	Partner	Date	Signature
From	GIANLUCA BALDASSARRE	ISTC-CNR	14/11/2007	
Verified	GIOVANNI PEZZULO	ISTC-CNR	14/11/2007	
Approved by	RINO FALCONE	ISTC.CNR	14/11/2007	

Files

Software Products	User files
MS Word™ File compressor (.rar)	MindRACES-Deliverable_D6.3. DOC 071115MindRACES-AnnexDeliverableD6_3OnIntegration.rar

Project information

Project acronym:	MindRACES
Project full title:	MindRACES: 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

1 Summary

This report illustrates the outcome of the *integration* research work carried out during the implementation of the EU funded MindRACES project. The scientific work of the project was organised around three thematic Work Packages: WP3: Attention and monitoring; WP4: Goal directed behaviour and analogy; WP5: Emotions. The integration work carried out within the project has the goal to implement cognitive architectures and models that integrate anticipatory cognitive capabilities along various dimensions: (1) within single Work Package; (2) between different Work Packages; (3) between different models and architectures investigated by different project partners. Initially, the report gives an overview of the integrated models produced by the consortium, emphasizing the different anticipatory cognitive capabilities they integrate, the partners involved, and the eventual robotic implementation of the models. Then it presents a description of the single models, aside with a clear indication of the anticipatory cognitive functions they integrate, and refers to specific published and in-preparation papers for details. The work illustrated in the report aims to furnish not only an overview of the integration work carried out within MindRACES, but also a sample of the state-of-the-art research carried out in the fast developing field of anticipatory cognitive systems.

2 Introduction

In the second review meeting, the referees of the project appreciated the work done within the project, and pointed out various directions for improvement, in particular they suggested: (1) to refine the research work carried out to that moment, and in particular to spend most efforts in integrating the models developed in the first two years of the project; (2) to carry out collaborative research work involving more partners to develop integrated models and to have common publications; (3) to spend most "writing efforts" in publishing refereed papers (e.g., vs. producing reports); (4) to collect the research efforts of the project into "project books". The consortium tried to do its best in following such indications. In particular it produced several models that integrate various anticipatory cognitive functions, involved different partners and collaborative work in the development of the various models, as illustrated throughout the report, and published the models in refereed papers., so This report presents an overview of the models, emphasising in particular the anticipatory cognitive functions that they integrate, also indicating specific papers which provide the complete description of the models. These papers have to be considered as an important annex of the report (note that these are only the consortium's papers that integrate various cognitive various anticipatory cognitive functions: during the last year the consortium also produced other papers that aimed to improve single anticipatory functions).

The indication to produce collective books has been followed, too. In particular, Butz et al. (2007) is a post-workshop book that collects the best papers of the third edition of the "workshop on Anticipatory Behaviour In Adaptive Learning Systems (ABiALS 2006)" organised by two partners of the consortium (UW and ISTC). This book not only clearly shows that the project partners played an important role of "catalyst" within the European and International research field of "cognitive anticipatory systems", but it also allowed publishing some theoretical papers produced by the project (e.g. Butz et al., 2007a, see section 3.1). This experience is being followed by a second edited "MindRACES book" fully produced by the consortium partners (Pezzulo et al., in preparation). This book will present various review chapters, which can be considered position papers regarding some important sub-sets of anticipatory functions, together with some other chapters which focus on some models produced during the project and that illustrate the paramount role that anticipation plays in various cognitive functions, from vision to attention, from goal-driven action to planning, and from motivations to emotions. The provisory table of contents of the book in preparation is included in the annex (see file PezzuloFalconeEtAlInPreparation.pdf).

In synergy with this forthcoming book, this report gives an overview of the integrated models produced by the consortium in the last year of the project. In this respect, Table 2 furnishes a broad overview of the main features of the integrated models reviewed in the report. The table lists the models (to which a conventional name has been assigned to ease referencing within the report), and for each of them indicates (a) the anticipatory cognitive functions it integrates, (b) if the model was tested in a real robot, and (d) the partners that contributed to develop it. The table attempts to list the models by following a (rough) order of increasing cognitive complexity, from models involving low-level cognitive functions to those involving high-level cognitive functions.

The table allows classifying the models in broad categories. To this purpose, it is useful to recall that the scientific work of the project was organised around three thematic Work Packages which can be considered as three items of a broad taxonomy of anticipatory cognitive functions:

1) WP3: Attention and monitoring
2) WP4: Goal directed behaviour and analogy
3) WP5: Emotions

This taxonomy allows us to group the contributions/models reported in the table in the following broad categories:

1) *Theoretical contributions*. This are not strictly models but either theoretical analyses of cognitive functions (THEORETICAL ANALYSIS) or presentations of general principles/paradigms which can be used to shape the architectures of anticipatory cognitive systems (IMP and TOTE).
2) *Models integrating cognitive functions within single WPs*. These models integrate more cognitive functions within the same WP, in particular functions involving perception and attention (PAM ATTENTION MODEL), action (RL-SURE_REACH) and planning/reasoning (OBJECT INTERCEPTION MODEL).
3) *Models integrating cognitive functions between WP1 and WP2*. These models integrate functions involving both perception/attention and action. Among these, one focuses on attention and the control of a robotic arm trained with reinforcement learning (RL) (ATTENTION ROBOT ARM), another focuses on active perception and RL using a camera mounted on a mobile robot (CAMERA ROBOT), another focuses on a mobile robot which has to produce motor control and at the same time monitor the outcome of action so as to correct action in the case of a hardware failures (HOLONOMIC ROBOT), while a last model focuses on the interplay between visual/attentive processes and analogy-based predictions and is tested in an Aibo robot (IKAROS-AMBR ROBOT).
4) *Models integrating cognitive functions between WP2 and WP3*. This category includes a model that integrates anticipatory emotional processes and reasoning functions (EMOTION-BDI AGENT).
5) *Models integrating cognitive functions between WP1, WP2 and WP3*. These category includes a model that integrates high-level planning cognition with low level sensorimotor processes, and at the same time integrates them with motivations generated from drives and goals (INTENTIONAL SENSORIMOTOR ARCHITECTECTURE). The category also includes a last model which can potentially be used to tackle in a formal way (machine learning point of view) the most complex human manifestations of anticipatory cognition, such as the capacity for "discovery" and "creativity".

The rest of the report is organised as follows. Section 3 illustrates in detail the features of the various integrated models and architectures. Each sub-section of this section refers to a single model: it initially indicates the project partners that participated to the development of the model and points to the reference(s) of the publication(s) in which the model is described in detail, then reviews the model functioning and the results of their test, and finally clearly indicates the cognitive functions integrated by the model. Section 4 closes the report by drawing conclusions and giving some indication for future work.

Table 1: Schematic overview of the integration work carried out within MindRACES. The table attempted to list the models in a rough ascending order of cognitive complexity, from models involving low-level cognitive functions to those involving high-level cognitive functions. From the left to the right column, the table reports the name of the models, used here to ease referencing, the number of the section of the report that gives the details of the model, the anticipatory functions integrated by the model (with reference to the Work Packages they tackle), the eventual test of the model with a real robot, and finally the project partners that participated to build, implement and test the models. Fields in italics refer to functions involving standard reinforcement learning algorithms (RL): these functions, although involving emotional aspects of behaviour, did not aim to propose particularly innovative results and models within the field of emotions (see text).

		Integrated anticipatory function involving various Work Packages' themes				Partner participation							
MODEL	**Sect**	**WP3 Attention, monitoring**	**WP4 Goal directed behaviour, analogy**	**WP5 Emotions**	**ROBOT**	**ISTC**	**LUCS**	**UW**	**NBU**	**IST**	**OFAI**	**IDSIA**	**NOZE**
THEORETIC. ANALYSIS	4.1	Overview of anticipatory functions	Overview of anticipatory functions	Overview of anticipatory functions		**V**		**V**					
Theoretical analysis: IMP and TOTE	4.2	- IMP: learns action-state models - TOTE: goal-based monitoring to stop/act	- IMP: selection of goal-based action selection - TOTE: acts to reduce state/goal mismatch			**V**		**V**					
RL-SURE_REACH	4.3		- Goal-based control of redundant arm - Flexible selection of goals on the basis of RL	*RL of arm's targets*		**V**		**V**					
PAM ATTENT. MODEL	4.4	- Bottom-up attention for finding cues - Top-down attention for finding target		*RL based on salient events (e.g. target discovery)*		**V**	**V**						
ATTENTION ROBOT ARM	4.5	- Bottom-up attention - Top-down attention	Goal-based motor control of robotic arm	*RL based on salient events*	**V**	**V**	**V**						
CAMERA ROBOT	4.6	Anticipatory vision, attention	Navigation, model-based RL	*RL (model-based)*	**V**							**V**	
HOLONOMIC ROBOT	4.7	Attention to hardware failure in view of behavioural change	Chooses actions based on own model		**V**							**V**	
IKAROS-AMBR ROBOT	4.8	Attention to collect information in view of reasoning	Analogy-based anticipation and reasoning		**V**		**V**		**V**	**V**			
OBJECT INTERCEPT. MODEL	4.9		- World modelling - Planning								**V**		
EMOTION-BDI AGENT	4.10		Goal-directed behaviour and reasoning	Emotivectors and emotions for appraisal of unexpected events and for deliberation		**V**				**V**			
INTENTIONAL SENSORIMOT. ARCHITECT.	4.11	Monitoring of mismatch of effects of actions with predictions	Expectations used to select goals and plans, and these activate sensorimotor schemes	Integration of low-level and high-level motivations (drives and goals)		**V**							**V**
CURIOUS AGENT	4.12	Curiosity-driven exploration	Adaptive prediction	*RL based on anticipation progress*								**V**	

3 Contributions

3.1 THEORETICAL ANALYSIS (UW, ISTC)

Butz et al. (2007); Butz et al. (2007a) *(In cooperation with:* Olivier Sigaud, external partner Animat Lab, LIP6, Paris, France)

Summary and main results

The matter of anticipatory behaviour in adaptive learning systems is steadily gaining interest, although many researchers still do not explicitly consider the capabilities of their systems under the framework of anticipation. Given the importance of anticipation, in 2006 some partners of the consortium (UW and ISTC) organised the third edition of the "workshop on Anticipatory Behaviour In Adaptive Learning Systems (ABiALS 2006)", held on the 30th of September 2006 during the "Simulation of Adaptive Behaviour Conference (SAB 2006)".

This workshop was followed by a post-proceedings book (Butz et al. 2007). The workshop and the book contributed to show that anticipatory cognitive mechanisms for behaviour and learning strongly overlap among researchers from various disciplines, including the whole interdisciplinary cognitive science area, and for this reason there are great opportunities in treating them in an integrated fashion. In this respect, the workshop was perfectly in line with the MindRACES project that contributed to show yet again that there are striking similarities and common underlying principles in different anticipatory functions studied in diverse cognitive systems.

Thus, further conceptualizations of anticipatory mechanisms seem mandatory. The introductory chapter of the book (Butz et al., 2007a) therefore proposes a taxonomy of how anticipatory mechanisms can improve adaptive behaviour and learning in cognitive systems. During the workshop it became clear that anticipations are involved in various cognitive processes that range from individual anticipatory mechanisms to social anticipatory behaviour. This book reflects this structure by first providing neuroscientific as well as psychological evidence for anticipatory mechanisms involved in behaviour, learning, language, and cognition. Next, individual predictive capabilities and anticipatory behaviour capabilities are investigated. Finally, anticipation relevant in social interaction is studied.

Anticipatory behaviour research on cognitive, adaptive systems aims at exploiting the insights gained from neuroscience, linguistics, and psychology for the improvement of behaviour and learning in artificial cognitive systems. However, this knowledge exchange is expected to become increasingly bidirectional. That is, the insights gained during the design and evaluation of different anticipatory cognitive mechanisms and architectures may also provide insights into how anticipatory mechanisms can actually shape, guide, and control natural brain activity. This book reveals many interesting and thought-provoking connections between distinct cognitive science areas. We strongly hope that these connections do not only lead to a deeper understanding of the functioning anticipatory processes but also enable a more effective, bidirectional knowledge exchange and consequently more effective scientific progress in the natural and artificial cognitive systems research disciplines.

3.2 IMP and TOTE (UW, ISTC)

Pezzulo et al. (2007)

Summary and main results

An important theoretical chapter (Pezzulo et al., 2007) of the aforementioned book on anticipation (Butz et al., 2007), tackles a fundamental issue for anticipation, namely the nature and role of *goals*. Goals are at the core of several cognitive anticipatory functions as they represent *desired future states*, that is future states of the world that can guide agent's action so that it will eventually make them happen. The concept of goal is a fundamental pillar of anticipatory behaviour aside *predictions*, which are anticipated future states the world will assume either spontaneously or under the effects of the agents' action. Goals differ from predictions in that they have a "theleonomic charge" that allow them to direct agents' actions. Indeed, while goals *cause* actions, predictions *are caused* by action. As we shall see throughout the report, the concepts of goals and predictions play a central role in cognitive functions.

How can goals be represented in natural and artificial systems? How can they be learned? How can they trigger actions? The paper describes, analyses and compares two of the most influential models of goal-oriented behaviour: the ideomotor principle (IMP), which was introduced in the psychological literature, and the ``test, operate, test, exit" model (TOTE), proposed in the field of cybernetics. This analysis indicates that the IMP and the TOTE highlight complementary aspects of goal-orientedness. In order to illustrate this point, the paper reviews three computational architectures that implement various aspects of the IMP and the TOTE, discusses their main peculiarities and limitations, and suggests how some of their features can be translated into specific mechanisms in order to implement them in artificial intelligent systems.

Anticipatory cognitive functions integrated

The paper contrasts the two important approaches to goal-oriented behaviour, the IMP and the TOTE, but also shows that both integrate important aspect of anticipatory behaviour. In particular, both principles have the premise that anticipatory behaviour starts with a goal representation. However, the IMP focuses on learning aspects and less how the goal triggers the behaviour and when the goal is achieved. In doing so, it stresses the importance of the execution of action and also the monitoring of their consequences which allow the system to form action-effect associations. On the other side, TOTE focuses more on the "Test" aspect of the achievement of goals, that is on the monitoring of the effects produced by actions with the purpose of persevering action execution or stopping it. The paper shows how various architectures, further developed within the project (see the following sections on the RL-SURE_REACH and ATTENTION ROBOT ARM), incorporate various principles underlying the IMP and TOTE.

3.3 RL-SURE_REACH (UW, ISTC-CNR)

Herbort et al. (2007)

Summary and main results

This work produced a developmental neural network model of motor learning and control, called RL-SURE_REACH, capable of controlling a simulated robotic arm by activating suitable goals for it (goal-based action selection). The RL component of the system allowed it to reach a rewarded target while at the same time avoiding penalties. The selection process was realized with an actor-critic neural-based RL model for arm control developed by ISTC within the project MindRACES. This was integrated with an anticipatory goal-directed behavioural model for redundant-arm control (the SURE_REACH architecture) studied and improved by UW within the project.

In a "childhood phase", a motor controller for goal directed reaching movements with a redundant arm develops in an unsupervised fashion action-effects association in accordance with the Ideomotor Principle. In subsequent task-specific learning phases, the neural network acquires goal-modulation skills. These skills enable RL-SURE_REACH to master a task that was used in a psychological experiment by Trommershauser, Maloney, and Landy (2003). This task required participants to select aim points within targets that maximize the likelihood of hitting a rewarded target and minimizes the likelihood of accidentally hitting an adjacent penalty area. The integrated model RL-SURE_REACH showed to adapt similarly to tested humans on the aim-pointing selection task. In particular, it exhibited an accurate pointing behaviour for rewarded target areas and, at the same time, a varying pointing distance from punished regions next to the target, with the distance positively related to the level of punishment.

Anticipatory cognitive functions integrated

RL-SURE_REACH integrates two major anticipatory functions of motor control: (1) Flexible goal-directed behaviour based on redundant anticipatory representations of alternative goal postures and paths to the goal; (2) Flexible selection of goal locations based on an actor-critic RL technique, which resulted from predictive reward representations.

3.4 PAM ATTENTION MODEL (ISTC, LUCS)

Ognibene et al. (in press).

Summary and main results

How can visual selective attention guide eye movements so as to collect information and identify targets potentially relevant for action? Many models have been proposed that use the statistical properties of images to create a dynamic bottom-up *saliency map* used to guide saccades to potentially relevant locations. Since the concept of saliency map was introduced, it has been incorporated in a large number of models and theories.

These bottom-up mechanisms have been enhanced with top-down processes in models that learn to move the eye in search of the target on the basis of foveated objects. In many of these systems, top-down attention is guided by task-related information that is acquired through automatic learning procedures.

The paper proposed a novel model that improves on this type of top-down mechanisms by using an eye-centred *potential-action map* (PAM). The PAM keeps track of all the potential locations of targets based on the information contained in a sequence of fixation. In this respect, the PAM works as a short term memory for potential target locations. Each fixation suggests potential locations for targets or other relevant cues and the evidence for each possible location is accumulated in the PAM. Overall, the PAM makes up an efficient mechanism for accumulating evidence for potential target locations in a action-oriented compact format readily usable for controlling eye movements.

Anticipatory cognitive functions integrated

The model integrates two important anticipatory functions within attention: (1) Bottom-up attention processes, based on various filters, that detect information-rich portions of the image at a gross grain: this allows the system to anticipate where to concentrate fine-grain visual processing; (2) Top-down attention processes, based on the PAM, which encode the chances of finding the targets in various regions of the image; (3) a RL components that learns to use image information to activate the PAM.

3.5 ATTENTIONAL ROBOTIC ARM (ISTC, LUCS)

Ognibene et al. (in preparation).

Summary and main results

This model will be presented in the final robotic demo of MindRACES. The model is a neural-netework architecture resulting from the integration of the PAM ATTENTION MODEL and the controllers of the robotic arm developed during the project by ISTC. The PAM ATTENTION MODEL component functions as illustrated in the previous section. The second component is composed by a goal-based action selection controller for robotic arms which functions on similar basic principles as the RL-SURE_REACH system: (1) "childhood's" learning of actions-effects associations (according to the ideomotor principle); (2) "dynamic competition" to select goals for the arm (implemented with a dynamic neural network); (3) triggering of actions on the basis of the activation of goals (here corresponding to desired arm's end-points).

The two components of the system contributed to the functioning of the overall system in an integrated fashion. In particular, the attention components identify potentially highly informative regions of the image and learn to scan the image in search of the target. Moreover, attention can learn to stay on the target if this is rewarded. The attentional system signals to the arm controller the currently foveated region, and this signals is used to bias the "dynamic competition" of the arm controller. The result is that if the attentional system finds and fixates the target, the posture of the arm corresponding to it will be triggered, and the arm will reach the target. Note that in the current implementation of the architecture, the attention component is directly trained on the basis of its capacity to foveate the target. However, in future implementations the arm movement will produce a reward when it results in a contact with the target. The reward obtained by the arm will be used to train the attention system by RL, so closing the loop between the two components of the architecture.

Anticipatory cognitive functions integrated

The architecture integrates various functions both within the WPs themes and between the WPs themes. In particular, the attention system integrates various anticipatory aspects of bottom-up and top-down attention (see previous section on the PAM ATTENTION MODEL). Attention scans the images so as to collect the information needed to guide the arm. The arm controller uses a goal-based action selection mechanism and interacts with the environment so as to obtain rewards. Finally, the actor-critic RL components of the system use information on past rewards to anticipate future rewards and drive the learning processes of the system.

3.6 CAMERA ROBOT (IDSIA)

Bakker (2006). Zhumatiy et al. (2006);

Summary and main results

A robot equipped with a color camera is placed into a room. The task is to find and move to a unique colored cup randomly placed in the room. The camera is mounted in front of the robot and looks a bit downwards. It has a very limited field of view in relation to the room. Therefore, the robot has to find the cup before it can move to it on the basis of other items spread in the room.

The controller of the robot translates sensory input data to the robot's movement commands. It is trained with different reinforcement learning methods. In Zhumatiy (2006) the mean position of all camera pixels in a specific color range of the target object is used as input for the reinforcement learner. To reduce the huge amount of memory for the policy a Piecewise Continuous Nearest-Sequence Memory (PC-NSM) algorithm is used for general metrics over state-action trajectories. In Bakker (2006) the visual information from the camera is preprocessed into a 5x4 binary grid, which represents the position of the cup in the camera image, if the cup is visible. To reduce the general long training time for reinforcement learning algorithms for real robots a probabilistic world-model is learned from less real robot experiments. This world-model is then used to make mental experiments on this model to train the controller with Prioritized Sweeping, a modified version of the standard Q-Learning algorithm. The policy is applied with a high repetition rate during the learning process of the mental model and with a real time repetition rate in the physical world.

Anticipatory cognitive functions integrated

The system uses a probabilistic world-model to estimates future rewards and stats by mental experiments, and to speed up learning. The system integrates active vision and attention, Goal directed behaviour, and action control and monitoring.

3.7 HOLONOMIC ROBOT (IDSIA)

Rojas et al. (2006).

Summary and main results

Omnidirectional robots learn to correct inaccuracies while driving, or even learn to use corrective motor commands when a motor fails, both partially and completely. A feed forward neural network with historic and current information of robot's poses is used for learning the robot's response to the commands. The learned model can be used to predict deviations from the desired path, and take corrective actions in advance, thus increasing the driving accuracy of the robot. The model can also be used to monitor the robot and assess if it is performing according to its learned response function. The paper shows that even if a robot loses power from a motor, the system can relearn to drive the robot in a straight path, even if the robot is a black box for the controller and we are not aware of how the commands are applied internally.

Anticipatory cognitive functions integrated

The robot anticipates how its actions influence the environment. It uses this knowledge to choose the best action fulfilling its intention in the future. The robot also observes itself to detect drifting effects of its actions and adapt its own world-model. The model integrates action control (it choose actions bases on its own self-model), attention (in case of hardware failures it change its behaviour) and monitoring (it observes actions to detect deviations from expected effects).

3.8 IKAROS-AMBR (NBU-LUCS-IST)

Kiryazov et al. (in press); Petrov et al. (2007).

Summary and main results

The papers outline an approach to building robots with anticipatory behaviour based on analogies with past episodes. Anticipatory mechanisms are used to make predictions about the environment and to control selective attention and top-down perception. An integrated architecture is presented that perceives the environment, reasons about it, makes predictions and acts physically in this environment. The architecture is implemented in an AIBO robot. It successfully finds an object in a house like environment. The AMBR model of analogy-making, developed by NBU, is used as a basis for top down perception and selective attention and it is extended with new mechanisms for anticipation related to analogical transfer. The bottom up visual processing is performed by the IKAROS system for brain modelling, developed by LUCS. The chapter describes some experiments performed with the AIBO robot controlled by the architecture, and demonstrates the usefulness of the analogy-based anticipation approach. The papers also report a comparison with human data to ensure psychological validity of the anticipatory aspects of the model.

NBU also conducted some studies, at the theoretical level, for including emotions in the DUAL architecture, in particular to use some mechanisms of the *FearNot* model. This study showed that emotions can influence AMBR mechanism of analogy making basically in two ways: (1) Emotions can be attached to the description of episodes, thus past episodes with similar felt emotion will be retrieved with a higher chance; (2) Emotions ‘felt’ by the model modify some of the parameters of the architecture which influence the volume of working memory, the speed with which agents enter/leave WM, etc.

Anticipatory cognitive functions integrated

The paper presents a model which provides mechanism of anticipations by analogy. The robot makes a prediction about the environment based on analogy with past episode of its experience. Then it acts in it based on the prediction. The model also uses anticipation in scene representation building. Relations and properties of the objects in the world are predicted via the mechanisms of analogy making. The perception system also checks if predictions are correct. A theoretical study also identified possible ways for using emotions to improve memory processes for episode recall and to regulate on the fly important parameters of the architecture.

3.9 OBJECT INTERCEPTION MODEL (OFAI)

Lewandowski A. (in preparation).

Summary and main results

This paper proposes a robot with the ability to anticipate the location of reappearance of a moving target, usually a ball, which is temporarily hidden behind a wall. The images taken by the robot are simplified and transformed into a smaller number of sector views, whereby each sector is assigned to one of the states "target", "wall" or "background". Based on the current observed sector view an action sequence is started, which is either followed until a different sector view has occurred or a pre-defined number of steps have elapsed. The goal for the robot is to lower the distance to the object under a given threshold in the smallest number of steps. During training rewards are derived from the average number of steps needed to reach a new sector view and a matrix of estimated transition probabilities is updated with every transition. Bumping into walls is penalized. Accessible states which belong to views with the so far observed maximum number of target sectors are declared as goal states. We alter the transition and the reward matrices to allow the application of known optimization algorithms to find a path to the goal states. We validate the algorithm with simulation experiments.

The model learns transition probabilities from observed state to observed state if a certain action sequence is performed. Together with dynamically generated goal states and a given forecast horizon H (e.g. H=20), "best actions" are recursively derived, starting with the end of the horizon t+H and going backwards until we arrive at the current time t. With this construction the robot is able to deduce which action it should choose for every possible and already known state for every time point between now and the horizon, based on the so far observed data.

The system only follows the advice for the current time as, after each action execution, the transition matrix is updated and a new plan is calculated. Importantly, the system can choose an action at t which maximizes an expected reward during the time span from t to t+H and hopefully will lead to one of the goal states as quick as possible. The probabilities of future events are coded in the transition matrix and can be estimated for every assumed chain of action sequences. Assuming that our model is a good description of reality, we are able to predict the outcome of an arbitrary chain of action sequences. The outcome is here a probability distribution over known states including the goal states.

Anticipatory cognitive functions integrated

With experience, the model learns a model of the world in the form of transition probabilities between states. The model of the world is used to plan movements and to trigger actions. Planning involves prediction the consequences of future action sequences on the basis of the world model.

3.10 EMOTION-BDI AGENT (ISTC, IST)

Piunti et al. (2007).

Summary and main results

ISTC-CNR and IST worked on the integration of their models of emotions; the integration aimed at improving and integrating the ISTC's BDI architecture with some of the anticipatory and affective components developed by IST. In particular, the integration involved the "emotivector" component, developed by IST and the goal directed architecture developed by ISTC-CNR. The resulting integrated architecture was used for guiding agents and orienting reasoning processes according to expectations both on low and high cognitive level. The agents were capable of performing deliberation and practical reasoning and of exploiting emotions to affect decisions and to dynamically adapt to dynamic environments. Moreover, they appraised external events as opportunities or threats and recruited additional resources to pro-actively respond to them. The architecture was tested with experiments using a dynamic scenario.

Anticipatory cognitive functions integrated
The emotivector component was used for: 1) building predictions and expectations of future states on the basis of a monitoring signal; 2) cognitively appraising mismatches between the predicted-expected signal and the effective perceived input; 3) producing emotions in order to modulate expectations and affect deliberation and future decisions.

3.11 INTENTIONAL SENSORIMOTOR ARCHITECTURE (ISTC, NOZE)

Pezzulo et al. (2007); Pezzulo et al. (submitted).

Summary and main results
We studied and implemented an integrated architecture to control a simulated agent. The architecture is composed of two “layers”, an intentional layer, that is able to plan and deliberate actions by reasoning on declarative knowledge, and a sensorimotor layer, that is able to execute the plans by means of sensorimotor schemas. The aim of our study was understanding how intentional actions (formulated at a quite abstract level) can be implemented by means of situated motor actions. For example, how the distal goal of preparing a coffee is first deliberated and then actually achieved by chaining the right sequence of proximal actions.

In the work we focussed on: (1) the relations between two different levels of the intention-to-action hierarchy, i.e. the intentional layer (that includes beliefs, plans and goals) and the sensorimotor layer (that includes schemas for situated action); (2) how representations and processes at the intentional layer are grounded in sensorimotor activity; (3) how an agent architecture can deliberate and achieve its goals while remaining responsive to opportunities offered by its environment (i.e., the dynamics between top-down and bottom-up motivations and 'pressures'); (4) how low-level and high-level motivations interact, respectively, in terms of drives and goals.

Both layers of the system exploit anticipation, although in two different ways. At the sensorimotor layer, expectations are produced by the sensorimotor schemas (that include a forward model able to predict the next sensed stimuli, given the actual state and motor command) that serve to control action and to select the right context/action. At the intentional layer, these expectations are used to ground beliefs and other epistemic states. Expectations about the (distal) results of the agent's actions are used to select goals and plans, which in turn activate the appropriate sensorimotor schemas.

The results of the tests of the architecture show that it outperforms both traditional deliberate agent architectures (such as BDI) and situated agent models such as behaviour-based architectures (wich lack the intentional level), in complex environments in which the ability to achieve both proximal and distal goals is required.

Anticipatory cognitive functions integrated
The layered architecture integrates the intentional and sensorimotor levels of action control. There are two main aspects of such integration: (1) Declarative knowledge at the high level (e.g. beliefs that are used for reasoning and selecting among multiple possible goals) is grounded in the patterns of activity of sensorimotor schemas that run either in the real world or 'in simulation' (i.e., a re-enactment is required in which predictions take the place of real stimuli); (2) Plans selected at the intentional level can activate or inhibit sensorimotor schemas which permits to achieve long-term planned goals that require a sequence of actions.

3.12 CURIOUS AGENT (IDSIA)

Schmidhuber (2007); Schmidhuber (2007a)

Summary and main results
The papers postulate that human or other intelligent anticipatory agents function or should function as follows. They store all sensory observations as they come - the data is holy. At any time, given some agent's current coding capabilities, part of the data can be anticipated and is compressible by a short and hopefully fast program / description / explanation / world model. In the agent's subjective eyes, such data is more regular and more "beautiful" than other data. It is well-known that knowledge of regularity and repeatability may improve the agent's ability to plan actions leading to external rewards. In absence of such rewards, however, known beauty is boring. Then "interestingness" becomes the first derivative of subjective beauty: as the learning anticipatory agent improves its compression algorithm, formerly apparently random data parts become subjectively more regular and beautiful. Such progress in compressibility is measured and maximized by the curiosity drive: create action sequences that extend the observation history and yield previously unknown / unpredictable but quickly learnable algorithmic regularity. We discuss how all of the above can be naturally implemented on computers, through an extension of passive unsupervised learning to the case of active data selection by anticipatory agents: we reward a general reinforcement learner (with access to the adaptive compressor) for actions that improve the subjective compressibility of the growing data. An unusually large breakthrough in anticipation and compressibility deserves the name "discovery". The "creativity" of artists, dancers, musicians, pure mathematicians can be viewed as a by-product of this principle. Several qualitative examples support this hypothesis.

Anticipatory cognitive functions integrated
The system integrates various anticipatory functions: (1) an adaptive anticipatory predictor plays the role of history compressor; (2) anticipation progress is measured and serves to compute the reward of the curious exploring agent; (3) a significant improvement of anticipatory power can be interpreted as a "discovery"; the "creativity" of artists, dancers, musicians, and pure mathematicians can be explained as a consequence of the intrinsic reward obtained for such discoveries.

4 Conclusions

This report presented an overview of the research work carried out within the EU project MindRACES directed to design and implement architectures that integrate various anticipatory cognitive functions. In particular, the report presented a number of models developed in the course of the project that integrated cognitive functions both *within* the fields of vision/attention, action, and emotion, and *between* them. Most integrated models were the outcome of collective research work involving at least two partners of the consortium. Moreover, most of them were validated by the peer-review process involved by the publications with which the models were diffused among the scientific community.

Overall, we believe that the theoretical work carried out within the project, and the models developed by the consortium, fully achieve the goals of the MindRACES project. In particular, this work not only contributes to the advancement of the state of the art of the existing specific models of anticipatory cognitive functions, but it also proposes innovative ways of integrating these functions in the form of both theoretical analyses and whole integrated architectures.

5 References

Note: a '*' marks the papers included in the annex of the report

* Bakker, B., Zhumatiy, V., Gruener, G., Schmidhuber, J. (2006). Quasi-online reinforcement learning for robots. *Proceedings of the International Conference on Robotics and Automation (ICRA06)*.

Butz, M. V., Sigaud, O., Pezzulo, G., Baldassarre, G. (Eds.). (2007). *Anticipatory Behaviour in Adaptive Learning Systems: From Brains to Individual and Social Behaviour*, LNAI 4520 (State-of-the-Art Survey). Berlin, Springer Verlag.

* Butz, M.V., Sigaud, O., Pezzulo, G., Baldassarre, G. (2007a). Anticipations, brains, individual and social behaviour: An introduction to anticipatory systems. In Butz, M.V., Sigaud, O., Pezzulo, G., Baldassarre, G. (Eds.), *Anticipatory Behaviour in Adaptive Learning Systems: From Brains to Individual and Social Behaviour,* pp. 1-18. LNAI 4520. Berlin: Springer Verlag.

* Herbort, O., Ognibene, O., Butz, M.V., Baldassarre, G. (2007). Learning to select targets within targets in reaching tasks. In Demiris Y., Scassellati B., Mareschal D. (Eds.), *The 6th IEEE International Conference on Development and Learning (ICDL07).* E1-6. London: Imperial College.

* Kiryazov, K., Petkov, G., Grinberg, M., Boicho, K., Balkenius, C. (2007). The interplay of analogy-making with active vision and motor control in anticipatory robots. In Butz, M. V., Sigaud, O., Pezzulo, G., Baldassarre, G., (Eds.), *Anticipatory Behaviour in Adaptive Learning Systems: From Brains to Individual and Social Behaviour.* LNAI 4520, Berlin, Springer Verlag.

* Lewandowski, A. (in preparation). *Learning to anticipate a temporarily hidden moving object.*

Ognibene, D., Balkenius, C., Baldassarre, G. (in preparation). Integrating attention and goal-based action selection for controlling a robotic arm.

* Ognibene, D., Balkenius, C., Baldassarre, G. (in preparation). A reinforcement-learning model of top-down attention based on a potential-action map. In Pezzulo, G., Falcone, R., Castelfranchi, C. (eds.). *The Anticipatory Approach.* Berlin, Springer Verlag.

* Petkov, G., Kiryazov K., Grinberg, M., Kokinov, B. (2007). Modeling Top-Down Perception and Analogical Transfer with Single Anticipatory Mechanism. In *Proceedings of the Second European Cognitive Science Conference.* Greece.

* Pezzulo, G., Baldassarre, G., Butz, M.V., Castelfranchi, C., Hoffmann, J. (2007). From Actions to goals and vice-versa: theoretical analysis and models of the ideomotor principle and TOTE. In Butz, M.V., Sigaud, O., Pezzulo, G., Baldassarre, G. (Eds.), *Anticipatory Behaviour in Adaptive Learning Systems: From Brains to Individual and Social Behaviour,* pp. 73-93. LNAI 4520, Berlin, Springer Verlag.

Pezzulo, G., Calvi, G. (submitted). Grounding intentional action in a layered agent architecture.

* Pezzulo, G., Calvi, G., Castelfranchi, C. (2007). DiPRA: Distributed Practical Reasoning Architecture. *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence,* pp. 1458-1464.

* Pezzulo, G., Falcone, R., Castelfranchi, C. (in preparation). *The Anticipatory Approach.* Berlin, Springer Verlag.

* Piunti, M., Goncalves, J., Martinho, C. (2007). Modeling expectations for affective agents. In *Proceedings of 2nd International Conference on Affective Computing and Intelligent Interactions (ACII07).* Lisbon.

* Rojas R., Förster, G. A. (2006). Holonomic control of a robot with an omnidirectional drive. *KI - Künstliche Intelligenz,* vol. 20, n. 2. BöttcherIT Verlag.

* Schmidhuber, J. (2007). Simple algorithmic principles of discovery, subjective beauty, selective attention, curiosity & creativity. In Corruble, V., Takeda, M., Suzuki, E. (Eds.), *Proceedings of the 10th International Conference on Discovery Science (DS07),* pp. 26-38. LNAI 4755, Berlin, Springer Verlag.

Schmidhuber, J. (2007a). Simple algorithmic principles of discovery, subjective beauty, selective attention, curiosity & creativity. In Hutter, M., Servedio, R. A., Takimoto, E. (Eds.), *Proceedings of the 18th International Conference on Algorithmic Learning Theory (ALT07)* p. 32. LNAI 4754, Berlin, Springer Verlag.

Zhumatiy, V., Gomez, F., Hutter, M., Schmidhuber, J. (2006). Metric state space reinforcement learning for a vision-capable mobile robot, *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS06).* Tokyo.