

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/249606299>

Endowing Artificial Systems with Anticipatory Capabilities: Success Cases

Article · July 2008

DOI: 10.1007/978-3-540-87702-8_11

CITATIONS

0

READS

59

12 authors, including:



Giovanni Pezzulo

Italian National Research Council

326 PUBLICATIONS 9,098 CITATIONS

[SEE PROFILE](#)



Martin Butz

University of Tuebingen

301 PUBLICATIONS 4,707 CITATIONS

[SEE PROFILE](#)



Cristiano Castelfranchi

Italian National Research Council

398 PUBLICATIONS 11,125 CITATIONS

[SEE PROFILE](#)



Rino Falcone

Italian National Research Council

156 PUBLICATIONS 3,231 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Hybrid Cognitive Architectures for Artificial Agents [View project](#)



Human Flexible Goal-Directed Manipulation of Internal Representations [View project](#)

Chapter 11

Endowing Artificial Systems with Anticipatory Capabilities: Success Cases

Giovanni Pezzulo, Martin V. Butz, Cristiano Castelfranchi, Rino Falcone, Gianluca Baldassarre, Christian Balkenius, Alexander Förster, Maurice Grinberg, Oliver Herbort, Kiril Kiryazov, Boicho Kokinov, Birger Johansson, Emilian Lalev, Emiliano Lorini, Carlos Martinho, Maria Miceli, Dimitri Ognibene, Ana Paiva, Georgi Petkov, Michele Piunti, and Vin Thorsteinsdottir

11.1 Introduction

This book has provided various theoretical perspectives on anticipatory processes in natural and artificial cognitive systems. Advantages have been proposed and confirmed in various detailed case studies, which may have given the reader detailed insights into anticipatory processes and their importance in various cognitive systems tasks. To wrap up these advantages and give a concluding overview of various current anticipatory process advantages, this final chapter highlights a concise collection of precise success stories of anticipations in artificial cognitive systems. We survey fourteen case studies, which were developed during the EU project **MindRACES**¹. In these studies, simulated or real robots were tested in different environmental tasks, which required advanced sensorimotor and cognitive abilities. These abilities included the initiation and control of goal-directed actions, the orientation of attention, finding and reaching goal locations, and performing mental experiments for action selection. All the studies have shown advantages of anticipatory mechanisms compared to reactive mechanisms in terms of increased robot autonomy and adaptivity. In some cases, anticipations even caused the development of new cognitive abilities, which were simply impossible without anticipatory mechanisms. For each case study, we indicate relevant associated publications, in which the interested reader may find further details on the relevant computational architectures, the involved anticipatory mechanisms, as well as on the analytical and quantitative results.

While the book as a whole has laid out the theoretical principles and design methodology for such advancements, this final chapter thus provides various possible starting points for further developments in both the surveyed system architectures and the presented solutions to the cognitive tasks addressed.

¹ MindRACES: From Reactive to Anticipatory Cognitive Embodied Systems; funded under grant FP6-511931 under the “Cognitive Systems” initiative from the EC.

11.2 Flexible Goal-Directed Arm Control: The SURE_REACH Architecture

SURE_REACH (a loose acronym for *Sensorimotor, Unsupervised, REdundancy-REsolving control ArCHitecture*) is a hierarchically structured, self-supervised learning control architecture (Butz et al., 2007a; Herbort and Butz, 2007). Initially, SURE_REACH explores the interactions of its associated body with the environment by means of random motor babbling. During these explorations, it forms associative models of motor-dependent sensory correlations, which is referred to as a sensorimotor model, and of bodily inverse kinematics. The developing knowledge of SURE_REACH about its body and environment is represented by population-encoded spatial body representations, and associative structures that correlate the body representations with each other.

SURE_REACH has been applied to the control of a three-degree of freedom arm in a two-dimensional environment. Due to this setup, each potential target position can be reached with various arm goal postures and on various paths from the current posture to these postures. Given a goal location, SURE_REACH consequently determines not only the most suitable target posture but also the currently optimal path to that posture online. Due to the population encoding, the associative correlations, and the goal-directed, anticipatory behavior approach, it has been shown that SURE_REACH is able to represent multiple problem solutions implicitly in parallel and thus is able to adjust its behavioral policy highly flexibly to current task constraints and changing environmental circumstances.

In the original setup (Butz et al., 2007a; Herbort and Butz, 2007), SURE_REACH represents an *extrinsic hand space*, which encodes hand locations (x-y coordinates) with a uniformly distributed, partially overlapping 2-D array of neurons. Additionally, an *intrinsic posture space* similarly encodes arm postures with a uniformly distributed, partially overlapping 3-D array of neurons (shoulder, elbow, and wrist angles). SURE_REACH learns to correlate these two spaces with each other in a *posture memory*, which associates hand with posture space neurons—effectively encoding an inverse kinematics model—and a *sensorimotor bodyspace model*, which associates postures with each other action dependently. The resulting system is able to predict which posture is reached given a current posture and chosen motor command. More importantly, though, it is also able to deduce the posture that *preceded* a given posture given some action was executed. This latter capability enables the system to choose motor commands when given a current posture and a goal posture. If the goal posture(s) are not in the immediate vicinity, then dynamic programming can be used to generate potential fields in the representation that are able to guide the arm to the goal posture by means of closed-loop control.

The representation together with the goal-directed behavior control processes enable highly flexible and adaptive anticipatory behavior control. Essentially, due to the anticipatory, redundant encoding of behavior alternatives, the system is able to **initiate goal-directed actions highly effectively and context-dependently** (Butz et al., 2007a). Essentially, the anticipatory architecture learns to reach goal postures as well as hand goal locations highly reliably. Moreover, the system can flexibly

adjust its behavior to various task constraints: it can avoid obstacles that block the shortest path to the goal; it can compensate for broken joints and prefer particular joint movements over others; it can combine multiple goal constraints—such as a hand goal location with a particular position of a joint (as long as the hand goal is still reachable). These capabilities were achieved by simple multiplicative influences on the architecture.

Moreover, the system can also account for future goal priorities while executing reaches to current goal locations, essentially preparing for a **faster and smoother reach** to a subsequent goal (Herbort and Butz, 2007). Additionally, in collaboration with other MindRACES partners, the architecture was coupled with goal state selection mechanisms. These selection mechanisms were based on reinforcement learning and particular the actor-critic method. In the examples investigated, the goal selection mechanisms chose to reach goals inside a goal region that were the further away from a punishment region the closer the punishment region and the stronger the punishment. Essentially, these experiments showed **improved decision making** due to the anticipatory representations and the implicitly anticipatory goal location choice, which emerged from the reinforcement learning architecture. Besides the general effectiveness of the action choice, the data also closely reproduced data from psychological experiments (Herbort et al., 2007).

In conclusion, SURE_REACH is a control architecture that effectively stores alternative (redundant) behavior means and goals. When a particular goal is desired, it effectively constrains the representation and issues control commands that yield the behavior that is maximally suitable given the current goal and (optionally) further current constraints. The achieved behavioral flexibility is only possible due to the effectively encoded sensorimotor model and the anticipatory, goal-dependent control structure.

11.3 Learning Cognitive Maps for Anticipatory Control: Time Growing Neural Gas

The time-growing neural gas (TGNG) approach builds spatial representations linking sensory codes time- and motor-dependently (Butz et al., 2008b). Essentially, TGNG learns sensorimotor spaces from scratch via random motor babbling. It forms a neural network representation of the experienced space and the connectivity within that space. Distances in the space are represented by the motor activity necessary to reach one spatial location from the other. Sensory proximity between close spatial representations is not required. The exploration essentially leads to the generation of a ‘cognitive’ map of an environment, which is represented by the growing neural network. The representation is related to place-cell and head-direction cell encodings found in the hippocampus in rats (Wiener et al., 2002). Each node in the growing network essentially encodes a place cell, which is activated by appropriate sensory input. Each edge, which connects two place cell nodes, associates the average motor command that was executed to move from one node to the next one. TGNG is complementary to SURE_REACH: it may replace the currently uniformly

distributed bodyspace encodings in SURE-REACH by the flexibly developing spatial encodings.

TGNG was used to control a robot vehicle in a maze environment, which the robot initially explores by random movements. The developing map enables the invocation of goal-directed control commands that move the vehicle to currently activated goal locations. This behavior is initiated by the activation of the network nodes that correspond to the goal locations, the subsequent propagation of that activity via dynamic programming principles, and the final closed-loop movement to the goal. Behavior is controlled by activating those motor commands that are associated with the edge that leads to the next higher activated node from the node best representing the current robot location. The spatial anticipatory encoding showed to yield **highly effective and context-based action initiation** and **fast and smooth behavior execution**. Particularly, the system reached goal locations reliably and effectively and it was able to flexibly adapt its behavior given additional task constraints, such as preferred movement directions (Butz et al., 2008b).

In sum, TGNG showed that cognitive maps can be learned by simply associating growing perceptual nodes motor-dependently. As long as the problem is Markov—so that each perceptual node maps to a unique state in the environment—the resulting representation is very suitable to generate flexible, anticipatory, goal directed control commands. The advantage of connecting states in time is that states can be connected that may strongly differ in how they are perceived as long as they are easily reached from each other. As a consequence, the cognitive map is rather independent of sensory proximity. It represents proximity directly motor dependently, that is, which motor effort is necessary to reach a location in space from another one. The consequence is that direct anticipatory, highly flexible behavioral control is possible.

11.4 Learning Effective Directional Arm Control: The Evolutionary System XCSF

The XCSF classifier system is well-known in the evolutionary computation community for its robust capability of iteratively learning function approximations accurately and reliably (Wilson, 2002). The system has been applied to approximation tasks of up-to seven dimensions and it has shown to be machine-learning competitive in several respects (Butz et al., 2008a). In behavioral tasks, the system has also been applied to the problem of learning generalized Q-value functions in real-valued domains (Lanzi et al., 2006).

Recently, during the MindRACES project, XCSF was further developed to yield directional, anticipatory control structures. In a first application of the resulting architecture, a directional control structure was learned for an arm with three degrees of freedom in a two-dimensional environment. The architecture mapped the arm posture space by evolving a population of overlapping, piece-wise (linear) classifier approximations. Each classifier encoded the control commands necessary for a directional movement in a particular arm posture subspace of the environment.

During learning, XCSF partitioned the posture space of the simulated arm to accurately predict how motor actions affect hand movements. The inversion of the predictions enabled goal-directed, closed-loop control of reaching movements. The system reached remote hand locations accurately, reliably, and effectively. Moreover, it was shown that the learning approach did not rely on the particular sensory inputs nor on a linear mapping. In fact, the evolving control map in XCSF is inherently non-linear. Due to the predictive learning approach and the consequent, inverse, piece-wise linear control approach, the system yields **fast and smooth behavior execution** patterns (Butz and Herbort, 2008).

In sum, an evolutionary approach was used to learn a forward-inverse piece-wise linear mapping of an arm control space, which was represented by a population of neural classifiers. Only the employed forward-inverse anticipatory representations of the classifiers coupled with goal-based, directional closed-loop control enabled the effective invocation of fast and smooth behavioral control.

11.5 Anticipatory Target Motion Prediction

To solve the task of predicting the movement of a visual target, a learning linear associator with memory, embedded within a Kalman filter was developed. While the Kalman filter takes care of the prediction of the target location the linear associator learns the model of the target motion (Balkenius and Johansson, 2007; Balkenius and Gardenfors, 2008). The memory component stores previous observations and allows the associator to train on a large number of observation in each iteration. This technique effectively emulated some of the advantages of batch-training methods within an on-line learning system.

The learning system has been applied to a number of task in which a tracking component is necessary, including the tracking of moving balls in partially occluded situations and the modeling of pursuit eye-movements. By combining the learning predictor with an inverse model of a robotic arm with three degrees of freedom, it became possible to catch a plastic toy fish that moved along a regular circular path. The system learned to predict the location of the target using color tracking in combination with the associator described above. The implemented system had a delay of approximately 500 ms from camera image to motor control. The prediction mechanism was essential for **successful tracking** as well as for the **manipulation of moving objects**. It also allowed for **faster and smoother behavior execution** since actions can be directed toward the future location of the target.

Prediction is an important ability that is useful as a component in many different applications. The results showed that motion prediction can be effectively included in many different tasks. Furthermore, it was investigated how different learning systems can be adapted for prediction by delaying inputs, outputs, and training data in different ways. A general conclusion is that any learning system can be adapted for prediction tasks in the way outlined and that predictive learning can be vary fast in simple situations.

11.6 Anticipatory Spatial Attention with Saliency Maps

A new approach of learning saliency maps was formulated that allows standard reinforcement learning techniques to be used in a number of attention tasks (see chapter 4). The approach is based on a novel and compact formulation of a saliency map, which allows many types of visual information to be combined in a coherent way. In a current implementation, feature-based and spatial attention was combined in a seamless way. The central idea is that the saliency map can be seen as an approximation of a value function for reinforcement learning. Unlike the standard action-value function in reinforcement learning, there is no state in this formulation. Instead, each location in the image corresponds to an individual action that directs attention to that location. Since all different sources of attentional signals eventually lead to attention that is spatially focused, the approach provides a common language for all such processes.

The mechanism has been applied to selective attention as well as priming in sensory processing. The mechanism is general enough to be used in any system that includes any form of sensory selection. In particular, the mechanism was used to select targets for visual tracking. It can also be used to **improve top-down attention** by tuning it to external reinforcement. Moreover, it **improves information seeking** by allocating larger processing resources to input data that resembles previously rewarded stimuli.

The new mechanism shows how it is possible to add reinforcement learning also to systems that are not used to control actions directly and suggests a general strategy for the marriage between reinforcement learning and perceptual processing (Balckenius and Winberg, 2008). Moreover, the proposed mechanism can be used as an important part of a complete reinforcement learning architecture to select stimuli as well as actions. To the best of our knowledge, this is the first computationally efficient implementation of a mechanism first suggested in (Balckenius, 2000).

11.7 Behavior Prediction in a Group of Robots

A combination of several techniques was used to anticipate the future behavior of a group of robots (Johansson and Balckenius, 2007). Kalman filters were used for short term prediction and correction of tracking data. Associative anticipatory attention mechanisms were used to learn where robot will reappear after they disappeared behind obstacles and to produce epistemic actions in the form of directed attention to gain optimal information about the behavior of other robots. The system also used internal simulation based on internal models of the other robots to anticipate how they will behave, in order to select appropriate actions in relation to the other robots. The robots also used a form of primitive joint attention through communication about their observations.

The combined mechanisms were used to investigate how a robot can control its own behavior depending on the anticipated behavior of another robot. A hiding scenario was implemented using a multi-robot set-up. There were four robot thieves and two guards that patrolled the environment in a regular fashion. The task of

the thieves was to hide from the robot guards while navigating to certain places in the environment. The implemented system showed **improved decision making** by simulating the other robots' behavior. **Faster and smoother behavior execution** was achieved by employing anticipatory mechanisms at three levels: at a low level of motor control, at an intermediate level for avoiding collision with other robots, and at a higher level for reaching desired goals without interference with other robots.

The implementation shows how anticipation at a number of levels can be made to work together within a unified agent architecture and addresses many of the difficulties that arise as complex anticipatory systems are built. It is believed that this is the first architecture that combines this many aspects of anticipation in a coherent system and successfully controls the individual anticipatory behavior of a robot as well as the emergent interactions between a group of robots with similar or conflicting goals.

11.8 Enhanced Adaptivity in a Predator-Prey Scenario

Another study investigated how anticipations can enhance the adaptivity in a predator-prey scenario, in which a predator (usually the dog like Sony AIBO™ robot) is supposed to catch a prey (such as a simpler robot with a reactive behavior, or in more sophisticated scenarios, a second AIBO™). First of all, the magnitude of the benefits by implementing anticipatory behavior for the predator depends on its physical abilities like velocity and agility, but seen with respect to the analogue capabilities of the prey. If, for instance, a predator can navigate much faster than the prey, then the simple behavior of heading towards the prey and then approaching it will usually be a sufficient strategy once the prey has been detected. This strategy can still be filed under reactive behavior, as the necessary anticipatory capabilities of learning the effects of necessary movements can be achieved quite easily.

A certain form of anticipation is necessary, if prey and predator possess comparable navigation abilities and both operate in the open field. The predator has obviously some advantages if it is able to predict the trajectory of the other robot. If additionally obstacles are present, which might occlude the opponent, then anticipatory capabilities deliver substantial advantages and it is easy to construct scenarios in which pure reactive behavior rarely succeeds, for example, if the prey is only visible for a short amount of time, which is not sufficient for a successful access.

Two approaches have been developed, one is based on Markov models (Lewandowski, 2007) and the second uses artificial immune systems (De Castro and Timmis, 2003). Both approaches operate in the space of observed sensor values and do not try to estimate robot locations within a world map. For the Markov approach, the current camera image is transformed into one of a final number of possible views. The elementary building blocks are the estimated probabilities of transitions between the current and the following state, given the current action, which are updated continuously while maneuvering. The ability to construct 1-step predictions could be used now to form chains of actions and to predict possible future outcomes, or to be more exact, to predict probability distributions over possible states for a

given time point. The system knows in its representation every so far encountered state.

Actually, the algorithm implements goal-oriented behavior in a different way. A desired state is one already experienced view, for which the prey has been seen from the nearest so far observed distance. Multiple desired states are allowed to exist. A plausibility check for a necessary condition ensures that the desired goal states are probably achievable from the current state with the acquired knowledge of transitions. A backward induction algorithm (Puterman, 1994) is then applied in a recursive manner that finally a current action can be chosen that, based on the history of transitions, one of the goal states will be reached as fast as possible. The algorithm allows that the prey is temporarily hidden, and therefore it is not mandatory that each step reduces the distance to the prey.

The building blocks and systemic operations of the AIS approach are elements and procedures that are in line with the immune system metaphor. Simple atomic elements that code for a condition, an action and an expectation, in terms of the possible outcome of an action, correspond to the antibody that reacts to a certain degree matching environmental stimulus, the antigen. Sensor inputs and therefore the epitopes of the antigens are coded as strings and are compared to the condition parts of the artificial immune systems elements. The best corresponding element is chosen and its action is executed. The expected outcome of this single action can be predicted and this in consequence allows the anticipation of the outcome of future action sequences in relation to the current situation and therefore the prediction of the outcome after multiple time steps.

After each successful run, the population of elements is dynamically updated according to certain interdependency patterns that closely follow Niels Jerne's immune network theory (Jerne, 1974), with the consequence that the concentration of elements belonging to successful runs will grow higher and their elements will be chosen more likely to produce genetically varying offspring than less effective antibodies, which are subsequently suppressed. Within AIS, the goal to catch the prey is therefore not explicitly given, but is implicitly coded in the distribution and concentration of the elements of the network's population. A main goal of the experiments with artificial immune systems was the evaluation of their capabilities in controlling a successful predator (and in some cases prey) in match to other control strategies. Among these were genetically evolved static strategies like differently complex subsumption architectures, as well as other evolutionary computation algorithms, including classifier systems (XCS) and even anticipatory classifier systems (Sigaud and Wilson, 2007). In all scenarios, from the simplest without any obstacles, to the most complex, where both players are controlled by anticipatory control mechanisms and a hiding place was introduced, the artificial immune system approach was able to solve the task of catching the prey (or escaping) very well. Additionally the behavior exhibited to the human observer in many cases appeared like a reasonable anticipatory strategy that one would observe also in biological hunter and prey scenarios.

In conclusion both approaches can be successfully applied, if the behavior of the prey under similar circumstances remains stable. When the presence of obstacles is

allowed, they provide **substantial advantages over purely reactive models**. The behavior of the prey does not need to be deterministic, but it should not change abruptly. Thus, the algorithms in their current form are best applicable, if the prey's behavior is stationary.

11.9 Adaptive Navigation and Control with Anticipation

In a real robot task an omnidirectional robot learns to correct inaccuracies while driving, or even learns to use corrective motor commands when a motor fails, both partially or completely, to optimize the driving accuracy for soccer competitions (Rojas and Frster, 2006). 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 adapts its own world-model accordingly.

A feed forward neural network with historic and current information of the 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 takes corrective actions in advance, thus improving 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. It was demonstrated that even if a robot loses some motor's power, the system can **relearn to drive the robot** in a straight path, even if the robot is a black box and we are not aware of how the commands are applied internally. The robot controller framework integrates action control by choosing actions based on its own self-model. It integrates attention and monitoring by observing its own hardware quality to change its behavior and monitoring. Without these three anticipatory mechanisms, the robot would still fulfill its task to some extent, but not sufficiently (in the sense of robustness and accuracy) to win an international multi robot contest.

11.10 Mental Experiments for Selecting Actions

A robotic task consisting in finding and moving to a randomly placed unique colored cup in the room illustrates that anticipatory mechanisms are essential to fulfill the assignment (Bakker et al., 2006; Zhumatiy et al., 2006). The key idea to solve the task is to use a forward model to translate sensor inputs to robot movement commands. This module anticipates a probabilistic world-model to estimate future rewards by *mental experiments*. The results of the mental experiments are then used to select the most promising actions for the motor controller.

The robot is equipped with a color camera and placed in a room, which contains the colored cup. 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 the target position.

The controller of the robot translates sensor input data to robot movement commands. It is trained by various reinforcement learning methods. In (Zhumatiy et al., 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 et al., 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 long training time for reinforcement learning algorithms for real robots, a probabilistic anticipatory world-model is learned from comparatively few real robot experiments. This world-model is then used to conduct mental experiments to train the controller with Prioritized Sweeping, a modification 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.

The robot controller framework combines action control, active vision, attention, goal directed behavior, and monitoring. Removing one of these anticipatory modules makes it impossible to learn the whole task on a real robot within a reasonable time frame.

11.11 Anticipations for Believable Behavior

The emotivector (Martinho and Paiva, 2005) is an affective anticipatory mechanism situated at the agent sensory interface. Each emotivector is coupled with a sensor and generates affective signals resulting from the mismatch between sensed and predicted values. Inspired by the psychology of emotion and attention, our implementation of the mechanism showed how attention grabbing potential as well as elementary sensations can be automatically generated from the observation of the values flowing through the agent sensors, and be used to generate believable behavior for the agent.

Two tasks were performed to evaluate the effectiveness of the emotivector mechanism: one taking place in the virtual world, where Aini, a synthetic flower, helped users to perform a word-puzzle task; and another, taking place in the real world, where the iCat social robot played a game of chess against the user. In both tasks, the affective signals generated by the emotivector were used to directly control the affective expression of the synthetic character interacting with the human user. The results showed that behavior autonomously generated by the emotivector is perceived as believable and understandable by the users.

Anticipation (and the associated uncertainty) plays an important part in the generation of affective signals. In the emotivector mechanism, anticipatory mechanisms are used for: deciding whether a percept is salient; defining the quality of the elicited affective states; allowing the mechanism to be context-free, without any requirements related to the manual tuning of parameters. As such, besides being a crucial factor in the autonomous generation of attentive and affective signals, anticipations allow parsimonious design.

The research issue resulting from the gap between the scope of psychological theory and the engineering needs of the believable character community has only started to be addressed in a principled manner. The emotivector mechanism is located in this area of relevance and addresses the specific question of creating autonomous believable behavior to support the engineering of believable synthetic characters. The emotivector achieves this goal by fusing the fields of anticipatory computing and affective computing. By design, this approach is broadly applicable and provides practical means to significantly improve the capabilities of believable characters driven by such architectures.

11.12 Anticipatory Behavior in a Searching-for-an-Object Task

In section 9, the use of analogy making as a prediction mechanism was investigated. Analogy allows prediction making about the current situation based on one episode in LTM, which could be from a different domain. The analogy involved can be very superficial or very deep. The AMBR model of analogy making (Kokinov, 1994b; Kokinov and Petrov, 2000, 2001) was further developed and augmented with a transfer and evaluation mechanism, which allowed the implementation of a real robot scenario involving perception and action execution. These mechanisms allowed the usage of analogy as a selective attention mechanism and top-down perception mechanism, which directed attention of the robot to anticipated objects or properties. The scene representation was build gradually including only relevant objects and relations. The use of analogy as a basis for anticipatory mechanism is novel and proposed for the first time.

The task in which the model was tested was a ‘searching for an object task’. The AIBO robot had to find its hidden bone under an object in a room. The objects differed in shape and color. The robot analyzes the scene, makes a decision where the object could be, and goes to find it.

In such a task, using a non-anticipatory approach will lead to full search for the objects below all shapes. Although it cannot be guaranteed that the episode retrieved by analogy will lead to the correct solution, in many cases it does and the way the solution is found is unique and was only possible due to analogy making. The advantage of the analogy based prediction compared to other prediction methods is the ability to use just one positive trial in order to generate the prediction. The proposed methods for top-down perception and selective attention based on anticipation deal with the hard problem of processing visual input. The huge space of objects and relations is filtered, which allows the mind to handle only small but relevant aspects of the available information.

The analogy based anticipatory mechanisms seem very promising for finding solutions in some situations. They have to be further tested in richer environments in order to explorer their full potential and scalability. A promising further development seems to be the use of analogy based predictions as a basis for models of perception and action and work along this line is in progress.

11.13 The Role of Anticipation in Cooperation and Coordination

In section 10, the problem of cooperation as a basic principle underlying was investigated in the framework of Iterated Prisoner's Dilemma Game (IPDG). Results from simulations with a connectionist architecture of an anticipatory PD player allow the conclusion that anticipation is of high importance for cooperation to be present in 2x2 interactions as well as in simulated PD play in a society of agents. The architecture combines a simple recurrent neural network with an auto-associator and a forward looking evaluation mechanism (Lalev and Grinberg, 2007), implementing an anticipation-based decision making mechanism. In IPDG, the recurrent network processes the flow of available information: the structure of the PD matrix, the players' moves, and the payoffs obtained from the game. Due to the learning mechanism, the network correlates this information in time, whenever appropriate (for example, how players' moves correspond to gains from the game), and tries to infer information that is not available yet—such as the move of the opponent. In the case of a single couple of IPDG players, the model managed to reproduce results from experiments with human subjects by (Hristova and Grinberg, 2004). Anticipation, rather than backward-looking reactive behavior, was responsible for the cooperation of the model against the simple-strategy computer opponent also used in the experiment with human subjects. Manipulation of anticipation forward-looking parameters also revealed that the anticipatory properties of the model's decision making contribute most to the observed dependence of cooperation on the structure of the payoff matrix (the so-called cooperation index).

With instances of the validated model architecture, simulations of IPDG playing in small societies were conducted. The aim of the simulations was to investigate the role of anticipation in a society of payoff-maximizing agents on cooperation and coordination among them. For this purpose, small groups of agents with different parameterization of the model played IPDG in simulated social environments. The parameters were chosen to have five groups of players with increasing anticipatory capabilities. The analysis of the processes in each society was based on the overall level of cooperation, mean payoffs, as well as cooperative coordination. It turned out that the level of cooperation in the simulated IPDG societies grew with anticipation, starting from 5% in the first society and reaching up to 30% in the fifth society. Corresponding to their anticipation, the intermediate types reached intermediate levels of cooperative interactions.

A tendency of increased mean number of mutual cooperation cases per simulation was observed with increased anticipatory properties of the agents in the societies. The opposite was valid for the mean number of double non-cooperative choices (mutual defection) per simulation, as this number increased with diminishing anticipations in the societies. Higher mutual cooperation is considered an advantage as long as this is the most profitable outcome for the society in the long run. The summary payoffs, which were gained in each society, were also positively correlated with forward-looking capabilities: the higher the anticipation within a society, the higher the payoffs obtained by the members of the corresponding society.

As a first measure for the level of coordination between the agents, the mean number of mutual cooperation games played in a series per IPDG session was used. The longest mutual cooperation coordination was observed in the societies with highest anticipation. Although the sequences are not very long, the influence of anticipation is considerable.

Cooperation and coordination play a positive role in a society and represent a decisive advantage. In IPDG, for example, bilateral coordinated cooperation would result in higher gains for both players. On the level of a society, cooperation and coordinated actions will lead to high overall productivity and benefits. These simulations showed that anticipation is decisive for high levels of cooperation and higher coordination. According to the results, the higher the anticipatory capabilities are, the higher the cooperation rate and the coordination in cooperation between agents. As human cooperation in IPDG is close in rates to the cooperation of our anticipatory agents, the prediction is that coordination series among human subjects may be in close ranges to those observed in the simulations.

11.14 Anticipatory Effects of Expectations and Emotions

Recent computational models in the context of cognitive systems are providing simple affective states in terms of their functional effects on agent's behavior. Their roles are argued to enable adaptive and situated cognition and span from reactive methods of control (similar to those employed in primitive biological organisms) to the control of computational resources, attention, and decision making processes. Systems based on appraisal theory stressed different relations between emotions and cognition, arguing emotions as a causal precursor for the mechanisms to detect, classify, and adaptively cope with significant events and environmental changes. Typically, emotions are modeled as cognitive mechanisms to monitor goal pursuit in terms of functional appraisal of action achievement and failure. Besides, emotion signals may rule intelligent resource allocation, improve situated cognition, and generate goals, which are translated into purposive behavior.

Moreover, an approach was investigated that promotes the anticipatory effect of emotions as a main breakthrough. In so doing, we envisage to address either the route from emotion to anticipation, or the reverse one, from anticipation to emotion.

Emotions and Anticipation in Goal Directed Agents Our theoretical model of emotions have been developed by using mathematical and logical tools which have been developed in the field of decision theory and applied modal logic (Castelfranchi, 2005; Lorini and Castelfranchi, 2007). This level of specification has been a first foundational step towards the design of computational architectures for affective and anticipatory agents.

In realizing a computational model for emotions and anticipations, a multifaceted approach was adopted by distinguishing different processes behind goal oriented behavior. In so doing, practical reasoning, epistemic reasoning, and situated reasoning

were distinguished and basic principles at the basis of emotions in terms of their cognitive ingredients were investigated. This promoted a clear disambiguation between slow, decisional processes, processed devoted to deal with knowledge and processes related to cope with situated events.

Whereas typical approaches in modeling cognitive agents are oriented at including graded primitives and temporal dimensions (i.e. belief on the future), a cognitive approach was adopted, which introduced expectations as emerging attitudes coming from epistemic and motivational states. A particular use of expectations enhancing problem solving and learning abilities has been modeled in the deliberation and goal selection phases. In addition, by pointing out the subjective character and the functional role of expectations, as intrinsic cognitive ingredients of many basic emotions (i.e. surprise, hope, relief, disappointment), a further kind of interaction between emotion and anticipatory activities was considered that consisted not only in predicting future events, but also in anticipating future emotions.

In (Piunti et al., 2007c,b,d), the quantitative influence of expectations upon the terms of a rational decision was investigated. In so doing, expectation driven decision making was introduced (Piunti et al., 2007c), which enabled agents to proactively take decisions either on the basis of anticipated events (i.e. trends of monitored signals) or on the basis of ongoing needs and desires.

As in appraisal-inspired models, emotions and mental states were provided to coordinate different computational and physical components required to effectively interact in complex environment. A clear methodological separation of concerns allowing the modeler by breaking down the work into two separate and independent activities was promoted: while the former was defined referring to the goal overview in the problem domain and clearly involves decisional processes (i.e. deliberation of alternative courses of actions), the latter can be defined through control frames to improve situated behavior. The following step was made to reinstate the two approaches by taking into account the correlations and the relative interactions occurring in system execution model. This allowed to integrate the low-level, situated reasoning to be used to inform higher decisional processes.

The emergent nature of affective states enables agent to adopt a mental frame while both expectations and emotions are conveyed to inform reasoning for redirecting resources and adopt long term strategies once a disturbing event is detected. To this end the contribute of Mental States is twofold: from the one side they can relieve the deliberative and the attentive processes from the burdens to process weakly relevant information in decision processes, excluding action alternatives that are likely to be less promising or have vanishing likelihood to be achieved. On the other side, Mental States provide ready to use action selection and resource allocation strategies that may relieve agent's need for resource-demanding and meta decision processes. An additional effect of modeling emotions through mental states is for agent's intention reconsideration. Traditional reconsideration strategies indicate an agent to abandon an intention when a related goal is achieved, when a goal become infeasible or when the agent relieve some inconsistencies between the world state and the external conditions necessary for goal achievement. Our model allows basic emotions to elicit an *interruption* on normal cognitive processes when unexpected

events require servicing. Once based on expectations of future states, intention reconsideration becomes anticipatory and can be used to coordinate behavior with prediction of future states.

A further ability based on expectation processing is to allow agents to modify their courses of action in order to anticipatorily coordinate with the other agents behavior. At this stage, agents were investigated that are able to thwart some expected events if the expectation is threatful or to promote them if the expectation is promising (with respect of the ongoing goals) (Piunti et al., 2007a).

Surprise Signals as Filters for the Update of Relevant Beliefs Taking into account the above mentioned model, a novel approach in epistemic reasoning and active perception was proposed. A *surprise driven* belief update processes introducing a notion of information relevance based on goal processing was presented by (Lorini and Piunti, 2007). By considering a proactive and an anticipatory perceptive process, the proposed model implements a novel strategy for epistemic reasoning according to which agents search and filter information from their environment not by monitoring nor perceiving all the retrievable data, but according their ongoing needs, desires, concerns, thus filtering and assessing only what is expected to be relevant for pursuing their goals.

The surprise based filter mechanism allows agents to consider useful for their belief updates only those information related with their goals and expectations. Raised from a mismatch between agent's knowledge and his perceived facts, a surprise signal is sent back to the control system in order to trigger a belief update process. The filter mechanisms is then responsible for (1) signaling the inconsistency between beliefs and an incoming input which is relevant with respect to the current task and (2) the revision of Beliefs and Expectations on the basis of the incoming relevant information.

The proposed filter allow agents to realize as useless and unnecessary those additional costs spent for data processing. Hence, surprise governed attention enable agents to process and filter the perceived data according to the ongoing expectations, modeled on the basis of a knowledge model (the belief base) coupled with subjective goals importance (related to subjective desires, purposes and concerns). In so doing they acquire the capability to divide the overall set of perceivable data in a *relevant* and *irrelevant* subsets.

Our experimental analysis measures the costs for perception and belief updates in agents engaged in dynamic environments. The results show that to higher environment dynamism, the greater costs sustained for epistemic activities are not compensated by an enhancement of achieved tasks. This elicit an important, general strategy exploitable by all those agent engaged in information rich worlds, with big sized information sources to be reported in their belief base. In these conditions, only the relevant information was filtered from the environments that was guessed to be a critical issue for forthcoming cognitive systems (consider for example agents engaged in an information retrieval task in the context of open system applications).

11.15 On-Line and Off-Line Anticipation for Action Control

In a series of studies, the effectiveness of reactive and anticipatory control architectures in predator-prey scenarios were studied. In these scenarios, it was necessary to satisfy competing drives, such as hunger and the avoidance of predators (Pezzulo and Calvi, 2006a; Pezzulo, 2008b).

An important aspect of the undertaken investigations focused on the trade-off between accurate control of action and time spent to form predictions. Anticipatory mechanisms used on-line with action make it more accurate and permit forecasting possible dangers arising from it. At the same time, prediction is a costly operation that requires time, and in principle it can make situated agents less effective or less responsive to dynamics in their environment such as dangers. This trade-off is further complicated when anticipatory mechanisms are used off-line with action to predict multiple and/or distal events (for example, the long-term outcome of several alternative potential courses of actions). Here the (computational) costs for engaging in ‘imagination’ and ‘planning’ are higher and can prevent effective situated action. Due to the failure of the earlier AI systems to deal with this trade-offs, it is believed by several researchers that belong to the ‘novel AI’ (Brooks, 1991) that situated agents should better act than reason.

Advantages of Anticipatory Mechanisms used On-Line with Action In a first study (Pezzulo and Calvi, 2006a), the performance of two schema-based agent architectures were compared in a predator-prey scenario involving multiple entities (predators and preys), obstacles, and moving objects. The first agent architecture included multiple perceptual and motor schemas (for detecting and escaping predators, detecting and catching preys, etc.), having one inverse and one forward model (Wolpert and Kawato, 1998) each. These two (coupled) internal models were used for determining the motor action and for predicting its sensory effects respectively. Prediction errors (of the forward models) were used for action control and schema selection: schemas generating reliable predictions (and related to the current active drives) were selected for controlling action. The second agent instead lacked the internal forward models.

The results indicated that the first agent architecture demonstrates a better adaptivity, since it was better able to satisfy its multiple drives (Pezzulo and Calvi, 2006a). This shows that in such dynamical and demanding environments it is advantageous to use on-line anticipation for action control and schemas selection, despite the costs of running forward models in real time.

Overall, contrary to the view that situated agents need to be reactive (Brooks, 1991), this experiment indicates that internal modeling is highly advantageous for situated agents if it is done on-line with action and produces representations whose format is compatible with the agent’s sensorimotor loop (as in the case of internal forward models).

Further Advantages of Off-Line Simulations In a second study (Pezzulo, 2008b), it was investigated how the same anticipatory mechanisms exploited in the first study can be exploited off-line to anticipate several steps in the future (an internal, ‘mental’ simulation of behavior) for the sake of (1) preventing dangers and (2) planning goal-directed action (that is, mentally generating and selecting sequences of actions to realize further). In this case, the possible advantages or disadvantages of engaging in ‘imagination’ during situated actions were tested.

Again, two agent architectures were compared, with and without the capability to re-enact schemas in simulation. In the first agent architecture, the same sensorimotor schemas adopted in the first study were used, but now they were allowed to run off-line in *simulation mode*, to predict the long-term sensory consequences of their motor commands. In simulation mode, motor commands were inhibited (not sent to the actuators), but fed as sensory inputs to the forward models, which then produced new sensory predictions that were used by the inverse models for generating new motor commands ‘as if’ the agents actually sensed the predicted future. The loop between forward and inverse models allowed generating long-term predictions for an arbitrary number of future steps. Two additional mechanisms were responsible for (1) stopping the current action if its predicted outcomes are evaluated as dangerous (a kind of ‘somatic marker’ mechanism, Damasio, 1994), and (2) evaluating, selecting, and activating the better ‘plans’ (that is, sequences of ‘simulated’ actions). The second agent architecture used forward models in the on-line control of action, but lacked the ability to run in generation mode.

The two agent architectures were tested in a task consisting in exploring a simulated house for finding a ‘treasure’ without being captured by guards (as in the first study, each agent architecture had concurrent drives to satisfy). Our results have shown that the first agent architecture, using anticipation both on-line and off-line (that is, in generation mode), outperformed the second one in terms of adaptivity thanks to its capability to predict possible future dangers and to plan from time to time (Pezzulo, 2008b). This happened despite the costs of simulating and planning.

Consistently with recent simulative theories of cognition (Damasio, 1994; Grush, 2004; Hesslow, 2002), the results of our experiments indicate that *mental simulation* is an effective strategy for avoiding dangers and planning in dynamic environments despite the fact that ‘imagination’ can in principle make an agent less efficacious in its current sensorimotor interaction. As argued in (Grush, 2004; Pezzulo and Castelfranchi, 2007), it is believed that off-line mental simulation is a suitable, embodied alternative to ‘reasoning by symbol-crunching’ of traditional AI systems, since it permits internal manipulation of (anticipatory) representations without losing grounding and situatedness. Thanks to anticipation, artificial systems can engage in mental operations that when performed by ‘ungrounded’ AI systems typically determine a poor performance in situated activities.

Overall, these studies have the potential to shed light on the role of mental simulations, how they enable increasingly complex cognitive capabilities, and how they open the passage from present-directed to goal-directed, purposive actions (Pezzulo, 2008a; Pezzulo and Castelfranchi, 2007). See also (Pezzulo et al., 2005; Pezzulo and Calvi, 2005, 2007a) for related studies.

11.16 Conclusion

We have surveyed fourteen case studies, which were developed during the EU project **MindRACES**. In each study, anticipatory mechanisms of (simulated or real) robots determined an advantage in terms of behavioral flexibility, adaptability, reliability, or social interaction. The presented case studies, however, are only few among a growing number of examples of how anticipatory mechanisms can enhance the cognitive and behavioral capabilities of artificial systems or can even develop novel capabilities that may be impossible without anticipatory mechanisms. Overall, these studies illustrate that our design methodology, based on a systematic endowment of artificial systems with anticipatory capabilities, determine huge advantages in terms of increased autonomy and adaptivity.

We hope that the design methodology we have illustrated throughout the book—and exemplified in this chapter with success cases—can inspire future research and real-world product development of cognitive system architectures. Taking the passage ‘from Reactive to Anticipatory Cognitive Embodied Systems’, this book may serve as the guideline to successfully design and develop further truly flexible, adaptive, effective, and interactive cognitive agents.