FP6-511931

Mind RACES

from Reactive to Anticipatory Cognitive Embodied Systems

DELIVERABLE D3 (D2.1)

Scenarios Specification and Problem Finding

Due date of deliverable:
31 / 05 / 2005

Actual submission date:
13 / 07 / 2005

Start date of project:
01 / 10 / 2004

Duration:
36 month

Organization name of lead contractor for this deliverable:
ISTC-CNR

Revision:
4

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

Document identifier:	**DELIVERABLE_WP2_N_1**
Date:	**13/07/2005**
Work package:	**WP2**
Partner(s):	*IDSIA, IST, ISTC-CNR, LUCS, NBU, NOZE, OFAI, UW-COGSCI*
Lead Partner:	*ISTC-CNR*
Document status:	*Approved*
Deliverable identifier:	**WP2_D2.1**

Delivery Slip

	Name	Partner	Date	Signature
From	LUCA TUMMOLINI	ISTC-CNR	27/06/2005	
Verified	RINO FALCONE	ISTC-CNR	04/07/2005	
Approved by	RINO FALCONE	ISTC-CNR	08/07/2005	

Files

Software Products	User files
MS Word™	D2_1.DOC

Project information

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

TABLE OF CONTENTS

PART 1 – Management Overview

1 Document Control

This document is a co-production of all the partners mentioned above. First, all the partners have contributed to the forum based discussion posting their contributions and acting as discussants (from M4 to M6). With these results the partners have focussed on sub-problems and agreed on the final scenarios presented in the deliverable.

2 Executive Summary

The goal of this deliverable is to identify two common environments (the HOUSE and the ROOM ENVIRONMENTS) and the three final scenarios, with the aim of including the main issues proposed by the partners during the discussion phase of the project.

The identification of the environments and the scenarios that will be presented is based on the consideration that the common scenarios should also be considered as a means to achieve integration among the different approaches characterizing the partners.

In Section 4 the results of the forum-based discussion are presented and discussed. To facilitate the presentation and the groupings of these different proposals the original contributions have been simplified and organized in a table format. The original proposals with the involved research problems are attached to this Deliverable in the APPENDIX section.

In Section 5 the two environments are individuated and discussed and in Section 6 common constraints on the robots are envisaged.

Finally, in Section 7 on the basis of the six theoretical scenarios, the final three scenarios (FINDING AND LOOKING FOR, PREDICTING IN A DYNAMIC WORLD, GUARDS AND THIEVES) are described and discussed. These final scenarios are the main achievement reported in this deliverable representing the common framework shared by the Partners that will be used in the next phases of the MindRACES project.

3 Terminology

The following table summarizes the working definitions used throughout the document.

Robot:	A real or simulated agent having specific sensors (e.g. camera, infrared/ultrasound sensors), actuators (e.g. two wheels, a three-segment arm), and a body (e.g. a cylinder, three rigid segments).
Mechanism:	Specific architecture and algorithms (=structure + functioning) of a model/controller.
Environment:	A particular real or simulated arena with specific features (i.e. dimensions, walls, type of "terrain"), containing particular objects (i.e. balls, boxes, doors, lights), and containing robots with specific features (i.e. sensors, actuators and bodies).
Scenario:	A set of tasks that share a common portion of an environment.
Task:	The specific goal one robot or group of robots have to accomplish in a scenario.

Table 1 Terminology adopted in the document.

PART 2 – Deliverable Content

The main goal of this deliverable is to individuate six general challenging scenarios that show the importance and the role of *anticipation* for cognitive systems both for theoretical and practical purposes. Such scenarios have been first presented by each partner individually. Then they have been grouped according to similarities based on the shared portion of environment and the involved cognitive functions. On this basis, in the last section three scenarios are presented that will be implemented during the next phases of the MindRACES project.

4 Scenario-based Discussion

To achieve the objective of individuating the scenarios for the theoretical discussion, three parallel thematic forums on the Project's portal (http://www.mindraces.org) have been created reflecting the thematic Project's Work-packages (WP3, WP4 and WP5). This activity lasted for three months (M4-M6).

The goals of the forum-based discussion have been:

1) The identification of cognitive capabilities dealing with anticipation and the specification of their function within the scenario.
2) The identification of other functions for the cognitive capacity beyond those presented in the scenario.
3) The proposal of models or analysis of that capacity with pointers to relevant literature.
4) The description of mechanisms that can be used to implement that function.

Partners have based these discussions on the proposal of specific scenarios. For each *theme* there has been three *cycles* of discussions, each based on one scenario. The scenario had the goal of triggering and being a concrete base to support the discussion on relevant issues.

During each cycle, the partners have played one of these three roles:

1) *WP Coordinator*: he has coordinated the discussion relative to a given theme during all the three cycles of discussion. The WP coordinator had the role to moderate and stimulate the discussion, keeping it on the topic arose by the proposed scenario. The WP leaders acted as the coordinators for the themes.
2) *Proponent (of the scenario)*: each partner has prepared a *scenario document* that has been used as basis for the discussion. In the scenario document, he has proposed a concrete scenario and a discussion focused on the four goals listed above.
3) *Discussants*: Each partner has discussed at least two different scenario documents preparing detailed *reply documents.* In these documents, partners gave structured replies, criticisms and additions to the issues raised in the proponent's scenario document.

This activity has produced the individuation and the discussion of 27 different scenarios, raising many scientific and technological issues that are more deeply discussed in concomitant deliverables (D3.1, D4.1, D5.1).

The first section of this deliverable reports the results of these discussions and schematically illustrates the proposed scenarios (see the APPENDIX for the original Partners' contributions).

4.1 Scenarios for WP3 (Attention Monitoring and Control)

The discussion in this WP has covered all the situations where a cognitive system selectively perceives and attends to its environment to improve the knowledge on its state. The emphasis is on selective attention and knowledge representation *viz.* not on the use of knowledge *for* the guidance of action (to achieve pragmatic goals).

The general goal of the cognitive system explored in this WP is to "understand" the world around it by focussing on the relevant details, hence the cognitive system is mainly motivated by an *epistemic goal.*

However:

- The use of *actions* as necessary for reaching this epistemic goal has been explored (such actions are called epistemic actions).
- Even if the focus has been on the epistemic goal, it is important also to consider the top-down influence of the *context* of active pragmatic goals on these processes.

Besides the different approaches to perception, attentive control and knowledge representation characterizing the Partners, the major goal in the discussion has been to unravel the role of *anticipation in these cognitive functions*.

To frame the discussion inside the two most promising state-of-the-art approaches to perception here are two major strands emphasizing anticipatory abilities:

- *Active perception*: treat perception as a *temporally* extended pattern of activity (a dynamical process). Main focus in *low-level* perception. The cognitive system is able to perceive only by means of both sensors and actions that produce as effects new sensory signals (i.e. self-modification of the body orientation towards the stimuli). The cognitive system is *attuned to* changes in the sensory inputs. The possibility of perception is linked to the capacity to "implicitly" understand (implicit practical or sensorimotor knowledge - skill) the effects of movements on sensory stimulation. The main problem is to temporally integrate a sequence of sensory stimuli.

- *Constructive perception*: is focussed on *high-level* perception where the sensorial stimulation is structured or organized through the application of schemas (recognition). The percepts are constructed through the schematization of the sensorial stimulation (sensations). The schemes are seen as procedures to *construct and interpret* the percepts. The sensorial stimulation is seen as a "sign" of the percept that the cognitive system has the goal to construct. The cognitive system *sees* the sensorial stimulation *as* a token of a specific type (or schema) and in this sense guesses what there is out there, makes hypothesis (to be verified) about the reality. The interpretation of the sensorial stimuli is then oriented by the already available schemas for interpretation that influences the way in which new information is gathered and actively pursued (epistemic actions) to fill such schemas. Besides the schema *activates more knowledge* of the object than what is present in the current stimulation. In front of an apple, the cognitive system sees it *as* an apple beginning from a particular stimulation i.e. by focussing on its shape and colour. Emphasis is given also to the influence of *top-down* influences in perception (background assumptions, active pragmatic goals, external context).

More broadly, the discussion has addressed the role of anticipation in the:

1. Acquisition of information from the world
2. Selection of what is *relevant* in the sensorial stimulation (filtering the information flow; attentive mechanisms)
3. Organization and structuring of the sensorial stimulation
4. Selection of different interpretations for identical or similar sensorial stimulation in different contexts
5. Development of an interpretation for a new sensorial stimulation

N.	Partner	WP	Environment	Tasks	Cognitive Functions	Agents
1	ISTC-CNR	3	Several rooms with walls of different dimensions No objects The rooms are contiguous with open passages Each room has different dimensions	The robot is pre-programmed to follow the walls leading to different rooms The robot learns to detect temporal regularities in its perceptual flow	<u>Event detection</u> by predicting the next primitive percepts (primitive prediction) <u>Hierarchical categorization</u> of events that can be used to predict at different time scales (abstract prediction) Providing <u>pseudo-feedback</u> to be matched with real feedback Generation of <u>surprise</u> triggering <u>model update</u>	1 wheeled robot (Pioneer) with a camera or infrared or ultrasound sensors
2	ISTC-CNR	3	One room A static object graspable by the robot's gripper	The robot detects the object in the room and tries to grasp it The camera learns to recognize the action of the robot and its goal	<u>Macro-action understanding</u> (categorization of biological motion) Prediction of the <u>next micro-action</u> by observation only of previous micro-actions Recognition of the <u>intended result</u> of the action (goal)	1 movable camera and 1 Robot (Pioneer) with a camera and a gripper
3	LUCS	3/4	One room An object with internal dynamics: multiple moving targets	The robot looks at the object and learns to predict the behaviour of multiple targets The robot learns to grasp the targets at specific time points (Fish game)	Event detection Learning spatial and temporal attention: to focus attention at a particular point in time and space Generalization of the learnt model from different spatial perspectives	1 wheeled robot with a movable camera and a gripper
4	LUCS	3/4	One room One marble and different blocks that can be arranged in different shapes	The robot looks at the specific arrangement and learns to predict how the marble will move The robot manipulates the blocks to obtain a specific motion of the marble matching a desired goal state (Marble run game)	Learning causal relations between events Use a causal model to predict a future state	1 wheeled robot with a movable camera and a gripper

5	LUCS	3/4	One room A ball Possibly some occluding obstacles	The robot looks at the ball rolling towards him from different perspectives The robot learns to predict the motion of the ball (Roll the ball game)	Tracking the trajectory of a dynamic object Prospective grasping	1 wheeled robot with a movable camera and a gripper
6	LUCS	3/4	One room The three previous set-ups are present contemporaneously	The robot learns the three games at the same time and is able to select the appropriate predictions	Detection of context Use of context in the selection of action	1 wheeled robot with a movable camera and a gripper
7	LUCS	3/4	A start room and a goal state room The rooms are contiguous with open passages The passages do not allow more then one robot at a time	The robots, while mutual monitoring, have to arrive first at the goal state room avoiding collisions (Race game)	Macro-action understanding (categorization of biological motion) Prediction of the the next micro-action by observation only of previous micro-actions Recognition of the intended result of the action (goal) Anticipatory coordination and strategic reasoning Tradeoff between monitoring actions and practical actions	At least two wheeled robots and 1 overhead camera or many simulated agents
8	LUCS	3/4	One room No objects	The robots have to avoid collision between each other and to be touched by the chaser The robot that is touched by the chaser becomes the chaser and has to touch the others (Game of tag)	Turn taking behavior and role understanding Pretence and deception	At least two wheeled robots and 1 overhead camera or many simulated agents
9	LUCS	3/4	Several rooms Several objects that can occlude the visibility of the robots in different ways, ie containers or tables or boxes	One robot has to detect possible spots in the rooms that affords different hiding strategies while the other has to seek for the other agent as much efficiently as possible (Hide and seek)	Selective attentional scan of the rooms Anticipation of the spatial perspective of another agent Categorization of objects according to their occluding properties Anticipation of the attentional strategy of the other	Two wheeled robots and 1 overhead camera
10	IDSIA	3	One room Several objects that co-occur frequentially or are semantically related, ie a table, a bottle and a cork	The robot has to find a target object (the cork) in the room by producing a sequence of saccades or other movements until the target is centered in the visual field	Selective attention Categorization and schema development Sequential search for informative inputs by the anticipation of the information gain	One wheeled robot with a movable camera or other directional sensors

4.2 Scenarios for WP4 (Goal Directed Behaviour, Pro-activity and Analogy)

The discussion in this WP has covered all situations where an agent uses *models of the world that take into consideration the consequences of its own action*, that is models of the world that predict the future state of the world on the basis of current state and the planned action (also known as "forward models").

Some of the problems that arose from the discussion are:

- The role of forward models in different cognitive functions, e.g.: feedback control when the feedback takes too long to arrive, and in planning.
- The evolutionary passage from reactive systems to systems endowed with forward models.
- The routinization of planned behaviours: how planned behaviours are routinized, that is "compiled" into reactive anticipatory behaviours.
- The shifts between the reactive and planning control systems.
- The role of "social forward models" ("theories of mind") for the prediction of the behaviours of other agents on the basis of the actions.
- The appropriate "format" of expectations in order to be matched with sensorial stimuli and/or goal states.
- The different kind of expectations (implicit, explicit, at different temporal abstractions) that can be used for different tasks.
- The different kinds and functions of *goals* and *goal states* (internal motivational states, expected reinforcement, etc) and their different representations.

Two main functions of anticipations have been distinguished:

- *Anticipation for deciding what to do next*: predictions are compared with a goal state not with a world/sensorial state; and are used for action/plan selection. E.g.: two choices are possible and the cognitive agent generates predictions for both; one action is selected (e.g. because the corresponding prediction satisfies a goal); the corresponding prediction is used successively for action monitoring; the other one will never "met the reality".
- *Anticipation for action monitoring*: expectations are compared with the actual state and the match/mismatch information is used e.g. for adjusting and tuning actions, generating surprise, shifting from a routine to a deliberative control, "compile" behaviours, etc.

N.	Partner	WP	Environment	Tasks	Cognitive Functions	Agents
11	ISTC-CNR	4	Several rooms The rooms are contiguous with open passages There is a target room that is central in the environment with different entrances In the other rooms there are static objects that afford hiding strategies	The goal of the robot is to plan a defensive strategy, monitoring the other rooms, and an offensive one actively looking for possible intruders (Watchdog)	Learning the model of the environment to support navigation Reasoning about the predicted consequences of his actions for planning a complete path Routinization of the plan preserving the expected consequences of actions Monitoring and adjustment of actions Surprise Shift from routinary to deliberate control Prediction of biological motion with categorical reasoning Goals management Active search in the environment looking for possible hiding spots	One wheeled robots or a simulated agent
12	ISTC-CNR	4	Several rooms The rooms are contiguous with open passages There is a target room that is central in the environment with different entrances In the other rooms there are static objects that afford hiding strategies	The same goal as in the previous but involving a group of collaborating robots	Helpful behaviour Delegation by trust	At least two wheeled robot or simulated agents
13	OFAI	4	One room A ball and an occluding obstacle, ie a wall or a tunnel	The robot has to reach a rolling ball moving after the occluding obstacle	Learning object continuity Predict where and when an occluded moving object will reappear Prospective reaching	An AIBO robot
14	OFAI	4	One room A ball and an occluding obstacle, ie a wall or a tunnel An obstacle behind the wall	The robot has to reach a rolling ball that has been blocked behind the wall The noise made by the bumping ball is used to predict that the ball is behind the wall	Learning object continuity Surprise Predict that the object will not reappear Curiosity	An AIBO robot
15	UW-COGSCI	4	One room Several objects	The robot has to search and retrieve a specific object	Learn a mental model of the room Recognition of objects by affordances, shapes, occlusion properties Search and collect	A wheeled robot with moving camera and a gripper
16	UW-COGSCI	4	One room Several objects	The robot has to search and retrieve specific objects at different points in time	Balance between present motivation and future ones Grounding of abstract representations	A wheeled robot with moving camera and a gripper
17	UW-COGSCI	4	One room Several objects	The robot has to search and retrieve a specific object after an abstract request	Matching the abstract representations with the sensory grounded ones	A wheeled robot with moving camera and a gripper

18	UW-COGSCI	4	One room Several objects	The robots have to search and retrieve a specific object (either cooperating or competing)	Understanding the actions of others Anticipatory coordination	At least two wheeled robots with moving cameras and grippers
19	UW-COGSCI	4	One room Several objects	The robots have to search and retrieve a specific object (either cooperating or competing) and are able to recognize each other	Understanding the actions of others Trust	At least two wheeled robots with moving cameras and grippers
20	NBU	4	Several rooms The rooms are contiguous with open passages Some transitions are regulated by signals (ie semaphores) that can also indicate obligated directions The passages are such that only one robot at a time can pass	Starting from an initial state room the robot has to move to a goal state room as fast as possible to find an object. There is more than one possible way to reach the target There are three signals whose state the robot has to consider in its choice Three possible conditions in the states of the signals: homogenous, non-homogenous but periodic and non deterministic (Taxi)	Learning the signals sequence Anticipation of the next signal Analogy at different complexities in the different conditions	One wheeled robot or AIBO
21	NBU	4	Several rooms The rooms are contiguous with open passages Some transitions are regulated by signals (ie semaphores) that can also indicate obligated directions Several occluding objects (ie boxes) Some target objects	The robot has to find target objects that are distributed in the environment and are hidden behind occluding objects Once the object has been found, the robot has to remove the occluding objects and point or grasp the target one There can also be guardian robots that have to block the robot or to change the location of the target (Tomb Raider)	Anticipatory coordination based on analogy	At least two robots (AIBO or Pioneer) or many simulated agents
22	NBU	4	Several rooms The rooms are contiguous with open passages Some transitions are regulated by signals (ie semaphores) that can also indicate obligated directions There are resources (ie food) that periodically appear and disappear in the environment	The goal of the robots is to survive by finding resources that regularly appear in the rooms. There is a cost in moving and trespassing forbidden passages (signals) and a benefit picking resources. The robots have local knowledge of their environment but to survive need to move to new environments.	Event detection Anticipatory coordination Transfer of knowledge from the local domain to new domains (analogy)	At least two robots (AIBO or Pioneer) or many simulated agents

4.3 Scenarios for WP5 (Emotions as Anticipations in Computational Architectures)

Beyond the general problem of modelling emotional responses, the discussion in this WP has been devoted to the identification of all the different relationships between *anticipation and emotions*, or better, between anticipatory behaviours or representations and mechanisms, and emotional responses.

Emotions form complex, hybrid subjective states because they integrate somatic, cognitive and motivational components. Such basic constituent elements of emotions are *beliefs, evaluations, goals, arousal* - i.e. somatic activation and its *proprioception* - the “tendency towards action” or *conative* component, and the *expressive* component.

In particular, a meaningful model of emotions should account for the fact that they are ‘felt’ and such feeling is integrated with the other representational and motivational components. To have genuine emotional responses, the cognitive system needs to have a body, and no real body is possible without modelling also the ‘feeling’ (the body sends sensorial signals about its current state).

The body can autonomously react to the external stimuli, without high-level (belief-based) evaluation and forecast, and its ‘reaction’ (*motion*) is perceived and interpreted by the control system. However emotions also need higher-order cognition, to interpret and attribute received bodily responses to categorized external events.

Emotions are a fuzzy set; not only a general agreement about the boundaries of this set is lacking (if certain affective or mood states are to be considered emotions or not; which is the distinction between ‘emotions’ and ‘feelings’; etc.), but emotions are related to other felt experiences and motivational states (like thirst, hunger, sexual excitement, or arousal, surprise, relief, etc.).

In this Project *emotions* are those *bodily motions that are not simply information about bodily processes and states, but are ‘about’ an event or imagination, have a subjective qualitative value (are pleasant/unpleasant), and represent an implicit felt ‘evaluation’ of that event or idea (not based on reasons)*. So one might consider close to emotion an exciting or painful surprise, an intense relief, and disappointment; while perhaps one should leave aside mere arousal or relaxation. Anyway, the discussion was oriented towards the modelling of the specific (bodily, motivational, cognitive, architectural..) components that those phenomena partially have in common.

Broadly, three basic different relationships between emotion and anticipation have been identified:

i. Emotions eliciting an anticipatory (preparatory) behaviour.
ii. Emotions due to Anticipatory Representations:
   a) Emotion now as a response to the predicted future event;
   b) Emotion at the very predicted moment due to the previous expectation and its (mis)match with reality.
iii. Anticipating Future Emotions

N.	Partner	WP	Environment	Tasks	Cognitive Functions	Agents
23	ISTC-CNR	5	Several rooms The rooms are contiguous with open passages There are several objects but some of them are 'dangerous' for the robot (ie fire) The dangerous objects leave signs in the environment (ie smoke)	The robot has to explore the environment navigating through the rooms to reach a target room The robot is able to monitor his internal environment (body) and feels the internal reactions to the dangerous objects (ie a feeling of fright) The robot learns to avoid dangerous object reacting to the precursor 'signs'	The bodily motion elicits an anticipatory avoidance behaviour The reaction to a stimulus is appropriate to react to a forthcoming event No mental anticipated explicit representation of the dangerous event Primitive form of fear or fright Appraisal and feeling	A wheeled robot or a simulated agent
24	ISTC-CNR	5	Several rooms The rooms are contiguous with open passages There are several objects but some of them are 'dangerous' for the robot (ie fire)	The robot has to explore the environment navigating through the rooms to reach a target room In planning its course of action, it is able to have explicit representations of possible dangers It feels fear and use this reaction to plan a safe path The robot can perceive a dangerous situation and feels a sense of 'anxiety' and increase his epistemic exploration of possible dangers	Emotion is caused by an explicit prediction (expectation) Feeling of the bodily reactions to future prospects Hope Fear Anxiety Worry Reaction to the feeling with expressive or impulsive behaviour Surprise Disappointment Relief	A wheeled robot or a simulated agent
25	ISTC-CNR	5	Several rooms The rooms are contiguous with open passages There are several objects but some of them are 'dangerous' for the robot (ie fire)	The robot has to reach a target room as quickly as possible but avoiding dangerous objects There are three possible paths to the target. The robot has conflicting goals or motivations (be quick, avoid danger, avoid feeling fear, feel pleasure or joy)	Prediction of future emotions Use of this anticipated emotion in decision and in decision and planning Change of preferences	A wheeled robot or a simulated agent
26	ISTC-CNR	5	One room No objects	One robot predicts the emotional reaction of another and act in order to evoke this emotional state (Frightening Game)	Prediction of future emotions Use of this anticipated emotion in decision and in decision and planning	At least two wheeled robot or simulated agents
27	IST	5	One room One ball and several distractors (objects with similar color or shape) Several obstacles (ie boxes)	The robot has to reach the target ball that is thrown in the environment Along the planned path, the robot's attentive system is distracted by object with similar shape or color	Expressive behaviour Curiosity	An AIBO robot

5 The two environments

Drawing on the similarities among the scenarios, two common environments that could support all these tasks have been identified. In this section the kind of *arenas* and *objects* composing these environments are detailed, while we the *robots'* properties are discussed in a separate section since they are common for all the environments.

5.1 First environment: the "House"

This environment has been inspired by the scenarios proposed by the partners that require that the robots navigate in closed multiple arenas, and by the "Artificial City" scenario of NBU (see page 64 for details).

The features of this environment are as follows:

- Several adjacent rectangular rooms surrounded with walls, having a flat floor, and connected by door-like passages. The size of the rooms is not specified.
- The walls of rooms might possess surfaces to aid the functioning of infrared and ultrasound sensors (e.g. they might have high-reflection tapes/covers on them).
- The objects described in the "game room" environment reported below, and some coloured lights (see below) might be placed on the walls as landmarks to aid robots in distinguishing different rooms, and in order to orient and navigate inside them.
- Coloured lights or leds (yellow, red, blue, green, white, black) might be placed on the sides of the doors to represent: traffic lights (pass/do not pass rules), rewards and goals, open-closed doors, and landmarks (sub-groups of partners might decide to use non-marked open/closed real doors in order to have more realistic tasks).

5.2 Second environment: the "Game Room"

This environment has been inspired by the scenarios requiring the robots to interact with *different objects* having different properties (manipulation and object interactions tasks).

The features of this environment are as follows:

- A rectangular room surrounded with walls with a flat floor. The size of the room is not specified.
- The walls might possess surfaces to aid the functioning of infrared and ultrasound sensors (e.g. they might have high-reflection tapes/covers on them).
- The objects described below, and some coloured lights (as those used in the "house" environment described above) might be placed on the walls as landmarks to allow the robots to orient and navigate inside the room.
- The room might contain a number of balls and cubic boxes.

Balls will have one of the following properties:

- Size (diameter): 2.5, 5, 10, 15, 20, 30, 40, 50, 60, 70, 80, 90, 100 cm.
- Colour: yellow, red, blue, green, white, black

- Weight: free (sub-groups of partners tackling same tasks should converge on this)
- Material: free (sub-groups of partners tackling same tasks should converge on this)

Cubic boxes will have one of the following properties:

- Size: 2.5, 5, 10, 15, 20, 30, 40, 50, 60, 70, 80, 90, 100 cm.
- Colour: yellow, red, blue, green, white, black
- Weight: free (sub-groups of partners tackling same tasks should converge on this)
- Material: free (sub-groups of partners tackling same tasks should converge on this)

Some boxes might have one or two missing faces to allow other objects to be put in them, and pass through them, and to allow robots to hide inside them and pass through them.
Notice that boxes with specific textures (e.g. stripes or lattices) might be obtained by putting together several boxes. “Complex” objects, e.g. a bottle with a cork on it, might be approximated by putting together several of these standard modules (boxes and balls).

6 The robots

Convergence on the properties of the robots used by the partners is needed to ease the work on the same scenarios (= environment + tasks) and proceed towards integration. Convergence is of course more feasible in simulation. Unfortunately, it is impossible to converge on the same real robots since partners develop on top of different robots. However, luckily we envisage the possibility of converging on a restricted number of sensors and actuators, or at least on restricted classes of them having the same properties at a particular level of abstraction, as indicated here:

Sensors:

- Camera (eventually motorised, see below). We leave open the scope (height/width angles, e.g. a webcam or a 360° panoramic camera), definition (number of pixels), and sensitivity (luminosity and contrast) of the cameras used by partners (partners tackling the same tasks should converge on this).
- Proximity sensors: infrared and ultra-sound sensors. We leave open the scope, distance reached, number, and placement on the robots of these sensors.

Actuators:

Locomotion actuators (wheels, legs in the case of Aibo): as mentioned above, since all partners are interested in displacement of robot in space, and not in the details of movement, we can abstract over details by referring to quite abstract displacement actions, e.g.:

- Desired displacement speed (space/time covered by the barycentre of the robot)
- Desired rotational velocity (degrees/time that the robots turn with respect to the arena)
- Notice that these two dimensions of displacement are the most commonly used commands in the APIs of real wheeled robot's controllers. They should be also suitable for controlling legged robots.
- Motors to control active cameras: pan and tilt motors. The commands will be in terms of desired position of the camera (in degrees).

7 The Three Final Scenarios

In this section the three final scenarios (FINDING AND LOOKING FOR, PREDICTING IN A DYNAMIC WORLD, GUARDS AND THIEVES) are presented and discussed.

7.1 Grouping the scenarios

In what follows the previous 27 scenarios have been grouped in six general scenarios that have been used to focus the discussion and foster the convergences between the partners. At this stage, each general scenario is still composed by a set of tasks proposed by the partner. The tasks have been re-described and simplified to be adapted to the three common environments. The numbering reflects the original scenario descriptions.

The "House" scenarios

N.	Partner	WP	Environment	Scenario 1 (four tasks)	Cognitive Functions	Agents
7	LUCS	3	Several rooms	Race: Competing robots have to arrive first at a goal place avoiding to collide other robots them (robots know where to go)	Anticipating the actions of others to avoid collisions	Several wheeled robot with proximity sensors, and cameras
8	LUCS	3		Game of tag: A chaser robot has to touch other robots. The robot that is touched by the chaser becomes the new chaser	Anticipating the actions of others to avoid being touched	
				Collective retrieval: Robots have to cooperate or compete to collect particular objects	Modelling others; trust	The same robot with a gripper
18,19	UW-COGSCI	4	Several objects to be collected	Survive: several robots collect energy objects, regularly appearing in different places, while avoiding energy subtraction places	Detection of events' timing and cooperation based on analogy	
22	NBU	4	Several lights/objects signalling either energy increasing or energy decreasing places			

N.	Partner	WP	Environment	Scenario 2 (four tasks)	Cognitive Functions	Agents
9 11	LUCS ISTC-CNR	3 4	Several rooms connected by open door-passages Several boxes for hiding	Hide and seek: 1 defender robot has to spot one or several attacking robots (so blocking them) that have to achieve a central goal spot without being seen by the defender	Selective attention, simulation of perceptual perspectives and attention of others, categorisation of hiding places	1 (or several) wheeled defender robot with proximity sensors and cameras
11	ISTC-CNR	4		Collective hide and seek: Several defenders and several attackers	Modelling of environment, planning, routinization of plans, routine	1 (or several) wheeled attacking

					monitoring, routine-deliberation shift	robots with proximity sensors and camera
21	NBU	5	One treasure object	Tomb rider: 1 attaching defender robot has to search a treasure object in the rooms: the object might be occluded by other objects; there are several defenders	Anticipatory coordination based on analogy	
26	ISTC-CNR			A seeking robot has to frighten and cause an hiding robot to escape	Acting on the basis of anticipation of others' emotions	

N.	**Partner**	**WP**	**Environment**	**Scenario 3 (two tasks)**	**Cognitive Functions**	**Agents**
23 24 25	ISTC-CNR	5	Several lights on doors, regularly on and off, signalling that they are open/closed or perceived as dangerous	Reach a target room avoiding dangers causing different emotional reactions	Planning with anticipatory behaviours and emotions	1 wheeled robot with a camera
20	NBU	4		Taxi: Achieving assigned goal places in the shortest time	Goal achievement and planning on the basis of analogy (e.g. applied to traffic lights' regularities in time)	wheeled robot with proximity sensors and camera

The "Game Room" scenarios

N.	Partner	WP	Environment	Scenario 4 (three tasks)	Cognitive Functions	Agents
5	LUCS	3	1 Ball Several boxes	The robot has to hit the ball or avoid to be hit by it	Anticipation of trajectories of the ball; preservation of the anticipation capabilities notwithstanding changes of perspective	1 wheeled robot with a camera
27	IST	5	1 moving ball Several obstacle balls and boxes	The robot has to reach the target ball notwithstanding the distractors and obstacles	Emotional control of attention	1 Aibo dog (camera + movement)

N.	Partner	WP	Environment	Scenario 5 (three tasks)	Cognitive Functions	Agents
1	ISTC-CNR	3	The House Environment with several rooms connected by open door passages	A robot, that executes wall following in the room, learns to categorise and anticipate walls, corridors, corners, etc., and uses this capacity to reach particular landmarks (target objects) in the room	Extracting regularities in time at different time scales; using them to anticipate future events	1 wheeled robot with proximity sensors
10	IDSIA	3	Several target objects Several boxes	The camera has a limited vision: the robot has to spot the target object as quickly as possible in a partially observable environment	Orienting, Active perception, Selective attention	1 motorised camera mounted on a wheeled robot
15,16,17	UW-COGSCI	4		A robot with different internal drives has to search and retrieve one or more target objects	Present/future motivations, world modelling, perceiving objects' affordances, grounding of internal representations	1 motorised camera mounted on a wheeled robot with gripper

N.	Partner	WP	Environment	Scenario 6 (three tasks)	Cognitive Functions	Agents
3	LUCS	3	The Room Environment with several objects with intrinsic dynamics.	Fish Game: The robot has to anticipate when the lights became red; it has also be capable of doing this while moving (i.e. with a changing perspective)	Selective attention, anticipation of rhythmic dynamic processes, preservation of anticipatory capabilities with changing perspectives	1 wheeled robot with a camera
4	LUCS	3		Marble run: The robot has to anticipate the movement of the ball. The robot has to develop compositional prediction capabilities in order to generalise prediction when the obstacles of the cell game are rearranged	Understanding causality. Compositionality of models of the world	1 camera robot
5	LUCS	3		The learns the two tasks at the same time and is able to select the appropriate predictions	Use of context in the selection of action	
13, 14	OFAI	4		The robot has to anticipate that a ball passing behind an obstacle will eventually come out of it; if the ball hits a second obstacle while occluded (causing noise) the robot should anticipate that the ball will not appear	Object persistence, sensor fusion, curiosity, surprise	1 Aibo (camera + movement)

7.2 Specification of the final scenarios

7.2.1 The Game Room and the House Environments

The *Game Room* environment is placed in a large room with a flat floor.
It can be transformed into the *House* environment composed of multiple rooms with differently coloured walls. The colours should differ in different rooms in order to be used as landmarks by the robots. The size of the walls has to be fixed so that they are higher than the robots and prevent them from seeing in other rooms from the room where they are. The walls of the rooms can have special surfaces to efficiently support the functioning of infrared and ultrasound sensors (e.g. high-reflection wallpapers will be put on them).

To help the robots avoid collisions with the walls and fixed objects some boundary lines might be used either with the same colour in all the rooms and corridors or with a colour specific to each room. Different objects of a limited set of shapes and sizes (e.g. small, middle and big cubes, cylinders and parallelepipeds) can be placed in the rooms.

Gates allowing the passage of only one robot connect the rooms. These gates should have a simple signalling system – indicating the passing state (open or closed). The signals are of two kinds: visual and acoustic. Each of them can be positioned on the wall next to the gate. The duration of each state and the type of signal (visual or acoustic) may vary across the gates. To simplify the signalling system, two lights (red and green) and two different sounds can be used. The occurrence of such 'events' introduces an interesting temporal dynamics in the environment.

7.2.2 Scenarios

In what follows the three final scenarios are presented and their tasks discussed.

Finding and Looking For

[IDSIA, IST, ISTC-CNR, NBU, UW-COGSCI]

This scenario consists of two main varieties: the first based on the assumption that the robot has a map of the House and the second based on the assumption that the robot is constructing such a map on the fly.

Each variety includes search, optimization of the search (rejecting unlikely cases), recognition and report for the object found. The tasks do not imply moving of the object searched.

The agent used to this purpose is either a pan-and-tilt camera mounted on a fixed base, or a pan-and-tilt or fixed camera mounted on a mobile robot (the robot is eventually equipped with other sensors to aid navigation and obstacle avoidance, e.g. a compass, ultrasound sensors, etc.). The camera might have a simulated fovea, with higher resolution at the centre.

- ***Finding a specific object (Game Room)***

The purpose of this task is to find a specific object in the environment (e.g. a red cube). The degree of detail in the description must be sufficient to define unambiguously a single object, not a class of similar ones. For example "red cube" is to be used in the case when there is a single red cube, and "big red cube" if there are several red cubes with different sizes and only one of them is big. The target localization tasks can be divided into several levels of complexity:

- *Finding an object in an eye-catching colour in a clean room.*
- *Finding an object among similar ones: more objects with different colours and similar/same shapes and/or dimensions, or objects with the same colour as the searched one, but with different shapes and dimensions.*
- *Finding an object that can be partially or fully occluded by other objects.*

The prediction/anticipation concerns the possible object's position and can be based on similar past trials, generalisation, positions already visited, regularities in the spatial relations among objects (e.g. the red cube has a certain probability of being above yellow cubes, and yellow balls tend to be near big blue objects).

- ***Finding members of a class of objects by class description (Game Room)***

The purpose of this task is to find any object matching some general or partial description (for example "find a cube" or "find a red object"). As in the previous case, prediction or anticipation can be based on previous experience, recurring spatial relations, etc.

This task involves approaching the following problems as:

- *Object recognition*
- *Detection of relations among objects*
- *Selective attention (only the information relevant to object description is to be processed)*

- *Relation extraction and encoding (including recognition of semi-visible objects behind other objects)*
- *The agent must at the same time memorize where it has looked before*
- *The agent might generalize from previous experiences which actions to take are best.*
- *The agent must be able to (at least implicitly) categorize observed objects, whether only parts are observed or the whole thing at once.*

➢ ***Looking for an object in the House (House)***

This task is placed in the House environment. Coloured light signals might be positioned above/aside passages between rooms. These lights signal if the passage is open or closed, and might have periodic behaviours. In this task the robot's goal is to find an object that is hidden in one of the rooms in the shortest time or using the shortest way. The level of detail in the object's description may vary. In the case of class-definitions of the target, the purpose of the robot is to find any object that matches the given description.

In some conditions, the target's location is probabilistically biased towards certain locations (e.g. red cubes tend to stand on yellow cubes, although not always, or to stay in some rooms).

These task is much more complex than the previous ones. It might imply solving the following problems:

- *Learning/making inferences in order to estimate and anticipate the regularities of the environments.*
- *Performing exploratory behaviours by anticipating relevant information gain and executing those actions that lead to more cues about the target's location.*
- *Decision making about the rooms to visit involving prediction and anticipation.*
- *Learning that the gate is signalling the permission to pass (e.g. by reinforcement learning).*
- *Transferring knowledge about a gate behaviour (rate of the signals, light or sound, etc.) to other gates.*
- *The agent might build models of the world structure and/or regularities between objects on the basis of experience, and use this to find the targets.*

➢ ***Looking for an object in a "dangerous" House (House)***

In this task the robot is looking for a target object in a House where there are also dangerous objects. The three subtasks are designed to explore three specific relations between emotions and anticipation.

- *While looking for the object, the robot meets dangers that cannot be fully avoided (say fire). This danger leaves signs in the zone around its location (like smoke or smell) that when encountered by the robot elicits some internal motion (or* appraisal*) in its body (i.e. a feeling of 'fright'). The robot learns to anticipatorily detect the danger just by conditioning an avoidance behaviour not to the danger (Ev) but to its precursor sign (St).*
- *While planning to get to the room where the target object is located, the robot foreseen a given scene; this event is bad for it, is a threat, a danger. It feels 'fear' and changes its*

path (avoids the danger) or escapes away if the danger is moving and arriving. Later, while perceiving a possible danger or an unsafe zone or situation the robot, feeling a sense of anxiety, might multiply its investigating attitude and be more cautious but active for knowing about actual dangers or successes.

- *While planning to get to the room where the target object is located, the robot choose a particular path because it expects to feel pleasure and joy there, although the other path would be shorter.*

- ***Fetch that object!***

This is a human-robot interaction task focussed on believability. In the room there are several crates lie scattered around, acting as obstacles between Aibo and its searched target.
The human throws a red ball into the next room, then turns to an Aibo robot and says: "Fetch!" The robot should run into the room and designs a plan to find the red ball. While searching the space, its attention is drawn to a small handkerchief whose colour is just as the ball it is searching for. With its ear pointing forward, Aibo starts running, waving its tail and barking in anticipation. However, as soon as the robot realizes it is a mere handkerchief, its ears drop back and its tail falls between its legs. With a disappointed face, Aibo starts moving back, its gaze wandering across the room...

PREDICTING IN A DYNAMIC WORLD

[LUCS, OFAI]

This scenario involves prediction of objects that are characterized by an intrinsic dynamics. The first three tasks involve looking at two games with moving objects and the task is to predict the movement of one or several objects. The last one deals with the prediction of a rolling ball with increasing levels of complexity.

- ***The fish catching game (Game Room)***

In the fish catching game, the movement of the targets is very regular but there are two types of predictions that can be made:

- *the path of the fish*
- *the time when it will open its mouth.*

When the scene is viewed from different angles, the system need to predict the movements of the fish regardless of from where it is looking at it. Ideally, the learned model should allow for quick relearning (or reparameterization) when the viewing angle changes.

- ***The marble run game (Game Room)***

In the marble run, the movement is again very regular, but the different components of the game can be rearranged to produce different paths for the marble. These scenes combine the continuous dynamics of the ball with a compositional structure. This allows for generalization between different configurations of the elements of the run.

- ***Learning the two games at the same time (Game Room)***

To add some complexity to the previous tasks, the cognitive system could simultaneously look at and learn the different games. This makes the learning context sensitive. It also makes it possible to study how the current game can be used to prime the relevant features of the visual scene that should be used for anticipation. Ideally, the system should learn that there are two different games by itself by detecting the relevant contexts. The only given goal of the system will be to anticipate the state (e. g. location and velocity) of some predefined objects in the scene. By simulating a delay in the perceptual system (as would result if a robot was used), it becomes necessary to predict the behaviour of the moving object for tracking to occur.

There are a number of problems that can be studies within this scenario:

- *Anticipatory model based target tracking. Since the position data available is delayed, some form of prediction is necessary to accurately track the target.*
- *Tracking in a structured environment. In the marble run, the structure of the environment determines the movement of the tracked object. This allows for the study of many types of interaction between the target and the environment.*
- *Tracking of partially occluded targets. Since the target is sometimes occluded, the tracking system must be able to maintain the position of the target even when it is not visible. By anticipating where the target will reappear, a better performance will be possible.*
- *Generalization of target model from one viewing condition to another and between the different games. This stresses the importance of different coordinates systems and the separation between the viewer and the tracked object.*
- *Context dependent selection of target, predictive models and tracking strategy. The type of game, the configuration of the different elements in the scene and the viewing position are all contextual factors that influence the position of the target.*
- *Simultaneous modelling, recognition, prediction and estimation of viewer pose. Although each of these problems can be addressed separately, the scenario also makes*

it possible to simultaneously study these different tasks.

- *Optionally different types of anticipation based interaction with the target. In particular in the fish catching game, actually catching the ball may be a possible objective.*

➢ ***Tracking the rolling ball (Game Room)***

This task is a bottom-up scenario, inspired by the idea of interactivism. The scenario is divided into three developmental stages:

In the first early developmental stage ("how to" development), the robot starts basic interactions with its environment (like walking around, looking at things, poking them), driven by basic instincts and motivations. After a specific amount of time and training the robot learns through reinforcement which interactions environmental features, or to be more specific, objects allow, thus in the first stage the robot acquires concepts of objects and the world itself, if not to say affordances of objects.

- *The robot through interactions gains knowledge about the concept of objects and the world around it.*
- *The main developmental achievement is to acquire object generalisations and by that also certain concepts (one could even say affordances of objects such as the affordance of an object to be moved).*
- *As shown in the figure below, one achievement of this developmental stage will be e.g. to learn to expect where a moving ball will reappear after being occluded by an obstacle (e.g. a wall).*

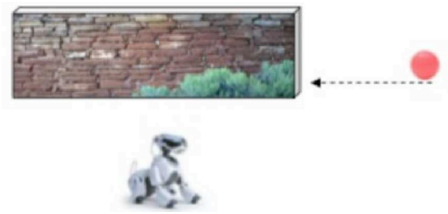

In the second developmental stage, after having acquired certain basic how-to knowledge about interacting with the world, the robot learns generalisation. The "how to" knowledge which has been acquired in the first level is anchored in a process of supervised learning.

- *Now the robot should develop more sophisticated concepts, such as object continuity – e.g. it learns to anticipate that, depending on the speed, the ball will reappear on the other side of or remain behind the wall.*
- *If the ball moves behind the wall with a low speed, the robot should learn to go looking for the ball on the right side of the wall, if the speed vector increases it should learn to go looking for the ball on the other end of the wall.*

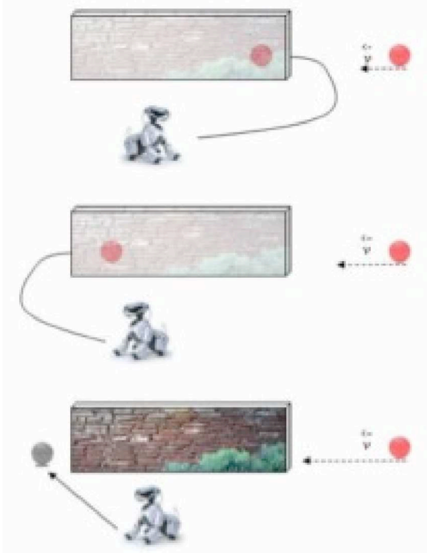

In another setup, the wall might be blocked at the end, and a ball, coming in with a high speed, normally reappearing on the other side, now does not reappear, and there is a sound (the ball bumping against the wall) coming from behind the wall.

- *The robot should then find out, that something has changed and that it needs to go and search for the ball from the other side again.*

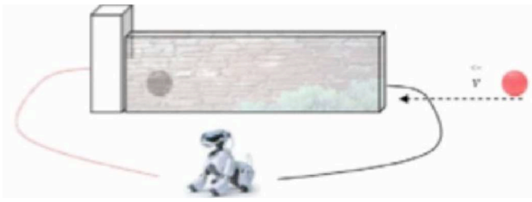

In the final stage, as the robot has seen the ball disappearing behind several "hiding places" (walls, obstacles), it shall now learn how to find the ball again and move around, looking for the ball, anticipating it to be in one of the observed "hiding places".

- *The final goal is to instantiate a hunter-prey – sub-scenario, where prediction of behaviour based on the capabilities acquired in stage one and two is implemented.*

A second robot can be added to the scenario (i.e. KURT3D, for further details next section). This task will converge on the following scenario. By observing the intruder, the AIBO shall "hunt" the intruder by simply intercepting it, or after some experience, going to a place where it anticipates the intruder will go, realising an offensive tactic.

- *The robot thus performs epistemic actions such as look if, look for. This offers a bridge to constructive perception, what means that in general, what is seen is interpreted by the means of what is expected.*
- *Expectation then can lead to "ask questions to the world".*

GUARDS AND THIEVES

[ISTC-CNR, NBU]

In this scenario, robots or agents can have two different roles (guards or thieves). Some objects are considered to be valuable and the thief's aim is to find and pick them all. The thief has a store where it places all the valuables it succeeds to take away. The goal of the guard is to protect the valuables. It could restrict access to them – either by blocking the entrance of the room when a thief is nearby, or by blocking access to the place where the valuable is.

In some cases, only one of the two kind of agents (thieves and guards) will be modelled as an anticipatory system, while the other kind will be a simple, routine-based one (used as a baseline adversary). In the more complex examples, both the kind of agents will be anticipatory agents.

➢ ***Conflict in accessing the valuables - simple (House)***

This is a social task involving two agents – one thief and one guard. In the beginning several valuables are hidden in at least two different places or there are several accesses to the hidden place, in order to make the guard's task non-trivial. The session ends either when the thief has collected or found all the valuables or when the guard has arrested the thief either by blocking him or by touching him.

This task includes the solution of the following problems:

- *Recognition of the adversary among the moving (or moveable) objects*
- *Prediction/anticipation of the adversary behaviour (avoiding/intercepting)*
  This can be done for instance by using counterfactual reasoning ("if I were the adversary I would …") or on the basis of previous experience about the opponent's behaviour. The thieve needs to find the optimal and safe route to a valuable and take it while the guard must find the optimal guarding route to prevent the latter.
- *Finding objects of the right kind (see previous scenario)*
- *Integrating different levels of action control (e.g. routinary, reasoning), based on different kinds of expectations (e.g. implicit, explicit) and being able to arbitrating them by shifting level of control or by mediating.*
- *Being able to transform the representations used for the different levels of control; e.g. learning as routinization of behaviours that are first adopted in a deliberative way; or abstracting concepts that are first learned in a trial-and-error way.*
- *Having two or more conflicting goals (e.g. protect two places), possibly conflicting, and arbitrating between them.*

➢ ***Conflict in the access to valuables - complex (House)***

This is a social task involving several agents – several thieves and a guard. The session ends either when all the valuables have been collected or found (no matter by whom) or when the guard has arrested (caught) all the thieves as described in this scenario.
In addition to all the problems listed before this task implies that the thieves should be able to distinguish between guards (danger) and rivals/fellows (competition/cooperation).

This task includes the solution of the following problems:

- *Help and critical help by anticipating other's needs, actions or capabilities, e.g. by removing obstacles or doing part of other's work.*
- *Delegating by trusting; e.g. an agent can explicitly "ask another one for help".*

- ***Coordination in accessing the valuables - several thieves (House)***

This is a social task involving several agents – several thieves (at least two). Some (types of) objects are considered to be valuable and each player aims to find them all. Thus the participants have to play the roles both of the thief and the guard from previous tasks. If one thief blocks another (the way the guardian could block the thief) the first takes the valuable from the second if currently it is carrying any. The session ends once a player has collected/found all the valuables or after some fixed amount of time.

One variety of this task is to coordinate the access to the valuables by means of a prediction of what the other agent "owns". The agent should be designed to support the understanding of what one has the right to access exclusively. *The agents are capable (and possibly learns) to understand that a resource or a territory (one or more rooms) are property of other agents and base their prediction on this knowledge and not on statistical means (who was accessing what in the past).*

Problems to solve in addition to the already listed:

- *Recognition of institutional relations between the agents*
- *Coordination based on anticipation*
- *Different bases for predicting the behaviour of others*

Anticipation of the adversary behaviour (avoiding/intercepting), including the more complex situation involving the combined "steal other's valuables" and "keep the stolen" behaviours.

7.2.3 Additional Scenarios Specification

In the course of discussion additional requirements on the robots that will be adopted and on the initial specification of the scenarios are emerged.

IDSIA

To approach the aforementioned tasks IDSIA will exploit the camera image as the sole sensor of the agent/robot. IDSIA plans to simulate a moving fovea, with possibly higher resolution in the centre and coarser resolution towards the borders.
Actions are abstract driving commands for the robot and velocity commands for the optical fovea. IDSIA does not intend to employ structures for explicit knowledge representation so long as it is possible to solve the tasks without them. For implicit representation of the environment, IDSIA will use neural networks, e.g. LSTM, that learn to "react" in a supervised or reinforcement-learning-type way to sensors and past experiences to generate actions.

Figure 1 The Robertino Robot

IDSIA uses a fully-autonomous Robertino robot (see Figure). It is equipped with a holonomic three-wheel drive, and has a PC-103 (industry standard) on board and communicates through WLAN. Its sensors consist mainly of an omnidirectional FireWire camera, on which we simulate the fovea. Other sensors will not be exploited. The actuators are the three wheels and the simulated fovea, and abstract commands for direction and velocity.

IDSIA will implement directly a simple simulator, because more realistic tools like ODE are overly complex and slow. The simulator is based on a physics simplification of the Robertino, and OpenGL for vision.

ISTC-CNR

One of ISTC-CNR main topics of interest are the high level aspects of cognition, such as practical reasoning, and the high level cognitive constructs such as beliefs, goals and explicit expectations. ISTC-CNR plans to approach these themes with many instruments: theoretical and conceptual analysis, logic formalization, simulations. With respect to these themes, ISTC-CNR architectures will be able to explicitly represent and reason about current and future states of affairs, choose the current behaviour on the basis on their expected consequences, deliberate over goals on the basis of supporting beliefs, about the past, the present and the future; perform epistemic actions with the goal of obtaining information and confirming or rejecting hypothesis.

Our models and architectures are designed to manage complex representations, not only related to sensory data but at different levels of abstraction; some of them will be even available for high level, symbolic reasoning. ISTC-CNR plans to explore a range of instruments and approaches: the BDI (belief, desire, intention) model of practical reasoning; a parallel and distributed control schema inspired by behaviour-based robotics; fuzzy systems; neural networks and evolutionary computation.

In order to tackle the scenario of object finding, ISTC-CNR will use a Pioneer 2 robot (see Figure below), produced by ActiveMedia. The robot will be equipped with the following sensors (these will be used depending on the different versions of the tasks tackled): a) bumper sensors (used for robot's safety); b) 16 front-and-rear ultrasound sensors (used for navigation and objects recognition); c) camera (used for object recognition and eventually navigation). The robot will be controlled with a portable computer set on top of the robot.

Figure 2 Robot Pioneer 2 (ActiveMedia), with and without camera

The group will also use customised simulators of the robot, and eventually software libraries to simulate vision. The controllers will be first developed and studied in simulation, and then will be tested or redeveloped on the real robot.
The group will use neural-network controllers that will be either evolved through genetic algorithms, or trained with learning algorithms (e.g., Hebbian learning, reinforcement learning). The group will use both monolithic and modular neural networks. Recurrent networks/architectures will be used in the case prediction will be exploited to guide control.

LUCS
To approach the tasks, LUCS will record movie of the two scenes (fish and marbles) from at least five different angles and coded in a number or formats:

1. MPEG at a resolution of 640x480 pixels, 25 frames per second. This is the raw format to use when the complete visual recognition and anticipation task is addressed.
2. MPEG at a lower resolution of 320x240 pixels, 5 frames per second. This format is used as reference for the raw tracking data when a lower bit rate is desired.

3. Raw tracking data (coordinates and state) for the target object in each movie at 25 values per second together with a static description of the scene.

The raw tracking data for the fish game will consist of the x and y coordinate of the target fish in the image and a third component that identifies whether the mouth of the fish is open. This data will be coded in ASCII files with four columns of numerical data (See table below). Since the scene is cyclic, the each data file will contain one cycle with typically lasts less than 5 seconds.
To allow generalization between different views the scene will be recorded five times from different viewing angles. The location of the camera relative to the center of the game will also be supplied.

Col	Data	Range	Unit
0	Time	≥ 0	Ms
1	X	0…1	Fraction of width
2	Y	0…1	Fraction of height
3	state	0…1	0=closed; 1=open

Table 2 Data file format for the fishing game tracking data

The raw tracking data for the marble game will consist of the x and y coordinate of the marble in the image and third component that identifies whether the mouth of the fish is open. This data will be coded in ASCII files with four columns of numerical data (See table below). The tracked data will contain the position of the marble from the time that it enters the scene until it disappears.

Col	Data	Range	Unit
0	Time	≥ 0	Ms
1	X	0…1	Fraction of width
2	Y	0…1	Fraction of height
3	visibility	0…1	0=occluded; 1=visible

Table 3 Data file format for the marble run tracking data

An additional file will contain a description of the elements of the scene separate from the position of the marble in the format specified in the following table. The two coordinates for each element indicate the start and end of the marble run through the element. For elements without clear locations of this kind, both coordinates code the centre of the element.

Col	Data	Range	Unit
0	type	≥ 0	Index of the element type

1	X_0	0…1	Fraction of width
2	Y_0	0…1	Fraction of height
3	X_1	0…1	Fraction of width
4	Y_1	0…1	Fraction of height

Table 4 Data file format for the marble run scene description

There will be five different scenes with different arrangements of the elements.
Two of these scenes will contain elements partially occluding the pathway of the marble. Each scene will be recorded from five different visual angles. The location of the camera relative to the centre of the game will also be supplied.

Each system that is tested with the data set is first trained for one or several iterations on the data set and then tested on the same or a different set of data for the same scene.

The system is trained and tested on delays of 100 ms, 250 ms, and 500 ms. For the fish game the following situations are of interest:

- Train on one data set, test on the same data set.
- Train on four data sets, test on the fifth data set.

For the marble run game, the following situations can be studied:

- Train on one data set, test on the same data set.
- Train on four views of one data set, test on the fifth view.
- Train on four configurations, test on the fifth configuration.

A combined situation is also possible in which the system is trained on four views of the fish game and four views of the marble run and then has to predict the movement of the target in on of the scenes without knowing which one.

NBU
The NBU's robots/agents (both virtual and real) should be able to perceive the environment (visually and auditory) and to move inside it and safely reach a target location. The robots should be able to catch and manipulate objects of a suitable size adapted to the available grippers. That is the robots should be equipped with an arm and/or a gripper or some functionally equivalent equipment.

The robots that will be used are:

- One Pioneer 3DX endowed with a 5 DOF arm, a gripper, and a visual system
- Several dogs AIBO capable of grasping and transferring small objects (aibones)

A simulated environment similar to the described above can also be built by using the open source ODE and Webot platforms.

The robots should be able to deal with the following relations:

- Spatial
- Temporal
- Causal

The spatial relations recognized should be of two kinds:

- Global – e.g. “North”, “West”, “South”, “East”, Above, Below, Angle (0°-North, 90°-West) etc.
- Ego-centric – e.g. Ahead, Left, Behind, Right, Above, Below, Angle (0°-Ahead, 90°-Left), In touch.

Example: “object #1 is in front of me” or “object #2 is to the left of object #1”, where the between-object relation is determined from the perspective of the observer.

In some cases, spatial relations could have a more precise quantitative specification, e.g. by making use of some coordinate system, which would allow the determination of objects’ size and position. The representation of the world by the robot may also include fuzzy spatial relations like “close”, “far”, etc.

The robots/agents need to identify and remember cause and effect relations and to be able to use them in decision-making. The challenge here is to recognize such relations in the environment either by statistic data accumulation or by event (episode) analysis and evaluation.

The temporal relations are related to:

- Temporal ordering – like “before”, “simultaneously” and “after”. They can be used to encode or learn cause-and-effect relations or event occurrence time;
- Duration – like “longer”, “shorter”, “same” etc.

The robots must be able to perceive and encode the following types of features:

- Shape – e.g. cubes, parallelepipeds, cylinders, aibones etc.
  To simplify shape recognition and thus object identification the number of shapes used should be limited and some colour code used. Depending on their shape the objects can be easily movable like balls or cylinders or more static like cubes. The latter can be used to control the complexity of the prediction/anticipation tasks to be executed in the environment.
- Size – e.g. “small”, “middle” and “big”.
  Size’s evaluation involves some fuzzy classification and depends on spatial relations.
- Colour
- Distance/position
- Emitted light/sound
- Time ordering and duration

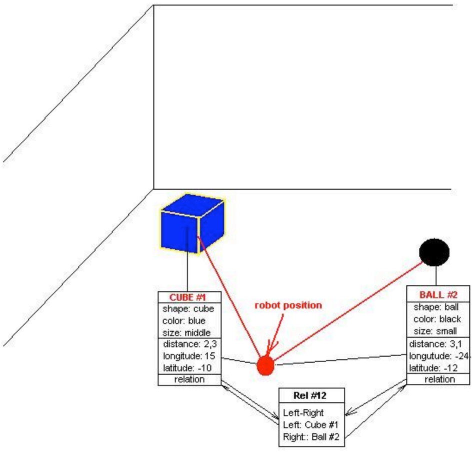

Figure 3 An example of relations recognition

OFAI

For the main experiments, OFAI is using the Sony AIBO robotic platform.

Figure 4 The AIBO robot

Although AIBO was created initially as an entertainment robot for the home, it has been embraced by many academics and researchers looking for a low-cost programmable robot platform. AIBO is completely programmable at a variety of different levels and is an excellent platform for research as well as education. There exists a family of different programming kits for AIBO, suitable for a wide variety of developers.

45 multi-color LEDs allow ERS-7M2 to express emotions. In addition, these LEDs indicate the status and function of ERS-7M2. Illume-Face (using 24 of these LEDs) provides a completely new way for ERS-7M2 to show when it is happy, sad, angry, surprised, etc. Tactile touch sensors on the back, head, and chin allow for more organic interaction, and also contribute to ERS-7M2's growth development.

AIBO is used in a wide range of university courses. Whether at undergraduate or postgraduate research level, it is used to support courses including, introductions to artificial intelligence, automated systems, decision-making processes and behaviour, voice recognition, image processing and software programming.

The main advantage of working with Sony's AIBO is that it is a complete and stable development platform. In addition, it features state of the art hardware and free and downloadable software-programming tools. This enables universities to fully gear resources and focus to programming in the area of Artificial Intelligence.

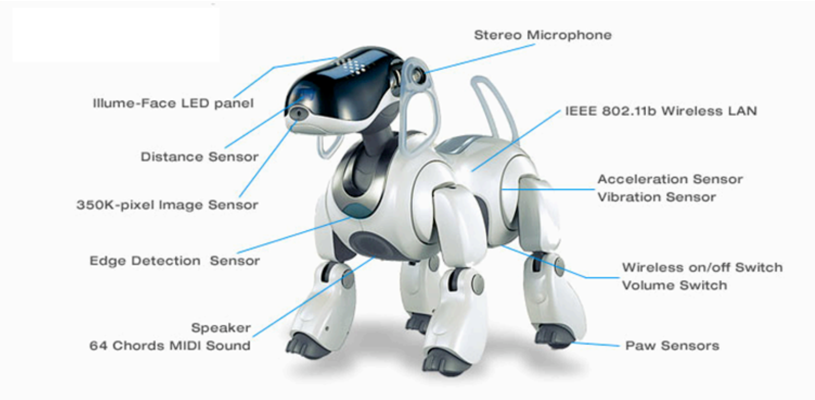

Figure 5 AIBO Features - Front
[source: Sony AIBO-Europe homepage: www.aibo-europe.com]

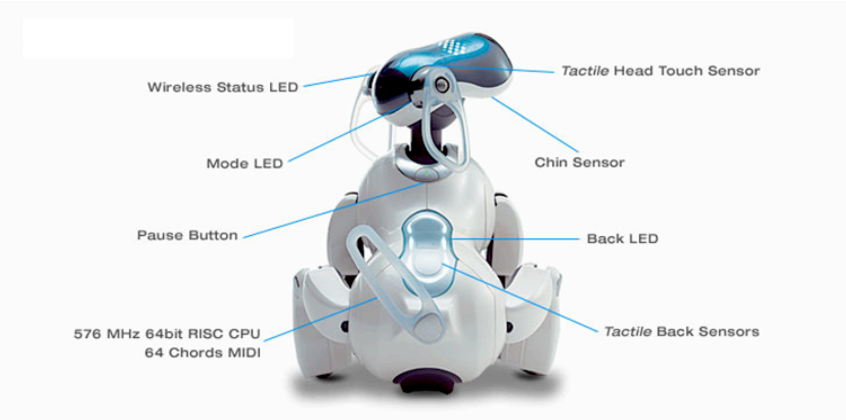

Figure 6 AIBO Features – Back
[source: Sony AIBO-Europe homepage: www.aibo-europe.com]

Figure 7 Sony AIBO Illume-Face capabilities
[source: Sony AIBO-Europe homepage: www.aibo-europe.com]

Sony ERS-7 features summarised:

- Components: Body, Head, Leg x 4, Tail
- CPU: 64bit RISC processor
- Main Storage: 64MB SDRAM Program Storage Medium Memory Stick™
- Input/Output: PC Card Slot Type 2 In/Out, Memory Stick™ Media Slot In/Out,
- AC in Power Supply Connector Input
- Image Input: CMOS Image Sensor (300K pixel)
- Audio Input: Miniature Stereo Microphones
- Audio Output: Miniature Speaker
- Built-in Sensors: Temperature Sensor, IR Distance Sensor, Acceleration Sensor, Pressure Sensors (head, face, back, legs and tail), Vibration Sensor
- Power Consumption approx. 9W (standard operation in autonomous mode)
- Battery Charging Time approx. 2 hours LCD Display Time, Date, Volume, Battery Condition
- Operation Temperature 5 - 35 degrees Celsius (41 - 95 F.) Operation Humidity 10 - 80%
- Dimensions: 180 x 278 x 319mm (w x h x l)
- Mass approx. 1.65Kg (including battery and Memory Stick™ media)

The Objects that the AIBO should be able to recognize have been specified are the AIBall, which comes with the AIBO equipment, that meets the requirements specified in Section 5. Additionally OFAI intends to add the AIBone because AIBO can grab the object with the mouth.

Figure 8 AIBall and AIBone as two examples for the objects used in the Game Room.

In the third stage of the OFAI scenario, a hunter-prey scenario has been suggested, as part of the Game Room scenario. For the realisation of this task, the AIBO robot will be used as hunter, after learning the required capabilities and as predator the KURT3D robot will be used.

Figure 9 KURT3D at the Fraunhofer laboratories
[source: http://www.ais.fraunhofer.de/ARC/kurt3D]

KURT is an experimental robot platform for sewerage inspection, hence the name, which is a German acronym for sewerage inspection robot ("*KanalUntersuchungsRoboter Testplattform*"). KURT3D is an autonomous mobile robot. The dimensions are 45 cm (length) x 33 cm (width) x 26 cm (height) and an approximate weight of 10.4 kg.

The robot carries an IBM ThinkPad T42p (1.8 Ghz, 512MB RAM, 2kg) and a 3D laser range finder (based on a Sick LMS, +7.0 kg) that increases the height to 51 cm and the weight to totally 22.6 kg. KURT2 operates for about 4 hours with one battery charge (28 NiMH cells, capacity: 4500 mAh). The core of the robot is a Pentium-III-600 MHz with 384 MB RAM and Real-Time Linux. An embedded 16-Bit CMOS microcontroller is used to control the motor and lower sensors (Phytec. Minimodul C167, inc. flash rom). The current maximum velocity the robot can be controlled with using the control architecture and the laser range finder (about 70 scans/sec, 181 values in 180 degrees) is 4.0 m/s (14.4 km/h) and is reached by two 90 Watt Maxon motors (transmission 1:14) .

The most important specifications summarized are (according to the KURT3D homepage):

- 90W motor
- power supply: 38V
- Main Sensor: 3D laser scanner based on a Sick LMS, 181 values in 181 degrees in 13 ms, 24V extra power supply
- Wheel encoders
- Maximal possible speed 5.4m/s
- Maximal controlled speed 4.0 m/s. (Controlled speed means that the robot avoids humans and other obstacles.)

- Weight: 3D laser: 7kg, Laptop: 4.2 kg, KURT2 body: 8.6 kg, cover: 2.8 kg (the cover includes batteries for the 3D scanner)
- Additional sensors: 2 cameras.

UW-COGSCI
Since UW-COGSCI does not work with robots directly, it is dependent on collaboration with other partners in this respect. However, for the next year, it is expected that work in simulated environments will be sufficient for the intended approaches to visual processing, attention, decision making and control. For this, UW-COGSCI already has implemented several simple arm simulators and intends to implement several other simulators suitable for the environments described above. As IDSIA, UW-COGSCI agrees that simpler, hand-coded simulators are sufficient for the investigation of base behaviors and learning mechanisms. Software, such as ODE, might be more realistic, but too time consuming to run and learn with. Moreover, UW-COGSCI intends to use the purely visual data provided by LUCS for the visual processing and attention-related challenges ahead.

UW-COGSCI intends to approach the aforementioned scenarios from two sides: (1) the sensory processing side and (2) the action control side. Both processing sides will be structured hierarchically and interconnected where appropriate. On both sides, UW-COGSCI intends to implement basic pre-processing units by hand. On the action control side, basic gradient fields will be provided. These fields may be controlled, that is, excited and inhibited, appropriately from higher control structures. Similarly, on the sensory processing side, UW-COGSCI will provide simple pre-processing stages that process sensory input, either on the low-level resulting in a lateral geniculate nucleus simulator effectively whitening (or decorrelating) sensory input, or, on the high-level side, providing a feature-based input representation. UW-COGSCI will try to learn the latter representation, though, expanding the available hot research topic on symbol grounding.

UW-COGSCI intends to approach these challenges with a combination or neural-based processing units and rule-based units. Neural processing units will include LSTM-type, gated units and hierarchical, predictive coding units in the style of Rao and Ballard (see D4.1 for details).These will be learned using neuroevolution techniques as well as gradient-based, Hebbian or backpropagation techniques. The rule-based processing units will use the principles derived from the XCS classifier system as well as from anticipatory learning classifier systems (ALCSs).

Both approaches are intended to be developed according to the principles of anticipatory behavior control, outlined by Hoffman in his original book. The major emphasize in this respect lies in the continuous comparison of predicted and actual behavioral consequences and the learning from the differences of these consequences. These principles are realized in the framework of Rao and Ballard as well as in the XCS and ALCS systems.

PART 3 – APPENDIX: The original partners' proposals

The following theoretical scenarios have been developed and posted for discussion by each partner during M4-M6 of the MindRACES project. They represent the initial basis on which the final scenarios illustrated in Section 7 have been agreed and specified.

8 WP3: Attention, Monitoring and Control (ISTC-CNR)

SCENARIO 1: Anticipation of physical worlds. *A robot endowed with a camera and/or infrared sensors and/or ultrasound sensors navigates an office formed by few rooms, on the basis of a stereotyped wall-following behaviour. The task of the robot is to predict the next "primitive" percept at t+1, or some primitive percepts at t+x in the future; or the task is to predict at a more abstract level, by categorising sequences of primitive percepts and by predicting on the basis of these categories.*

Notice that the robot's action has not a role in the task since it is generated by a stereotyped behaviour: the task would be the same if the robot is mounted on a trolley that follows a close-loop track, or if the robot sits on a chair and simply "watches" a video-recorded film.

Here we present the cognitive capabilities required by this scenario, the possible function that they might play within cognitive systems, the possible mechanisms used to implement them, some references to the relevant literature.

1. *Predict the next "primitive" percept at t+1, or some primitive percepts at time t+x in the future*
   - Using past primitive percepts to predict future percepts: The robot has to build a model of the world that captures regularities in time within the experienced perceptual sequence.

This cognitive functionality might play an important *role in a number of higher level cognitive functions* (this same functions can be played by the capacity, discussed below, to predict at a more abstract level):
- At very fine time scales, this capability might allow the generation of a pseudo-feedback simulating the feedback from the motor plant if the real feedback is too slow.
- Prepare action (or own state) in order to suitably react (or be in a desired state) when the world will assume certain anticipated states in the future.
- Generate surprise: when the world does not match the predictions, this can be used as an indication that the world has changed, or that the model is not complete (to the extent that one can distinguish between surprise due to ignorance and surprise due to truly unpredictable events). On its turn, curiosity can have different cognitive functions.
- Predicting as the basis for decision making: e.g. in lotteries, economy (stock exchange), or other situations involving decisions highly based on the capacity to anticipate future events.

The *mechanisms* that can be used to implement this functionality can be based on algorithms and architectures suitable to implement function regression. Example of this are:

- multiple-layer neural networks learning on the basis of gradient descent algorithms (error-back propagation and delta rule): the agents learns to associate to the experienced sequence of input patterns {st-n, st-(n-1), st-(n-2), …st-2, st-1, st} one or more of the following future input patterns {st+1, st+2, st+3, st+4, st+5, …}:
- Elman J.L.(1990). Finding Structure in Time. Cognitive Science, 14, pp. 179-211.

- Long-short term memory:
- F. Gers, N. Schraudolph, J. Schmidhuber. Learning precise timing with LSTM recurrent networks. *Journal of Machine Learning Research* 3:115-143, 2002.
- S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735-1780, 1997

- It would be interesting to investigate if Hebbian learning can be used to build the models of the world. These algorithms allows having a learning process that is robust with respect to possible variations of times in which the events take place:
-Abbot, L.F., Blum, K.I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex*, Vol. 6, pp. 406-416.

2. *Categorising sequences of primitive percepts*
   - Recognise/categorise events, that is sequences of primitive percepts, on the basis of their temporal structure (i.e. sequences of primitive percepts with particular order in time).

The capacity to categorise events might play several different *roles in cognitive systems*. What is relevant here is that this capacity allows predicting the world at a more abstract, coarse level (see next point n. 3).

Some of the *mechanisms* that can be used to categorise events (sequences):
- Kohonen-like networks that can be used to categorise compact representations of sequences:
- Nolfi S., Tani J. (1999). Extracting regularities in space and time through a cascade of prediction networks: The case of a mobile robot navigating in a structured environment. *Connection Science*, vol. 11(2), pp. 129-152.
- Neural networks based on units implementing 'counters' can be used to categorise events since they can count time delay between events' portions:
- F. Gers, N. Schraudolph, J. Schmidhuber. Learning precise timing with LSTM recurrent networks. Journal of Machine Learning Research 3:115-143, 2002.
- Nolfi S. (2002). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science* (14) 3:231-244

3. *…at a more abstract level…predicting on the basis of these categories*
   - Predict at an abstract level, that is at the level of the categories mentioned at the previous point. The system might categorise past events, and then use this categorisation to predict events that will take place in the future.

Predicting at a more abstract level might play the same *roles in cognitive systems* as the capacity to predict at the primitive level. However, predicting at a more abstract level might allow the system to focus on a time scale more suitable to achieve/satisfy its goals/needs (say: the time scale of minutes/hours vs. the time scales below one second). Moreover, some events are better or only predictable at a more abstract level. This means that sometimes it is easier (possible) to predict the category of the future events, but not the precise events in the future (e.g. the category of word, as a verb, that will follow some word such as "I will…."; the type of things I will see by going to "the park", but not the details of them).

This capacity might be based on the following type of *mechanisms*:
The same architectures and algorithms used to predict at a primitive level might be used to predict at a more abstract level:
- Nolfi S., Tani J. (1999). Extracting regularities in space and time through a cascade of prediction networks: The case of a mobile robot navigating in a structured environment. *Connection Science*, vol. 11(2), pp. 129-152.
Mechanisms capable of predicting at a symbolic level might be used to predict in terms of sequences of categories instead of primitive percepts? For example, consider the symbolic prediction based on "fractal machine":
-Tino P., Dorffner G.: Building predictive models from spatial representations of symbolic sequences, in Solla S.A., et al.(eds.), *Advances in Neural Information Processing Systems* 12, MIT Press, Cambridge/Boston/London, pp. 645-651, 2000.
-Tino P., Dorffner G.: Predicting the future of discrete sequences from fractal representations of the past, *Machine Learning*, 45(2)187-217, 2001.

SCENARIO 2: anticipation of intelligent (teleonomic) systems. *A robot (a camera) stands in front of another robot (camera + arm) or a human, observes its actions (that is the execution of different sequences of micro-movements) and: 1) recognises the other agent's actions; eventually the recognition takes place on the basis of the observation of the execution of only the first portion of the action sequence; 2) recognises the intended effect on the environment of a particular action executed by the other agent, that is its goal (e.g. grasping a particular object); eventually the recognition of the gaol takes place on the basis of the observation of the execution of only the first portion of the action.*

Here we present the cognitive functionalities involved by this scenario, the possible roles they might play within a cognitive system, the possible mechanisms used to implement them, some references to the relevant literature.

1. *recognises the other agent's actions; eventually the recognition takes place on the basis of the observation of the execution of only the first portion of the action sequence.*
    - *This cognitive capability allows the agent to categorise (i.e. recognise) actions lasting for some time*. If the agent is capable of recognising the action by observing only its first portion, this capability can be used to predict the remaining not-yet executed micro-movements composing the action.

With respect to the *roles that this and the following cognitive capabilities might play in cognitive systems*: these capabilities are very important in both cooperative and competitive set-ups, for example: building things together, aiding others to do things, hunting together a prey, avoiding to be caught, fighting against another agent.

Mechanisms:
The same mechanisms used to categorise perceptual sequences produced by physical systems can be usable to implement this cognitive functionality (see above). Question: what differences does it make if the systems to predict are teleonomic systems or physical systems?

2. *recognises the intended effect on the environment of a particular action executed by the other agent, that is its goal (e.g. grasping a particular object); eventually the recognition of the gaol takes place on the basis of the observation of the execution of only the first portion of the action.*
    - *This cognitive capability allows the agent to categorise the type of effect that the other agents is attempting to cause in the environment*. This capability implies having a model of the other agent's action *and* a model of the way the environment reacts to it. If the agent is capable of recognising the action's goal by observing the execution of the first portion of the action, this allows it to predict the relevant change that the action will cause in the environment.

Mirror neurons indicate that monkeys are capable of recognising complex actions (reaching, precision grasping, strength grasping, tearing, etc.), eventually on the basis of the execution of the first part of them:
- Fogassi L., Gallese V., Buccino G., Craighero L., Fadiga L., Rizzolatti G., Cortical mechanisms for the visual guidance of hand grasping movements in the monkey: a reversible inactivation study. *Brain* 124:571-586, 2001.
There are some preliminary models of the mirror-neuron system:
- Arbib M.A., Billard A., Iacoboni M., Oztop E. (2000). Synthetic brain imaging: grasping, mirror neurons and imitation. *Neural Networks*, Vol. 13, pp. 975-997.
- M. Ito and J. Tani: "On-line imitative interaction with a humanoid robot using a mirror neuron model", *Proc. 2004 IEEE Int. Conf. on Robotics & Automation* (ICRA2004), New Orleans, USA, pp.1071-1076, 2004.

9 WP3: Attention, Monitoring and Control (LUCS)

The following two scenarios address many of the problems that we hope the MindRACES project will study. Although we have tried to stay within the topic of perception rather than action control, we feel that these processes are highly interwoven which means that we could not avoid some overlap with the other topics.

When searching for real-world situations that could form the basis for our scenarios, we eventually came up with the idea of looking at children's games. Many of these appear to train children in exactly the skills that are the topic of the MindRACES project.

We have continued the numbering of scenarios with 3 and 4.

9.1 Scenario 3: Visual Prediction

A robot with a movable camera (and possibly with an arm and ability to move around) is looking at different games where the task is to predict the motion of an object. We suggest that there are three such games. (a) A fish catching game (see figure), (b) a marble run (see figure), and (c) an ordinary ball rolling toward the robot. In all three cases the task is to predict the movement of one or several objects but the situations are also quite different.

(a) In the fish catching game, the movement of the targets is very regular but there are two types of predictions that can be made: the path of the fish and the time when it will open its mouth. There is an interesting generalization situation here if the robot is allowed to move where it has to predict the movements of the fish regardless of from where it is looking at them. Ideally, the learned model should allow for quick relearning (or reparameterization) when the robot moves.

(b) In the marble run, the movement is again very regular, but the different components of the game can be rearranged to produce different paths for the marble. We want to study how a robot can learn to use the layout of the different components of the game to predict how the marble will move. This appears to be a perfect area to study not only anticipation, but also generalization of learned anticipations to new situations.

(c) Finally in the "roll the ball" game, the ball takes different paths every time as a human rolls it toward the robot. The robot has to learn to predict the motion locally based on learned behavior of the ball, for example that the ball is likely to move along a straight path at a certain speed. These are prediction in "ball-centric" coordinates rather than world coordinates as in the first two games.

Figure 10 The Fishing Game and the Marble Run Game

To add some complexity to the scenario, we would like the robot to learn the three games simultaneously in such a way that it can switch between them at any time. This makes the learning context sensitive. It also makes it possible to study how the current game can be used to prime the relevant features of the visual scene that should be used for anticipation. Ideally, the system should learn that there are three different games by itself by detecting the relevant contexts. The only given goal of the system should be to anticipate the state (e. g. location and velocity) of some predefined objects.

In all three cases, there is an obvious connection to actions if we add the requirement that the robot should catch the fish, marble or ball. (At least in the marble game, it would be great if the robot could lift the marble to the starting position on its own during training).

This scenario builds on our previous studies of learning mechanism in visual attention (Balkenius, 2000, Balkenius, Åström and Eriksson, 2004) and the detection of context (Balkenius, 2003, Balkenius, and Morén, 2000, Björne and Balkenius, 2004) and the use of context in the selection of actions (Balkenius and Björne, 2001, Balkenius and Winberg, 2004).

There is also a visual categorization components in this scenario which mainly lies outside the topics of the MindRACES project, but where we already have the required algorithms available (e. g. Balkenius, 1998, Balkenius and Kopp, 1997, Kopp, 2003).

9.2 Scenario 4: Anticipation of Group Actions

Our second scenario consists of a group of robots that play different children's games, e. g. (a) Race, (b) Tag, and (c) Hide-and-Seek.

(a) In the racing game the goal is to move from any initial location to a goal location. This in itself is a trivial action - what makes it interesting is that there are many robots trying to do the same thing at the same time. Unless they are able to predict the movements of the other robots they will all collide.

The game can be made more interesting by including obstacles in the environment that the robots must avoid. This makes the task for each robot more complicated, but the prediction task also becomes more complex. This situation also entails many of the aspects of the first scenario proposed by ISTC on anticipation in perception. There is also the possibility of studying 'recursive anticipation', i. e. that the a robot anticipates that the other robots will anticipates its movements and move accordingly.

To some extent, the work on flocking behaviour in video games and computer graphics have studied this problem, but usually in a situation where all agents have complete knowledge or are centrally controlled. In this scenario, we assume that each robot has limited attentional resources that must be allocated wisely to collect as much information as possible to predict the locations and movements of the other robots while not looking so much that they forget to move toward the goal. It is clear that although the game is extremely simple, the anticipatory abilities that can be studied within it are very complicated.

(b) In the game of tag, the goal is for one of the robots to catch any of the other robots which then becomes the chaser (It!). Like in the race game, each robot has to anticipate the movements of the other robots to avoid colliding. In addition, they also have to anticipate the movement of the chaser to avoid it. The dynamics in tag is the opposite of the race, Instead of having all robots heading toward the same location, in tag they are all heading away from the location of the chaser.

Again, the type of predictions that can be made can range from the trivial to the very complex. For example, a robot can use the anticipated movements of another robot to avoid the chaser and make that robot the target instead. For the chasing robot, the task is similar although the goal is different.

(c) In hide-and-seek the goal of each robot is to find a location in the environment where it is not visible for the robot that is looking for it. The seeking robot on the other hand must predict where the other robots are and look for them as efficiently as possible.

The game can be made more interesting by allowing the robots to move while the seeking robot is looking for them. In this case, they may look at the seeking robot and try to anticipate where it will look next and avoid being there to bee seen. In this case the game approaches the watch-dog scenario proposed by ISTC.

All games includes context sensitivity in the trivial sense that the task of robot differs depending on its current role in the game. There can also be different amount of learning from only the anticipatory component of the game to learning the actual rules through trial and error.

We envision a set-up with a number of small radio controlled robots with easily trackable markers on top which are monitored by a single overhead camera. Each robot is controlled individually and only has access to a small part of the complete image that corresponds to its visual field. This allows the sensory abilities of the robots to be manipulated in a simple way and also allows the behavior of the robots to be recorded for evaluation purposes.

References

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

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

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

Balkenius, C., Åström, K. and Eriksson, A. P. (2004). Learning in visual attention. In *ICPR '04 workshop on learning for adaptable visual systems* (LAVS).

Balkenius, C., and Björne, P. (2001). Toward a robot model of attention-deficit hyperactivity disorder (ADHD). In Balkenius, C., Zlatev, J., Kozima, H., Dautenhahn, K., and Breazeal, C. (Eds.), *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund University Cognitive Studies, 85.

Balkenius, C., and Kopp, L. (1997). Elastic template matching as a basis for visual landmark recognition and spatial navigation. In Nehmzow, U., and Sharkey, N. (Eds.), *Proceedings of AISB workshop on Spatial reasoning in mobile robots and animals.* Manchester: Manchester University, Department of Computer Science, Report number UMCS-97-4-1.

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

Björne, P., and Balkenius, C. (2004). Cognitive development in context: Learning to pay attention. In *Proceedings of ICDL '04*.

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

10 WP3: Attention, Monitoring and Control (IDSIA)

A (possibly but not necessarily mobile) robot with camera or other directional sensors has attention-shifting actions such as "turn sensor right by 10 degrees."

Its goal: find some object in the visual scene, through active perception, by producing a sequence of saccades and other actions until the object is centered in the visual field.

Arbitrary degrees of difficulty are possible through complex visual scenes, partial observability, partial occlusions etc.

The robot may profit from sequential search for informative inputs that increase the probability that subsequent actions will bring even more promising search areas into view. For example, if the goal is to center in the visual field the cork of a wine bottle, then better first find a bottle, because once you have a bottle in view the remaining attention-shifting actions needed to focus the camera on the cork may be easy. And bottles are more likely to be located on the table, not on the floor, so a rational attentive system should first focus on the table, etc.

In the supervised case a teacher whose input is exactly the input of the robot (no global view for the teacher) presents several successful trajectories leading to the desired input; the robot learns to imitate the teacher, given the sequences of sensory inputs. The essential tests take place in previously unseen situations outside the training set: can the robot generalize and still find the target object?

In case a recurrent network is used for supervised learning: Certain activation patterns across the internal units of the trained network may be viewed as embodying the results of constructive perception. For example, an "apple detector" might switch on for a while when an apple has been observed through a sequence of actions leading to a sequence of low-level inputs, to be compactly represented by the high-level apple detector.

In the intuitively harder reinforcement learning (RL) case there is no teacher but only reward once the target object is centered, and some sort of RL algorithm will be necessary.

Selective attention as illustrated above implies the anticipation of information gain: the robot should perform action sequences that will increase its knowledge about the present surroundings, such that it can expect to quickly identify the target object.

11 WP4: Goal Directed Behavior, Pro-activity and Analogy (ISTC-CNR)

SCENARIO 1: The watchdog*: a robot that patrols a house and defends it by an intruder. The house has three doors (that can not be spotted from the same location) that can be open or close; there are trees and rocks all around. Some more dynamic elements can be added: e.g. night and day, or moving elements.*

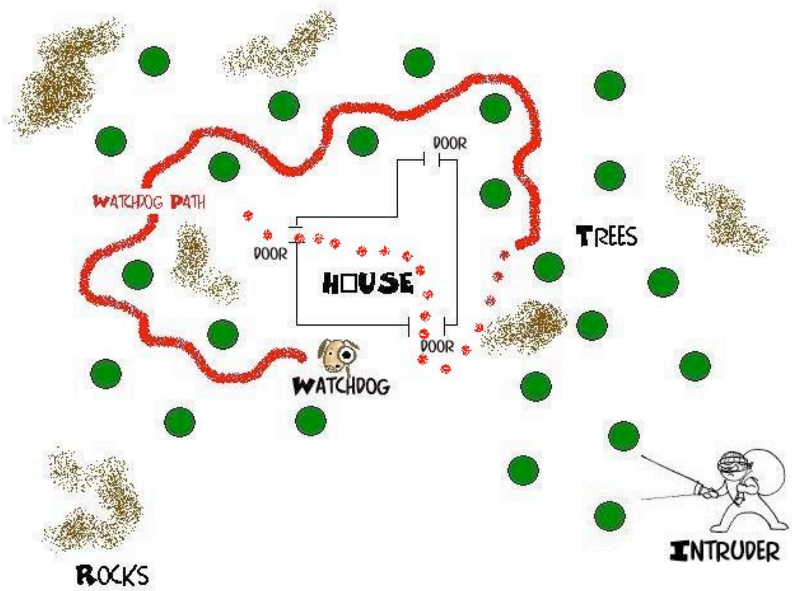

1. *The Watchdog learns to follow a good "patrol path" all around the house that traverses all the doors and many possible hiding places (Defensive Tactic).*
   - <u>Using past percepts and actions to predict future percepts and actions</u>: The Watchdog needs a (predictive) model of its world that affords (at least) navigation, both in a routinary and in a deliberative way; this *model building activity* is preliminary to many other capabilities. For example, the Watchdog can learn a good path (e.g. by reinforcement) by "attempting" different actions/directions and by generating appropriate expectations. Different kinds of anticipations, used for many cognitive functions, can be built:
     - Implicit (i.e. embedded into the S-R structure, just because Responses are selected to be adapt to the future state of the world with respect to Stimulus);
     - Explicit type 1 or Reliability (embedded as an anticipatory part in a 3-parts schema, e.g. the E of anticipatory classifiers S-R-E);
     - Explicit type 2 (these are real, explicit expectations; explicitly represented and used for deliberate control and reasoning; e.g. a belief in BDI).

[Many references (e.g. about "Predictive learning") are in the Annex 1 (pagg. 13-22), including predictive representations of state (PSRs); observable operator models (OOMs); reinforcement learning (especially with respect to temporal issues); anticipatory classifiers; BDI etc.]
[Some more interesting pointers: In the Drescher's schema mechanism objects are not primitives; they are "synthetic items", represented as a sets of expected interactions with the environment.

Drescher, G. L. (1991). Made-Up Minds: A Constructivist Approach to Artificial Intelligence. Cambridge, Massachusetts: MIT Press. A more extended and very interesting approach, based on schema theory, is proposed by *Deb Roy. Grounding Language in the World: Schema Theory Meets Semiotics. (Submitted to Artificial Intelligence)*].

- <u>Build an expectation as a consequence of each action of its action repertoire</u>. If the actions are represented as 3-parts schemes, the expectation can be extracted (thus become available in an explicit form). *Evolving explicit expectations from implicit ones* is a central evolutionary issue; how it is possible to pass from implicit to explicit forms of anticipation? Which are the specific advantages? This capability is a building block of many others, e.g. used for *deciding what to do next* (planning, reasoning about the possible consequences of each actions).

- <u>Building macro actions and plans from a repertoire of micro actions</u> (e.g. by concatenating them). Actions can be assembled into more coarse-grained representations, that are more abstract and general; for example, the action "reach point x" can be activated from different starting points and it is able to adapt on-line to the context. Even expectations can be produced at a coarse-grained level (e.g. only expectations about important intermediate steps, or only about the final results); this simplifies the monitoring cycle but makes the system more vulnerable to errors.

- <u>Planning (and proto-planning).</u> Fine or coarse-grained representations such as macro actions and plans can be exploited for new behaviors; for example, if required the Watchdog is able to go from point x to point y by explicitly planning a path even if has not a routine from x to y. It needs explicit expectations (of the consequences of its actions/plans) as well as a way to combine micro and macro action to reach a goal state. It can re-plan if some expectations are violated.

[You find in the Annex 1 some references to the family of Dyna-based models; including the architecture of Baldassarre Baldassarre G. (2002). *Planning with Neural Networks and Reinforcement Learning.* PhD Thesis. Colchester - UK: Computer Science Department, University of Essex.]

- <u>Action monitoring:</u> the expectation (e.g. embedded in a 3-parts schema) is used for on-line control by matching it with the perception. The match-mismatch information can be used for a monitoring activity that leads to actions tuning and adjustments. A serious mismatch can even generate surprise.

- <u>Generate surprise: when the world does not match the anticipations.</u> during *action monitoring*, expectations are always matched with perception ("if I do x, y will happen"). If there is a mismatch (i.e. an expectation is violated), this generates surprise. Surprise can be generated both in automatic and deliberate control of actions and used as a building block of a number of cognitive functions (such as the following two).

- Routinization/compilation of sequences/patterns of actions by surprise. The Watchdog can start by explicitly planning paths (slow and requiring many resources); after a while it "routinizes" a certain path (that no longer needs deliberate control); some expectations are used in the routines for on-line action monitoring. This is in fact a model of *"skill learning" as "compilation" of sequences/patterns of actions*. The routinization mechanism is based on surprise: when a planned sequence of actions no longer generates surprise, it can be safely routinized. However, the expectation is not lost, but embedded into the routine (we call it reliability) and used for on-line, automatic action monitoring (e.g. using anticipatory classifiers).

- Passing from routine to deliberate control by surprise. The converse operation (un-compiling skills) uses surprise, too. When a surprise is generated in the automatic control of the routine the system can pass anew from the routine to the deliberate control (e.g. activate a goal or a plan). This happens e.g. when the Watchdog encounters an un-expected situation in its path (e.g. a new obstacle, a door is close, another agent passing). The routine is stopped, many resources and/or attentive control are raised, etc.

- On-line adjusting/tuning of plans and routines; combine deliberative planning with reactive plan execution (including reactive local plan optimization during execution). If the Watchdog sees an intruder it can build a plan to reach it; but the intruder moves, it can adjust the plan on-line. Minor expectation violations (requiring fine-tuning of actions) should be allowed without "special" mechanisms such as surprise; they simply show that the system is robust and quite fault-tolerant.

[Aaron Sloman describes a demo of an hybrid planning-reactive mechanism combining planning and on-line adjusting of actions (implemented in his Simagent toolkit). Available online on: http://www.cs.bham.ac.uk/research/poplog/figs/simagent/#hybriddemo]

2. *The Watchdog learns the behavior of the intruder in order to actively search it (Offensive Tactic).*
   - Prediction of behavior. For example, the Watchdog can spot some (hiding) places because he found there the intruder in the past, (e.g. by reinforcement); or exploit categorical reasoning (e.g. intruders always hide behind trees; intrudes always come by night); or analogy (this place is similar to the other I found the intruder; it always follows a certain strategy).
   - Goals Managing. The watchdog can have two or more competing patrolling Goals (e.g. "stay always close to the house" and "when you see the intruder, follow it". It needs an on-line effective mechanism for "deciding what to do next". There can be many possible strategies: automatic activation of one goal, with inhibition of the other one; explicit decision to fulfill one of them (e.g. by means-ends calculus); "mixing up" of the two goals (the resulting trajectory will emerge from some constraints of both), etc. Both automatic and explicit decisions about goals mainly depend on their expected results. A related point is exploiting opportunities: a goal can be preferred (and an action/plan executed) because it opens some possibilities to other goals, even if these results are not explicitly intended in the

goal selection phase (e.g. "come back to the house" produces exploitable results for "control if the main door is open").

3. *The Watchdog "hides" and "ambushes" the intruder by going to a place where it anticipates the intruder will go (Offensive Tactic).*
   - The watchdog exploits explicit expectations for deciding what to do next ("the intruder will pass from this point") in order to find a good location for waiting the intruder.
   - Hiding (and understand that it can not be spotted in a certain place) is really complex and challenging. It involves not only advanced perceptive and attentive capabilities, but at least a minimal "theory of mind" of the other agent (and it is thus "social"). A more sophisticated (social) extension of this scenario is having many (collaborative) intruders and one Watchdog that hides; or the converse.

4. *A POSSIBLE BRIDGE TO WP3, NOT TO BE DISCUSSED HERE: The Watchdog actively searches for intruders: traces and events (Offensive Tactic).*
   - Perform epistemic actions such as look if, look for. The Watchdog actively searches traces and indices for locating the intruder. This involves driving perception and attention (search into some spots) and performing actions that have epistemic goals (e.g. illuminating a spot).
   - This offers a bridge to constructive perception. In general, *what is seen is interpreted by the means of what is expected*; expectation can lead to "ask questions to the world", an abductive process that is driven by hypothesis and expectations, that are used to search/select perceptive data to be matched. This case is similar to action monitoring, but the action is actually an epistemic action. A surprise can thus lead to an epistemic rather that an action revision (e.g. revise a background hypothesis about what is happening).
   - Moreover, the Watchdog can interpret what it sees in terms of its searching activity (e.g. if it founds a tree with a broken branch it can interpret it as the passage of the intruder; if it founds a closed door it can form the hypothesis that the intruder closed it); objects that afford different actions (a tree can afford hiding and repairing from rain) are categorized my the means of the current activity/goal.

SCENARIO 2: The watchdogs*: many watchdogs that collaborate*

The scenario is the same but there is more emphasis on the collaborative side.

5. *They can learn to have "areas of influence" and dynamically split the patrol activity*
   - Help and critical help by anticipating other's needs, actions or capabilities. Many possibilities for cooperative behavior; for example, a watchdog can "help" another by providing it an information (or performing a certain action, e.g. removing an obstacle; or by letting an olfactive track) only if it knows that it is relevant/necessary to the other in order to achieve its goals. Or they can learn to search in a given direction that they know no other watchdog can spot. All this involves a (more or less complex) understanding/theory of mind of other's behaviors and goals. It has to be noticed that there can be both explicit and implicit "division of labor".

[You can refer to many works by Castelfranchi and Falcone, e.g.: *Castelfranchi, C., Falcone, R., (1998) Towards a Theory of Delegation for Agent-based Systems, Robotics and Autonomous Systems, Special issue on Multi-Agent Rationality, Elsevier Editor, Vol 24, Nos 3-4, , pp.141-157.*]

6. *Each Watchdog can be specialized for a given function (search traces, see far, chase); each has a given availability (depending on its current tasks);*
   - Delegating by trusting. A Watchdog can explicitly "ask another one for help", choosing the best Watchdog to ask depending on trust (i.e. evaluating its competence, availability, possible environmental obstacles, past interactions, etc).

[You can refer to many works by Castelfranchi and Falcone, e.g.*: Falcone R., Castelfranchi C. (2001). Social Trust: A Cognitive Approach. In Castelfranchi C., Tan Y.H., Trust and Deception in Virtual Societies (pp. 55-90). Kluwer Academic Publishers.*]

A Final note: in these scenarios we model the Watchdog and not the Intruder side. In the Kickoff meeting LUCS proposed a scenario for modeling the Intruder: we have the opportunity to integrate the two scenarios and let the two agents interact/compete.

12 WP4: Goal Directed Behavior, Pro-activity and Analogy (OFAI)

The following section represents a brief (partial and open) discussion about the background of the OFAI scenario.

The “artificial life” route towards artificial intelligence (Steels and Brooks, 1994) argues that behaviour-based robots can function within dynamic environments, but they fail to go any further and reach higher levels of intelligence, because they lack representation. Brooks and others, e.g. (Brooks, 1991), have argued strictly against representation, what applies on a purely stimulus-response-based level.

Representations acting as stand-in for objects in the world (external representations, (Gombrich, 1963) imply a conceptualisation of the world, which is expressed using the properties of the medium and the set of emerging conventions in the group. Piaget showed that the ability to construct and interpret (external) representations could be seen as crucial in the development of children. (Piaget, 1970). If it does not happen, it is an early indicator of mental retardation. These findings suggest that representation-making might be a crucial bootstrapping device for higher mental function. At some point, external representation-making becomes internalised to form the basis of thinking, in the sense of inner dialogs or mental imaging (Steels, 2003). Internal representations co-evolve with external representations on a natural way. It is important to ask the question – why should the robot do something in the first place? Why should it patrol the house and consider doors as something to be worth to be protected? Let’s consider that, according to Steels, representations should be seen as organisers of activity rather than abstract models of some aspect of reality. Steels argues that the intervention of a cognitive agent is essential for a material structure to become a representation. The structure must trigger categorisation and then action selection which depends on the outcome of this categorisation. (Steels, 2003)

As a side note, McFarland’s ethological research in animal behaviour (McFarland and Bosser, 1994) suggests, that robots’ behaviours need to be motivated. This is well known and established in behaviour-based robotics. Basic motivations in the OFAI scenario could be: e.g. the drive to wander around, to discover an intruder (moving object), to look for intruders, hunger (recharge), etc.

Based on these questions and insights, the scenario proposed by OFAI will be a bottom-up approach, descending from the ISTC scenario.

In the ISTC scenario a watchdog sees an intruder running towards a tree, then disappearing and hiding behind the tree. The first expectation would be that the intruder appears again from behind the tree after some time. If this does not happen, the dog needs to approach the tree and look for intruder. In the following scenario, the intruder is being represented by a moving object; the tree is being replaced by a wall. The first main goal is to learn continuity and then learn to expect where the ball (intruders) reappears.

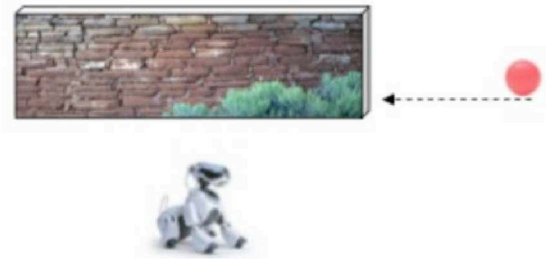

Figure 11 AIBO watching the ball disappear behind the wall.

The task for the watchdog is now to learn to anticipate, that depending on the speed, the ball will reappear on the other side or remain behind the wall. If the ball comes in with a low speed, the watchdog should learn to go looking for the ball on the right side of the wall, if the speed vector increases it should learn to go looking for the ball on the other end of the wall.

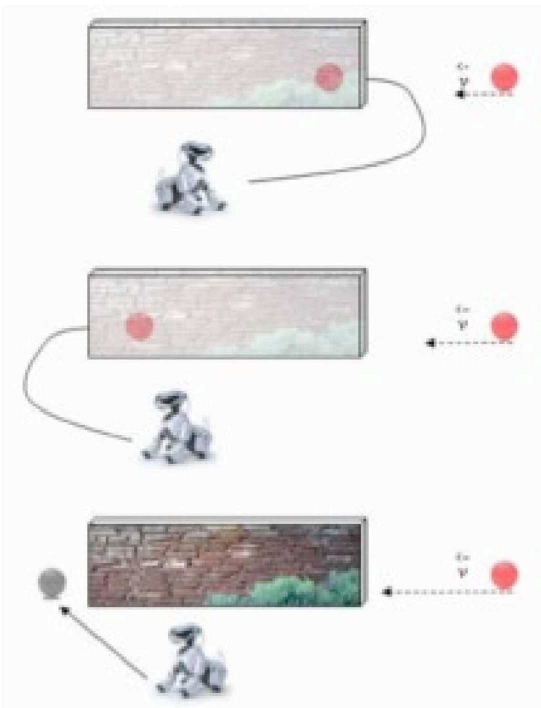

Figure 12 AIBO learns to anticipate, where the ball is "hiding"/reappearing and learns to go "looking for the ball".

In another setup, the wall might be blocked at the end, and a ball, coming in with a high speed, normally reappearing on the other side, now does not reappear, and there is a sound (the ball bumping against the wall) coming from behind the wall. The watchdog should then find out, that something changed and that it needs to go and search for the ball from the other side again.

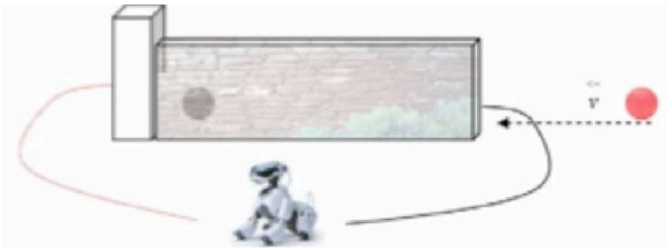

Figure 13 The wall is suddenly blocked at the end; AIBO needs to find out that the ball is still there, but can only be reached from the right side.

The next step could be a "hide and seek" scenario, where the robot has seen the ball disappear behind several "hiding places" (walls, obstacles) and then learns how to find the ball again and moving around, looking for the ball, anticipating it to be in one of the observed "hiding places".

References

Bargh J. and Chartrand T. (1999). The Unberable Automaticity of Being. *American Psychologist*, Vol. 54, No. 7, 462-479.

Brooks, R. (1991). New approaches to robotics. *Artificial Intelligence,* 47:139–159.

Gombrich, E. (1963). Meditations on a Hobby Horse and other essays on the theory of art. Phaidon, London.

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

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

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

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

Steels, L. and Brooks, R. (1994). The "Artificial life" Route to "Artificial Intelligence". Building Situated Embodied Agents. Lawrence Erlbaum Ass, New Haven.

Steels L. and Kaplan F. (2001). AIBO's first words. The social learning of language and meaning, in: H. Gouzoules (Ed.), *Evolution of Communications*, Vol. 4, No. 1, Benjamin. New York.

Steels L. (2003) Intelligence with representation. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, Vol. 361, p.2381-2395.

13 WP4: Goal Directed Behavior, Pro-activity and Analogy (UW-COGSCI)

After a rather high-level initial scenario on the watchdog and a rather low-level scenario on the moving object tracking and consequent activity, in this final scenario discussion session we propose several increasingly complex scenarios all under the general topic of search and retrieve.

13.1 Search and Retrieve Scenarios

The general scenario takes place in an assumably finite area (e.g. room), in which the real or simulated robot is located. The robot may be a simple robot arm combined with a camera or a mobile robot with a mounted gripper device. It may be assumed that the robot knows how to perform basic movements and manipulative actions. Many suggestions in the following proposal are inspired by Roy's (2004) experiences with their robot Ripley – in which nearly all features and suggested transitions between perceptual and internal representational modalities are hard-coded. One of the big challenges for the Mind RACES project should be to change the hard-coded capabilities of Ripley to adaptive, cognitive capabilities using and studying the advantages of anticipatory mechanisms.

Task 1: *Plain locate and collect scenario: Learn a mental model of the environment parsing the observed visual scenes into an internal, abstract mental model that can be used to detect and retrieve certain objects (first very few, potentially only one object such as the ball or train in the previous discussion session that may be occluded at times as proposed in the previous discussion cycle). Use the mental model to locate and retrieve certain objects efficiently.*

Although this first scenario may still be solved purely behavior-based, the suggested internal model should enable the system to achieve the task faster and more directed. Additionally, the internal representation will be increasingly relevant in the subsequent scenarios. Several challenges need to be faced in this still quite simple scenario:

First, objects need to be recognizable. Hereby, the feature of object persistence (particularly visually) may be used to identify objects and integrate them into the mental model. However, in the spirit of Gibson (1979), other properties of an object may be utilized such as its affordance, its characteristic shape, its occlusion properties, etc. These features need to be searched for by the chosen adaptive learning mechanisms in order to reliably build internal (predictive) object representations.

Second, objects need to be searched and found. Thus, the system needs to have a search algorithm that causes the robot to "look around". Dependent on the environment and the mobility of the robot, this may be accomplished by a simple camera movement. However, also anticipatory search mechanisms can be incorporated that should lead to an active exploration of occluded or unseen areas in the environment.

Third, an identified object should be collected. In this first, simple scenario, it would be interesting to have the robot collect all objects (that are graspable) and transport them to a certain location.

Initially, search and collection may be hard coded subroutines. A motivational module may be necessary that links detected objects to a decision making module that then triggers the collection routine.

Task 2: *Motivationally-guided scenario: Add a more elaborate motivational module that requires the retrieval of certain objects at certain points in time. Properties of objects need to be distinguished and linked with internal motivations of possessing or consuming certain objects.*

In this case, homeostatic variables need to be added to control current needs such as the energy level. Balances and resulting urgencies need to be accounted for in the motivational module. Moreover, each homeostatic variable may encounter a specific increase upon the retrieval of objects dependent on the objects' properties. For example, there may be a homeostatic variable for red objects. In this case, the need for the color red may be linked to red objects, which in turn would trigger the retrieval of a red object (if available and if there is nothing currently more urgent to do). The anticipatory agent should also account for the difficulty and expected time until retrieval of the current target object potentially projecting the levels of the other variables into the future (especially the one for energy). Essential for successful behaviour in the scenario is the grounding of an abstract (towards symbolic) object representation (Roy, submitted; Steels, 2003a) linking this representation with the proposed homeostatic variables. Certainly, it would be most desirable, if the abstract representations can evolve on their own.

Task 3: *Helper scenario: Act as a support system that searches for and retrieves requested items.*

In the third scenario, we propose the addition of a parsing system that is capable of processing requests, which may be actual linguistic utterances or simpler commands mediated for example via the keyboard. With respect to the previous scenario, such a request may be integrated into the motivational module in that, for example, a request always has highest priority until it is satisfied.

More challenging is the necessary linkage of object identities (or properties) with the actual request (that is, understanding the specifics of the request). It would be most appealing if the robot was able to learn to link specific requests with concurrently relevant objects. In a teacher scenario, a teacher may present one object at a time concurrently uttering its name. In a reinforcement-based scenario, requests may always be understood as retrieval requests. Initially then, the robot may start to retrieve all object candidates. Reinforcement learning may then lead to further distinctions of which object is referred to by which request. Even more important than in Task 2, symbol grounding is necessary to have candidate objects and object properties available (Roy, submitted; Steels, 2003a). Additionally, the symbolic representation needs to be associated with the 'request' input modality. The addition of mirror capabilities may enhance comprehensive capabilities estimating the most likely meaning of the request projecting own potential needs on the hypothesized speaker's needs (Arbib, 2002).

Task 4: *Simple social scenario (offensive or cooperative): Act with multiple other robot agents in the environment to achieve a certain task or also simply to continuously survive. In the simple* *scenario, the robot's capabilities may not be sufficient to distinguish between other individual robots restricting its beneficial interactive capabilities.*

In the simplest case, robots may continuously search and retrieve objects in the environment. Robots with anticipatory capabilities can be expected to retrieve items faster by (1) anticipating the retrieval before or after another robot and by (2) being more selective and directed in which objects to retrieve (dependent on the motivational module). Such a scenario may be interesting to study from an artificial life perspective in the same vein as the much simpler ECHO system (Hraber et al, 1997). The competition for the scarce resources should lead to the evolution of anticipatory agents, which provably have advantages over reactive agents (Davidsson, 1997, 2003). Also, simple forms of one-on-one trade may be imaginable, if different agents need different resources, as long as no immediate fraud is possible, potentially leading to symbiotic systems.

Task 5: *Complex social scenario: Interact with other robots recognizing other individuals consequently developing a form of trust. The task may be the same survival task as in the previous section.*

The capability of recognizing other individuals should lead to additional social, anticipatory capabilities such as the formation of trust worthiness of another individual and the consequent implications for encountering in trade or other forms of interaction. Note the relation of this and the previous scenario with the iterated prisoners dilemma – if the individuals cannot recognize each other, trust cannot be established and the interaction is prone to fraud. In the more complex scenario, a trust measure may be formed allowing beneficial interaction (for a simple general overview, see Ridley, 1996). The capability of recognizing and trading with other individuals may again also be investigated in an artificial life scenario in which individuals of Task 4 may compete with individuals of Task 5, potentially quantifying the additional cognitive capabilities of identifying other individuals. Moreover, the interaction with other known individuals may make the formation of mirror capabilities advantageous, potentially setting the stage for an emergent development of an artificial language (Arbib, 2002). This could be started up by the further development of playing simple "language games" (Steels, 2003b).

Some additional comments:

- Although the hope of the proposed scenario succession is that each scenario may be added onto the previous one and the capabilities are increasingly enhanced, clearly, the latter and particularly the social scenarios may be studied in isolation with action primitives for locating and classifying objects, simple and complex models of other agents, etc. This seems particularly necessary in the proposed artificial life setup.
- The latter two scenarios are meant rather as an outlook than as an immediate challenge (particularly the language parts). However, the references show that research is in progress, albeit the systems in these studies are interacting with a very simplified, abstracted world.
- Note the relation of Task 4 to the initial watchdog scenario. In this case, objects could be the doors and the intruded is modeled as an offensive other agent. The watchdog robot needs then the motivation to "possess" all doors, consequently disallowing other robots to reach them. Hereby,

doors have the properties that they cannot be "retrieved" so that the watchdog scenario may be seen as yet another challenge that may also be integrated into the proposed scenario succession.

References

Arbib, M. (2002). The mirror system, imitation, and the evolution of language. In Dautenhahn, K., Nehaniv, C.L. (Eds): *Imitation in animals and artifacts.* MIT Press, Cambridge, MA.

Davidsson, P. (1997). Learning by Linear Anticipation in Multi-Agent Systems, *Distributed Artificial Intelligence Meets Machine Learning*, LNAI Vol. 1221 (pp. 62-72) Springer Verlag.

Davidsson, P. (2003). A Framework for Preventive State Anticipation. In Butz, M. V., Sigaud, O., & Gérard, P. (Eds.), *Anticipatory Behavior in Adaptive Learning Systems: Foundations*, *Theories, and Systems* (pp. 151-166) Springer Verlag.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Erlbaum.

Hraber, P.T., T. Jones and S. Forrest (1997). The ecology of Echo. *Artificial Life* 3(3):165-190.

Roy, D. (submitted). Grounding language in the world: Signs, schemas, and meaning. *Artificial Intelligence.*

Ridley, Matt (1996). *The Origins of Virtue: Human Instincts and the Evolution of Cooperation*. Middlesex, England: Penguin Books.

Roy, D., Hsiao, K.-Y., Mavridis, N. (2004). Mental imagery for a conversational robot. *IEEE Transactions on Systems, Man, and Cybernetics*, Part B, 34: 1374-1383.

Steels, L. (2003a). Intelligence with representation. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 361(1811):2381-2395.

Steels, L. (2003b). Evolving grounded communication for robots. *Trends in Cognitive Science*. 7(7): 308-312.

14 WP4: Goal Directed Behavior, Pro-activity and Analogy (NBU)

14.1 Environment

We suggest using a ***common environment*** that can be employed by various scenarios. We think that this environment might be used also by other groups to reformulate their scenarios in the same environment (e.g. the watchdog scenario).

The proposed environment is an ***artificial city***. It consists of a network of streets and signalized crosses (i.e. with traffic lights). The streets are only marked on the floor of a room, i.e. there is only one flat top but the robots are not allowed to cross the marked lines depicting the streets. The streets allow only one robot to pass at a single moment of time. It is possible to complicate the environment by adding obstacles and screens. There might be landmarks which allow the robot to position itself on the map. The robot may have a built-in map of the city or may learn it from experience. The signalization is simple – it can be in one of 4 signals: “cross”, “do not cross”, “turn right”, “turn left”. This is the *allowed* direction of movement in this particular moment. The robot has its own goals which may not coincide with the allowed direction and in this case it will wait until the signalization changes. Again the robot will have either built-in knowledge about the signals or will learn their meaning from experience. In the following scenarios we assume that the robot already knows the map of the city and the signalization meanings.

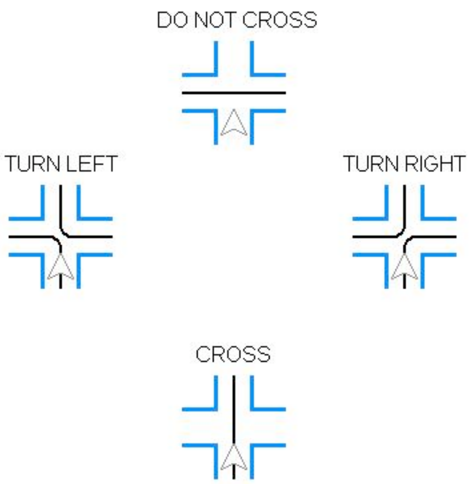

Figure 14 The Artificial City

14.2 Scenarios

Taxi (getting to the target in the fastest possible way)
Conditions:

- light signaling (e.g. four different levels of brightness: the lowest level means “stop” and the

highest level means “pass”)
- a robot having a cognitive map of the environment
- screens might be added outside the streets in order to prevent seeing some signaling from distance
- these screens might be dynamically changed in more complicated versions of the scenarios

Tasks:
- Positioning itself on the map
- Finding a way from the initial position to the target
- Getting to the target in the fastest way (minimizing the total waiting time at the crossroads)

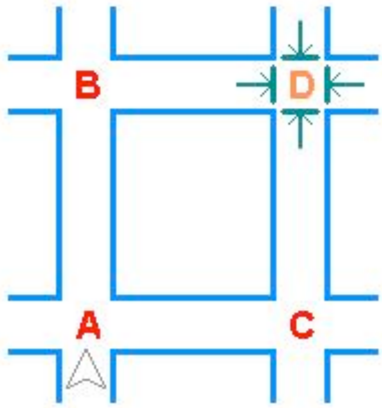

Figure 15 The Taxi scenario

Here is the simplest task: the robot has to move from A to D in the fastest possible way. The robot has to take into consideration the states of three traffic lights at the three cross roads – A, B, and C. There are three less or more complicated versions of this task:
- ***homogenous*** – the periodicity of the changes of the traffic lights is the same for all cross roads.
- ***nonhomogenous*, *but periodic*** – all traffic lights are different (have different sequencing of the signals), but the sequences are fixed for a given crossroad and do not change with time.
- ***nondeterministic*** – the sequencing of the signals are stochastic.

This task requires for the robot to anticipate the signal of the next traffic light (in the above case C and B while in A). In the case of remote goal it will have to anticipate many more states of the traffic lights along the roads.

Possible extensions of the scenarios:
- The traffic lights may break down and start signaling in a different sequence, or not changing the signal for too long time period, or the signal becoming unclear (unrecognizable). The question is whether the robot will be able to transfer the learnt knowledge to the new situation.

- At some point the signalization might be changed from traffic lights to traffic sounds (highly bright light is replaced by a sound with high pitch) – the robot should adapt to this new situation. The question is whether the robot can transfer the learnt knowledge to the new situation.
- We may introduce rotating bridges – they connect the two sides of the river and rotate over a certain period of time to connect the other two sides of the streets. Thus the robot will have to wait until it can get onto the bridge and then again until it can step down on the desired bank of the river. Suppose that the bridges rotate in a sequence equivalent to the traffic light signals. The robot has to predict the next bridge behaviuors in order to choose the fastest way.The question is whether the robot can easily adapt to the new task and use the learned sequencing, i.e. whether it can transfer the learnt knowledge to the new situation.
- We may introduce other robots navigating in the same city and then the robot will have to anticipate also their behavior in order to plan its own movements (since two robots cannot pass each other on a street).

Treasure-Hunter (Tomb Raider) – the task is to find a treasure which is hidden somewhere in the labyrinth of the city streets (in order to use the same environment) and which might be guarded by other robots.

Conditions:

- light signaling (e.g. four different levels of brightness: the lowest level means “stop” and the highest level means “pass”)
- none of the robots is allowed to violate the signal
- the treasure (or treasures) is located at the dead end of a street and is hidden – covered with a pile of cubes (the dead ends consist of such piles of cubes)
- we can have one or more treasure-hunters
- there is a treasure-hunter store (or stores) where the treasures should be collected
- there are one or more guardians of the treasures (if there are more than one guardians then they are staying in touch with and recognizing each other as players of the same team)
- there are no special signs to distinguish a guardian from a treasure-hunter
- the numbers of treasures is bigger than the number of guardians
- the guardians may put screens around the streets and dynamically change their position during the game (either transfering them or requiring the supervising humans to do so)

Tasks:

A single treasure-hunter with a simple tasks:

- to find all the piles that might potentially hide treasures
- to dig the treasure from the pile cubes (cube manipulation with arm(s))
- to move the treasure to the hunter’s store

One treasure-hunter and one guardian

- the treasure-hunter aims at finding all the piles that might potentially hide treasures
- the treasure-hunter aims at digging the treasures from the piles of cubes (cube manipulation with arm(s))

- the treasure-hunter aims at moving the treasure to the hunter's store
- the treasure-hunter aims at avoiding dangerous meetings with the guardian (see below)
- the guardian has no access to the hunter's store(s)
- the guardian aims at standing in the treasure-hunter's way to the treasure and thus stopping it from going there
- the guardian may put and move screens around the streets to obstruct the treasure-hunter's view
- the guardian may change the location of treasures inside the pile or the pile itself
- the guardian may block the treasure-hunter in a dead-end
- a blocked treasure-hunter leaves the game (is put in jail)
- the game continues until the treasure-hunter is arrested or until it expects no more treasures for grabbing

Several treasure-hunters and several guardians

- the treasure-hunters aim at finding all the treasures
- the treasure-hunters aim at digging the treasure from the piles of cubes (cube manipulation with arm(s))
- the treasure-hunters aim at storing the treasures at the hunters' stores
- the treasure-hunters aim at avoiding dangerous (see below) meetings with the guardians
- the guardians have no access to the stores
- the guardians aim at standing in the treasure-hunters' way to the treasure and thus stopping it from going there
- a guardian may put and move screens to obstruct the treasure-hunters' view
- a guardian may change the location of treasures inside the pile or the pile itself
- a guardian may block a treasure-hunter in a dead-end
- two guardians may block a treasure-hunter in any street section between two crossroads
- a blocked (arrested) treasure-hunter leaves the game (is put in jail)
- the game continues until all the treasure-hunters are arrested or until any of the free hunters expects no more treasure for grabbing

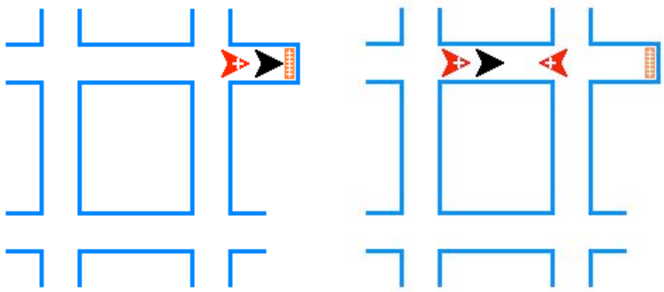

Figure 16 Two examples of "arresting"

These are examples of blocking (arresting). The treasure-hunters search, move to, and examine the piles. In first case a guardian has succeeded to block the offender in a dead-end. And in the second case two guardains have blocked the offender on a street section. The speed of both types of players is the same, thus anticipation and surprise are crucial.

Possible extensions of the scenarios:
All treasure-hunters form a coalition

- they warn each other for a guardian's presence or guardian's approaching; the warning is perceived by everybody
- they inform each other about the location of a treasure
- the treasure-hunters may spread deliberately to mislead or block the guardians

Treasure-hunters competition

- each treasure-hunter has its own store and they compete for the tresures
- if a treasure-hunter blocks another one into a dead-end section of a street it grabs its treasure but the robbed one survives

Dynamic formation and breaking of treasure-hunters coalitions

- negotiation between treasure-hunters
- to remain free is a common interest among the treasure hunters
- to become the richest is a competing individual interest among the treasure hunters
- dynamicly forming and breaking of treasure-hunters coalitions in correspondence with the above mentioned interests
- the treasure-hunters belonging to a coalition warn each other for a guardian's presence or guardian's approaching; the warning is perceived by everybody
- the treasure-hunters belonging to a coalition may spread deliberately to mislead the guardians
- each treasure-hunter has its own store
- if a treasure-hunter blocks another one in a dead-end it grabs its treasure but the robbed one survives

Presence of tourists

- the tourist just walk on the streets
- the tourists have no special sings to be distinguished
- the tourists and the guardians know a special password
- the tourists go close to the treasures but do not rob them
- the guardian may try to arrest a tourist but actually the guardian only loses some time to get the password from the tourist and leaves it free

Dangerous cubes

- some of the cubes (specially signed) explode when dropped down on the floor
- some of the cubes explode if put next to a triggering neighbour (both specially signed)
- if a cube explodes next to a robot (no mater guardian or a hunter) the robot leaves the game

Survival
Conditions:

- several cities divided by screens
- light or sonor signaling of the street crosses
- different signaling in each city (different color values; different sound values)
- plants bear fruits with a constant (for all the cities the same constant) periodicity
- several robots; by default each in a separate (its “native”) city
- by default each robot knows the map of its native city
- the robot’s main characterization is its vitality – a non-negative number (0 means dead)
- the vitality decreases linearly over time (breath expense); this expense is inevidable
- further more the vitality decreases linearly with moving (moving expense)
- the vitality may be increased by picking plants’ fruits (in a short interval between ripe and overripe)
- the vitality decreases with a constant on each non-allowed street cross (distress expense)
- all the plants of a city are insufficient (bearing in mind the ripe period) for a robot to survive
- the robot may move from one city to another and find plants there, however, it has to transfer its knowledge about signaling, plants, ripeness, etc. from its native environment to the new environment.

Tasks:
Each robot has the task to survive as long as it can. It may explore foreign cities in search for food. It may cross on “red light” if this will pay back (it grabs a fruit before the rival).

14.3 Mechanisms of anticipation involved

What are the mechanisms of anticipation needed to perform each of these scenarios?

Taxi
Getting to the target in the fastest way (minimizing the total waiting time at the crossroads)

- Learning the traffic lights sequence and anticipation of the next signal
- In the case of homogenous city the robot anticipates by analogy that the familiar sequence of one traffic light will transfer to another one; this can possibly be performed by simple generalization and transfer in a neural net
- In the nonhomogenous, but periodic case the robot may falsely anticipate the same behaviour of the next traffic light, but then may be able to transfer by analogy a higher-level structure (periodicity, the time periods, etc.). We suppose that this kind of transfer will be more difficult for simple NN techniques, but it might possibly be modelled by the AMBR mechanisms of analogy-making (Kokinov, 1994, 1998, 1999, Kokinov & Petrov, 2001, Kokinov & Grinberg, 2001).
- In the nondeterministic case the anticipation won’t help much, but still the robots are expected to try to predict the signalling by analogy with a recent traffic light, or with a previous episode with the same traffic light, here the behavior will be very context-sensitive and under different circumstances will use different old episodes as bases for analogy. Again, this might be possibly modelled by AMBR mechanisms, and it might be difficult for NN approaches. It will, however, require active perception for perceiving the relevant features of the environment to be used as context-cues.

Broken traffic lights

- Anticipating the signal or completing it by analogy with previous cases of the same traffic light or with previous cases of broken traffic lights. Again AMBR might be used as a basis for modeling.

Signal type change (e.g. from lights to sounds)

- anticipating the message by analogy – this is an interesting case of rather abstract analogy and again AMBR might be useful for modeling it, but would not be easy.

Rotating bridges

- analogy between the rotation of the bridges and the traffic light allowed direction
- anticipation of the behavior of the bridges by analogy with the learnt sequence of signals of the traffic lights

These analogies are again very abstract and remote and not easy, but we can try to model them by AMBR.

Streets crowded with robots

- anticipating the movements of the fellow-citizens by analogy with its own behavior or other robots already met before. This can be possibly modeled by NN or by AMBR.

Treasure-Hunter (Tomb Raider)
The guardians anticipate:

- the treasure-hunters to direct to the piles
- the treasure-hunters to avoid meeting them
- the treasure-hunters to direct to a store when they got the treasure
- a repetition of each specific treasure-hunter's behaviour (an analogy)

The treasure-hunters anticipate:

- presence of treasures in the piles in general (up to some moment)
- presence of treasures in some specific piles (analogy based on similarity with other treasure piles)
- analogous placement of treasures into different piles
- analogous behaviour of the cubes that have already been classified (some explode, some do not – see Dangerous cubes)
- misleading the guardians (coalitions, tourists)
- the guardians to chase them

Analogies can be widely used in this anticipatory behaviours: a robot may act by analogy with its own behaviour in a specific previous case (which it found analogous), a robot may predict the behaviour of another robot (hunter or guardian) by recognizing it and based on an its behaviour in the previous encounter, or finally, the robot may predict the behaviour of another robot based on analogy with another previously met robot and its behaviour. All these cases might be attempted to be modelled by the mechanisms of AMBR.

Survival
Each robot

- calculates its chances to survive after certain period of time (or to die)
- calculates its chances to increase its vitality after consuming a specific fruit
- anticipates that plants bear fruits periodically
- anticipates the presence or lack of fruits in a given city depending on the other robots' behaviour (whether they enter or leave that city)

Using analogies with the known cities each robot may

- calculate the length of ripe fruit period in a given city
- recognize the traffic signals in a given city (again they can vary from different traffic light signals to sound signals, use different periodicity, etc.)
- evaluate the food reserve of a given city based on other robots' behaviour

Again these analogies might be attempted to be modeled by AMBR.

We have emphasized the need for analogies in these scenarios, however, we believe that a variety of mechanisms will be needed for the successful modelling of the behaviour of the robots in this environment and we suggest trying to converge on this environment (or similar ones) and bringing together the approaches of all groups in achieving the desired behaviour. This environment may be initially emulated and various algorithms tested and then implement it with real robots and exploit their perception-action strengths.

References

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

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

Kokinov, B. (1999). Dynamics and Automaticity of Context: A Cognitive Modelling Approach. In: Bouquet, P., Serafini, L., Brezillon, P., Benerecetti, M., Castellani, F. (eds.) Modeling and Using Context. Lecture Notes in Computer Science (Lecture Notes in Artificial Intelligence), vol. 1688, Springer Verlag.

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

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

15 WP5: Emotion as Anticipation in Computational Architecture (ISTC-CNR)

SCENARIO 1: *A robot is moving around in an environment with dangers that when met cannot be fully avoided (say fire). This danger leaves signs in the zone around its location (say: smoke or smell) that when encountered by the robot elicits some internal motion (or appraisal) in its body (i.e. a feeling of 'fright'). The robot learns to anticipatorily detect the danger just by conditioning an avoidance behaviour not to the danger (Ev) but to its precursor sign (St).*

In this scenario there is no need of an explicit 'mental' representation of the future dangerous event. It is just a case where **emotions elicit an anticipatory or preparatory behaviour**.

A Stimulus St is exploited (thanks to selection or learning) as the precursor and the 'sign' of a following event Ev, and it is adaptive for the organism to react immediately to St with a behaviour which is in fact just the 'preparation' to the forthcoming event.

The advantage in this case is that the organism is 'ready', 'prepared to' the event. But this does not require a 'mental' anticipated explicit representation of Ev, that is the prediction, or better the 'expectation' that Ev will occur.

In this reactive process (St ➔ preparatory R), the emotional reaction plays a mediating role; that is the stimulus elicits an emotional response and the emotion activates an impulsive reaction that is 'preparatory' to a possible event.

St ➔ Emotional Response ➔ preparatory behavioural R

This seems to be for example the case of primitive forms of 'fear' or, better, '<u>fright</u>' in the scenario where not a prediction or a belief about the future but simply the stimulus itself (like a noise or a sudden movement) elicits the emotion.
In 'fright' in fact the reaction of the body (automatic retraction, reduction and hindering, stress, other?) looks 'conative' for escaping or for avoidance or appeasement towards a possible forthcoming danger or aggression.

We can ascribe this kind of emotion-based anticipatory behaviour to animals like rats or perhaps birds, but not to insects that only have merely reactive S➔R.

The problem of this rather simple model and scenario is that it is not at all clear what really the 'emotion' is in this case, what's its function, and why it shouldn't be simply skipped in a more simple mechanism like: St ➔ preparatory R.
Is the emotional 'mediation' just a trick, an empty and superfluous postulation?

Emotion is <u>the internal response of the body</u> to a given stimulus (or representation), felt by the control system.

Why should the body have an internal reaction (that could also be reduced to the elicited motor behavior) and especially why (for doing what) should we perceive, be informed about the reaction of our body to the events?
What does this add to the mere adaptive reactive behaviour (S ➔ R)?

First of all, the emotional response has a qualitative dimension, its experience is pleasant or unpleasant, and the feeling of this dimension represents (provides the organism with) an implicit 'evaluation' of the St (and of the Ev) as good or bad, as positive or negative for the organism. Such appraisal provides a sort of categorization of that kind of St/Ev for future uses that allows analogies and generalizations. For example next time along the same path the robot might remember the felt fear and avoid this area without any sign of danger (smoke); just on the basis of the associated and evocated emotional experience.

Second, the emotional response plays a role of reinforcement, reward, in learning process (in a non clear way). It seems that this provides (independently of the success) an internal measure of the importance of the rule reinforcing it: the stronger the emotional activation the more reinforced the rule and the greater the probability that it will be activated next time in similar circumstances.

Anyway, in order to model something less empty, and closer to a bodily 'motion' and 'felt' emotion, we should model a reactive variation of some internal bodily state (activating an external behavioral response) and then a signal of this reaction for memory, learning, etc.

SCENARIO 2: *The robot foreseen a given scene; this event is bad for it, is a threat, a danger. It feels 'fear' and changes its path (avoids the danger) or escapes away if the danger is moving and arriving. Later, while perceiving a possible danger or an unsafe zone or situation the robot, feeling a sense of anxiety, might multiply its investigating attitude and be more cautious but active for knowing about actual dangers or successes.*

This scenario focuses on emotional responses that are caused by anticipatory representations, by predictions.

In particular the robot feels its **bodily reactions to endogenous representations of future events**. The mental prediction of a given event elicits Hope, Fear, Anxiety, Worries.

The emotion is the response here and now to the anticipation of a future event. Bad events, threats, elicit unpleasant emotions, while expected positive events elicit pleasant ones.

The system should be able to:

- make predictions;
- evaluate them against goals;
- react to this cognitive appraisal with a bodily reaction ('motion');
- feel it;
- use this feeling for something;
- react to all this with an expressive or impulsive behavior.

The function of this faculty should be an associative, intuitive, fast, experience-based appraisal of future events and situations impacting on current motivation, maintaining commitment against difficulties and procrastination, discouraging us, or also activating other motives like reducing ignorance, acquiring additional information, being prepared.

A different relationship of the emotional response to the anticipatory representation is when the robot is **facing the confirmation of disconfirmation of its expectations**.
The fact that there was a given prediction (mental representation of the future) where the organism was interested and concerned (that was important for its goals) and the fact that this expectation is invalidated or realized, elicit specific 'affective' states.

Expectation (Prediction, possibly + Goal) & Ev ➔ Emotion ➔ …

Function: This is the area for the theory of 'Surprise' (non-anticipatory systems cannot be 'surprised'), 'Disappointment' and 'Relief'.
As we have said, not everybody would consider 'surprise' a really 'affective' state or an emotion. But more people would agree that intense Relief or Disappointment are emotions. Anyway, they should be modeled in Cognitive Systems dealing with the future.

The function of surprise seems to lie in the mobilization of resources for coping with 'abnormal' events (arousal), in particular processing/cognitive resources: attention.
It seems that it is important for learning and changing habitual assumptions and rules: after a 'surprise' one cannot continue in its routine behavior, must be aroused or careful.

Our hypothesis about the function of 'disappointment' is that it is learning not to be too optimistic in predictions, adjusting beliefs about predictability of the world, standards of sources reliability, self-confidence, etc.

SCENARIO 3: *The robot X avoids a path P1 (although perhaps this path is the shortest), because predicts – on the basis of the retrieval of previous experience and associated memory of felt emotions (see previous scenarios) – that it will feel fear, and does not want feel fear again (goal of not feeling fear).*
The robot prefers path P2 to P3 because it expects to feel pleasure and joy there, although the other path would be shorter.
Robot Y predicts X's emotion in a given case (for example if Y does action A) and decides to do A in order to produce that emotion (say fear) in Y (frightening-game).

In this scenario **the emotion is the object of the anticipatory representation**, not its effect (I predict to feel guilt, or regret, or joy, or embarrassment) and it is taken into account in current reasoning or decisions.

Function: The claim is that the ability to anticipate possible emotions affects current decisions in various ways, and in particular (since emotions are positive or negative, i.e. one searches for them or wants avoid them) changes preferences about foreseen scenarios.

References

Bechara A, Damasio H, Tranel D, Damasio AR. (1997) Deciding advantageously before knowing the advantageous strategy. *Science*, 275(5304), 1293-5

Camille, N., Coricelli, G., Sallet, J., Pradat, P., Duhamel, JR, and Sirigu, A. (2004), The involvement of the orbitofrontal cortex in the experience of regret, *Science*, 304 (5674), 1167-1170.

Frijda, N.H. The place of appraisal in emotions. *Cognition and Emotion*, (7), 357-387.
Lazarus, R.S. *Emotion and Adaptation*. New York: Oxford University Press, 1991.

Miceli, M., Castelfranchi, C. (2002)The Mind and the Future: The (Negative) Power of Expectations. *Theory and Psychology*, (12), 335-366.

Ledoux, J. (1996) *The emotional brain: the mysterious underpinnings of emotional life*, New York: Simon & Schuster.

Öhman, A. (1986) Face the beast and fear the face: animals and social fears as prototypes for evolutionary analysis of emotions. *Psycophisiology* (23). 123-145.

Ortony, A, Clore, G.L., Collins, A. *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press, 1988.

Snyder, C. R., Harris, C., Anderson, J. R., Holleran, S. A., Irving, L. M., Sigmon, S. T., Yoshinobu, L., Gibb, J., Langelle, C., & Harney, P. (1991). The will and the ways: Development and validation of an individual-differences measure of hope. *Journal of Personality and Social Psychology*, (60), 570-585.

Tooby, J., Cosmides, L. (1990) The past explains the present: emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology* (11), 375-424.

Zeelenberg, M., van Dijk, W. W., Manstead, A. S. R., & van der Pligt, J. (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition and Emotion*, (14), 521-541.

16 WP5: Emotion as Anticipation in Computational Architecture (IST)

16.1 The Zen of Anticipation

When designing an agent system, we generally provide it with the ability to search the space into which it is integrated and devise an optimum plan allowing it to reach its goals while minimizing a cost function. Under this point of view, achieving a goal is one if not the most important concern of the agent and, of course, anticipation and anticipatory affect do play an important part in such design, as several proposed scenarios will support. Our main concern, however, follows a complementary approach: that the journey towards achieving the agent goal is as important as achieving the goal itself. This "zen" approach is specially relevant when designing believable synthetic character systems.

Consider the following example: Lucia throws a red ball into the next room, then turns to Aibo, the dog, and says: "Fetch!". Aibo runs into the room and designs a plan to find the red ball. While searching the space, its attention is drawn to a small handkerchief which color is just as the ball it is searching for. With its ear pointing forward, Aibo starts running, waving its tail and barking in anticipation. However, as soon as Aibo realizes it is a mere handkerchief, its ears drop back and its tail falls between its legs. With a disappointed face, Aibo starts moving back, its gaze wandering across the room...

From the planning algorithm point of view, Aibo may have found itself in a local minimum, however, from the user point of view, much more had happened, orthogonally to the search plan. When designing a system in which believability is a key factor defining the *qualia* of the interaction, the path can become more important than the goal itself. Our scenario aims at providing an evaluation to this "zen" approach.

16.2 Context

Synthetic characters have proved to be an affective medium to enrich and enhance the interaction between the user and the machine, be it from the usability point of view, be it from the entertainment point of view. A critical yet subjective concept to account for when defining the quality of the machine-to-user side of the interaction with a synthetic character is *believability*. By believable character, we mean a digital being that "acts in character, and allows the suspension of disbelief of the viewer" (Bates, 1994).
Disney animators have been dealing with the creation of believable characters since the dawn of the last century, and have developed a set of guidelines to help in the creation of such believable characters (Thomas and Johnson, 1995). The general principle is to "display the internal state of the character to the viewer". This simple principle strives to make the character *aware* of its surrounding environment by consistently making the character react emotionally to what happens around it. In other words, even the characters are not "real", even the environment where they evolve is not "real", the relations between them are!

The concept of awareness can be further developed into what we call the *behavior loop*. Consider the following example. Aibo, the dog, is laying down near the fire when Lucia enters the room.

Aibo should respond by looking at Lucy (Aibo's attention focus is on Lucia). Furthermore, Aibo should clearly express an emotional reaction (perceived as caused by Lucia). The same principle should be applied to all intervening characters, including Lucia. This behavior loop increases the believability of characters (Martinho, 2003).
Our work researches which mechanisms are suited to control both the focus of attention and the emotional reactions of a synthetic character, to increase its believability through the behavior loop.
Furthermore, this work strives to make such control as autonomous as possible from the agent processing, in the attempt to extend the base agent architecture with a module designed to provide support for believability in synthetic character creation. And this is where anticipation enters.

16.3 Architecture and Anticipatory Independence

Our agents are implemented as software agents (Russel and Norvig, 1995). To make the control as independent as possible from the agent processing, we provide the agent with an autonomous module: the salience module (see figure 1).

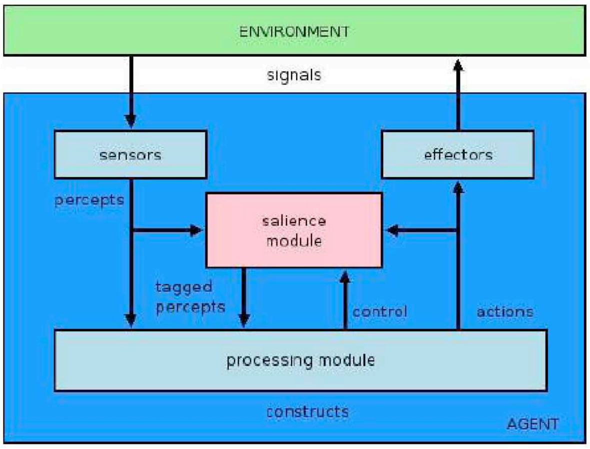

Figure 17 Extended Architecture

The salience module performs a semantic-independent monitoring of the percepts flowing from the sensors to the processing module as well as the action-commands flowing from the processing module to the agent effectors, in the attempt to capture “the feeling of what happens” (Damasio, 1999). This monitoring is possible since the code of the information flowing through the agent is usually consistent, in the sense that it is the repeated measurement of a specific internal or external aspect of the agent on a same scale over time.

The salience module is composed of several *emotivectors* (Martinho, 2004), each one associated with a sensed dimension. An emotivector is a module that keeps a limited record of a signal history and possesses mechanisms to anticipate the next expected value based on this history. By confronting the expectation with the sensed value, and using an anticipatory affective model based on the one described in (Martinho, 2005), the emotivector computes the sensor salience, and the percept is tagged with information providing both its attention focus potential as well as its emotional potential. The salience module also possesses a set of strategies to manage all the emotivectors together, that will also have to be assessed by the scenario.

The tagged percepts reaching the processing module of the agent are meant to be used as a guideline towards human-like behavior. Specially in the case of our scenario, they will be used as parameters for an autonomous mechanism controlling the synthetic character behavior loop.

16.4 Scenario

Our scenario takes place in a household environment where Aibo, the synthetic dog, "lives". As a starting scenario, we are aiming at an environment alike to a small warehouse, where several crates lie scattered around, acting as obstacles between Aibo and its searched target. Several distractors, will be added to difficult the task and provide with opportunities for Aibo to "play in character", following the same principles as the "Commedia d'ell Arte" improvisation directives.

We will strive to get a running simulation with a simple physics support that will demonstrate the impact of anticipatory affect in synthetic character design. Aibo the dog "digital body" will be modeled according to the real Sony Aibo specification, and will be as faithful as possible to its "real-life counterpart". This digital actor will be used to evaluate the possibilities of the real robot.
However, some factors, as the response speed of the robot effectors will be tweaked in the attempt to measure the adequacy of both the anticipatory affective approach as well as the Aibo Sony robot real-time support mechanism. Aibo's "digital mind" will integrate the different aspects of the "Anticipatory Continuum" and be a testbed of the benefits/disadvantages of such an approach.

References

(Bates, 1994) The Role of Emotions in Believable Agents.
(Damasio, 1999) The Feeling of What Happens.
(Martinho, 2003) Synthetic Emotension: Building the Behavior Loop.
(Martinho, 2004) Synthetic Emotivectors.
(Martinho, 2005) The Zen of Believability.
(Thomas and Johnson, 1995) The Illusion of Life.