Grant agreement for a large-scale integrating project (IP)

Annex I – Description of Work

Project acronym: *CogX*
Project full title: *Cognitive Systems that Self-Understand and Self-Extend*
Grant agreement No: *215181*

List of Beneficiaries

Beneficiary			Country	Date of entry	Date of exit
No.	**Name**	**Short name**			
1(coordinator)	University of Birmingham	BHAM	United Kingdom	1	50
2	Deutsches Forschungszentrum fur Kunstliche Intel-ligenz GmbH	DFKI	Germany	1	50
3	Kungliga Tekniska Hogskolan	KTH	Sweden	1	50
4	Univerza v Ljubljani	UL	Slovenia	1	50
5	Albert-Ludwigs-Universitat	ALU-FR	Germany	1	50
6	Technische Universitat Wien	TUW	Austria	1	50

Contents

Part A

1 Overall budget breakdown for the project

Grant agreement Preparation Forms

EUROPEAN COMMISSION

7th Framework Programme on Research, Technological Development and Demonstration

Collaborative Project

A3.2:
What it costs

Proposal number (1)	215181	Proposal acronym (2)	CogX

ONE FORM PER PROJECT

Participant number in this project	Organisation short name	Estimated eligible costs (whole duration of the project)				TOTAL A+B+C+D	Total receipts	Requested EC contribution
		RTD / Innovation (A)	Demonstration (B)	Management (C)	Other (D)			
1	BHAM	1,964,690.00	0.00	279,336.00	67,320.00	2,311,346.00	0.00	1,820,173.50
2	DFKI	2,114,475.00	0.00	34,848.00	48,279.00	2,197,602.00	265,906.00	1,668,983.00
3	KTH	1,357,387.00	0.00	25,400.00	60,160.00	1,442,947.00	0.00	1,097,250.00
4	UL	796,640.00	0.00	16,800.00	74,560.00	888,000.00	0.00	688,840.00
5	ALU-FR	866,016.00	0.00	2,000.00	16,000.00	884,016.00	0.00	667,512.00
6	TUW	1,062,880.00	0.00	25,468.80	34,560.00	1,122,908.80	0.00	857,189.00
TOTAL		**8,162,088.00**	**0.00**	**383,852.80**	**300,879.00**	**8,846,819.80**	**265,906.00**	**6,799,947.50**

Figure 1: Overall budget breakdown.

Comment

The partner DFKI plans to use resources made available by third parties (i.e. the states of Rhineland-Palatinate and the state Saarland). The resources include salaries of professors and researchers paid by the governments, as well as equipment, infrastructure and services paid by the governments. The total amount of such receipts will be 265.906 Euro and charged as percentage of own personal costs.

2 Project summary

Rationale: The goal of the Cognitive Systems programme call was to develop cognitive systems able to work in open ended, challenging environments, dealing with novelty, uncertainty and change. The starting point for CogX was the recognition that two key elements necessary for such cognitive systems have made progress in recent years. The first is that the range and power of available learning methods has advanced significantly, with specific advances based on these in certain domains, such as computer vision, or human augmented mapping. The second is that we have made progress in the engineering science of building integrated robotic systems that incorporate multiple modes of sensing and acting. Examples of the latter include a raft of robotic systems in the past four years that combine mapping, manipulation, vision, language, planning and learning. What is missing, and thus what we must build upon these advances, is a clean way of thinking about how agents such as robots should understand their own abilities and knowledge (self-understanding), and choose which of their abilities or knowledge to extend at any one time (self-extension). If we had a framework for self-understanding and self-extension for a robot or cognitive system then we would have taken a significant step along the road to achieving the aim of building systems that can handle the types of environments specified in the call.

Aims: The high level aim of the project is to develop a unified theory of self-understanding and self-extension with a convincing instantiation and implementation of this theory in a robot. It is important to note that the aims of the project are neither to produce a complete theoretical framework for self-understanding and self-extension that remains unimplemented in a robot, nor to merely create further specific algorithms for robot learning. Instead our goal is to build a bridge between these two, by creating a framework that is convincingly instantiated and studied in robot systems. This will require specific advances in the areas of mapping, language, and manipulation. But the overall way we will tie these advances together within a framework is as important.

Key Innovations: Our basic insight is that for any class of representations (e.g. logics, probabilistic models, dynamical systems) that a cognitive system might use it is possible to represent uncertainty or incompleteness in the agent’s knowledge that is encoded in those representations. There has been much effort in developing these, notably in work on epistemic logic or probabilistic learning and reasoning. What we will do is to develop methods for generating and reasoning about such representations for a robot with a number of specific sub-systems (language, spatial modelling and mobility, manipulation, vision), and to develop methods for combining the uncertain beliefs from each subsystem, so that the robot can for example reason about what information or knowledge gathering activities it should next engage in so as to achieve its task or learn about the world. This could be, for example, whether it should next learn about how to grasp the new cup in front of it, or explore the room next door to complete its map, or ask a person questions about the properties of the objects it can see. Only by being able to reason about the effects on its knowledge state can the robot plan what it should do to acquire new knowledge. We will both pull existing strands of work on representing beliefs and uncertainty in knowledge together, and make advances in each of the domain specific areas. These will be as important as the framework that connects them.

Benefits: The principal benefits of this project will be that it will be another step on the long road to cognitive systems that will inhabit our everyday world. Progress is slow but steady, and many of these technologies are now being exploited. With each advance a wider range of applications becomes possible. The primary impact of the project will be scientific, but we will make as much of our work as possible publicly available to encourage commercial exploitation.

Part B

1 Concepts, Progress, Methodology and Workplan

1.1 Concept and objectives

The challenge is to understand the principles according to which cognitive systems should be built if they are to handle situations unforeseen by their designers, other forms of novelty, and open-ended, challenging environments with uncertainty and change. Our aim is to meet this challenge by creating a theory — grounded and evaluated in robots — of how a cognitive system can model its own knowledge, use this to cope with uncertainty and novelty during task execution, extend its own abilities and knowledge, and extend its own understanding of those abilities. Imagine,

> a cognitive system that models not only the environment, but its own understanding of the environment and how this understanding changes under action. It identifies gaps in its own understanding and then plans how to fill those gaps so as to deal with novelty and uncertainty in task execution, gather information necessary to complete its tasks, and to extend its abilities and knowledge so as to perform future tasks more efficiently.

One way to characterise such a system's behaviour is to say that it is a system that has self-understanding. In this characterisation we use the terms 'understanding' and 'self-understanding' in a limited sense. By 'understanding' we mean the system's collection of models of the environment, and the way they are used by the system to achieve its tasks. These models could be of an enormous variety. In this project we will mainly study: models of action effects; maps; observation models; and type hierarchies (or networks) that organise information about objects, their properties and relations, and the actions that can be performed upon them. By 'self-understanding' we therefore refer both to models of these models of the environment and also to the ability to learn and reason about them. In other words to have self-understanding the system must have beliefs about beliefs and use these to model and reason about how its actions will change its beliefs. Such a system should be then capable of identifying, planning how to fill, and then filling gaps in its models of the environment; in other words, capable of self-extension. Note that in using the term self-extension we are not referring to systems that just learn, but to systems that can represent what they don't know, reason about what they can learn, how to act so as to learn it, execute those actions and then learn from the resulting experience. Our planned work is therefore based on the idea that self-understanding is necessary for self-extension (as we have defined the terms). We can summarise the aim of the project as to develop:

> a unified theory of self-understanding and self-extension with a convincing instantiation and implementation of this theory in a robot.

The technical challenges that arise from this aim are significant. Different types of information (e.g. stemming from various sensory modalities) require different kinds of representations, and thus representations of the accuracy or completeness of those representations (i.e. beliefs about beliefs) will also vary with the type of information being modelled. This means that devising a unifying framework for the representation of beliefs about beliefs will be challenging. In addition to this, methods for efficiently reasoning about beliefs, and planning in the kinds of belief spaces that we will consider, will require significant progress beyond the state of the art.

1.1.1 The role of robotic implementations in the research

We will evaluate our theory in robots. This is because they present a challenging domain where the benefits of a particular solution are clear. Our theories should, however, be applicable more generally. Specifically we will ground the evaluation of our theory in a single scenario centred around a mobile robot with manipulation abilities. The robot's environment will be an office or home, and it will be required to find, fetch, grasp and perform simple manipulation on a range of everyday objects. We assume that there are natural lighting conditions, and some degree of scene clutter. The robot will be able to interact with humans, with an emphasis on dialogue. This communication will be important in that the robot will have to reason about the beliefs of other agents, communicate its own ignorance, establish common understanding with them, and use dialogue to help achieve goals and extend its understanding. Achieving natural communication with humans is therefore central to the scenario.

The tasks in our scenario can be posed with varying degrees of difficulty. We will pose them in forms where uncertainty, novelty or incomplete knowledge is present. In the later stages of the project we will handle all of these together. This will include handling incomplete maps of the building, and underspecified descriptions of objects or locations. In manipulation, mobility, dialogue and sensing the robot will have to handle uncertainty in action and observation, and deal with interruptions to task execution. It is important that the aim of the project is not to produce a working system for a narrow domain (although we will do this), but to produce a theory able in its most general form to support systems for a wide range of domains. It is, however, important to note that because we have chosen a particular robotic domain it will heavily influence many of the specific representations and algorithms we use and devise. This is where much of the technical challenge lies. Reasoning, learning and planning under uncertainty, and reasoning about beliefs have been tackled before, but dealing with them in the context of dialogue, vision, manipulation and mobility in a unified manner is new.

1.1.2 An example task in our scenario

A specific, if very simple example of the kind of task that we will tackle is a domestic robot assistant or gopher[1] that is asked by a human to: “Please bring me the box of cornflakes.” There are many kinds of knowledge gaps that could be present (we will not tackle all of these):

- What this particular box looks like.
- Which room this particular item is currently in.
- What cereal boxes look like in general.
- Where cereal boxes are typically to be found within a house.
- How to grasp this particular packet.
- How to grasp cereal packets in general.
- What the cornflakes box is to be used for by the human.

The robot will have to fill the knowledge gaps necessary to complete the task, but this also offers opportunities for learning. To self-extend, the robot must identify and exploit these opportunities. We will allow this learning to be curiosity driven. This provides us, within the confines of our scenario, with the ability to study mechanisms able to generate a spectrum of behaviours from purely task driven information gathering to purely curiosity driven learning. To be flexible the robot must be able to do both. It must also know how to trade-off efficient execution of the current task — find out where the box is and get it — against curiosity driven learning of what might be useful in future — find out where you can usually find cereal boxes, or spend time when you find it performing grasps and pushes on it to see how it behaves. One extreme of the spectrum we can characterise as a task focused robot assistant, the other as a kind of curious robotic child scientist that tarries while performing its assigned task in order to make discoveries and experiments. One of our objectives is to show how to embed both these characteristics in the same system, and how architectural mechanisms can allow an operator — or perhaps a higher order system in the robot — to alter their relative priority, and thus the behaviour of the robot.

[1]Gopher as in to “go for”.

1.1.3 Summary of objectives

Our overall aim, as stated previously is to develop a unified theory of self-understanding and self-extension. We have broken down this aim into the following measurable objectives:

1. A unified framework for representing beliefs about representations of action effects, observation models, incomplete information and categorical knowledge. [WPs 1,4,5]
2. Specific representations of beliefs about beliefs for the specific cases of dialogue, manipulation, maps, mobility and some types of vision. [WPs 2,3,6]
3. Representations of how actions will alter the belief state of the cognitive system, and those of other agents, as represented in the first two objectives, i.e. models of the effects of actions on beliefs about space, categorical knowledge, action effects, dialogue moves etc. [WPs 1,2,3,4,5,6]
4. A theory of how to reason, plan, act and interact using such representations of beliefs, and beliefs about beliefs, to achieve a task in the face of incomplete information, uncertainty and novelty. [WP 4]
5. A theory of how to use these representations to identify learning opportunities, plan and execute plans in order to learn so as to perform future tasks more effectively and efficiently. [WPs 4,5]
6. Methods for perception and manipulation of objects that enable a robot to actively explore objects, to extend its manipulative skills, and its understanding of these. [WP 2]
7. Methods for perception and spatial modelling that enable a robot to identify gaps in its spatial models (e.g. maps) and to extend them so as to support natural communication with humans. [WP 3]
8. New representations and algorithms to allow a robot to extend its categorical knowledge by identifying gaps and learning the relationships between different modalities (e.g. vision and language). [WP 5]
9. Methods that enable a robot to represent and reason about its beliefs and those of other agents to support natural dialogue and to extend its own abilities and understanding. [WP 6]
10. A theory of how a cognitive system can trade-off task driven and curiosity driven activity. [WP 1]
11. A robotic implementation of our theory able to complete a task involving mobility, interaction and manipulation, in the face of novelty, uncertainty, partial task specification, and incomplete knowledge. [WPs 2,3,6,7]
12. Within the same implementation the demonstration of the ability to plan and carry out both task driven and curiosity driven learning goals. [WP 1,7]

1.1.4 The example revisited: explore, explain, extend

To illustrate how we will address these objectives we return to our example based on the instruction, “Please bring me the box of cornflakes.” Assume that the robot only has to deal with the second knowledge gap listed in Section 1.1.2: ignorance of the box’s location. We assume it has prior knowledge of: some categories of objects; physical, dialogue and information processing actions; and locations. Different versions of the scenario can be generated by varying its prior knowledge. All our work will assume some innate knowledge: our robots will not learn tabula rasa. In our framework, in order to complete the task while engaging in curiosity-driven learning, the cognitive system must engage in three broad types of activity, which it may cycle through several times. We refer to these as exploration, explanation and extension. These are merely the main stages, and there are several other activities in the cycle as shown in

Figure 2. Initially the robot must plan how to achieve the task. If it doesn't know where the cornflakes box is, it must reflect in order to realise it must fill that gap, and perhaps identify general-purpose knowledge it can use to do so. If it has a category system, for example, it could use information about the super-type of cornflakes boxes (cereal boxes) — e.g. that cereal boxes are often found in kitchens or dining rooms. Following this, in the planning stage the robot will need to plan how to find out which location is correct, and then retrieve the box. If there are many courses of action to achieve this it must assess their relative worth. Asking a question, for example, may be quicker than an exhaustive search for the box, but if the robot needs to find someone to ask a question of, then it may be slower. To devise a truly efficient plan the robot should consider such trade-offs quantitatively. Thus planning will require representations of action effects on both the physical and belief states of the system, and represent the non-deterministic effects on each of these in both a qualitative and probabilistic manner.

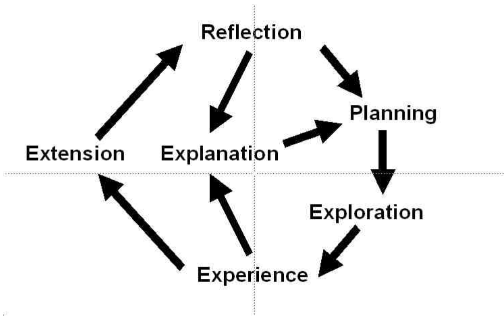

Figure 2: The structure of the explore, explain, extend cycle.

In executing the plan the robot will have to continually monitor execution and re-plan. We call this the exploration stage, in which it acts in the world, gathering information as it goes. If the plan unfolds in a benign way, with little true novelty encountered then the task will be accomplished. Perhaps the robot asks the human, "Is the cornflakes box in the dining room or the kitchen?"; is told it is probably in the kitchen; goes to the kitchen; sees the box on the sideboard after some visual search; grasps and returns the box to the person. This benign execution is extremely unlikely, even for such a simple task. There are many things that can go wrong: recognition of the box will be unreliable; the box may be in the kitchen cupboard rather than conveniently placed on the table; the robot may try to grasp the box and have it slip from its gripper.

1.1.5 The importance of surprises

These unexpected outcomes need to be dealt with. The first part of this process is the explanation stage in our schema. Each outcome has a number of possible explanations, and the robot has to choose which one to act on. The resulting action might assume the correctness of the explanation, or constitute an attempt to confirm it. For example if the robot has a model that tells it the perceptual routine it ran was unreliable, then it may explain away the unseen box as perceptual failure. In this case it should run a different algorithm on the image, or look at the scene from a new viewpoint in an attempt to find the box. Alternatively it may explain the unseen box as being due to the fact that the box is not in the kitchen after all, but in another room. Finally it may use reflection to hypothesise that the cornflakes box is in a cupboard, and that this explains its visual absence. These explanations must be generated using models of the world, and representations of how reliable those models are. Thus explanation is a process of reflection to explain the experience arising from exploration. There are several deep technical challenges here. If action outcomes are uncertain how can it distinguish between unreliability that cannot be improved upon, and a genuinely unusual outcome that can be explained in such a way that in future it will be more predictable? If for example, the robot grasps the box, which then slips from its fingers. Is this due to poor localisation of the object or because the packet is heavier than supposed? In the first case perhaps it should move to get a new view and try again. In the latter case it might hypothesise that it should apply a greater grasping force. How to distinguish these cases, let alone how to generate and then reason about possible explanations, is a truly difficult set of problems.

1.1.6 The importance of self-understanding in learning

Once the robot has generated hypotheses it must produce and execute a plan that allows it to test them. It may need to ask a human, if present, whether the cornflakes are in the cupboard. Or it may have to execute a new grasp. This leads the robot through another phase of planning, exploration and experience. Finally if the hypothesis is correct the robot has to evaluate whether a learning opportunity has arisen. Having been told that the cornflakes are in fact in the cupboard it may spot a learning opportunity by hypothesising that cereal packets in general are found in cupboards. This will require the ability to spot the potential knowledge gap, hypothesise a way of filling it, plan and execute a test of the hypothesis, and adjust its knowledge. This kind of reflective, hypothesis-driven learning activity is the final phase in our schema, the extension phase. Being able to represent and reason about learning opportunities in robotic domains will be the most challenging of all the objectives we have set ourselves. This is where we will encounter the problem of choosing between many explanations or experiments, and between curiosity driven learning and task driven activity.

The task of obtaining a box of cornflakes is rather mundane. However, in unpacking it we have shown both its underlying richness, and the way we will frame our work. We will use the framework to create a unified testable theory of self-extension and self-understanding. When tested on a number of scenarios of the sort we have described this will produce what we believe will be a convincing demonstration of its worth.

1.2 Progress beyond the state-of-the-art

In the previous section we outlined a focussed, but ambitious set of objectives. In this section we will describe how achieving these will require us to move beyond the state of the art in a variety of areas. In the past four years there has been significant progress in several fields related to the planned work. There are however important restrictions on the abilities of integrated systems that make use of these advances. In this section we describe briefly the contributions we will move beyond the state of the art in a number of areas and themes that cut across the project.

1.2.1 Project baseline

Five of the six participants, and a majority of the PIs have participated in the CoSy project (FP6-IST-004250-IP). It is thus important to set out the scientific baseline on which we are building. During the past four years, we and others, have made significant progress in a number of areas. In computer vision part based recognition and categorisation of objects has improved; the abilities of human augmented mapping systems have increased; we have created architectures (e.g. CAS) suitable for integration of multiple modalities of sensing and action; we have implemented the first simple cross modal learning systems in this framework; and in planning we have made some advances in posing dialogue planning as a multi-agent planning problem. The consortium, also has experience of working together, building integrated robotic systems, both for manipulation and mobility. We have not yet combined these, however, and we still have very simple visual abilities when it comes to recovering object shape, or predicting object behaviour.

1.2.2 Relation to behaviour based robotics

Our approach to learning and extension is clearly heavily representational, and in this sense is rather different from behaviour-based approaches to robotics. Behaviour based robotics typically eschews representations or models of the world state, including hidden state, relying on the current sensory signal to determine. It also divides the controller according to task. In this second aspect there are similar problems to ones that we will face, since in our architectural approach we have modality and motor specific modules. The control outputs of these may need to be synchronised, giving either parallel action, or sequenced action. This need to integrate behavioural output is dealt with in behaviour-based systems in a variety of ways: by using subsumption, linear combination, or voting mechanisms. We will plan action, but utilise task achieving behaviours to achieve these actions where appropriate. Thus while our philosophy is quite different to that of behaviour based systems we have a pragmatic attitude, and recognise the need for short feedback loops at lower levels in the system. Where we will thus make contributions is in trying to adapt insights from behaviour based systems about how selected behaviours are integrated and coordinated during their expression.

1.2.3 Self-reflection and meta-cognition for robots

We will develop representations for knowledge producing actions within the CAS architectural schema. This will mean providing linked Bayesian and logical representations of the belief state both within modalities, and globally. We will represent the quantitative and qualitative effects of information gathering actions (such as sensing, dialogue or information processing) on the belief state. This will be challenging given that we will have very different forms of knowledge gaps and uncertainty in different sub-domains such as manipulation, perception, dialogue, spatial reasoning and categorical knowledge. We will use these representations in WP4 to both plan courses of action and introspect on them to provide explanations and hypotheses. Finally we will implement this theory for our scenario in a robot.

1.2.4 Self-motivation

We will create methods that allow multiple, modality specific sub-systems (vision, language, manipulation) to throw up opportunities for action asynchronously and concurrently. We will create representations of intrinsic motives, and methods to alter relative priority. Using these we will create methods to generate specific desires and intentions from the intrinsic motives, and the internal and external state of the robot. A key task is to clearly explicate the different steps in this process with reference to a system architecture. We have done this for human motive processing [4, 41], but the implementation was limited to a simulated domain [40].

1.2.5 Cognitive robot architectures for cross-modal processing

We will use explicit models of uncertainty in knowledge, and how beliefs change under action to develop new algorithms for cross-modal processing. We will model incompleteness and uncertainty at a variety of levels of abstraction. We will use these to develop algorithms able to select information gathering actions (such as clarification requests); to create more sophisticated models of cross-modal attention; and to perform cross-modal inference in the face of uncertain and missing information. We will use the CAS architectural schema [18] to support all this work. Currently CAS performs mediation based on categorical understanding [3, 14, 27]. We will also develop methods to deal with cases where the categorical information itself is uncertain or incomplete. In summary we will explain how information can be fused and processed across modalities more effectively by using algorithms that employ models of uncertainty in knowledge. The result of this work will be robots that more rapidly, efficiently and effectively build up a coherent picture of the world in the face of uncertain information.

1.2.6 Acquiring models for object manipulation

In the planned work we will go beyond the state of the art, both with respect to modular mo-tor learning and robotic manipulation where we will tackle one of the unsolved problems in robot grasping, namely the generation of feasible and stable grasps for un-modelled objects in unstructured scenarios using two- and three-fingered robot hands. Most robotic systems today suffer from insufficient perception and do not employ suitable methods to represent the extracted knowledge that also allows for learning. One of the most important contributions of this project is therefore the intended interplay between perception and manipulation. We will thus create visual models of object shape which can be refined by exploring the shape. Visual input will also be used to execute some of the early grasping hypotheses regarding both two- and three-fingered grasps. Once a successful grasp is obtained, controlled movement of the object will be performed in order to extract additional information about the object structure. With respect to modular motor learning for robot manipulation we will create and incorporate a visual channel into the MOSAIC model [39, 17] that explicitly models the object shape, and the changing contacts under simple manipulation, as well as controlling exploration of the object to learn, and adding new contexts to the model as required.

1.2.7 Life-long acquisition of conceptual knowledge

Extended learning on the basis of physical and social interaction will make use of mechanisms from all parts of the project, and will also contribute to most of them by helping to drive their learning. Moreover, the integration of all these mechanisms must be supported by an adequate architecture with flexible self-extension and self-reflection capabilities (Section 1.2.3). This architecture will be based on our earlier work on CAS (cf. Section 1.2.5 and [18]). Attempting to develop new forms of learning based on self-understanding in combination with other competences will provide a demanding challenge for the architecture and tools.

1.2.8 Situated dialogue processing for human-robot interaction

The project will provide an approach to establishing common ground in understanding the environment, sometimes in a learning context, as discussed in Section 1.2.7. This will, like most of the project, depend on integrating diverse competences in a common architecture (e.g. since visual perception, and planning and other tasks, will provide some of the context required for language understanding). It will also make use of the robot's self-understanding capabilities insofar as it can tell which aspects of the context it is using and detect that there is a communication problem because it lacks some context, which can lead to a decision to acquire some more information to help with understanding an utterance, e.g. by looking somewhere, by manipulating something, or by asking for clarification or help. This will need novel mechanisms for processing dialogue about learning in dynamic environments (WP6 in connection with WPs 1,2,3,5), and mechanisms supporting interaction between planning, language, and spatial knowledge (WP6 in connection with WPs 3,4,5). The work will also contribute to continual learning by devising novel methods for acquiring mappings between non-linguistic knowledge and linguistic meaning: a situated model of language acquisition [5].

1.2.9 Continual planning of situated action

Our work will build on decades of research on planning and decision making under uncertainty, including continual planning with interleaved planning and execution, planning with stochastic and non-deterministic actions, planning with partial state knowledge and active learning. In our case it is significant that the planner will need to respond in real time, so we will combine both decision-theoretic and symbolic planners using bounded rationality arguments [33] to decide when each is appropriate. The real-time requirement is a serious challenge as many approaches have exponential worst cases. Our approach is to give up optimality and produce plans that are 'good enough'. We will explore approximation techniques, the use of abstraction, and the continual planning approach of planning a small way ahead, and then executing the partial plan to remove uncertainty.

The combination of continual and decision-theoretic planning by switching between them depending on the situation will be a novel contribution of the project. Previous mixed-mode planners include robotic architectures such as 3T [6], where a high-level symbolic planner sends actions to a lower-level sequencer. In our domain this will not work because uncertainty can occur at the highest abstraction levels, and ignoring it can lead to inefficient or unusable plans. For example in the 'fetch the cornflakes' task a continual planner can plan to ask about the location, and then continue planning after asking, but decision-theoretic arguments are needed to reason about the possible outcomes of asking before the action has been executed, for example to decide where to go to find someone to ask. Deciding when to switch planning mode on the basis of self-knowledge and knowledge about the situation is a requirement for successful performance in our robot, and is presently well beyond the state of the art.

1.2.10 Qualitative models in spatial cognition

We will move beyond the state-of-the-art in several ways. Previous work using only artificial objects [13] or a very limited set of natural objects [11, 25] in working memory will be extended by incorporating more advanced strategies for visual search and more principled ways of combining conceptual and spatial models; and also investigating in more detail relations between and the use of short-term and long-term memory. We will model not just the spatial environment, but also the gaps and uncertainties in the robot's model of it, and the ways that its exploratory actions will generate new information. The key is to build up a picture not just of where objects and places are, but what they are, and what can be achieved with them. This means generating what we call a functional understanding of space.

1.3 S&T methodology and associated work plan

1.3.1 Overall Strategy and General Description

The shape of the workplan is summarised in Figure 3. There are three domain specific packages, related to object perception and manipulation (WP2), spatial cognition (WP3), and dialogue (WP6). These will concentrate on developing new representations and algorithms for self-understanding and self-extension that are essentially domain specific. There are then three workpackages that cut across the domains. These are WP1 on motivation and reflection, WP4 on planning, and WP5 on cross-modal learning. In each of these we will have to find and use representations and algorithms that link or unify the domain specific representations used in WPs 2, 3 and 6. Since representations are the key in our view to these connections we have devised some tasks in some workpackages that run for all 50 months. These are tasks that work on how to represent the gaps (incompleteness or uncertainty) in particular kinds of knowledge. These project long tasks for representation will link WP5 (gaps in cross-modal and categorical knowledge), WP2 (gaps in knowledge of objects and manipulation) and WP3 (gaps in spatial knowledge) with a task that will bring their results together in WP1 (representations of beliefs and models of belief change under action). There will thus be regular feedback between these tasks.

The integration for the project will be driven by the robot scenario for which a simple example task was described in Section 1.1. Each year we will aim at integration to achieve another stage of the explore, explain, extend schema. The first two years will be dominated by the explore stage, with the second year focusing on relaxing the environmental assumptions so as to introduce rather more uncertainty than in year 1. These integrated experimental platforms define the project wide milestones, measuring the overall project progress:

- Milestone 2 (Month 15): Task driven exploration & learning.
- Milestone 4 (Month 27): Task driven exploration under uncertainty.
- Milestone 6 (Month 39): Explanation with limited extension.
- Milestone 8 (Month 50): Full curiosity driven extension.

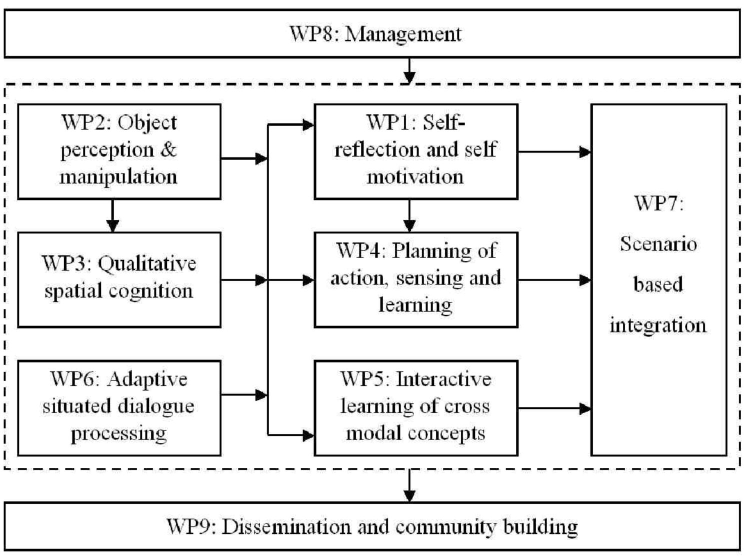

Figure 3: Work package relationships.

The whole project is constructed in roughly six month cycles — although the first cycle is longer because of the ramp up period, and the last is slightly shorter because of the wind-up period (months 49-50). Progress on theories and mechanisms that make the parts of our theory (WPs 1-6) will be summarised half way through each year, and the results reviewed by the General Assembly. At this point the precise tasks for integrated systems that form experimental platforms will be decided upon, and integration will proceed in parallel (WP7) with more theoretical work. This integration phase will last for six months leading up to the annual project review. In the six months following the review the integration package will be used to experimentally analyse the integrated systems developed. Therefore, the workplan strategy is to have regular information exchange coupled with well defined periods for integration. By this approach we aim to achieve both a systems-level and a component-level understanding of how self-understanding and self-extension should work. We now detail the objectives and activity within each of the work-packages, including those related to management, and dissemination and community building. A detailed view of the timings and dependencies between the work package tasks can be seen in the Gantt chart in Figure 4 (page 20).

The Workpackage summaries give detail on the tasks into which the workpackages are currently decomposed, together with their time, and our current estimates of the total and partner effort for each task. The time estimates in particular are necessarily very loose, and are not to be read as binding. Typically we might normally expect the partner effort within a task to vary from the specified figure by anything up to double or down to a half of the number of person months listed. The total effort per task should deviate by less, however, all these figures are intended to be indicative, rather than binding. We will modify the effort distribution and allocate new tasks if necessary.

1.3.2 Scenario-based integration

It is absolutely central to our aim and method that we study self-understanding and self-extension in integrated robotic systems. We need to evaluate our theory, not just in terms of its components, but at a system level. This is where the most difficult challenges will arise. The other important aspect of integration is the degree of open-endedness, uncertainty and variety with which we pose these tasks. This means systematically defining for each integrated system the degree of situational variation with which the robot should be able to cope. We will develop a series of four integrated systems, one every 12 months, and use these as experimental platforms to empirically study the robot's complete behaviour as we change aspects of either its design, or the environment. These

will be based around our scenario of a home/office robot with mobile manipulation. We have said that we will aim for a single robot capable, by the end of the project, of two types of behaviour: task focused, and curiosity driven. We refer to these behaviours as the robot assistant/gopher and the curious child-robot scientist.

We set out one detailed example from our scenario in the first section. It is very important to note, however, that integration work in the project will be driven by the overall scenario and not by this specific example task. The following are tasks that may be set either by a human, or by the robot itself:

- Fetching a particular item, e.g. "get me my coffee mug".
- Fetching items specified at the category level, e.g. "get me a mug from the kitchen".
- Fetching items specified by function e.g."get me something to drink from".
- Searching for people and giving instructions or messages to them.
- Finding out information about where people or objects are.
- Learning the association between visual features of an object and its linguistically expressed properties.
- Learning the causes of the dynamic behaviour of an object under pushing activity,

and generalising this to predict the behaviour of other objects based on their shape.

- Learning dialogue strategies for efficient information gathering.
- Learning maps of space that contain explicit representations of objects and their relations.
- Learning grasps for objects, that exhibit a degree of generalisation.
- Learning the qualitatively different poses that an object can be in relative to another.

We can characterise each of the systems we will develop as moving us further along the explore, explain, extend process. At the same time we will gradually increase the degree of environmental uncertainty (situational variation with which the robot can cope). We characterise the scenarios that we will work towards below. As part of that description we give three illustrative tasks within our scenario:

- ”Please bring me a box of cornflakes.”
- ”Please fetch me my mug.”
- ”Find Aaron and tell him we are meeting in room 225 now.”

It is extremely important to note that we give these illustrations in the spirit of being as informative and detailed about our aims as possible. These are not necessarily the precise instantiations of our scenario that we will test and demonstrate. It is not possible to predict the state of the art in five, or even two years to the level of detail required to guarantee that the precise descriptions we give will be the systems we demonstrate. Thus we reserve the right to alter the demonstrated tasks within the scenarios (using slightly different objects, lighting conditions, or vocabularies etc) while seeking to satisfy a majority of the generic targets for any given year.

Task driven exploration & learning (Month 15): The first integrated system will concentrate on the exploration stage of our schema in a fairly controlled setting. The robot will be mostly task driven, and will be able to fill simple knowledge gaps (such as in which room an object is) using dialogue, fixed visual search routines, or by learning. These gaps will be defined, and mostly driven, by a human. It will represent non-determinism qualitatively rather than in a quantitative manner. The robot will fail semi-gracefully in that it will try to return to the human and state when it has failed, but not explain why. The lighting will be controlled and there will be no scene clutter, objects to grasp will be quite simple (e.g. boxes of varying dimensions), and humans will be present and helpful. Reflection will be limited to identifying the knowledge gaps that need to be filled to complete the task, and how to partially fill them using categorical knowledge. Vision will be limited to instance recognition, and modelling of simple shapes. Finally the robot will be able to verbalise when it fails to understand a situation or command, and to initiate clarification dialogues.

Illustration: When asked for a box of cornflakes the robot has a trained identification model for a specific box of cornflakes. It uses this together with a map of the house to search for the cornflakes. It will have a likely set of places the cornflakes could be, and a model of human belief that assumes that humans are essentially omniscient with respect to the beliefs it can represent and reason about. It will have both a map of the domestic space that explicitly represents objects and the places they are in, and when it finds objects it can place them in its map. Thus the robot can ask where the cornflakes are, or search for them visually. In the first year we will use controlled lighting conditions. There will always be a human present in each room. The robot will know how to grasp the box, but if it fails to grasp an object it will be able to ask the human to help it. The object will sit on a visually un-textured surface, with no object clutter, in a reachable position. We will vary the location of the object within and between the rooms, but there will be no variation in the object to be grasped.

Task driven exploration under uncertainty (Month 27): In the second year we will incorporate our switching planner to allow the robot to reason quantitatively about trade-offs between actions. The robot will perform in the same task-driven manner as the first year system, but will have to cope with a greater degree of uncertainty; limited variation in target objects and lighting conditions; and a little scene clutter. The robot will be able to plan algorithmic actions at the qualitative level (e.g. knowing it needs to run a recognition algorithm to find an object). It will incorporate the more sophisticated models of space developed in WP3 by now, and plan active visual search or simple manipulation to refine the segmentation of an object against textured background. The grasping strategies will vary for a small range of objects. The robot should also now be opportunistic with respect to fulfilling its task, so that if the target appears in an unexpected place it will exploit this serendipity. In dialogue the robot will be able to learn dialogue strategies that are efficient from its experience. Finally the robot will be able to perform some cross-modal learning.

Illustration: The robot is asked for the mug of the person it is interacting with it. It has several stored persons with corresponding stored identification models for their mugs. It has a model that suggests that if the owner's mug is not found a generic coffee mug will do. It may encounter the specific (distinctive) mug and other mugs, as well as other objects. Lighting conditions may vary slightly between the rooms. It may not find a person in every room, but if a person is present they will helpfully identify themselves. In this case the grasp model will be for a number of specific objects, but for simple objects (e.g. boxes) there will be some variation between them. Humans may not know the answer to all questions. The objects will again be placed on un-textured surfaces, but these may now contain other objects. There may be some occlusion of the desired object from various places in a room, but there will be viewing locations with unobstructed views. The robot will thus have to reason about how to view the object so as to increase its confidence in the identification or categorisation.

Explanation with limited extension (Month 39): In the third year we will build on the earlier systems to achieve a robot that can reflect on the cause of failure in plan execution, and provide via dialogue an explanation for its failure for domain models where abduction can be performed. It may be able to plan very limited extension activities to test this hypothesis, or at least to express it in some formal manner. The robot may now also be able to grasp novel objects by planning novel grasp points. It will demonstrate the ability to generalise across objects when pushing with respect to shape and when grasping with respect to weight by using a modular motor learning strategy. Active segmentation will now work more reliably, with the robot able to represent uncertainty about the object segmentation or shape, and act to minimise this using explicit models of uncertainty. It will now be able to spot opportunities for self-extension, with an emphasis on generating questions about objects to acquire categorical knowledge. In dialogue the robot will also be able to plan and handle spatial referencing to maximise understandability. Finally we will demonstrate active curiosity driven learning for simple manipulation (pushing) and cross-modal learning. The system will switch between being task focused, and curiosity driven, but won't mix these.

Illustration: The robot is asked to tell Aaron about a meeting. It may be able to identify Aaron using vision (unreliable) or dialogue. It has an incomplete map of the building, and so may have to ask questions about where locations as well as people are. It doesn't know where Aaron is, and so asks, and then searches in the likely locations. It also plans to go to locations where it can ask people it believes know where Aaron is. When Aaron is not in the place expected it can revise its beliefs and provide an internal explanation for the failure (he is somewhere else). It attempts to explain perceptual failure, and if Aaron cannot be found at all, will attempt to execute a plan in which it tells the original person it could not find him. This may involve reasoning that that person is now likely to be in the specified meeting room. If the robot doesn't have an initial visual model of Aaron it could attempt to learn one when it meets him, thus extending its knowledge.

Full curiosity driven extension (Month 50): In the final year we aim to be able to demonstrate a mix of task driven behaviour and curiosity driven learning. The planner will now be able to plan active learning for self-extension, and to plan activities to test its hypotheses from the previous system about the causes of unexpected events. The robot will be opportunistic, in that it will be able to interrupt its activities on the fly to pursue curiosity driven learning. It will also now be able to generate hypotheses due to reflection that enable it to test theories about the qualitative behaviour of objects under simple manipulation, and will be able to learn and extract models of an object's behaviour that capture both its continuous and qualitative changes in pose. The robot will now also have the capacity for curiosity driven dialogue about its environment. The robot will also be able to exploit its functional models of space.

Illustration: While engaging in the task of finding the coffee mug, as described in the second illustration, the robot fails to grasp it, and must revise its grasp model for that object. Later it is given a new object to grasp, such as a beaker, and has to adapt its grasp model to this new object category. Also, while searching for the object it encounters other objects with which it is unfamiliar (e.g. a ball). It is able to identify this new object, pokes it to learn a motion model under action, and when it encounters a human, asks what it is so as to extend its ontology.

It is very important that we give clear caveats regarding the set of integrated system targets:

- The illustrations do not constitute targets to which we commit ourselves. We may vary details of these, or use another instantiation of our overall scenario to test the theory if necessary.
- We do not plan to hit all the targets listed above in each year. That is why we describe them quite clearly as upper bounds on integrated performance. We will deem the target to have been achieved if a majority of the abilities described are integrated at each stage.
- We will often take early versions of systems to perform integration, and the integration will be used to feedback into the design of those subsystems.
- Our main target is to understand the principles along which such systems should be built, and the difficulties and trade-offs in doing so. We therefore consider failure analysis for systems that fail under certain types and degrees of environmental variation to be as informative as a complete success.
- The most important deliverables from integration will not be the systems, but the reports analysing system performance. The reports and systems will be spaced at six month intervals.

The role of our scenario in establishing the theory is crucial in that it is precisely a test of our framework for self-understanding and self-extension that we seek in this work. We can only achieve this by building complete systems and using them as experimental platforms within which to test the theory. In establishing a theory of self-understanding and self-extension for an embodied cognitive system we seek a convincing test of that theory. The scenarios described above test the ability of our theory to account for the following:

- How missing information of a variety of kinds, can be planned for and acquired.
- How this can be done when the effects of information gathering actions may be uncertain.
- How to distinguish between the different possible causes of surprise.
- How to generate both hypotheses and plans to test them in the real world with its uncertainty and change.
- How to connect a variety of models of incompleteness and uncertainty in belief that reside in domain specific sub-systems, so that their joint incompleteness can be reasoned about.

Specifically our scenario-driven integration allows us to meet objectives 11 and 12 as laid out earlier.

1.3.3 Management and Risk

Given the integration-focused approach of we work planned, and the ambitious nature of the target integrated systems described in Section 1.3.2, management has a central role to play in making the project a success. In addition to handling the administrative and financial aspects of the project, the management will be actively involved in facilitating and monitoring progress. Central to this process will be the task of ensuring timely exchange of information and results between workpackages, particularly when these relate to integration. In our experience, a project that fails to manage these aspects of integration (in addition to the software side) is one that fails to produce work that is truly integrated.

The other critical task of the management will be to monitor and measure the overall progress of the project. This monitoring, via regular progress reports, will allow the management to provide additional scientific direction over the project's lifespan. Part of this will inevitably include regular re-evaluations of whether the project's stated scientific aims are being met, and assessments of changes required to adjust the direction of the research or the nature of the targets being set across the work packages.

1.4 Work package tables

1.4.1 Work package list

WP no.	Work package title	Activ.	Lead partic. no.	Lead partic. short name	Pers. mon.	Start mon.	End mon.
WP 1	Architectures for self-reflection and self-motivation	RTD	1	BHAM	70	1	50
WP 2	Object perception and manipulation	RTD	6	TUW	146	1	50
WP 3	Qualitative spatial cognition	RTD	3	KTH	80	1	50
WP 4	Planning of action, sensing and learn-ing	RTD	5	ALU-FR	110	1	50
WP 5	Interactive continuous learning of cross-modal concepts	RTD	4	UL	98	1	50
WP 6	Adaptive situated dialogue processing	RTD	2	DFKI	68	1	50
WP 7	Scenario-based integration	RTD	2	DFKI	134	1	50
WP 8	Management	MGT	1	BHAM	28	1	50
WP9	Dissemination and community building	OTHER	1	BHAM	26	1	50
	TOTAL				760		

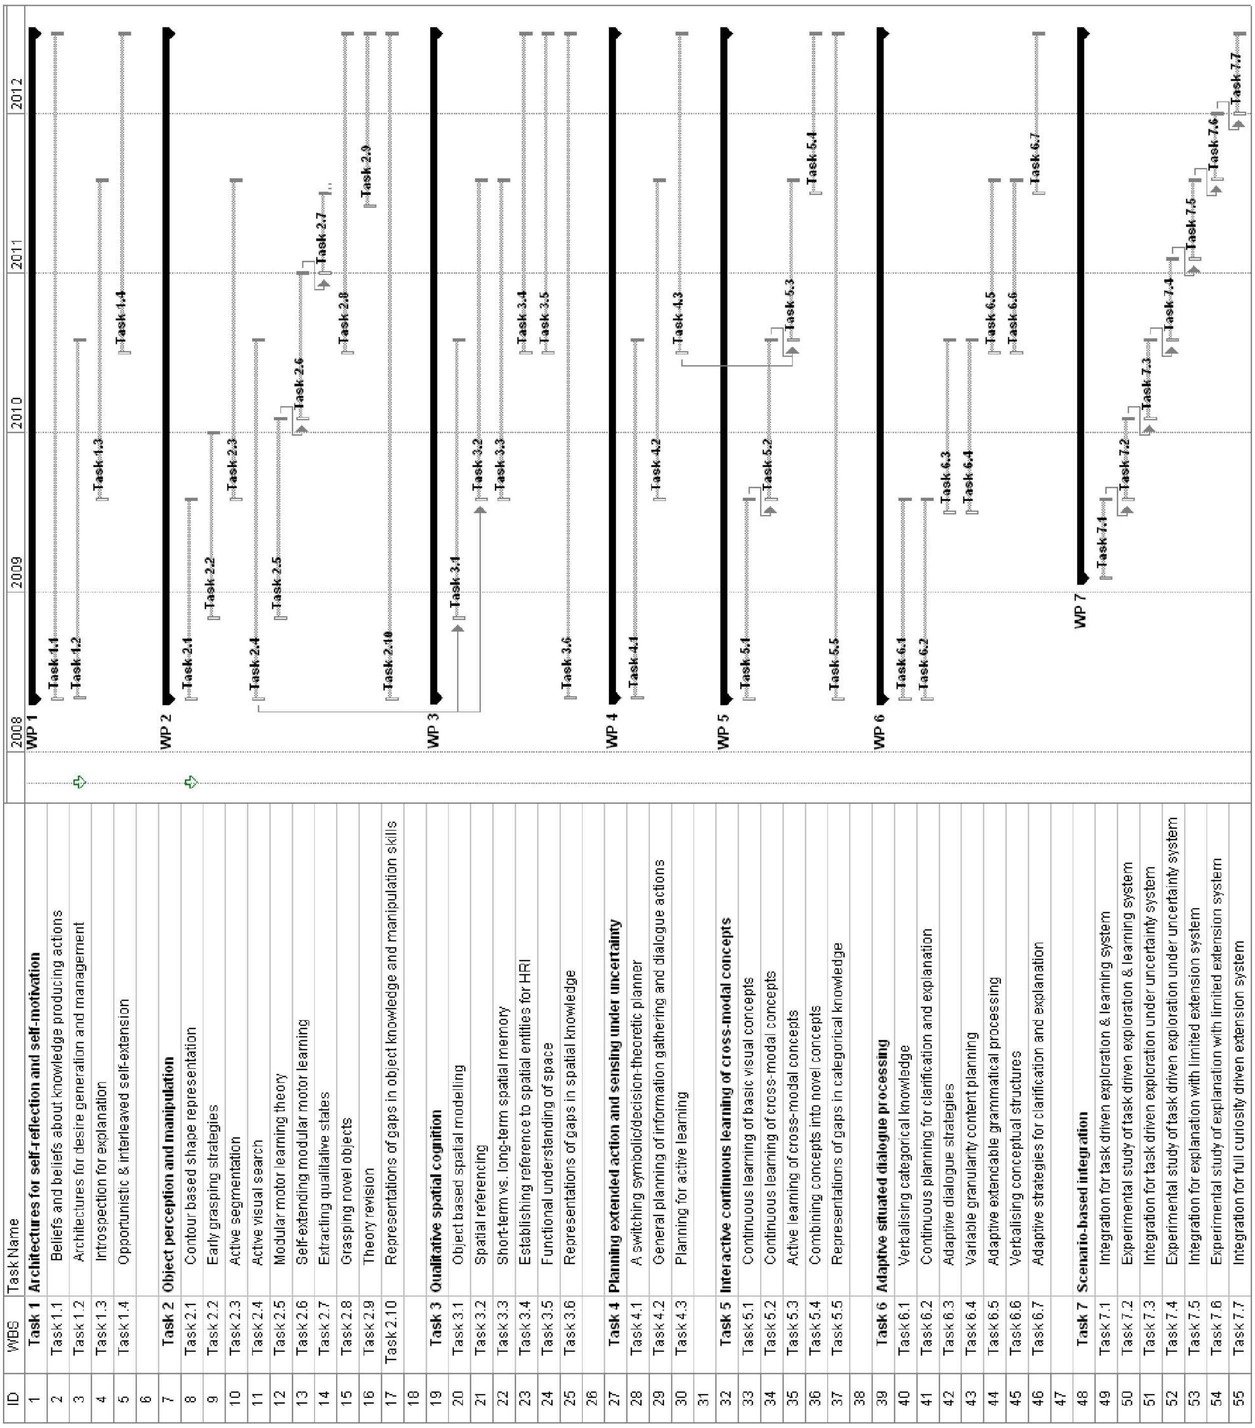

Figure 4: Gantt chart of work package tasks.

1.4.2 List of deliverables

Del. no.	Deliverable name	WP no.	Lead	Psn Mnths	Ntr	Diss. level	Deliv date
DR.1.1	Motive Management	WP 1	BHAM	14	R	PU	**15**
DR.1.2	Unifying representations of beliefs about beliefs and knowledge producing actions	WP 1	BHAM	14	R	PU	**21**
DR.1.3	Architectures and representations for introspection and motive management in a robot	WP 1	BHAM	14	R	PU	**39**
DR.1.4	Integrating intention changes into continual planning and acting	WP 1	BHAM	14	R	PU	**48**
DR.1.5	Unifying representations of gaps in knowledge	WP 1	BHAM	14	R	PU	**48**
DR.2.1	Representations of 3D shape for manipulation	WP 2	TUW	29	R	PU	**15**
DR.2.2	Active Vision, learning and manipulation	WP 2	TUW	29	R	PU	**27**
DR.2.3	Representations of gaps in knowledge about objects	WP 2	KTH	29	R	PU	**27**
DR.2.4	Manipulation of previously unseen objects	WP 2	KTH	29	R	PU	**39**
DR.2.5	Qualitative models of object behaviour, and grasping of novel objects	WP 2	KTH	30	R	PU	**48**
DR.3.1	Object based representations of space and gaps therein	WP 3	KTH	20	R	PU	**27**
DR.3.2	Spatial referencing and short-term vs. long-term memory	WP 3	KTH	20	R	PU	**39**
DR.3.3	Spatial entities for HRI and functional understanding of space	WP 3	TUW	20	R	PU	**48**
DR.3.4	Spatial knowledge and gaps therein	WP 3	TUW	20	P	PU	**48**
DR.4.1	Requirements and architectures for integrating existing symbolic and decision-theoretic planners	WP 4	ALU-FR	27	R	PU	**15**
DR.4.2	Planning for cognitive robots	WP 4	ALU-FR	27	P	PU	**27**
DR.4.3	Planning for knowledge changes	WP 4	ALU-FR	27	R	PU	**39**
DR.4.4	Planning for active learning of new actions and concepts	WP 4	ALU-FR	29	P	PU	**48**

DR.5.1	Continuous learning of basic visual concepts	WP 5	UL	20	R	PU	**15**
DR.5.2	Continuous learning of cross-modal concepts	WP 5	UL	20	R	PU	**27**
DR.5.3	Representations of gaps in categorical knowledge	WP 5	DFKI	20	R	PU	**27**
DR.5.4	Active learning of cross-modal concepts	WP 5	UL	20	R	PU	**39**
DR.5.5	Combining basic cross-modal concepts into novel concepts	WP 5	UL	28	R	PU	**48**
DR.6.1	Transparency in situated dialogue for interactive learning	WP 6	DFKI	14	R,P	PU	**15**
DR.6.2	Adaptive dialogue strategies supporting transparency	WP 6	DFKI	14	R,P	PU	**27**
DR.6.3	Adaptive extendable grammatical processing	WP 6	DFKI	14	R	PU	**39**
DR.6.4	Situated dialogue with adapting levels of vagueness and abstraction	WP 6	DFKI	14	R,P	PU	**39**
DR.6.5	Mixed initiative situated dialogue-guided curiosity	WP 6	DFKI	12	R,P	PU	**48**
DR.7.1	Analysis of a robot that achieves tasks under partial information	WP 7	BHAM	20	R	PU	**21**
DR.7.2	Analysis of a robot that acts under partial information and uncertainty	WP 7	DFKI	20	R	PU	**33**
DR.7.3	Analysis of a robot that explains surprise	WP 7	BHAM	20	R	PU	**44**
DR.7.4	Design methodologies for integrated cognitive systems	WP 7	BHAM	27	R	PU	**44**
DR.7.5	A curiosity driven self-extending robot system	WP 7	BHAM	47	R	PU	**48**
DR.9.1	CogX Website and Intranet	WP 9	UL	5	O	PU/CO	**1**
DR.9.2	Proceedings of Summer School	WP 9	UL	4	R	PU	**15**
DR.9.3	Proceedings of Summer School	WP 9	KTH	4	R	PU	**27**
DR.9.4	Proceedings of Summer School	WP 9	TUW	4	R	PU	**39**
DR.9.5	Final version of software toolkit	WP 9	BHAM	9	O	PU	**48**

1.4.3 Summary of staff effort

Partic. no.	Partic. short name	WP 1	WP 2	WP 3	WP 4	WP 5	WP 6	WP 7	WP 8	WP 9	Total person months
1	BHAM	42	42	0	42	0	0	28	16	6	176
2	DFKI	12	0	12	12	40	40	25	2	2	145
3	KTH	4	56	32	4	0	0	28	4	4	132
4	UL	4	12	12	4	48	4	12	2	6	104
5	ALU-FR	8	0	0	42	0	24	18	2	2	96
6	TUW	0	36	24	6	10	0	24	2	6	108
	Total	70	146	80	110	98	68	134	28	26	760

Mstne No.	Milestone Name	WP and Task No's	Lead	Deliv' Mon'	Means of verification
M1	Components for task driven exploration & learning	1.2, 2.1, 2.4, 4.1, 5.1, 6.1, 6.2	DFKI	**10**	Choice Point: technologies to adopt for integration for Month 15. These will be se-lected at the GA meeting.
M2	Task driven exploration & learning	1.2, 2.1, 2.2, 2.4, 2.5, 3.1, 4.1, 5.1, 6.1, 6.2	BHAM	**15**	Comparison to targets for in-tegrating scenario Month 15.
M3	Components for task driven exploration un-der uncertainty	1.2, 1.3, 2.2 – 2.5, 3.1 – 3.3, 4.1, 4.2, 5.2, 6.3, 6.4	DFKI	**21**	Choice Point: technologies to adopt for integration for year 2. These will be selected at the GA meeting.
M4	Task driven exploration under uncertainty	1.2, 1.3, 2.3, 2.4, 2.6, 3.1 – 3.3, 4.1, 4.2, 5.2, 6.3, 6.4	BHAM	**27**	Comparison to targets for in-tegrating scenario Month 27.
M5	Components for ex-plaination with limited extension	1.3, 1.4, 2.3, 2.6, 2.8, 3.2 – 3.5, 4.2, 4.3, 5.3, 6.5, 6.6	DFKI	**33**	Choice Point: technologies to adopt for integration for year 3. These will be selected at the GA meeting.
M6	Explanation with lim-ited extension	1.3, 1.4, 2.3, 2.7, 2.8, 3.2 – 3.5, 4.2, 4.3, 5.3, 6.5, 6.6	BHAM	**39**	Comparison to targets for in-tegrating scenario Month 39.
M7	Components for full cu-riosity driven extension	1.4, 2.8, 2.9, 3.4, 3.5, 4.3, 5.4, 6.7	DFKI	**44**	Choice Point: technologies to adopt for integration for year 4. These will be selected at the GA meeting.
M8	Full curiosity driven ex-tension	1.4, 2.8, 2.9, 3.5, 4.3, 5.4, 6.7	BHAM	**50**	Comparison to targets for in-tegrating scenario Month 50.

Table 1: Project Wide Milestone list. In this table we list the eight project wide milestones.

1.4.4 List of milestones and planning of reviews

Decisions on how best to proceed with the developments in various workpackages, and how and when to feed developing technologies into the integrated systems, will be taken at 6 monthly milestones. The milestones will be of two types: pre-integration milestones (odd numbered) and integration milestones (even numbered). At pre-integration mile-stones decisions will have to be taken about the exact nature of the next integrated system, based on the developments in the preceding one and the progress across the other workpackages. At in-tegration milestones, we will produce an integrated system and examine it, along with the general progress of developments across the project, to determine the best use of future effort. Table 1 reports how the workpackage tasks relate to the milestones. Tasks 1.1, 2.10, 3.6 and 5.5 will be reviewed at all milestones (as they run for the length of the project). Our expectations for the integration milestones are presented in the system descriptions in work package 7 (see Section 1.3.2). Finally we also detail, for verification purposes workpackage specific milestones in Tables 2 and 3, and their verification against deliverables where appropriate. Reviews in accordance with Annex II and project wide milestornes will take place after month 15, 27, 39 and 50.

Mstne No.	Milestone Name	WP and Task No's	Lead	Deliv' Mon'	Means of verification
M.1.1	Prototype system for Motive Management	**1**	BHAM	**15**	This will be verified by a description of the prototype system in deliverable DR.1.1. At this stage we will show how to assess and reorder desires using fast mechanisms as well as slow deliberative ones.
M.1.2	Architecture for Motive Management and Intro-spection	**1**	BHAM	**39**	This will be verified by experimental stud-ies to be described in DR.1.3. This will in-tegrate both introspection and the ability to assess and reorder desires.
M.1.3	Handling shifts between task and curiosity driven activity	**1**	BHAM	**50**	Handling shifts between task and curios-ity driven activity (Month 48) This will be verified by a system description and anal-ysis in DR.1.4. The system will show the ability to change to curiosity driven be-haviour at different task junctures.
M.2.1	Simple grasps	**2**	TUW	**21**	Detect shape of manipulable object on a table and perform grasp based on given grasping strategy. Experimental evalua-tion reported in DR.2.1.
M.2.2	Learned grasps	**2**	TUW	**30**	Use learned object behaviour under ma-nipulation to predict effects of pushes and grasps. Experimental evaluation in DR.2.2.
M.2.3	Novel grasps	**2**	TUW	**44**	Grasp previously unseen objects. Experi-mental evaluation in DR.2.5.
M.3.1	An object based spatial representation	**3**	KTH	**27**	Using the results from WP2 on active seg-mentation and visual search the system will be able to build a spatial representa-tion where objects play a central role not only as landmark for localization and navi-gation but also as a corner stone for higher level reasoning and as a basis for interac-tion with humans. Experimental verifica-tions will be reported in DR.3.1.
M.3.2	Spatial referencing and understanding of space	**3**	KTH	**50**	The system will be able to generate spatial references between objects in the model and relative to the robot to faciliate bet-ter human-robot interaction. The sys-tem should also be able to autonomously determine what of its knowledge should be treated as short-term information and what should go into the long-term mem-ory. Finally, the system should be able to reason about the function of space. Ex-perimental verifications will be reported in DR.3.2-4.
M.4.1	A switching/symbolic decision-theoretic plan-ner	**4**	ALU-FR	**27**	This will be verified by application to a planning domain model defined for one of the target scenarios. This will be reported in deliverable DR.4.2.

M.4.2	Planning of information gathering and dialogue	**4**	ALU-FR	**39**	This will be verified by generation of dialogue and information gathering plans for a planning domain defined for one of the target scenarios. The results will be reported in DR.4.3.
M.4.3	Planning for active learning	**4**	ALU-FR	**50**	This will be verified by tests on the robot system. The results will be reported in DR.4.4.
M.5.1	System for continuous learning of cross-modal concepts	**5**	UL	**27**	The system will be able to learn associations between automatically extracted features of different modalities and semantically meaningful concepts provided by a tutor through dialogue in a continuous, open-ended manner.
M.5.2	System for interactive continuous Learning	**5**	UL	**50**	The system will be able to detect its ignorance, plan and execute suitable actions (in interaction with the tutor and its environment) that extend its knowledge.
M.6.1	Situated dialogue for transparent, interactive learning	**6**	DFKI	**15**	The system, based on incremental situated dialogue processing, will be able to use clarification and explanation in a dialogue with a tutor, to learn more about the environment. To achieve transparency it can verbalize what it does and does not know (categorically).
M.6.2	Adaptive dialogue strategies supporting transparency	**6**	DFKI	**27**	The system is able to learn how to adapt the ways it communicates with a tutor in interactive learning.
M.6.3	Adaptive extendable grammatical processing	**6**	DFKI	**39**	The system will be able to extend its grammatical knowledge on the basis of newly acquired categorical knowledge. The purpose is to be able to verbalize this new knowledge, in the context of an interactive learning dialogue.
M.6.4	Situated dialogue with adapting levels of vagueness and abstraction	**6**	DFKI	**39**	The system will be able to adapt how it refers to objects.
M.6.5	Mixed initiative situated dialogue-guided curiosity	**6**	DFKI	**50**	The system will be able to initiate and drive situated dialogues for interactive learning based on its own curiosity.

Table 2: Workpackage Specific Milestone list. This table lists workpackage specific milestones grouped by workpackage for clarity.

1.5 Work package summaries

1.5.1 Summary WP 1: Architectures for self-reflection and self-motivation

Work package number:	**1**	Start date or starting event:				Month 1
Work package title	Architectures for self-reflection and self-motivation					
Activity type	RTD					
Participant number	**1**	**2**	**5**	**3**	**4**	
Participant short name	BHAM	DFKI	ALU-FR	KTH	UL	
Person months	**42**	**12**	**8**	**4**	**4**	

Objectives

This workpackage will contribute representations of self-knowledge, of actions that alter this knowledge, and of desires for self-extension and task-based action. These representations will be integrated into a computational account of motivation for a robot, produced with reference to a system architecture. Building on the work of [41], we expect our account to include the following features: beliefs about the system and its world, both specific and general, and models about how beliefs change under action, in various forms and distributed over various subsystems (i.e. beyond a purely logical database of beliefs); a small set of intrinsic desires that specify the general types of behaviours the system can engage in; a collection of semantically rich structures representing instantiations of the intrinsic desires from various subsystems, called desire instances; a set of concurrently, asynchronously active processes that generate, activate or reactivate desire instances based on intrinsic desires and the current context, called desire instance generators; a set of desire management processes that filter, inspect, assess and select desire instances for subsequent processing (including action); and a set of intentions which represent a set of selected desire instances that the system intends to act on. Although we are using the terminology of BDI systems, our heterogeneous architecture-based approach should be distinguished from the purely logic-based approaches of such systems.

Given the system-wide aims of this workpackage, it will interact with aspects of all of the objectives specified in Section 1.1, but it will be particularly concerned with the following objectives. For Objective 1 we will develop a unified framework for representing a system's intrinsic desires and desire instances where these can refer to both beliefs and beliefs about how beliefs change under action (both information processing and physical). For Objective 10 we will develop an architectural framework for the generation of desire instances from either task-driven or curiosity-driven intrinsic desires, along with approaches for managing competing and conflicting desires. Many more intrinsic desires could be considered in more general work, such as physical well-being (homoeostasis) or social pressures. For Objective 3 we will investigate how actions that stem from the adoption of desire instances into intentions will alter the belief state of the system and other agents. We will evaluate our work using implementations of our approaches. In the project's implemented systems the desire instance generator and management mechanisms will provide architecture-wide control. This will contribute towards Objectives 11 & 12.

In summary the objective of this WP is to do the following:

• Investigate representations of beliefs about beliefs and actions that can operate on them.

• Provide an architectural account of motivation of an intelligent robot.

• Support the generation of behaviour from the intrinsic desires to complete tasks or to extend, explore or explain.

• Evaluate the designs in working systems.

Description of work:

The implementation of this workpackage is separated into four tasks. They are concerned with representations for self-knowledge particularly when related to knowledge that can lead to exploratory or self-extending action, architectures and processes for desire generation and management, and reasoning methods for integrating these processes with behaviour. When referring to desires, we are particularly concerned with desires that refer to tasks that the system can carry out, and those that refer to opportunities to satisfy the system's drives to explore, explain and extend (i.e. its curiosity-based desires).

Task 1.1: Beliefs and beliefs about knowledge producing actions. We will examine how a system can represent, in a unified way, beliefs about incompleteness and uncertainty in knowledge. This will start with work on their representation that will feed into WPs 2, 3 & 4, and it will later unify the modality specific representations of incompleteness and uncertainty coming up from these packages. Representations of knowledge producing actions will utilise these to represent the preconditions and effects of knowledge producing actions. These knowledge action effects will be used in WP4 for planning information gathering and processing. This task will also support work on introspection. (Months 1 – 50) (BHAM (15 months), DFKI (4 months), ALU-FR (2 months), KTH (2 months), UL (2 months)).

Task 1.2: Architectures for desire generation and management. We will analyse the requirements for representations of desires in the system, and architectures to generate and manage them. Extending our previous work on motivation we will explore designs for multiple concurrent desire instance generators combined with filter mechanisms [41, 40, 4, 31], within the CoSy Architecture Schema [19, 18]. (Months 1 – 27) (BHAM (8 months))

Task 1.3: Introspection for explanation. In order to provide explanations for unexpected events we will have to perform introspection on our models. This will require the ability to analyse reasoning or prediction failures, and to identify possible additional state variables or rules that we may be able to use to augment a model to explain the unexpected event. We will develop methods and representations for doing this in other workpackages, and here we will attempt to unify these. We will develop methods for introspection and hypothesis generation in both logical and probabilistic models. (Months 16 – 39) (BHAM (11 months), DFKI (4 months), ALU-FR (4 months), UL (1 month) KTH (1 month))

Task 1.4: Opportunistic & interleaved self-extension. In this task we are interested in the interaction between selected intentions, mechanisms for planning behaviour and the mechanisms for executing behaviour plans. We will investigate how asynchronously generated desire instances can trigger new behaviours that must be interleaved with current behaviours, and other trade-offs that can be made. This will be particularly relevant for taking advantage of self-extension opportunities during an otherwise task-driven activity. (Months 27 – 50) (BHAM (8 months), DFKI (4 months), ALU-FR (2 months), UL (1 month), KTH (1 month))

Deliverables:

DR.1.1 Motive Management. Report. (Month 15) (BHAM)

DR.1.2 Unifying representations of beliefs and knowledge producing actions. Report. (Month 21) (BHAM, DFKI, UL, KTH, ALU-FR)

DR.1.3 Architectures and representations for introspection and motive management in a robot. Report. (Month 39) (BHAM, DFKI, UL, ALU)

DR.1.4 Integrating intention changes into continual planning and acting. Report. (Month 48) (ALU-FR, BHAM, DFKI)

DR.1.5 Unifying representations of gaps in knowledge. Report. (Month 48) (BHAM, ALU-FR, DFKI, KTH, UL)

Milestones:

M.1.1 Prototype system for Motive Management. (Month 15) This will be verified by a description of the prototype system in deliverable DR.1.1. At this stage we will show how to assess and reorder desires using fast mechanisms as well as slow deliberative ones.

M.1.2 Architecture for Motive Management and Introspection in a Robot (Month 39). This will be verified by experimental studies to be described in DR.1.3. This will integrate both introspection and the ability to assess and reorder desires.

M.1.3 Handling shifts between task and curiosity driven activity (Month 50) This will be verified by a system description and analysis in DR.1.4. The system will show the ability to change to curiosity driven behaviour at different task junctures.

1.5.2 Summary WP 2: Object perception and manipulation

Work package number:	2	Starting date or starting event:			Month 1
Work package title	Object perception and manipulation				
Activity type	RTD				
Participant number	3	1	6	4		
Participant short name	KTH	BHAM	TUW	UL		
Person months	56	42	36	12		

Objectives

The ability to manipulate novel objects detected in the environment and to predict their behaviour after a certain action is applied to them is important for a robot that can extend its own abilities. The role of this work package is to provide the necessary sensory input for the above by exploiting the interplay between perception and manipulation. We will develop robust, generalisable and extensible manipulation strategies based on visual and haptic input. We envisage two forms of object manipulation: pushing using a “finger” containing a force-torque sensor and grasping using a parallel jaw gripper and a three-finger Barrett hand. Through coupling of perception and action we will thus be able to extract additional information about objects, e.g. weight, and reason about object properties such as empty or full. To summarise:

- to develop representations that allow robust detection of objects in realistic environments,
- to provide methodologies for manipulation of known and novel objects,
- to learn predictive models of object behaviour from a small set of objects,
- to develop generalisable and extensible manipulation strategies for two and three fingered robot hands.

Description of work

Underpinning this work package will be a strand of work on how to represent objects so as to be able to detect objects in the environment, learn predictive models of object behaviour from a small set of objects, and then generalise our models of their behaviour under action to novel, previously unseen objects. Where our model fails to generalise successfully to a new object the system should, by introspection on the extracted sensory input and previously learned models generate hypothesised experiments that would provide the information about the new objects. Perception has two roles in this work package. Firstly we need to perceive object structure, as it is the object’s behaviour under robot actions we are ultimately interested in, and behaviour depends on structure. We will use a combination of contour based segmentation approaches and structure from motion techniques. Secondly we have to detect objects in cluttered scenes, from drastically varying view points and distances and with illumination changes, occlusions and real-time constraints. We shall follow the active vision paradigm where, instead of passively observing the world, viewing conditions are actively changed to improve vision performance. To this end we plan to use cameras mounted on a pan/tilt-unit and zoom-able cameras. We plan to combine global appearance based methods (for initial detection) and local feature-based methods (for verification). Regarding manipulation, we will focus on developing a theory of modular prediction of the effects of actions on objects. This will be based on the theory of modular motor learning. Based on 3D shape models, we will acquire representations of pushing and grasping strategies that generalise across object categories and allow extension to novel objects. Current models of modular motor learning are essentially uni-modal and only predict the effects of an action on variables describing the internal state of the manipulator (e.g. proprioception). We aim to extend the theory to allow input and output channels from position sensing, force and vision. In relation to grasping, we will deal with both two- and three-fingered hands and investigate how different shape representations can be facilitated to generate the input necessary for defining grasp strategies through combination of approach vector (where to place the hand with respect to the object) and preshape (what type of grasp to use in order to grasp the object).

Task 2.1: Contour based shape representations. Investigate methods to robustly extract ob- ject contours using edge-based perceptual grouping methods. Develop representations of 3D shape based on contours of different views of the object, as seen from different camera po- sitions or obtained by the robot holding and turning the object actively. Investigate how to incorporate learned perceptual primitives and spatial relations from WP5. (M1 – M15) (TUW (12 months), KTH (1 month))

Task 2.2: Early grasping strategies. Based on the visual sensory input extracted in Task 2.1, define motor representations of grasping actions for two- and three-fingered hands. The initial grasping strategies will be defined by a suitable approach vector (relative pose with respect to object/grasping part) and preshape strategy (grasp type). (M7 – M21) (KTH (8 months))

Task 2.3: Active segmentation. Use haptic information, pushing and grasping actions i) for interactive scene segmentation into meaningful objects, and ii) for extracting more detailed object models (visual and haptic). Also use information inside regions (surface markings, texture, shading) to complement contour information and build denser and more accurate models. (M16 – M39) (TUW (12 months), UL (2 months), KTH (2 months))

Task 2.4: Active Visual Search. Survey the literature and evaluate different methods for visual object search in realistic environments with a mobile robot. Based on this survey develop a system that can detect and recognise objects in a natural (possibly simplified) environment. (M1 – M27) (KTH (12 months), UL (2 months))

Task 2.5: Modular motor learning theory. Using 3D contour based shape descriptors plus haptic and proprioceptive information, extend the modular motor learning theory to predict and control object trajectories and contact relations. Extend object models with attributes that can be learned/detected only in contact with the object, e.g. weight. This will be investigated for both pushing and grasping actions. (M7 - M21) (BHAM (10 months), UL (2 months))

Task 2.6: Self-extending modular motor learning. Create a modular motor learner that adds new contexts to its model, based on assessing the quality of its predictions. We will use a probabilistic model to identify whether to refine existing modules, or add a new module, and to control exploration while doing so. (M22 – 32) (BHAM (7 months), UL (2 months))

Task 2.7: Extracting qualitative states. To be able to perform introspection on possible qualitative cause we require a model that has not only continuous states, but also qualitative states, with qualitative explanations for the transitions between them. We will use the notion of force-aspect graphs to devise a learning algorithm capable of partitioning the continuous configuration space of the modular motor learning predictions into sets of qualitatively similar stable states, plus their basins of attraction. (M33 – M39) (BHAM (6 months))

Task 2.8: Grasping novel objects. Based on our object models, we will investigate the scalability of the system with respect to grasping novel, previously unseen objects. We will demonstrate how the system can execute tasks that involve grasping based on the extracted sensory input (both about the scene and individual objects) and taking into account its embodiment. (M27 - M50) (KTH (18 months), TUW (6 months))

Task 2.9: Theory revision. Given a qualitative, causal physics model, the robot should be able to revise its causal model by its match or mismatch with the qualitative object behaviour. When qualitative predictions are incorrect the system will identify where the gap is in the model, and generate hypotheses for actions that will fill in these gaps. (M38 - M50) (BHAM (10 months), UL (2 months), KTH (5 months))

Task 2.10: Representations of gaps in object knowledge and manipulation skills. We will develop representations of the incompleteness of, and uncertainty about, models of objects.This is a prerequisite for reasoning about information-gathering actions and performing introspection. This task will feed into the unifying work on this in WP1. (M1–50) (KTH (10 months), BHAM (9 months), TUW (6 months), UL (2 months))

Relation to other work packages

While WP5 explores the interaction of perception and language, tutor-driven learning (with variable tutor involvement) and the formation of semantic concepts defined by language (such as perceptual properties, spatial relations), WP2 focuses on the interaction of perception and manipulation. Semantic concepts formed here are objects and manipulable object parts. Furthermore WP2 will deliver the visual search routines needed in WP3. Visual primitives and spatial relations learned in WP5 will be used in WP2 and vice versa. Higher level features, such as extracted 3D surface patches, will be provided for WP5 to use as learning input.

Deliverables

DR.2.1 Representations of 3D shape for manipulation. Report. (Month 15) (TUW, KTH)

DR.2.2 Active vision, learning and manipulation. This deliverable groups several reports on: Vision routines for active visual search (KTH, TUW); Modular motor learning (BHAM, TUW); and Grasping learned objects (KTH). Report. (Month 27) (KTH)

DR.2.3 Representation of gaps in object knowledge. Report. (Month 27) (BHAM, TUW, KTH, UL)

DR.2.4 Manipulation of previously unseen objects. This combines two pieces of work on self

extending modular motor learning (BHAM, TUW) and grasping of previously unseen objects (KTH, TUW). Report. (Month 39) (KTH, BHAM, TUW)

DR.2.5 Qualitative models of object behaviour, and grasping of novel objects. This bundles two

pieces of work on building qualitative models of behaviour (BHAM), and on grasping of novel objects (KTH). Report. (Month 50) (KTH, BHAM)

Milestones

M.2.1 Simple grasps. (Month 21). Detect shape of manipulable object on a table and perform grasp based on given grasping strategy. Experimental evaluation in DR.2.1.

M.2.2 Learned grasps. (Month 30). Use learned object behaviour under manipulation to predict effects of pushes and grasps. Experimental evaluation in DR.2.2.

M.2.3 Novel grasps. (Month 44) Grasp previously unseen objects. Experimental evaluation in DR.2.5.

1.5.3 Summary WP 3: Qualitative spatial cognition

Work package number:	**3**	Starting date or starting event:				Month 1
Work package title	Qualitative spatial cognition					
Activity type	RTD					
Participant number	**3**	**6**	**2**	**4**		
Participant short name	KTH	TUW	DFKI	UL		
Person months	**32**	**24**	**12**	**12**		

Objectives

Spatial models of the environment are at the very core for mobile robot systems. Robots are currently in the progress of moving out from the factories and into our homes and offices to, for example, run errands for us or otherwise assist us. Communication between robots and humans will therefore become increasingly important. To support this the robot does not only have to be able to perform navigation, but must go beyond that. It has to have an understanding of the environment that allows it to answer questions such as, for example, “what function does this part of the space fulfil?”, “where would I typically find object X?”, “what can be done with object X?”. That means that, on top of a quantitative model of space, the robot needs to have a conceptual model of spatial entities – rooms or other topological units, objects, their individual properties, and their relations.

If the robot is to perform its tasks in a man-made, human-populated environment, and, moreover, if the robot is to communicate about its environment with humans, then the conceptual spatial model needs to reflect how humans conceptualise their environment. Comparing how robots and humans perceive and represent the world shows that there is a large difference. For communication between the two to work, this gap has to be bridged. Robots have to start perceiving things closer to the way humans do, especially when interacting with real end-users and not the scientists that designed the robot system.

Our hypothesis is that objects play an important role when building a spatial model of a man-made environment for interaction with humans and answering questions like the ones above. Spatial models with objects as core building blocks would allow for a scalable representation and act as the basis for much of the high level reasoning. The spatial modelling and the search for objects are naturally coupled. When looking for a certain object, the search can be directed towards areas in the environment where such objects are normally found. For example, when looking for a box of cereals, the floor is an unlikely place to find it. Furthermore, since manipulation of object is of special interest in this project we will primarily work with objects that can be manipulated by the system and focus on objects located within graspable distance of the system.

Another important question in conjunction with the spatial modelling is the question of short-term versus long-term memory. Traditional robot systems often rely on a global long-term representation for the environment, the map, that allows the robot to stay localised and plan its way from one part of the environment to another. For local navigation another local but much more detailed map is often used. This local map acts as a detailed short-term memory that the robot will forget as soon as this particular area is left. The global model of the environment can be layered as in [25] to support different functions at different levels of abstraction, from low level navigation to conceptual reasoning. The main objectives of this WP are:

- Study how to best incorporate objects into the spatial model of the environment?
- Investigate how the object-based spatial models can be used to infer knowledge about type

and function of an area, typical placement of object classes, etc?

- Study how to perform spatial referencing between object in the model and relative to the robot?
- Find out what part of the environment should be captured in short-term representations and what goes into long-term memory and investigate if the same layering is necessary for the short-term memory as for the long-term memory?
- Study how to utilise and learn spatial relations in human-robot interaction?
- Investigate how to represent gaps in spatial knowledge?

Description of work:

Task 3.1: Object based spatial modelling. Develop a framework that allows for a hybrid representation where objects and traditional metric spatial models can coexist (Months 7 – 27) (KTH (9 months), TUW (4 months), DFKI (2 months))

Task 3.2: Spatial referencing. Investigate what objects and other entities in the map should be referenced to and how. (Months 16 – 39) (KTH (3 months), DFKI (1.5 months))

Task 3.3: Short-term vs long-term spatial memory. What knowledge goes where and how does it depends on the task? (Months 16 – 39) (KTH (4 months), TUW (2 months))

Task 3.4: Establishing reference to spatial entities for human-robot interaction. Investigate, in the context of human-robot interaction, how the robot can refer to objects based on their spatial relations and how to learn this. (Months 27 – 50) (UL (9 months), DFKI (4.5 months), KTH (2 months))

Task 3.5: Functional understanding of space. Investigate how to gain knowledge about the function of space by analysing spatial models over time. (Months 27 – 50) (KTH (6 months), TUW (12 months))

Task 3.6: Representations of gaps in spatial knowledge. How to represent beliefs about beliefs of spatial knowledge. (Months 1 – 50) (KTH (8 months), TUW (6 months), DFKI (4 months), UL (3 months))

Deliverables:

DR.3.1 Object based representations of space and gaps therein. Report. (Month 27) (KTH,TUW, DFKI)

DR.3.2 Spatial referencing and short-term vs. long-term memory. Report. (Month 39) (KTH, DFKI, TUW)

DR.3.3 Spatial entities for HRI and functional understanding of space. This includes two reports on establishing reference to spatial entities for HRI and to the creation of functional models of space. Report. (Month 48) (UL, DFKI, KTH, TUW)

DR.3.4 Spatial knowledge and gaps therein. Report. (Month 48) (KTH, TUW, DFKI, UL)

Milestones:

M.3.1 An object based spatial representation. (Month 27) Using the results from WP2 on active segmentation and visual search the system will be able to build a spatial representation where objects play a central role not only as landmark for localization and navigation but also as a corner stone for higher level reasoning and as a basis for interaction with humans. Experimental verifications will be reported in DR.3.1.

M.3.2 Spatial referencing and understanding of space. (Month 50) The system will be able to generate spatial references between objects in the model and relative to the robot to faciliate better human-robot interaction.   The system should also be able to autonomously determine what of its knowledge should be treated as short-term information and what should go into the long-term memory. Finally, the system should be able to reason about the function of space. Experimental verifications will be reported in DR.3.2-4.

1.5.4 Summary WP 4: Planning of action, sensing and learning

Work package number:	**4**	Starting date or starting event:				Month 1
Work package title	Planning of action, sensing and learning					
Activity type	RTD					
Participant number	**5**	**1**	**2**	**6**	**3**	**4**
Participant short name	ALU-FR	BHAM	DFKI	TUW	KTH	UL
Person months	**42**	**42**	**12**	**6**	**4**	**4**

Objectives

A cognitive system that is self-extending and needs to act in environments with uncertainty, change and lack of knowledge needs representations of the state of the world and its internal state, and must be able to plan actions based on this knowledge. These actions may change the external state, but they may also be sensing actions (including dialogue acts such as asking questions) or algorithmic actions such as running a vision algorithm on an image. These last two, which we will refer to as information-gathering, only change the internal state of the system, and can be treated together.

There are a number of competing requirements for the planning component of CogX. It must be able to build plans that include physical and information-gathering actions. These actions may be stochastic or non-deterministic, and the system must reason about their possible outcomes. Also it must plan in a world that is not known with certainty (otherwise information-gathering would be unnecessary). Finally, the system must make decisions quickly. To achieve these goals we design a system that can switch between a fast continual planner and a more computationally expensive decision-theoretic planner. In the “get the cornflakes” scenario a symbolic planner that operates over epistemic states can be used if the information needed to carry out the plan is available, or if it is easily obtainable, for example by asking questions. If all the information required is not easily available then an efficient plan will require reasoning about the possible outcomes of information-gathering actions when deciding what to do. To build this switching planning system we will have to extend the state of the art in both classical planning for epistemic states and in decision-theoretic planning, as well as developing a bounded rationality-based reasoning system to determine which to use.

As we have said, to be truly self-extending the cognitive system must be able to learn about the world and learn new actions. This may involve planning actions to learn new things, or to refine existing knowledge. This kind of planning requires reasoning about the system’s internal model and how it might be changed by future experiences. It might include trying out an action in a new situation to learn what its effects are, or planning to test a hypothesis. While model-learning approaches used in reinforcement learning address this challenge to some extent, there is relatively little work on active learning of representations needed for planning or reasoning at a high level. In summary the planner should have the following characteristics:

- It should operate continually, interleaving planning and execution.
- It should be able to reason about non-deterministic and stochastic outcomes of actions when building plans.
- It should be able to cope with state ambiguity and gaps in its knowledge. It should build plans that include information-gathering actions or conformant plans that achieve their goals without requiring the missing information.

- It should be capable of planning dialogue activities and reasoning about both its own mental state and that of others.
- It should be able to plan to change its internal model of the world, choosing actions to facilitate learning of new actions or concepts.
- It should be capable of doing all this in real-time.

Description of Work:

To accomplish all these planning challenges, the work-package is divided into four tasks corresponding to the major areas of research:

Task 4.1: A switching symbolic/decision-theoretic planner (Months 1-27) In this task we will look at how to combine these two approaches (symbolic and decision theoretic) into the switching planner discussed above. The result should be a planner that can do limited rea- soning about belief states (the representation in terms of epistemic operators is much less expressive than a full probabilistic belief-state representation) but still make good decisions, and that will operate in close-to real-time. After 15 months this task should deliver a report on how to combine the planners, the architecture of the planning system, and how the plan- ner operates within the overall architecture of the CogX system, and after 27 months, the switching planner that uses bounded rationality analysis to decide when to use each approach, and that uses the representations of world states and actions used by the rest of the CogX system. (ALU-FR (18 months), BHAM (20 months))

Task 4.2: General planning of information gathering and dialogue actions (Months 15- 39) The symbolic planner developed in Task 4.1 is limited in that it uses epistemic operators to represent beliefs, so it can only represent that a fact is known to be true, known to be false, or unknown. Better plans can be achieved by representing a much richer set of beliefs, for example by using probabilistic belief states. This allows the system to reason about the most likely states given its current knowledge, so it can for example begin driving towards the kitchen when sent to look for the cornflakes because its a-priori belief is that they are most likely to be in the kitchen. In Task 4.2 we will extend the planning system to allow arbitrary belief states to be reasoned about. The aim is to produce a planner capable of planning over arbitrary belief states, but specialised for the requirements of our domain. The task should deliver a preliminary report on how to achieve this at Month 27, and a planner for integration with the rest of the system at month 39. This task will build on work about representations of knowledge producing actions in WP1 and contribute to work in WP6 on continual dialogue planning. (ALU-FR (14 months), BHAM (11 months), UL (2 months), TUW (3 months), DFKI (6 months), KTH (2 months))

Task 4.3: Planning for active learning. (Months 27-50) If an agent has an ontology which captures its general knowledge about the world—such as the types of objects, actions and qualities, and the typical relations between them—then to extend this it must be able to represent beliefs about its incompleteness or incorrectness, and reason about the effects of possible actions on these beliefs. To extend its ontology the system needs to perform actions or create situations in which it experiences things it understands poorly. Here the agent is planning how to act so as to modify its theory of how the world works. We will tackle this problem by devising new representations for beliefs about ontological gaps, and gaps or errors in causal theories. This will therefore be related to work we will do in WP2 on object manipulation and understanding causation, and to work in WP5 on continuous learning. (ALU-FR (10 months), BHAM (11 months), UL (2 months), TUW (3 months), DFKI (6 months), KTH (2 months))

Relation to other work packages: Planning has connections with all the other packages in that it either causes the systems they develop to do things (WPs 2,6), or uses the representations they produce (WPs 1,3,5). In particular, WP1 provides goals for the planner and the archi- tecture within which it fits. WPs 1,3 and 5 provide representations of beliefs, spatial concepts, and cross-modal knowledge respectively, which become the state variables that the planner reasons about.

Deliverables:

DR.4.1 Requirements and architectures for integrating symbolic and decision theoretic planning. Report (Month 15) This will include both a requirements study for integrating existing symbolic and decision-theoretic planners, and study of a suitable architecture for the planning system and how it fits into overall CogX architecture. (ALU-FR, BHAM)

DR.4.2 Planning for cognitive robots. Report/Prototype. (Month 27) This will include a prototype and description of a switching planning system capable of limited reasoning about state uncertainty and stochastic/non-deterministic action effects. It will also include a study of how to extend the planner to allow richer state uncertainty for planning dialogue and information-gathering. (BHAM, ALU-FR)

DR.4.3 Planning for knowledge changes. Report/Prototype (Month 39). This will include a report on a prototype system capable of planning over belief states, and planning dialogues and information gathering; and a study of how to extend the planner to perform active learning. (ALU-FR, BHAM, DFKI, KTH, UL, TUW)

DR.4.4 Planning for active learning of new actions and concepts. This should be a fully capable planning system for the extensible gofer task. Software prototype. (Month 48) (ALU-FR, BHAM, DFKI, UL, TUW, KTH)

Milestones:

M.4.1 A switching/symbolic decision-theoretic planner. (Month 27). This will be verified by application to a planning domain model defined for one of the target scenarios. This will be reported in deliverable DR.4.2.

M.4.2 Planning of information gathering and dialogue. (Month 39). This will be verified by generation of dialogue and information gathering plans for a planning domain defined for one of the target scenarios. The results will be reported in DR.4.3.

M.4.3 Planning for active learning. (Month 50). This will be verified by tests on the robot system. The results will be reported in DR.4.4.

1.5.5 Summary WP 5: Interactive continuous learning of cross-modal concepts

Work package number:	**5**	Starting date or starting event:				Month 1
Work package title	Interactive continuous learning of cross-modal concepts					
Activity type	RTD					
Participant number	**4**	**2**	**6**			
Participant short name	UL	DFKI	TUW			
Person months	**48**	**40**	**10**			

Objectives

An important characteristic of a cognitive system is the ability to continuously acquire new knowl- edge and new skills in a life-long manner. We refer to such ever-present, life-long learning as continuous learning. Moreover, the learning process in an artificial cognitive system is inherently cross-modal. The environment is perceived through different sensors (e.g. visual, haptic), and is acted upon using different actuators (motor wheels, robot arms). The interaction with the environment and communication (verbal and non-verbal) with a tutor should significantly facilitate incremental learning processes. These processes might induce different levels of tutor involvement and different levels of robot autonomy. Our goal is to analyse these different types of interactive learning and to develop a system (in collaboration with WP 6) that would be able to seamlessly switch between different learning modes in a convenient way.

This workpackage will therefore focus on interactive continuous learning of cross-modal concepts and on detecting the gaps in categorical knowledge and filling these gaps by updating the corresponding concepts. The system should build its competencies incrementally. Initially, simple concepts (e.g., colours, shapes, etc.) will be built by grounding these concepts to sensory data (i.e. associating them with the features extracted from the sensory data) using a combination of pre-linguistic learning and learning based on communication with the tutor. With the benefit of these acquired concepts as a firm basis, it should continuously build new cross-modal and amodal concepts achieving a progressively richer ability to reason, plan and explain the environment and the robot's situatedness in this environment.

The main objectives of this WP are to:

- Develop a system for interactive continuous learning of cross-modal concepts.
- Investigate how to best incorporate prelinguistic learning of discrete sets of cross-modal concepts to aid the acquisition of linguistic concepts.
- Investigate how to best support different modes of learning by an interactive self-extending architecture.
- Integrate tutor involvement with an interactively learning system in a user friendly and flexible manner.
- Find representations of skills, concepts and experiences that can form the basis of knowledge boundary identification (i.e. ignorance identification).
- Advance the learning system to be able to actively plan and execute new actions that may increase the system's knowledge.
- Address the stability/plasticity dilemma by means of introspective management of desires and active cross-modal validation of concepts and skills.

- Find cross-modal representations appropriate for merging information stemming from vision, haptics, language, manipulation, planning etc.
- Find a mechanism for combining concepts into novel concepts at a higher level of abstraction.

Description of work:

Task 5.1: Continuous learning of basic visual concepts. Develop a learning mechanism for learning basic visual concepts grounded to signals. The system will be able to build associa- tions between features extracted from input visual data (colour and depth images) and visual attributes (e.g., colour, shape) and connecting them using language in a dialogue with the tutor. Adequate mechanisms for unlearning will be investigated as well. (Months 1– 15) (UL (9 months), DFKI (9 months))

Task 5.2: Continuous learning of cross-modal concepts. Extend the system to consider features of other modalities and to build cross-modal category systems. Analyse the trade-offs between unsupervised and supervised learning. (Months 16 – 27) (UL (9 months), DFKI (8 months))

Task 5.3: Active learning of cross-modal concepts. Increase the system's autonomy to en- able continuous detection of ignorance, and active planning and execution of knowledge pro- ducing actions enabling autonomous continuous self-extension. (Months 28 – 39) (UL (12 months), DFKI (7 months))

Task 5.4: Combining concepts into novel concepts. Develop a system that is able to com- bine concepts learned in the previous tasks into novel concepts; to learn complex concepts and hierarchies of concepts. (Months 38 – 50) (UL (9 months), DFKI (8 months))

Task 5.5: Representations of gaps in categorical knowledge. Investigate how a system can exhibit a certain level of self-understanding and self-criticism to detect the gaps in its knowl- edge and how to represent these beliefs about beliefs of cross-modal categorical knowledge. (Months 1 – 50) (UL (9 months), DFKI (8 months))

Deliverables:

DR.5.1 Continuous learning of basic visual concepts. Report. (Month 15) (UL, DFKI)

DR.5.2 Continuous learning of cross-modal concepts. Report. (Month 27) (UL, DFKI)

DR.5.3 Representations of gaps in categorical knowledge. Report. (Month 27) (UL, DFKI)

DR.5.4 Active learning of cross-modal concepts. Report. (Month 39) (UL, DFKI)

DR.5.5 Combining basic cross-modal concepts into novel concepts. Report. (Month 48) (UL, DFKI)

Milestones:

M.5.1 System for continuous learning of cross-modal concepts. (Month 27) The system will be able to learn associations between automatically extracted features of different modalities and semantically meaningful concepts provided by a tutor through a dialogue in a continuous, open-ended manner.

M.5.2 System for interactive continuous learning. (Month 50) The system will be able to continuously and autonomously detect its ignorance, and actively plan and execute suitable actions (in interaction with the tutor and its environment) that may provide novel information useful for increasing its knowledge.

1.5.6 Summary WP 6: Adaptive situated dialogue processing

Work package number:	**6**	Starting date or starting event:				Month 1
Work package title	Adaptive situated dialogue processing					
Activity type	RTD					
Participant number	**2**	**5**	**4**			
Participant short name	DFKI	ALU-FR	UL			
Person months	**40**	**24**	**4**			

Objectives

Situated dialogue is a means for a robot to extend or refine knowledge about the environment. For this, the robot needs to be able to establish with a human some form of mutually agreed-upon understanding – they need to reach a common ground. The goal of this WP is to develop adaptive mechanisms for situated dialogue processing, to enable a robot to establish such common ground. We will focus on dialogues for continuous learning. In continuous learning, the robot is ultimately driven by its own curiosity, rather than by extrinsic motivations. Therefore, we want to conceive of dialogue as peer-to-peer, mixed-initiative communication – either the robot or the human can be the one to be asked to clarify, explain, or perform something.

Establishing common ground requires the robot to be able to process (i.e. comprehend and generate) clarification requests and explanations. A clarification request is a request for information to help overcome a breakdown in communication, or in understanding a situation. A request can be a single utterance, but also an entire sub-dialogue – e.g., when a request needs to be refined, or rephrased. An explanation provides information about why an agent does something, or believes a certain fact to be true. To achieve common ground in dialogue for continuous learning, these strategies are to serve two related purposes: They help the robot and the human to achieve transparency in what the robot needs to learn about the situation, and then interactively set up an appropriate context (scaffolding) in which a clarification request, an explanation, or a task can be used to trigger an appropriate learning goal [36, 37].

This problem is challenging because the robot's knowledge is continuously being adapted and extended – and this requires dialogue processing to be adaptive and extendible, too. First, as categorical knowledge grows, the robot needs to learn online how to use context information to focus clarification and explanation on what is relevant. Categorical knowledge is an associative network of concepts, and only some may be relevant for the robot to talk about to get the answer it seeks. Second, because the robot acquires knowledge over a period of time with varying degrees of supervision/autonomy, we cannot assume that the robot's grammatical competence always has the adequate coverage to express newly acquired (categorical) knowledge. This means the robot needs an ability to learn how to linguistically convey new categorical knowledge.

The objectives of WP6 are thus as follows:

- Provide the robot with the basis for interactively establishing a mutually agreed upon common ground of the user and the robot.
- Provide verbalisation of the basic and combined perception-based concepts learned in WP5 (extension of lexico-grammatical knowledge).
- Use motivation systems to guide mixed-initiative communication. Investigate to what degree situated dialogue can make the robot transparent to the user in terms of intentions, knowledge

level and other internal properties.

- Active scene manipulation to create situations that support the communication of certain aspects of knowledge (e.g. to illustrate what a certain concept refers to).

Description of work:

Task 6.1: Verbalising categorical knowledge. (Months 1 – 15). The goal is to enable the robot to verbalize its own categorical knowledge (or lack of knowledge) relative to a situation, and understand situated references. We will extend existing methods for comprehending and producing referring expressions to cover verbalization of relevant information from singular visual categories (WP5), and contextual reference [9, 20, 23, 21]. (DFKI 6 months, UL 1 month)

Task 6.2: Continual planning for clarification and explanation. (Months 1 – 15). We will extend strategies for planning clarification- and explanation dialogues [32, 24] using a con- tinual planning approach [7]. This offers the necessary flexibility to adjust a plan when interactively setting up an appropriate context, and provides a model of common ground in dialogue of [16, 30, 15]. These methods will be based in means for grounding the information expressed by clarifications and explanations in situated understanding. (DFKI 6 months, ALU-FR 6 months, UL 1 month)

Task 6.3: Adaptive dialogue strategies. (Months 15 – 27) By year 2, the robot will be able to clarify and explain what it does or does not know (WPs 1,4,5), and learn (simple) cat- egories and their associations (WPs 2,5). In this task we will investigate how we can use forms of reinforcement learning to adapt dialogue strategies to optimize planning content for verbalization, clarification requests and explanation on the basis of dynamic (i.e. extending, altering) categorical knowledge [8, 12, 29, 38]. Learning feedback is obtained directly through dialogue (misunderstanding, requiring further clarification) or indirectly (number of turns, subdialogues required till answer) (DFKI 6 months, ALU-FR 4 months)

Task 6.4: Variable granularity content planning. (Months 15 – 27). We will extend content planning techniques to include the use of vagueness to express properties to varying degrees of granularity [28, 10, 22]. (DFKI 4 months, ALU-FR 2 months)

Task 6.5: Adaptive extendable grammatical processing. (Months 27 – 39) By year 3, we have an insight in the dynamics of acquiring categorical knowledge (WP 5), and the architec- ture has a sufficient degree of autonomy to enable self-driven curiosity (WP 1), leading to the acquisition of categorical structures that are based more on the robot's own categorization than on information provided by a tutor. We will focus on adaptive, extendable grammatical processing. Using the combinatory categorial grammar framework [35, 2], we will adapt and extend existing lexico-grammatical knowledge to cover novel categorical knowledge. We will develop methods for learning two types of mappings: a mapping relating a word's lexical meaning to a predicate-argument structure based on the associations of the category this meaning reflects [26, 5, 3, 14], and a mapping relating a word's predicate-argument structure to a syntactic family [1] that can express the structure. (DFKI 5 months)

Task 6.6: Verbalising conceptual structures. Provide a mechanism for verbalising categori- cal and associative structures of combined concepts (cf. Task 5.4). (Months 27–39) We will extend verbalization to cover conceptual structures, focusing on a category and its immediate associations with properties, and other categories. This extends the approach developed earlier in Task 6.1. (DFKI 4 months, ALU-FR 4 months, UL 1 month)

Task 6.7: Adaptive strategies for clarification and explanation. (Months 39 – 50). To- wards the end of the project, the robot's learning is primarily curiosity-driven. This is an advance in that it now actively needs to initiate dialogues, if it wants to interact with other agents. We therefore want to investigate (adaptive strategies for) clarification and explana- tion, more from the engagement-level [34], to address the issue of how to set the context for a clarification request (i.e. scaffolding it), to avoid "out-of-the-blue" behaviour. (DFKI 9 months, ALU-FR 8 months, UL 1 month)

Deliverables:

DR.6.1 Transparency in situated dialogue for interactive learning Report, Prototype. (Month 15) (DFKI, UL)

DR.6.2 Adaptive dialogue strategies supporting transparency Report, Prototype. (Month 27) (DFKI, ALU-FR)

DR.6.3 Adaptive extendable grammatical processing. Report. (Month 39) (DFKI)

DR.6.4 Situated dialogue with adapting levels of vagueness and abstraction. Report, Prototype. (Month 39) (DFKI, ALU-FR)

DR.6.5 Mixed initiative situated dialogue-guided curiosity. Report, Prototype. (Month 48) (DFKI, ALU-FR, UL)

Milestones:

M.6.1 Situated dialogue for transparent, interactive learning. (Month 15) The system, based on incremental situated dialogue processing, will be able to use clarification and explanation in a dialogue with a tutor, to learn more about the environment. To achieve transparency in why the system needs something clarified, the system can verbalize what it does and does not know (categorically).

M.6.2 Adaptive dialogue strategies supporting transparency (Month 27) The system is able to learn how to adapt the ways in which it communicates with a tutor in interactive learning. The system tries to use the situated context to try and find optimal ways in which to verbalize its own knowledge, how to pose clarification questions, and how to explain what it does (not) know. The purpose is to find the optimal amount of information that needs to be communicated (and how) to obtain an answer to a clarification request.

M.6.3 Adaptive extendable grammatical processing. (Month 39) The system will be able to extend its grammatical knowledge on the basis of newly acquired categorical knowledge. The purpose is to be able to verbalize this new knowledge, in the context of an interactive learning dialogue.

M.6.4 Situated dialogue with adapting levels of vagueness and abstraction. (Month 39) The system will be able to adapt how it refers to objects, and groups of objects. It will be able to use varying levels of vagueness (over material and comparative properties), and abstraction (over type) to construct referring expressions which uniquely identify the referent(s) while at same time being only as explicit as needed about properties of an object (in a given context).

M.6.5 Mixed initiative situated dialogue-guided curiosity. (Month 50) The system will be able to initiate and drive situated dialogues for interactive learning based on its own curiosity.

1.5.7 Summary WP 7: Scenario-based integration

Work package number:	**7**	Starting date or starting event:				Month 7
Work package title	Scenario-based integration					
Activity type	RTD					
Participant number	**1**	**3**	**2**	**6**	**5**	**4**
Participant short name	BHAM	KTH	DFKI	TUW	ALU-FR	UL
Person months	**28**	**28**	**24**	**24**	**18**	**12**

Objectives

The overall aim of this WP is to create experimental integrated systems, and to analyse them systematically. The specific objectives are:

- Integration of components from WPs 1,2,3,4,5,6 using the CAS architectural framework to create a series of experimental platforms.
- Specification of representations and interfaces and design processes necessary for creating integrated cognitive systems.
- Testing and demonstration of integrated robot systems.
- Empirical and formal analysis of the robot system, and its behaviour.
- Understanding appropriate methodologies to empirically and formally analyse the behaviour of integrated robot systems.

Description of work:

Task 7.1: Integration for task driven exploration & learning system. (Months 10 – 15) (All partners (19 months total))

Task 7.2: Experimental study of task driven exploration & learning system. (Months 16 – 21) (All partners (19 months total))

Task 7.3: Integration for task driven exploration under uncertainty system version. (Months 22 – 27) (All partners (19 months total))

Task 7.4: Experimental study of task driven exploration under uncertainty system. (Months 28 – 33) (All partners (19 months total))

Task 7.5: Integration for explanation with limited extension system. (Months 34 – 39) (All partners (19 months total))

Task 7.6: Experimental study of explanation with limited extension system. (Months 40 – 44) (All partners (19 months total))

Task 7.7: Integration for full curiosity driven extension system. (Months 45 – 50) (All partners (20 months total))

Deliverables:

DR.7.1 Analysis of a robot that achieves tasks under partial information. Report. (All partners) (Month 21)

DR.7.2 Analysis of a robot that acts under partial information and uncertainty. This will also include a report on methodologies for the analysis of robot cognitive architectures. Report. (All partners) (Month 33)

DR.7.3 Analysis of a robot that explains surprise. Report. (All partners). (Month 44)

DR.7.4 Design methodologies for integrated cognitive systems. Report. (All partners). (Month 44)

DR.7.5 A curiosity driven self-extending robot system. Report. (All partners). (Month 48)

Milestones:

M2 Task driven exploration, and learning. (Month 15).

M4 Task driven exploration under uncertainty. (Month 27).

M6 Explanation with limited extension. (Month 39).

M8 Full curiosity driven extension. (Month 50).

1.5.8 Summary WP 8: Management

Work package number:	**8**	Starting date or starting event:				Month 1
Work package title	Management					
Activity type	MTG					
Participant number	**1**	**3**	**2**	**4**	**5**	**6**
Participant short name	BHAM	KTH	DFKI	UL	ALU-FR	TUW
Person months	**16**	**4**	**2**	**2**	**2**	**2**

Objectives

The objectives of this WP are to:

- To ensure timely exchange of information between workpackages.
- To provide scientific direction for the project.
- To create appropriate forums for collaboration within the project.
- To monitor progress, assess risk and revise project targets where appropriate.

Description of work:

1. Generate 6-monthly management reports for presentation to the General Assembly and the commission, as well as reports in compliance with Annex 2 to the Grant Agreement.

2. Organise bi-weekly meetings for knowledge transfer between work packages.

3. Monitoring of progress on milestones and deliverables.

4. Administration of events, publicity, project office.

1.5.9 Summary WP 9: Dissemination and Community Building

Work package number:	**9**	Starting date or starting event:				Month 1
Work package title	Dissemination and Community Building					
Activity type	OTHER					
Participant number	**1**	**4**	**6**	**3**	**2**	**5**
Participant short name	BHAM	UL	TUW	KTH	DFKI	ALU-FR
Person months	**6**	**6**	**6**	**4**	**2**	**2**

Objectives

The objective of this WP is to build strong connections between the researchers within the consortium, and to maximise the dissemination and thus impact of the scientific and technical achievements of the project.

Description of work: All the following tasks to run months 1-50:

Task 9.1 CogX Open Days: Disseminate knowledge within and outside the consortium through open days for scientists, industrialists and the media.

Task 9.2 CogX Web Site and Intranet: set up a web site and intranet facility to provide the most up to date results on the project to partners, researchers and the general public.

Task 9.3 Summer Schools: annually for up to 45 students and researchers.

Task 9.4 Specialist Workshops: seek external sponsorship and organise in conjunction with major conferences, European and national events.

Task 9.5 Dissemination to General Public: through website, science festivals and mass media.

Task 9.6 Develop Software Toolkit: maintain and document a public-ally available toolkit for prototyping cognitive robotic systems.

Task 9.7 Partner exchange: provide means for exchange of students and researchers.

Task 9.8 Scientific publications: publish papers, with an emphasis on joint work in leading conferences and scientific journals.

Deliverables:

DR.9.1 CogX Website and Intranet. (UL) (Month 1)

DR.9.2 Proceedings of Summer School. (UL) (Month 15)

DR.9.3 Proceedings of Summer School. (KTH) (Month 27)

DR.9.4 Proceedings of Summer School. (TUW) (Month 39)

DR.9.6 Final version of software toolkit. (BHAM) (Month 48)

2 Implementation

2.1 Management structure and procedures

The organisational structures for the CogX project are aimed at ensuring competent project management, both for day-to-day issues, and relative to the long-term project goals. The main elements are listed below. The formal powers and duties of these are specified in the consortium agreement.

- The General Assembly (GA)
- The project coordinator (PC)
- The project administrator who runs the project office (PO)
- The Executive Board (EB)
- The Scientific Advisory Board (SAB)

The General Assembly is the ultimate decision-making body of the consortium, and is composed of one duly authorised representative of each party with an equal voting right. The GA is responsible for matters related to the consortium agreement, budget allocation, and the general direction of the project. The GA will initially be chaired by the DFKI team leader, Geert-Jan Kruijff, DFKI.

The project coordinator (PC; Jeremy Wyatt, BHAM) is the single point of contact between the European Commission (EC) and the Consortium. The PC is responsible for the overall management of the project. He chairs the Executive Board, and prepares the meetings and records the decisions of the General Assembly and the Board. The PC is supported in his duties by the project administrator in the Project Office (PO).

The executive board (EB) is the executive committee of the consortium. The EB supports the PC in fulfilling obligations to the EC. It is responsible for coordinating the various activities for education, training and dissemination.

The Scientific Advisory Board (SAB) will advise the EB, helping in identifying risks and further potentials of the research activity. The SAB will consist of three internationally respected senior scientists from research fields relevant to the project.

2.1.1 Risk Management and Contingency Planning Mechanism

The risks with this project are numerous. We are attempting to make significant progress beyond the state of the art in areas that have consistently made slow progress over the past forty years. We are also promising to integrate the resulting components into working systems that we will use for experimentation. At each stage either progressing the theories and technologies for the components may prove to be harder than we anticipate, or they may prove to be unintegrable. The way we will handle these risks is captured in the structures of the tasks within workpackages. All tasks are synchronised to complete at six month intervals. These coincide with the meetings of our General Assembly, and we will use these mechanisms to internally review at 10, 21, 33 and 44 months whether certain areas of the project have not progressed significantly enough to warrant effort at that point for their integration. In those instances we will roll back to the previous algorithmic and representational solutions for those parts. While this does not ameliorate all risk, we believe it represents a sensible strategy to ensure that progress on the system level is not compromised too badly by failure to progress at the component level.

2.2 Beneficiaries

2.2.1 University of Birmingham (BHAM)

Organisation: Participating in this project, the School of Computer Science has a faculty of 35, and 103 research staff and doctoral students. It has a prominent research group in Artificial Intelligence and Natural Computation of which the Intelligent Robotics Laboratory is a part. This lab houses four senior academics, four research fellows and seven PhD students. It conducts basic research in cognitive robotics, cognitive architectures, object manipulation, decision theoretic planning, statistical machine learning, and automated diagnosis.

Contributions and experience: We have recently developed a new architectural theory called CAS which has been developed to enable rapid prototyping of robots with multiple modes of sensing and action. This will form the basis for work in WP1, and the current open source implementation (CAST) will provide a mechanism for integration in WP7. For WP2, we will bring experience in robotic manipulation from the CoSy project to bear together with background work with psychologists on the theory of modular motor learning. The final area of involvement is WP4 and our relevant expertise is in decision theoretic planning.

Staff: *Dr. Jeremy Wyatt* is a senior lecturer and co-director of the Intelligent Robotics Laboratory. He obtained his PhD in artificial intelligence from the University of Edinburgh (1996) on the topic of active learning and sequential decision-making, and has since published over forty refereed journal and conference papers in robot learning, sequential decision-making under uncertainty, human action recognition, cognitive architectures, and statistical machine learning. He has held grant awards from the Royal Society, the EU, the Leverhulme Trust, the British Council and the Nuffield Foundation.

Dr. Richard Dearden is a senior lecturer and co-director of the Intelligent Robotics Laboratory. He gained his PhD on planning and decision making under uncertainty from the University of British Columbia. He subsequently became Acting Head of the Model-Based Diagnosis and Recovery Group at NASA Ames (2002-2004) where he worked on probabilistic fault detection, and automated daily activity planning for planetary rovers.

Prof. Aaron Sloman's numerous contributions to AI include work on forms of representation, architectures for intelligent systems, vision, varieties of affect, ontology development, learning about causation, robotics, and software tools. Has previously had grants from SRC, Joint research council, the EU, GEC, DERA, EPSRC, BT, IBM, and Leverhulme. Numerous honours include being elected a fellow of AAAI (second wave) 1991, AISB (first wave) 1997, ECCAI (first wave) 1999, Honorary DSc (Sussex 2006).

2.2.2 German Research Center for Artificial Intelligence (DFKI)

Organisation: Founded in 1988, DFKI today is one of the largest nonprofit contract research institutes in the field of innovative software technology based on Artificial Intelligence (AI) methods. DFKI focuses on the complete cycle of innovation, from world-class basic research and technology development through prototypes to product commercialisation. R&D is carried out in several research labs, including Language Technology (LT; Uszkoreit), Intelligent User Interfaces (IUI; Wahlster), and Robotics (RL; Kirchner).

Contributions and experience: Scientifically, DFKI contributes primarily to aspects of human-robot interaction (situated dialogue processing, WP 6), its integration into the overall system (WP 7) and in particular the connection between dialogue processing, spatial cognition (WP 3), learning and categorical understanding (WP 5), and planning (WP 4).

Staff: *Dr.ir. Geert-Jan M. Kruijff* will lead the CogX efforts at DFKI, where he is a senior researcher and project leader in the Language Technology Lab. He holds a PhD in computer science from Charles University in Prague (2001). His research focuses on developing cognitively motivated architectures that model the dialogue capabilities of a robot, and connecting dialogue and its interpretation with models of a robot's embodied experience. He currently leads DFKI's efforts in the EU IP "CoSy" and held leading positions in several national and international projects at Saarland University and Charles University. He has over 80 publications in the fields of computational linguistics, human-robot interaction, situated dialogue processing, abduction, and cognitive architectures, and regularly gives invited lectures worldwide.

Dr. Henrik Jacobsson received a BSc and MSc in computer science from University of Sk¨ovde, Sweden, 1999 and 2000 respectively and a PhD from the Department of Computer Science of the University of Sheffield, UK, 2006. His research focuses on automated analysis and modelling of dynamic systems. Within the CoSy project, his research focuses on cross-modal representations and prelinguistic learning. He has several publications in leading journals (such as Neural Computation).

2.2.3 Royal Institute of Technology (KTH)

Organisation: KTH is the leading technical university in Sweden. The Department of Numerical Analysis and Computer Science has about 60 senior faculty. Its Computational Vision and Active Perception Laboratory (CVAP), performs research in computational vision and robotics, was formed in 1982, has 5 senior faculty, 8 postdoctoral researchers and about 20 research students. CVAP is integrated with the Center for Autonomous Systems (CAS), an interdisciplinary center for research on all aspects of robotics but with a focus on different aspects of service robotics. CAS and CVAP have support from the Swedish Foundation for Strategic Research and The Swedish Research Council. The group has participated in the EU projects CogVis, Insight2+ and VIBES, and is now a partner of MOBVIS, Muscle and Pascal. The laboratory is also involved in the FP6 projects EURON II, Cogniron, Cosy, Neurobotics and PACO-PULS.

Contributions and experience: KTH will contribute primarily to the research on vision and manipulation (WP2), spatial modelling (WP3), and the overall system integration (WP7). The CVAP and CAS research groups have a strong background in robotics and computer vision. A central research theme is the development of artificial seeing agents capable of using vision in its interaction with the environment, for e.g. manoeuvring, navigating, grasping, and recognising things.

Staff: *Dr. Danica Kragic* received a MSc degree in mechanical engineering from the Technical University of Rijeka, Croatia and a PhD degree in computer science from the Royal Institute of Technology (KTH), Stockholm, Sweden in 1995 and 2001, respectively. She is currently an assistant professor in computer science at KTH and chairs the IEEE RAS Committee on Computer and Robot Vision. She received the 2007 IEEE Robotics and Automation Society Early Academic Career Award. Her research interests include computer vision, service robotics and human-robot interaction. She is involved in teaching of graduate and undergraduate courses and supervision of MSc and PhD students.

Dr. Patric Jensfelt received a MSc degree from the School of Engineering Physics at Kungliga Tekniska H¨ogskolan (KTH) in 1996 and the PhD degree from the department of Signal, Sensors and Systems in 2001. After having completed a PhD degree Patric worked as project leader in two industrial projects. One of these was with the Stockholm International Fairs were a robot system put into operation in August 2003 and has been running since then. His current research focuses on navigation, localisation and SLAM as well as systems integration. He is also the co-PI for KTH in the EU project "CoSy".

2.2.4 Univerza v Ljubljani (UL)

Organisation: The Visual Cognitive Systems Laboratory (ViCoS) is a part of the Faculty of Computer and Information Science, which is in turn the leading teaching and research institution in the field of computer science in Slovenia. Currently, the laboratory consists of one professor, one senior researcher, 2 post-docs, and 7 PhD students. It is involved in basic research on artificial cognitive systems, with an emphasis on visual learning, recognition and categorisation.

Contributions and experience: Recently, the members of ViCoS have been developing a general framework for continuous learning of visual concepts by learning associations between automatically extracted visual features and words describing the scene. The research has also focused on learning scalable representations suitable for recognition and detection of a large number of object categories. Within this framework, an approach was developed which learns a hierarchy of spatially flexible compositions in an unsupervised, statistics-driven manner. They will focus on these research topics in CogX as well. Most of their efforts will be devoted to the research and development of continuous learning mechanism that will accommodate different modes of learning of cross-modal concepts.

The members of the laboratory gained a lot of experience in related projects in the past. They have been actively involved in a number of EU funded research projects in FP6 (CoSy, MOBVIS, VISIONTRAIN, euCognition) as well as in FP5 (COGVIS and ECVision) and FP4 (Copernicus Project RECCAD). In addition, they have been participating in a number of national and bilateral (SLO-A, SLO-GR, SLO-CZ, SLO-USA) projects and have established collaboration with the leading research institutions worldwide.

Staff: *Prof. Ales Leonardis* is a full professor and the head of the Visual Cognitive Systems Laboratory with the Faculty of Computer and Information Science, UL. He is also an adjunct professor in the Faculty of Computer Science, Graz University of Technology. He was a visiting researcher, a postdoctoral associate, and a visiting professor in the GRASP Laboratory at the University of Pennsylvania, at PRIP, Vienna University of Technology, and ETH in Zurich, respectively. His research interests include robust and adaptive methods for computer vision, visual learning, and scalable representations for categorisation and recognition of objects. He is a (co)author of more than 130 papers published in journals and conferences and he coauthored the book Segmentation and Recovery of Superquadrics (Kluwer, 2000). He is an associate editor of Pattern Recognition. He has served on the program committees of major computer vision and pattern recognition conferences, and was as a program cochair of the ECCV 2006. He has received several awards. In 2002, he coauthored a paper, ”Multiple Eigenspaces,” which won the 29th Annual Pattern Recognition Society award. In 2004, he was awarded a prestigious national award for his research achievements.

Danijel Skocaj is a senior researcher at the Faculty of Computer and Information Science, University of Ljubljana. He received his PhD degree from the same institution in 2003. His main research interests lie in the field of cognitive vision and include automatic modelling of objects and scenes from visual information with the emphasis on robust, incremental, and interactive visual learning and recognition. These were also the main topics of his research in the CogVis project, and in the on-going CoSy project in which he is a current active participant. He serves as the president of the Slovenian Pattern Recognition Society.

2.2.5 Albert-Ludwigs-Universitat (ALU-FR)

Organisation: The Artificial Intelligence Group at the University of Freiburg Albert-Ludwigs-University consists of the research laboratories for Foundations of Artificial Intelligence, Machine Learning, and Autonomous Intelligent Systems. The laboratory for Foundations of Artificial In-telligence has expertise in the areas of knowledge representation, action planning, reasoning, and cognitive robotics.

Contributions and experience: ALU-FR is one of the world’s leading research groups in AI Planning. Based on this expertise, we will strongly contribute to WP4 both theoretically and practically, in particular by building on our existing planning systems. In the projects CoSy and DESIRE our planning technology has already been applied to dialogue planning: dialogue is re-garded as continual collaborative planning where communicative actions are planned and executed just as physical ones. This work will be extended in WP6 of CogX: The planner will be enabled to reason about the gaps in its own knowledge which can lead to plans for gathering the missing knowledge by means of communication. ALU-FR will also contribute to research on moti-vations and knowledge in WP1, building on our expertise in Knowledge Representation and Reasoning.

Staff: *Prof. Bernhard Nebel* received his PhD from the University of Saarland in 1989. Between 1982 and 1993 he worked on different AI projects at the University of Hamburg, the Technical University of Berlin, ISI/USC, IBM Germany, and the German Research Center for AI (DFKI). From 1993 to 1996 he held an Associate Professor position at the Computer Science Department of the University of Ulm. Since 1996 he is a full professor at University of Freiburg and head of the Laboratory for Foundations of Artificial Intelligence.

Among other professional services, he served as the Program Co-chair for the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92), as the Program Co-chair for the 18th German Annual Conference on AI (KI'94), as the General Chair of the 21st German Annual Conference on Artificial Intelligence (KI'97), and as the Program Chair for the 17th International Joint Conference on Artificial Intelligence (IJCAI'01). In 2008, he will serve as Co-chair of the International Conference on Automated Planning and Scheduling (ICAPS-08). In 2001, Bernhard Nebel was elected as an ECCAI fellow.

2.2.6 Technische Universitat Wien, Automation and Control Institute, Vision for Automation Laboratory (TUW)

Organisation: TUW is represented by ACIN - Automation and Control Institute, which employs 42 persons most of them researchers graduated in the fields of Electrical and Mechanical Engineering, Business Management, Physics, and Computer Science. TUW strongly emphasises the close co-operation with industry.

Contributions and experience: The expertise relevant to the project is the development of cognitive vision techniques specifically suited for robotic tasks to achieve a robot executing task such as "James, please bring me my cup". TUW is committed to provide the detection of structural elements and the grouping of elements into proto-objects (visual entities as candidates for objects) that can be annotated and that specify the shape of the object to enable the attribution of grasping points and linkage to servoing for grasping (contributing to WP2). Results from EU projects robots@home, MOVEMENT and ActIPret, and the national Cognitive Vision project are fundamental for finding ways to integrate multiple cues and modalities for robust perception, to handle occlusions (specifically of hand/object) and model the spatio-temporal relationships and grasp points between entities in three-dimensions (contributing to WP3). Furthermore, the immediate surrounding will be modelled for purposes of collision avoidance, where in cooperation with AMROSE collision-free grasping could be shown.

Staff: *Dr. Markus Vincze:* first degrees from TUW and in 1990 M.Sc. from Rensselaer Polytechnic Institute, USA. He finished his PhD at TUW in 1993. With a grant from the Austrian Academy of Sciences he worked at HelpMate Robotics Inc. and at the Vision Laboratory of Gregory Hager at Yale University. Presently he leads the research group "Vision for Automation" at TUW. With Gregory Hager he edited an issue on Robust Vision for IEEE and is (co-)author of over 100 papers. Markus has coordinated EU projects RobVision and ActIPret, and coordinates the Austrian Cognitive Vision Network. He was key scientist in EU projects FlexPaint, ECVision, FibreScope, MOVEMENT and contributes to the FET-open project XPERO and the Coordination Action euCognition. At present he is coordinating the Advanced Robotics FP6 EU Project robots@home.

2.3 Expertise Across the Consortium as a whole

	BHAM	DFKI	KTH	UL	ALU-FR	TUW
Manipulation	X		X			X
Mapping			X			X
Architectures	X	X			X	
Learning	X	X	X	X		X
Vision			X	X		X
Planning	X				X	
Language		X				
Robotics	X	X	X	X	X	X

2.4 Resources to be committed

The total requested EC contribution to the project budget is of the order of 6.9m euros, of a total project cost of 8.9m that is broadly decomposed as follows. Approximately half of the total project cost (4.3m euros) is the cost of the salaries of the researchers, at an average basic rate of 6k euros per person month for around 60 person years of effort (of which 75% is the EU contribution). This includes 8 person years of effort by a small number of senior researchers on the project. We will also spend two years of effort on management of the research programme, and specialist dissemination and training activities. The total person months are, however, toward the lower end of the range of total effort for large scale integrating project. The non-staff costs associated with the dissemination plan, community building, the purchase of necessary equipment, and management meetings are therefore a significant part of the overall cost (some 900k euros). The remainder of the total project cost is overhead specified under the FP7 rules, and totals 3.5m euros. We deal with each of the non-staff costs in turn.

2.4.1 Resources for Dissemination and Community Building

As detailed in the work package on dissemination and community building we have decided to embark upon an ambitious plan to ensure maximum dissemination and take-up of our results and tools while keeping the overall cost as low as possible, thereby maximising the scientific impact for the resource expended. We have budgeted an average of 35k euros per site for dissemination of the project results through academic conferences, an average of under 3k euros or two international conferences per person year. We have also scheduled three summer schools that will be used for community building within the consortium, but which will also double up to assist in the dissemination of the results of the project to scientists outside the consortium. For this purpose we intend to run long (9 day) summer schools each year, which will be attended by researchers within the project and up to 25 researchers from outside the project. In addition to talks by leading scientists, these summer schools will have 6 days of hands on work with robots. The cost of all the equipment to support these is 60k euros. This will pay for 10 low cost robots with manipulation abilities (Pioneer P3-DX) that will be used both during the year at each site, and at the summer schools. We will bring a further two robots of the same specification to support this work. Employing a common platform has important advantages in that a hands on summer school with a mixture of platforms is not possible, and it will enable simple integration tests to be carried out during the year while partners engage in more advanced research on separate platforms.

The summer schools will have subsidised rates for outsiders to encourage participation. Excluding the equipment the total cost for each nine day long summer school for 45 persons will be 38k including all travel and accommodation for CogX participants. To save costs the summer school will incorporate one meeting of the General Assembly each year.

As another novel dissemination activity we will also run a project open day in conjunction with our second General Assembly meeting each year. This has worked successfully in projects such as iCub. This second meeting will also provide the opportunity for our scientific board to meet and review the project progress. The total cost of supporting the open day, the scientific advisory board travel and the management meetings across the entire project will be 80k euros.

2.4.2 Specialised equipment and Partner Exchange

We have budgeted for an average of under 35k euros per site for specialised research equipment (this figure does not include the pioneer robots for dissemination and basic integration research mentioned above). This includes, for example, the cost of a single three fingered hand for research into multi-fingered grasping (50k euros), upgrading of an existing arm (25k), a couple of mobile platforms able to carry a reliable 6 DOF manipulator (80k), and a stereo head with vergence control to support mobile vision. The consortium will bring a far greater value of existing equipment to the project. This includes five high end mobile platforms such as BlueBotics, B21r or Powerbot; six high end manipulators, and three stereo vision systems. In addition some equipment will be purchased for the project as a partner contribution (up to 40k euros). Thus the equipment resource marshalled relative to the overall EU contribution will be high. The total requested EU contribution to equipment, including the robots to support summer schools and dissemination is about 260k euros. To support research integration we will also engage in significant partner exchange, at a rate of 2k euros per person year of effort, giving a total of just over 120k euros. This degree of travel for an integrated project is necessary to ensure full integration of partners by building solid working relationships between the individual researchers.

3 Impact

3.1 Strategic Impact

In CogX we aim to make significant progress in the science of *how to build complete artificial cognitive systems*. There has been good progress in the past few years in beginning to put back the pieces of AI together into complete cognitive robot systems. Despite this success these systems are still closely tied to their human operators. Mapping robots, for example, typically follow their guides waiting for the person to teach them about where they are. Robots for collaborative manipulation, while they can increasingly understand the connection between what a person says and what they see in the scene, essentially follow instructions to act or to learn, raising only simple queries in limited settings.

The project answers both requirements specified as the call's target outcomes[2]. It considers the issue of "achieving general goals" at several levels and addresses the issue of interaction between humans and robots through dialogue.

3.2 Dissemination and/or exploitation of project results, and management of intellectual property

3.2.1 Dissemination and exploitation

CogX is primarily a scientific research project rather than a commercial development project. So there will be no direct commercial exploitation of results during the course of the project. However, we will still take as many reasonable steps as possible to disseminate the results to potential exploiters. The channels we have chosen will ensure that the impact of the project not only during but after it has finished is as great as possible. Our dissemination plan uses a couple of innovative mechanisms in addition to the established ones for research projects. In addition, to be efficient we have combined some dissemination activities with community building within the consortium. The activities are:

- Scientific publications in conferences and journals: this will be the mainstay of the dissemination of our scientific results to the research community. The emphasis will be on joint publications in highly reputed journals and conference proceedings.

- Project website and intranet: we will create and regularly maintain a public website, with sections for researchers and the general public. This will be regularly updated with papers from each partner, with electronic copies of breaking work posted on the website prior to publication elsewhere. We will place particular stress on communicating our results in an understandable manner to a non-specialist audience. The project intranet will be used to disseminate work within the consortium. We will maintain shared code repositories which will be open source, and available for public use within the terms of the IPR agreement.

- Open days: in conjunction with one of our General Assembly meetings each year we will hold a project open day, with a poster evening and demonstrations. The scientific advisory board, researchers on other projects, industrialists and general members of the scientific community will be invited. We will also, where appropriate use these events to disseminate results through the mass media.

- Hands on summer schools: A central plank of our dissemination activity and of the community building work within the consortium will be our annual summer schools. These will last for 8-9 days each and will combine talks from the leading figures in relevant fields with hands on work (supported by tutorials) using a suite of 12 relatively low cost, flexible mobile platforms (P3-DX Pioneer robots). The hands on work will be partially supported by an open source toolkit for prototyping of cognitive robotics systems that we will develop as part of the integration

[2] ftp://ftp.cordis.lu/pub/fp7/ict/docs/ict-wp-2007-08_en.pdf

work package. The aim is to encourage widespread take up of the toolkit as a research tool within the community. The summer schools will thus also be an excellent field test of the toolkit, and we will revise it according to the feedback received.

- Exchange of students: this is the essential mechanism for community building within the project to ensure really integrated joint research effort, and also that researchers from the different disciplines within the project (AI, vision, robotics, linguistics) understand each other's research agenda. It is also the basis on which joint publications are laid.
- Software toolkit for researchers: The software toolkit used for integration within the project, and used at the summer schools will be released in open source form, and will be freely available. To ensure utility we will keep the publicly available toolkit lightweight and robust to ensure acceptance.
- Presentations and written material for the general public: apart from the open days we will use a variety of mechanisms to disseminate results through the mass media and through giving talks and demonstrations at science festivals.
- Specialist Workshops: we will organise specialist research workshops, co-located with major conferences, such as IJCAI, ECCV, CVPR, ICCV, to name but a few.
- Community building at large: we are committed to contributing to community building activities at large through collaborations with other projects in the Cognitive Systems, Robotics, Interaction domains, funded by the EC, and through participation in joint events, and activities (including roadmapping) such as those organised by the "Concertation Action" euCognition, under FP6.

3.2.2 Management of intellectual property

All partners in the consortium are committed to making the knowledge generated during the course of the project as widely and as freely available as possible. Our intention is that all results will be of a form that should be made freely available for subsequent research and development, subject to the consortium agreement on IPR. Thus we aim to fully implement the public domain recommendation as per the Call Background notes[3]. For this reason all of our research papers, technical reports and software will be placed in the public domain. Where appropriate, we will license the products of our research with open-source licenses such as the GPL or LGPL licences. This supports our aim of facilitating access to the results of our research for the largest possible base of users (particularly developers of intelligent robots), whilst retaining appropriate control of IPR. We shall attempt to avoid reliance on proprietary software, and where such reliance is essential and it is possible we will design our systems so as to allow subsequent reimplementation using non-proprietary infrastructure.

[3] ftp://ftp.cordis.europa.eu/pub/ist/docs/cognition/fp7-challenge2-background_en.pdf

References

[1] J. Baldridge. Lexically Specified Derivational Control in Combinatory Categorial Grammar. PhD thesis, University of Edinburgh, 2002.

[2] J. Baldridge and G.-J. M. Kruijff. Multi-modal combinatory categorial grammmar. In Proceedings of EACL'03, Budapest, Hungary, 2003.

[3] L. Barsalou. Perceptual symbol systems. Behavioral and Brain Sciences, 22:577–609, 1999.

[4] L. Beaudoin. Goal processing in autonomous agents. PhD thesis, School of Computer Science, The University of Birmingham, 1994. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#38.

[5] P. Bloom. How Children Learn The Meaning of Words. The MIT Press, 2002.

[6] R. P. Bonasso, D. Kortenkamp, D. P. Miller, and M. G. Slack. Experiences with an architecture for intelligent reactive agents. In Proc. of the Int. Joint Conf. on Artificial Intelligence, 1995.

[7] M. Brenner and B. Nebel. Continual planning and acting in dynamic multiagent environments. In Proceedings of the International Symposium on Practical Cognitive Agents and Robots, Perth, Australia, 2006.

[8] H. Cuayahuitl, S. Renals, O. Lemon, and H. Shimodaira. Reinforcement learning of dialogue strategies with hierarchical abstract machines. In Proceedings of IEEE/ACL Spoken Language Technology 2006, 2006.

[9] R. Dale and E. Reiter. Computational interpretations of the Gricean maxims in the generation of referring expressions. Cognitive Science, 19(2):233–263, 1995.

[10] D. DeVault and M. Stone. Interpreting vague utterances in context. In Proceedings of the 20th International Conference on Computational Linguistics (COLING'04), 2004.

[11] S. Ekvall, D. Kragic, and P. Jensfelt. Object detection and mapping for service robot tasks. Robotica: International Journal of Information, Education and Research in Robotics and Artificial Intelligence, 25(2):175–187, March/April 2007.

[12] M. Frampton and O. Lemon. Learning more effective dialogue strategies using limited dialogue move features. In Proceedings of ACL 2006, 2006.

[13] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J. Fern´andez-Madrigal, and J. Gonz´alez. Multi-hierarchical semantic maps for mobile robotics. In Proc. of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems (IROS), Edmonton, Alberta, Canada, 2005.

[14] A. M. Glenberg. What memory is for. Behavioral & Brain Sciences, 20:1–55, 1997.

[15] B. Grosz and L. Hunsberger. The dynamics of intention in collaborative activity. Cognitive Systems Research, 7(2-3):259–272, 2006.

[16] B. Grosz and S. Kraus. The evolution of sharedplans. In A. Rao and M. Wooldridge, editors, Foundations and Theories of Rational Agency, pages 227–262. Springer, 1999.

[17] M. Haruno, D. Wolpert, and M. Kawato. Mosaic model for sensorimotor learning and control. Neural Computation, 13(10):2201–2220, 2001.

[18] N. Hawes, A. Sloman, J. Wyatt, M. Zillich, H. Jacobsson, G. Kruijff, M. Brenner, G. Berginc, and D. Skocaj. Towards an integrated robot with multiple cognitive functions. In Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07), 2007.

[19] N. Hawes, J. Wyatt, and A. Sloman. An architecture schema for embodied cognitive systems. Technical Report CSR-06-12, University of Birmingham, School of Computer Science, November 2006.

[20] H. Horacek. An algorithm for generating referential descriptions with flexible interfaces. In Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics, Madrid, 1997.

[21] J. Kelleher and G. Kruijff. A context-dependent algorithm for generating locative expressions in physically situated environments. In Proceedings of ACL/COLING 2006, 2006.

[22] C. Kennedy. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. Linguistics and Philosophy, to appear.

[23] E. Krahmer and M. Theune. Efficient context-sensitive generation of referring expressions. In K. van Deemter and R. Kibble, editors, Information Sharing: Reference and Presupposition in Language Generation and Interpretation. CLSI Publications, Standford, 2002.

[24] G. Kruijff, H. Zender, P. Jensfelt, and H. Christensen. Clarification dialogues in human-augmented mapping. In Proceedings of the 1st Annual Conference on Human-Robot Interaction (HRI'06), 2006.

[25] G. Kruijff, H. Zender, P. Jensfelt, and H. Christensen. Situated dialogue and spatial organization: What, where... and why? International Journal of Advanced Robotic Systems, 4(2):125–138, 2007.

[26] G.-J. M. Kruijff. A Categorial-Modal Logical Architecture of Informativity: Dependency Grammar Logic & Information Structure. PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic, April 2001.

[27] G.-J. M. Kruijff, J. D. Kelleher, and N. Hawes. Information fusion for visual reference resolution in dynamic situated dialogue. In E. Andre, L. Dybkjaer, W. Minker, H. Neumann, and M. Weber, editors, Perception and Interactive Technologies: International Tutorial and Research Workshop, PIT 2006, volume 4021 of Lecture Notes in Computer Science, pages 117 – 128, Kloster Irsee, Germany, June 2006. Springer Berlin / Heidelberg.

[28] A. Kyburg and M. Morreau. Fitting words: Vague words in context. Linguistics and Philosophy, 23(6):577–597, 2000.

[29] O. Lemon, K. Georgila, and J. Henderson. Evaluating effectiveness and portability of reinforcement learned dialogue strategies with real users: the talk towninfo evaluation. In Proceedings of IEEE/ACL Spoken Language Technology, 2006.

[30] K. Lochbaum, B. Grosz, and C. Sidner. Discourse structure and intention recognition. In R. Dale, H. Moisl, , and H. Somers, editors, A Handbook of Natural Language Processing: Techniques and Applications for the Processing of Language as Text. Marcel Dekker, New York, 1999.

[31] B. Logan. A design study for an attention filter penetration architecture. In Proceedings of the AISB'00 Symposium on How to Design a Functioning Mind, pages 94–101. AISB, The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, April 2000.

[32] M. Purver. The Theory and Use of Clarification Requests in Dialogue. PhD thesis, King's College, University of London, 2004.

[33] S. J. Russell and E. Wefald. Do the Right Thing: Studies in Limited Rationality. MIT Press, Cambridge, 1991.

[34] C. L. Sidner, C. Lee, C. Kidd, N. Lesh, and C. Rich. Explorations in engagement for humans and robots. Artificial Intelligence, 166(1-2):140–164, 2005.

[35] M. Steedman. The Syntactic Process. The MIT Press, Cambridge MA, 2000.

[36] A. L. Thomaz. Socially Guided Machine Learning. PhD thesis, Massachusetts Institute of Technology, May 2006.

[37] A. L. Thomaz and C. Breazeal. Transparency and socially guided machine learning. In Proceedings of the 5th International Conference on Developmental Learning (ICDL), 2006.

[38] S. Wilske and G. Kruijff. Service robots dealing with indirect speech acts. In Proceedings of IROS 2006, 2006.

[39] D. Wolpert and M. Kawato. Multiple paired forward and inverse models for motor control. Neural Networks, 11(7):1317–1329, 1998.

[40] I. Wright. Emotional agents. PhD thesis, School of Computer Science, The University of Birmingham, 1977. http://www.cs.bham.ac.uk/research/cogaff/.

[41] I. Wright, A. Sloman, and L. Beaudoin. Towards a design-based analysis of emotional episodes. Philosophy Psychiatry and Psychology, 3(2):101–126, 1996. Repr. in R.L.Chrisley (Ed.), Artificial Intelligence: Critical Concepts in Cognitive Science, Vol IV, Routledge, London, 2000.