Information Society Technologies

FET Proactive initiative

Beyond Robotics

COGNIRON
Integrated project

FP6-IST-002020

The Cognitive Robot Companion

Review Report N°: 2

Covering period 01.01.2005 - 31.12.2005

Version 1.1

1 April 2006

Report Preparation Date:	1 April.2006		
Classification:			
Contract Start Date:	01.01.2004	Duration:	48 months
Project Coordinator:	Raja Chatila, LAAS-CNRS, Toulouse, France		

Place of the Review:

54 Rue Joseph II
Bruxelles, Belgique

Date of review:

14 February 2006

Reviewers:

Prof. Ruth Aylett
Heriot-Watt University, UK

Prof. James L. Crowley
Institut National Polytechnique de Grenoble, FR

Prof. Ian Horswill
Northwestern University, USA

Prof. Anthony Cohn
University of Leeds, UK

European Commission

Project officers

Pekka Karp **Philippe Reynaert**

1 OVERALL ASSESSMENT OF THE PROJECT

Three things set the Cogniron project apart from other projects in cognitive robotics:

1) The "scope" of the integration effort.
2) The grounding of semantics by a robot interacting with its environment and with humans, and
3) The companion aspect of the project.

The consortium should build on these strengths, and seek to become a visible success by redefining the state of the art in cognitive robotics.

The scope of the integration effort requires a serious, professional effort in integration. This review showed that the consortium had made progress in the area, but that more progress remains to be made. The effort towards defining a "conceptual architecture" is fine, but this should not be confused with a cognitive architecture, nor, more importantly, a software systems architecture. Systems architectures play an important role by allowing interoperability of components, and direct, *in situ*, comparison of competing techniques. The project can perform an important contribution to the technology of autonomous robotics by developing methods and architectures that facilitate systems integration.

The consortium is treading on new ground in this project, so it is not unusual to see emerging separate systems architectures for each of the Key Experiments. This is fine, but these must be accompanied by serious efforts in documentation, analysis and evaluation. At the end of the project it must be possible to predict the strengths and weaknesses of different systems architectures, as well as to apply them to other problems. The XML schema defined here for communication between modules (software services) is a good start. Publish these schema, and promote them as the basis for industry standard protocols.

Grounding semantics through interaction with humans and with the world is an important scientific challenge that the Cogniron consortium is well equipped to meet. The European Commission has given the consortium a unique opportunity here. We look forward to real progress here.

The companion aspect of the project also raises important scientific challenges. There is more to social interaction than avoiding annoying users. The consortium should meet these challenges by more fully exploiting social interaction for learning in the key experiments and considering the issues of proactivity and internal goal management much more seriously within them. Humans play roles when they interact. What roles can robots play? Can these enhance effectiveness of interaction with a human tutor? We look forward to seeing real progress being made here.

The consortium has obtained a high-profile, privileged situation and it is expected that they expand the state of the art in the area of cognitive robotics. The state of the art is defined in terms of objective tasks and objective performance evaluation. It is imperative to define such tasks and evaluation metrics and use them as a scale to define the state of the art. Set the target that the others have to meet.

2 PROJECT ACHIEVEMENTS and FUTURE PLANS

2.1 Work carried out in the previous reporting period

2.1.1 Overall Comments

The objectives of the Cogniron project are

- To achieve significant progress toward understanding and building artificial embodied cognition, and
- To understand and design Human-robot interaction schemes and methods for a robot ‘companion’.

Overall, the Cogniron project has continued to make good progress towards these objectives during the period of Month 12 to 23. However, the quality of the work has been somewhat uneven during this period.

In order to focus their efforts, the consortium has identified 4 sub-areas (identified as "scientific questions"). These areas are

1) Perception and Understanding
2) Decision making
3) Learning
4) Communication and Interaction.

These areas are, in some cases, developed in parallel. Results are brought together in the context of three "Key Experiments". These are:

KE1: Robot Home Tour
KE2: Curious Robot
KE3: Learning Skills and Tasks

In the area of <u>perception</u>, during the period, the consortium has defined concepts for

a) Representing space in terms of regions and places
b) Object-Centered Representations
c) Models for interaction
d) Human body perception.

The presentations and reports indicate that a solid conceptual foundation has been established for these problems. However, it was not apparent how these concepts will translate into objective, measurable performance. This should be addressed in the upcoming year.

In the area of <u>Communications</u>, effort has been directed to the question of adapting to human presence, to understand human behaviour, and to improving interaction. This includes the use of flexible dialogue structures; interaction-related states; dialogue style; object and location references to ground human-robot shared representations; interpretation of gestures, postures and activities; and integration of theories of human social spaces and in the process of motion planning in presence of humans.

In the area of <u>learning</u> the consortium has addressed learning manipulation skills and tasks by demonstration. They have proposed a generic framework for inferring goals and representations. They have examined the use of social context for solving the correspondence problem in imitation learning.

In the area of Decision Making, the consortium has focused on defining a "conceptual architecture", on developing systems to support joint plans and intentions, and systems for safe operation in the presence of humans. However, it is not clear what the cognitive architecture really represents. What are the meanings of the boxes and arrows? What specifically does this architecture explain or predict?

It is not clear how the conceptual architecture can take the consortium from a conceptual framework to a cognitive system that can learn to act and understand through interaction with humans and with the physical world.

With regard to progress on the key experiments, the consortium reports:

a) Design of scenarios for the key experiments
b) Integration of results and software into a few demonstrations
c) Mapping of the conceptual architecture into the KE software architectures
d) Implementation and integration tools and sharing of software.

The reviewers were impressed with progress made in KE1: The Robot Home Tour. The experiment showed that an integrated system had been achieved and was being used to pose and explore questions in multi-modal interaction and spatial orientation.

The reviewers were concerned with the lack of integration or coherent scenarios for KE2: Curious Robot. The current plans are to produce an integrated system at month 36 of the project, essentially too late for posing and answering many of the important scientific and technological questions. In addition, the work seemed to focus too much on geometric aspects of path planning and not enough on behaviours for acquiring information about the environment and responding dynamically to it. In particular, it is not clear why it is necessary to seek an articulated kinematic model of humans for this task. This is a hard problem in itself, and solution would appear to be only marginally useful in the curious robot scenarios.

The reviewers noted that some progress had been made in KE3: Learning skills and tasks. However, we were concerned that for both KE2 and KE3, there was a significant divergence between the experimental designs developed under WP7.1 and reported in appendix E of D7, and the actual experimental designs being used by the groups developing KE2 and KE3. In both cases, the experiments had been split into smaller experiments with different designs. And in both cases, the groups implementing the experiments seemed unaware during the presentation of the appendix E designs. There was also concern that the KE3 scenarios could largely be accomplished without the learning technology being developed in RA4, and so weren't well suited to evaluating that technology. This was apparently due to designing the KEs around what was already integrated, with the intention of incrementally modifying the KEs and more technology is integrated. We are concerned that this may not provide sufficient guidance to the implementers of component technologies to ensure they can be properly integrated in phase 4. More generally, we felt that a more serious effort was required in integration, and well as in documenting experimental design and evaluation.

The consortium should use the KEs to define objective performance metrics related to the KE Tasks (Robot Home Tour, Curious Robot and Skill Learning). Performance metrics define the state of the art. Once these are defined, it is possible to objectively measure progress and position with respect to competition.

It is recommended that the consortium use KE1 as a model for KE2 and KE3. The reviewers expect to see the same degree of integration, allowing the system to be used for experiments in robot behaviour and learning rather than simply the focus for systems building.

2.1.2 Response to recommendations from the previous year.

Overall, the reviewers were satisfied with the response to the recommendations from the previous review. Two formal recommendations were given

1) Reinforce work on both cognitive and systems architectures.
2) Consolidate the deliverables.

Concerning the first recommendation, the consortium has increased in manpower in this area. They have increased the number of meetings and workshops devoted to this problem. However, as noted above, much of this effort has been focussed on defining a "Conceptual architecture" relating perception/interpretation, learning, decision-making and interaction. What is needed is additional work on software architectures, integration tools and development methods, particularly for the KEs. In addition, a cognitive architecture should be defined to describe the core cognitive capabilities of the system, and how these fit together and interact to provide capabilities for tasks and social interaction.

Concerning the consolidation of deliverables, the consortium has satisfactorily implemented this recommendation, and the results are a set of deliverables that, while still large, allow more resources to be devoted to fundamental issues rather than superficial formalities.

In addition, the reviewers gave the following recommendations.

1) Be vigilant in avoiding the inappropriate use of anthropomorphic terms such as “companion” and “intention”. The consortium has abided by this. The establishment of a common glossary is an excellent effort. We encourage publication or dissemination of this glossary.

2) Be very explicit about the problems for which the consortium seeks to advance the state of the art. Use these planned innovations to define the experiments and the requirements for technology development. Use the requirements to define the technological work in the remaining areas.

The consortium reports that this has been addressed in workplans for phases 2 and 3. Each RA now clearly states its contribution to the state of the art. The inter-relation between the research areas and key experiments has been clarified. Nonetheless, this recommendation remains important especially as it was not really addressed in most of the review reports, which concentrated on what had been done without clearly evaluating the significance of the work against the state of the art. The consortium remains excessively concerned with developing component technologies in isolation, and insufficiently driven by performance or by the requirement for integrated systems in the Key Experiments.

3) Adopt professional methods for integration. Proper project management requires clear specification for components. Reliable integration and proper performance-evaluation experiments require component interoperability. This can only be achieved if components are assembled according to clearly defined architectural principles, using well defined component interfaces.

This recommendation remains in force. The component interoperability required for project success can only be achieved if components are assembled using well-defined component interfaces.

4) Reduce the number of deliverables. As noted above, this has been implemented.

5) Improve the clarity and details of the implementation plan. It is not clear if this has been achieved.

2.1.3 Comments on work within research areas.

2.1.3.1 RA1: multi-modal dialogue

Work during this period has addressed questions of representation and development-evaluation cycles as well as continued work in the areas of modality integration, evaluation, and enhancing dialog capabilities. During the second phase from M12 - M30, this Research Area has been composed of three work packages:

WP1.1 Adaptive multi-modal dialogue
WP1.2 Representation and integration of knowledge for an embodied multi-modal dialogue system
WP1.3 Evaluation methods

The goal in WP1.1 is to achieve natural, adaptive multi-modal dialog. The consortium seeks to achieve dialog with a flexible structure that allows mixed initiative and is independent of dialog domain.

The current state of the art approach is to structure dialog with a network of task related states. The partners UniBi, KTH and UH have explored an alternative approach based on explicit grounding. They use grounding to define a network of interaction-related states that incorporate local exchanges (e.g. presentation - acceptance) with global discourse (e.g. relation between exchanges). They have proposed a simple implementation mechanism based on an augmented push-down automaton.

An important criterion for success is to propose a metric for dialog success. The consortium has proposed such a metric based on the number of support exchanges that are achieved in the dialog. The definition of such metrics is an important advance for the scientific community. We regret not to have seen the use of this metric at the review or in the written report. We encourage the consortium to record such metrics and use them to define progress in dialog.

In the period M24 - M30, the consortium proposes to develop a miscommunication measurement using multi-modal cues, to develop user-type dependent dialog and adaptation strategies, to incorporate user models, to define cues for determining user type, and to incorporate explicit and implicit prompting and initiation of interaction.

The research goals for WP1.2 are to achieve interactive resolution of object and location references, and to verbally connect specified information to environment representation. The approach has been to develop a multi-modal representation for dialog, including Multi-modal grounding as well as an interface with the object resolution ability defined in CF-FOR, and to adopt a shared representation of referenced objects in scene model. The consortium has followed three essential lines of work:
1) the design of a mechanism to handle multi-modal input,
2) the interaction with an object attention system, and
3) a common representation for objects for both the dialog system and the object attention system.

The dialog system represents communications as Interaction Units (IU). Completion of this structure drives interpretation, and partial failure can trigger dialogs for recovery. This is a very powerful concept that should find use outside the project. The object-attention system receives queries from the dialog system and searches the current scene with the help of a gesture recognizer. If successful it assigns ID, image, name and other attributes to the gesture. If not successful, it directs the dialog system to use a recovery dialog. Object representation includes mechanisms for converting between modalities. This appears to be a powerful system that goes beyond the existing state of the art in dialog control.

WP1.3 concerns evaluation methods. For this, the consortium has continued development and annotation of a corpus. This corpus has been used to define 20 types of miscommunication and to examine their implications for dialog design. Although little was said about metrics, it would appear that this taxonomy would provide a good basis for defining objective measures of dialog performance.

In the next period the consortium proposes:
1) to use situation awareness for dialog
2) to develop interfaces with task planning
3) to develop methods to use monitoring for interpretation of communication situation to initiate interaction
4) to extend dialog models to include spatial representation

Overall, progress in this workpackage is excellent with the development of innovative concepts and techniques that are likely to advance the state of the art in dialog modelling and have many uses beyond the Cogniron project.

2.1.3.2 RA2: Detection & understanding of human activities

In RA2 the consortium seeks to develop techniques to detect and classify human activities based on observation and analysis of human motions using an enhanced articulated body model. In the laboratory, they observe and analyze human motions using an 85 Degree of Freedom (DoF) articulated body model that includes visual features and motion dynamics. This model is to be used to track and recognize human activity using a Vicon motion capture system employing 10 cameras and bio-mechanical markers. For the configuration embedded in the robot, they propose to reduce this to a 25 DoF articulated model that is driven from data from a stereo pair of cameras and a time of flight range sensor. In this period they report achieving:

- Articulated person tracking on board the robot
- Fusion of different tracking modalities for posture estimation

They argue that this will provide a basis for activity recognition, as well as a tool for resolving object references.

The activity is organized into three workpackages:
WP2.1 Detection and perception of body parts based on sensor features
WP2.2 Human body model: Integration and fusion
WP2.3 Context-based interpretation and classification of activities

WP2.1 is concerned with detection and perception of body parts. Four partners contribute to this research area:

- LAAS: With 2D Vision tracking and short and intermediate range observation of human limb configuration using a particle filter.

- University of Amsterdam: With 2D Vision tracking, coarse, long distance human tracking, and re-identification of humans.
- IPA: with 3D Sensor calibration, tracking and model initialization, and
- University of Karlsruhe: with 3D Time of Flight tracking, and articulated human model estimation using the Iterative Closest Point (ICP) algorithm.

Three specific 2D-tracking strategies have been investigated: 1) long-range tracking, 2) Intermediate-range tracking, and 3) short-range tracking in an active interaction context. Evaluation has been performed to measure identification error, failure rate and processing time under variations of illumination and appearance changes, occlusions and discontinuity in the temporal sequence. The complete results of this evaluation will be reported in the PhD thesis of Ludovic Brèthes. An excerpt of some results is presented as an appendix to the deliverable report for RA2. This is somewhat frustrating, as what we would expect in a deliverable report is that the experimental results be presented and discussed. In particular, the appendix lists graphs of results comparing 6 different techniques without any discussion of how the metrics were defined or the exact experimental conditions under which the measurements were made. The reader is left to draw his own conclusions about which techniques are most appropriate for the project. One supposes that this information is available, (no doubt in French) in the thesis and the effort would have been excessive to translate and present it to the consortium. However, this would indicate that perhaps this work was not really important to the Cogniron Project.

For the 3D human tracking, the work addressed in the past year concentrated on the implementation of fast and robust versions of existing techniques. The goal was to apply an Iterative Closest Point (ICP) algorithm to track a person at a 10 Hz frame rate, while accommodating natural limb and body movements. This method scales linearly with the number of joints, and does not require any background knowledge apart from kinematic constraints. However, with this approach, tracking can succeed only if the input data contains all necessary information to determine the human's posture, and no tracking hypothesis can be generated for temporarily invisible body parts or ambiguous configurations. Evaluation results indicate that only in 5% of the sequences the tracking had lost the subject completely and for 32% minor temporal deviations had occurred. However, the report gave only superficial discussion of the evaluation method as well as the result's consequences for the Cogniron Project.

WP2.2 is concerned with integration and fusion for maintaining an articulated human body model. This work integrates 3D tracking methods using range data from a ToF camera and a SICK scanner with 2D methods based on image appearance. The approach is to use an extension of the ICP algorithm. An ICP-based fusion and tracking technique named VooDoo was developed. This technique uses a new approach for fusion of different input sensors and cues for tracking a 3D articulated human body model using information from 3D sensors including a Time-of-Flight-cameras (ToF), and a stereo reconstruction technique, together with cues from 2D trackers using a single camera. The system is designed to work solely with sensors on-board the robot. It can perform tracking and fusion in real-time at about 10-14 Hz. This would appear to be the same system described in WP2.1.

Results include definition of importance and likelihood functions based on "shape", colour, motion and face detection. It is not clear how shape, colour or motion are defined. Evaluation on a set of images taken from the robot was used to evaluate discriminative power, precision, and computation time.

Much of the work on 3D articulated tracking and motion capture in WP2.1 and WP2.2 does not seem to be driven by the needs of Cogniron. In particular, the consortium has not made a case for why 3D articulated motion capture is necessary for Cogniron, or how it will be used. Given the great difficulty of such techniques, one can wonder if it is necessary for the consortium to invest in this research.

WP2.3 concerns context-based interpretation and classification of activities. The objectives during the second year stressed four aspects of human activity interpretation and classification
This includes work on

- Classification of gestures and activities (IPA, UH, UniKarl)
- 2D Communicative and commanding gestures (LAAS, IPA)
- Activity analysis and recognition (UniKarl)
- Focus of attention and dereferencing objects (UniBi), and
- Rule based activity interpretation (All partners).

For classification of gestures a methodology for identifying the intent of human gestures is proposed as the basis for understanding gestures in the context of human-robot interaction. Five gesture classes are proposed. For each class, the defining characteristics and associated intents are given.

Work on 2D communicative gesture interpretation has been devoted to the classification of hand configurations and of fronto-parallel motions in the video stream. A mixed-state particle-filter algorithm has been proposed to detect the most likely hand posture and canonical motion model, thus ensuring an automatic switch between multiple templates/motions in the tracking loop. The aim is to get a compact representation of a gesture through a sequencing of hand configurations and motions.

Activity analysis and recognition is based on the hybrid activity hierarchy. This approach is based on the assumption that each activity can be modelled by a set of characteristic features, which describe motion primitives of body parts, relations between body parts and objects, or relations between the motion/configuration and the observer. The activity recognition was performed by Feed-Forward Neural Network classifiers.

For focus of attention and dereferencing objects, an Object Attention System (OAS) has been developed that can resolve multi-modal references to objects. The recognition results from a pointing gesture are combined with information from the dialog in order to acquire visual information about the corresponding object.

For preliminary testing and evaluation of the activity recognition system, a test set of 11 activities were recorded. This list includes: balance, bow, call, clap, flap, handshake, kick, manipulate, sit, walk, and wave. For each activity, 10 example sequences were recorded for one male and one female subject. This gave a total of 220 sequences with together 21222 frames. 50% of the recorded data were used for training, the other 50% for testing.

For evaluation, the main focus was laid on how the different feature-selection mechanisms influenced recognition rates. Results indicate that ambiguities or missing input during human-robot interaction can be resolved by the system. Pointing to unknown objects can be used for learning the object. The experiments also raised a number of problems for the next period.

Overall, RA2 has produced interesting results for Cogniron and it is possible to say that they have achieved the target milestone of classification and first methods for recognition of composite human activities for human-robot interaction. While the interest of some of the work related to motion capture in WP2.1 and WP2.2 is not clear, the work in WP2.3 has provided important sensor modalities for interaction.

2.1.3.3 RA3: Social behaviour & embodied interaction

Research Area RA3 is organized in four workpackages:
WP3.1 Personal spaces and social rules in human-robot interaction
WP3.2 Posture and Positioning
WP3.3 Recognition of user intent
WP3.4 Models and algorithms for motion in presence and in the vicinity of humans

Three partners contribute most of the effort to this WP: LAAS are involved in the development of a motion planner taking into account human sensibilities of robot motion speed and direction; UH and KTH have both carried out trials investigating human-robot social spaces, robot behaviour styles and posture and positioning in HRI. Small amounts of effort from UvA and UniKarl seem to be related to liaison with other RAs. This WP was impacted by the withdrawal of VUB from the consortium: overall effort is much the same at 66pm in the second 18 month period as in the first (64pm) but the VUB effort has largely been redistributed to KTH (up from 8pm in the first 18months to 20 pm in the second 18 months) and to UH (up from 15pm in the first 18 months to 29 pm in the second 18 months). In addition EPFL and IPA have both withdrawn from activity in WP3 in this period.

WP3.1 seeks to develop social rules for implementing socially acceptable, non-verbal behaviour in robots. It is particularly concerned with social distance in the context of interactions in the key experiments. This is led by UH.

Data collected during 3 months of 2004 were analysed and documented during months M13-M18. Based on these results, new studies have been designed and conducted in the second project phase, involving in-depth user studies, with the aim of identifying the initial set of “social rules” for non-verbal robot behaviour related to social spaces using a less institutional space than in the first 12 months. There have also been experiments with the use of video rather than live robot data in relation to shortening the time required for and reducing the complexity of experiments.

Results were compared with other user studies conducted as part of the Home Tour Scenario (e.g. studies by KTH and Bielefeld in RA1 and RA3) in order to arrive at a common understanding of the role of social spaces in human-robot interaction in the specific scenarios addressed. Dissemination of the results of this WP has been particularly strong, with a number of papers published or accepted. Rules relating to human comfort have been derived and used in the motion planner of WP3.3. It is perhaps not so clear however whether this WP can move beyond the issue of not annoying users to more collaborative or ‘companionable’ movement.

WP3.2 is led by KTH and addresses the role of posture and positioning in task-oriented robot-human interaction. As with WP3.1, this work concerns use of space for socially acceptable interaction. The goal is to perform user studies using real, implemented robot navigation capabilities rather than simulated, tele-operated robot movements as during the first year. Work has been closely related to KE1 and the Home Tour and extends the static human position of WP3.1 to a situation in which both robot and human are moving. The results could have been more clearly reported in the RA3 deliverable, which talks about Hall’s ‘personal zone’ and Kendon F-formations without clearly explaining what these are or why they are important for the project. Work over the 12 months has involved quantitative analysis of the previous year’s experiment plus a small 5-subject pilot of a multi-room version of that rather simple study: this does not seem such a lot compared to the WP3.1 experiments and in relation to the KTH effort allocation to WP3. Unlike WP3.1 there seems to have been much less dissemination of this work and the two papers mentioned as due out in 2006 are not presented in the appendix so that there is actually no way of clarifying these issues.

WP3.3 studies models and algorithms for motion in the presence of humans. During the period, a navigation planner was finalised taking into consideration results of user studies performed in WP3.1, though this has so far only been run in simulation. The modification of motion planning to take social spaces into account is a useful result, but what may be missing here is attention to adaptive movement, for example if the generic assumptions about personal comfort zones need to be modified in individual cases. It is not clear either what performance metrics will be used when the planner is ported onto a robot in which modelling errors and dynamic sensor-processing may impact the real-world accuracy of the motion plans demonstrated in simulation and thus require such adaptive modification. Integration issues with RAs 2, 5 and 6 will also have to be addressed and are not mentioned at all yet.

WP3.4 was added in response to ’reviewer’s comments at the first year review. It investigates gestures to facilitate communication and facilitation between human and robot, and has carried out experimental work examining the relationship between perceived robot appearance type and gesture as well as analytical work on a coding scheme for human gestures. A strong element of this WP, led by UH, is collaboration with UniKarl on automatic classification of human activity, in particular gesture, and projected collaboration with UvA on an annotated video corpus of a robot following human subjects.

Milestones for RA3 include
M24: Design of user studies on spatial adaptation and interaction management
M24: Design of user studies on personal spaces and social rules in human-robot interaction
M24: Algorithms for replacement and task reasoning
M24: First results on identifying requirements for contextual interpretation of body postures and activities

Overall RA3 has produced worthwhile inputs to the project as a whole with a definite integration function. UH have produced an impressive amount of dissemination and seem to be involved in most of the so-far rather few joint publications between partners in the project. It is important that RA3 participants push forward in the next period beyond basic considerations of user comfort into consideration of more ‘companionable’ forms of social interaction

2.1.3.4 RA4 : learning Skills & tasks

Research Area 4 is concerned with learning skills and tasks. This area is fundamental to the development of the cognitive robot companion (and indeed to the field of cognitive systems and artificial intelligence more generally, so may have impact far beyond the immediate concerns of the project, though the focus is more on applying existing learning techniques rather than developing new ones).

Research area RA4 is composed of 3 workpackages:

WP4.1: Inferring the goals, representation and metric of the task
WP4.2: Social Context for the Correspondence Problem
WP4.3: Incremental Knowledge Acquisition of Complex Tasks

The deliverable for RA4 contains three reports:
1) A report on successful decomposition of the task into uncorrelated constraints and extraction of goals (EPFL)

2) A Report on studies and methods for solving the correspondence problem in context (UH)
3) A Report on building of and reasoning over complex task knowledge (UKa)

WP4.1 is led by EPFL with contributions from UH. In months 13 to 24 it addressed the problems of *what to imitate* in simple manipulatory tasks, which is closely coupled with the problem of *how to imitate* addressed in WP4.2. The approach looks for spatio-temporal invariants across several demonstrations of a task to be imitated (simple gestures), assuming a certain level of systemicity across the demonstrations. PCA is used to reduce dimensionality, reduce noise, and provide a more generic representation. A probabilistic learning technique, based on Hidden Markov Models is used to encode the data. The HMMs can then be used to recognise the extent to which a new input corresponds to an existing gesture. A metric for imitation performance is then formed from this representation, and subsequently used to produce a robotic reproduction of a learned gesture, which minimizes this metric subject to robot-body constraints – thus addressing the *correspondence problem.* The framework is generic and produced smooth trajectories on replay in the test scenarios.

Part of the work addressed here is *extraction of goals;* it was difficult to find a clear definition of what *goals* means in this context, except that they "encapsulate the important features of the task to be accomplished". (It is noticeable that the notion of "goal" being used here does not seem to correspond well to the notion being used in the cognitive architecture in WA6.) The *representation* adopted is a set of spatio-temporal movement constraints, extracted by analysis of spatio-temporal invariants – the *purpose* of a goal is left undefined at present. In some tasks, the goal might be that a spatial target is attained (e.g. reaching for something), in others the trajectory might be important, with absolute spatial location unimportant (waving goodbye). Success with both kinds of goal extraction are reported, though it seems that in general explicit goal preference has to be given in the absence of feedback or explicit teaching. The work reported here represents useful progress.

WP4.2 is led by UH with contributions from EPFL. The *correspondence problem,* already alluded to above, is the problem of imitating a gesture in a different kinematic system (e.g. with different degrees of freedom). The work conducted in months 13 to 24 has extended the JABERWOCKY system reported on in the first annual review, to produce solutions to the correspondence problem from a single input manipulation to multiple output platforms. The JABERWOCKY system has been released across the consortium. In the ALICE system, simulation-based work has looked further at the correspondence problem, for example in the context of using a learned library of correspondences and a changing embodiment over time; this work also included an interesting study of the effect of proprioception on learning which requires further study.

In the intended domain of Cogniron, the social context plays an important role (compared to much previous work programming robots by demonstration, for example for industrial tasks); a variety of interesting work has been carried out in this WP including: a study of the tradeoffs in dynamic vs static observation (i.e. whether the imitator may/must/may not move during the learning process); cultural transmission between robots (studied in the ALICE system); and an initial study into *moulding* and *scaffolding*.

Cogniron has recognised the importance of evaluating the performance of imitative systems and two evaluation mechanisms were studied and compared and gave good alignment between a quantitative system-centred assessment and a more qualitative human-centered assessment.

The work in WP4.2 continues to be of the high quality reported in months 1 to 12.

WP4.3 is led by UKA. In months 13 to 24 it built on the work from the first 12 months, investigating in particular the learning of tasks in a hierarchical manner, the use of multimodal input (particularly speech commentary) can facilitate task learning. A measure of subtask similarity was defined which is a vital prerequisite for higher-level task learning. Experiments have been successfully conducted in learning how to set a table place using task precedence graphs and vocal comments. Success in this task is simply the end position of the various objects involved (rather than any more complex spatio-temporal sequencing which remains an item for future work).

Milestones are

1) Method for automatically decomposing the task and expressing its metric into a linear set of constraints.
2) Further experiments on robotic imitation with more sophisticated target imitator platforms using effect and state metrics.
3) First implementation of building of and reasoning over task knowledge using precedence graphs and used problem solving strategies.

These milestones have all been satisfactorily met.

Although there appears to be input to each of the reports from at least two partners, all the scientific papers cited are authored from a single institution (in each case the lead institution for the WP).

2.1.3.5 RA5 : Spatial cognition & multimodal situation awareness

RA5 is composed of two workpackages:
WP5.1: Models of space, and
WP5.2: Models of objects.

This is, in fact, the result of a reduction from 4 WP's in the first workplan, which has emerged as appropriate given results during the first year. WP5.1 has been created to include work on hierarchical models of space, categorization of space, spatial concept learning by interaction with humans and temporal models. WP5.2. encompasses work on categorization of objects, interactions with humans and temporal relations. Such opportunistic reorganisation constitutes an excellent initiative on the part of the consortium that shows an ability to react to emerging research results. The two work packages are linked through the important research being conducted in representing space via relations between objects.

WP5.1 focuses on a number of research issues related to the modeling of space for a mobile robot in a domestic environment. It addresses the scientific questions:

- How can a hierarchical representation of space be attained?
- How can a robot achieve a consistent representation between 'local' representations and the categorical and semantic representations?
- Are the formed categories meaningful to humans and how can human supervision aid categorization?
- Can methods be developed to represent space by detecting objects?

Work in WP5.1 was carried out in three tasks:
WP5.1.1 Sensor-based methods for topological map building.
WP5.1.2 Integration of local representations and categorical representations
WP5.1.3 Space descriptions using object detection.

Three approaches using hierarchical representations of space have been addressed.

a) hierarchical representations by means of clustering techniques on the sensory data using appearance, (UVA / EPFL)
b) by means of object detections (LAAS, EPFL) and

c) space representations for interaction with humans (KTH, UvA, UniBi).

Hierarchical representations based visual appearance
The appearance-based approach is based on grouping images using image-level signals. Omnidirectional cameras are used in order to have a wide view of the scene. EPFL uses 'fingerprints' derived from sensory patterns to describe the current observation of the robot that can be used to label larger regions in space. UvA uses local invariant features based on the SIFT operator to describe an observation and cluster appearance into categories.

Object detection
This approach seeks to model the physical world in terms of objects and the way they relate to each other. Both spatial and semantic inferences are demonstrated. Object modeling and recognition capabilities together with methods to detect spatial and semantic relationships between objects have been developed using a probabilistic, feature-based, simple, object-recognition system (color, texture) as well as SIFT features.

Interaction with Humans
At KTH a graph-based model of the environment has been acquired, incorporating information that is given interactively by humans. Interaction is online, since the assumption is that a "guided tour" is an appropriate way to give the user the possibility to personalize the robot's general environment representation – in particular spatial relations can be naturally communicated using the shared environment. This model emphasises qualitative, high level descriptions of places, regions and objects, the former being described in terms of spatial relations between the latter; it deserves further work developing this approach.WP5.2 focuses on acquiring object models appropriate for passive and active learning about the objects.

Object models are represented in a unified object model structure that combines layers for symbolic description, physical interaction, and object appearance. A general learning cycle combines perception and action to learn a model that integrates knowledge of appearance and physical interaction through manipulation.

Three partners contribute to WP 5.2

- IPA: With trainable appearance-based detection and model-based pose estimation
- LAAS: With techniques for learning sensory-motor relationships from observation
- UKA, IPA and LAAS: Working on generation and refinement of object models

IPA has worked on trainable appearance-based detection and model-based pose estimation. An interactive learning scheme was developed to train an appearance-based object detection system by showing the object. Learning combines information from a range imaging sensor and a colour imaging sensor. Their method uses appearance-based recognition to determine the coordinates and pose of the object relative to the sensor head and uses superquadrics as geometrical models. The appearance-based recognition scheme combines scale- and rotation-invariant SIFT features with a Support Vector Machine for object discrimination. They refer to this technique as keypoint teaching. It can be remarked that a number of groups around Europe have explored a similar approach in the last year, with impressive results. This combination of SIFT and SVM would seem to constitute the new state of the art for such trainable recognition.

LAAS has investigated the use of a Pulsed Neural Network for learning sensory-motor relationships from observations using an appearance-based representation. As with the IPA system, neither the goal nor objects, actions or space representations are here known prior to learning. The method provides a hierarchical, structure-based method for scene recognition using object models learned from views.

For generation and refinement of object models, UKA and IPA have worked together to investigate a combined representation for data obtained using a SwissRanger sensor coupled with a colour camera. Techniques from UKA are used to construct a 3D super-quadric model to provide a first approximation of shape. This is then used as a skeleton for the appearance-based object detection method developed at IPA. The IPA appearance-based matching technique provides a pose estimate of the object. LAAS has adopted a Bayesian approach to provide active shape-based recognition of an object.

Overall, workpackage 5.2 of RA5 is advancing the state of the art in two important aspects: methods for learning integrated object models that combine 3D Form, Visual Appearance and Manipulation information, and in the combination of robust invariant features (SIFT) with discrimination recognition (SVM). In order to emerge a leader in this area, it is fundamental that the group carefully document experimental performance-evaluation for their methods.

2.1.3.6 RA6 : Intentionality and Initiative

Work in RA6 was carried on in four WPs according to the RA6 deliverable:
WP6.1 Generic architecture
WP6.2 Models and algorithms for cognitive robots supervision
WP6.3 Human-robot collaborative problem-solving
WP6.4 Intentionality attribution

However according to the Detailed Implementation Plan, it is being carried out in THREE WPs:
WP6.1 Generic architecture
WP6.2 Decision-making for Interactive human-robot task achievement
WP6.3 Intentionality attribution

Before we can approve the work in this Research area, the project team must justify this deviation. We note also that the effort figures for IPA in these WPs as reported in the Detailed Implementation Plan on p47 seem to be incorrectly totalled.) In this report we make the assumption that the Detailed Implementation Plan, as the more generic document, has the correct WPs for RA6.

One generic issue with these WPs is that all the milestones are very late in the project at M36 or M42, so that it is very hard to see how results can actually feed back into the project. Given that RA6 offers theoretical integration just as RA7 covers practical integration, the late milestones are particularly disappointing. We recommend that this be revisited and a milestone be set perhaps in six months time (month 30) for a first version of the architecture for a cognitive robot.

In general the amount of work reported for WPs 6.1 and 6.2 does not seem all that impressive given the allocation to RA6 as a whole of 77 pms over the second 18 month period, equivalent to over four people full-time. Moreover the relationship between the work reported for WP6.1 and 6.2 is not very clear, a problem if one assumes that the work of 6.2 should be encompassed by the generic architecture of 6.1, or in what sense is it generic?

WP6.1
The main output of this WP reported is the development of a conceptual architecture. We welcome this as a step towards the development of a unifying architectural framework for a cognitive robot. However the architecture is reported at such a high level of abstraction that it is difficult to answer some obvious conceptual questions, such as:

- What is the relationship between the interacting functional modules at the low level and the learning and interpretation boxes? There are no arrows on the lines between them which makes even the direction of data flows hard to determine.
- Does sensed data always have to be interpreted before it can be learned?
- What is the relationship between the 'learning' box on the architecture and the approach of RA4?

Given that interactive capabilities – which is a key distinguishing feature of a robot companion – are distributed invisibly over many boxes in the diagram, how far can this be said to be an adequate conceptual architecture? This is especially problematic given that the stated goals for WP6.1 include investigating "[t]he links between dialog, interaction and deliberation, which are central to the operation of a robot companion" and to "investigate the specificities and added issues imposed by … [the fact that] the robot … achieves most of its tasks in interaction with humans" (Updated DIP, p.97).

Given that this is the focus of an entire work package, it is distressing that the discussion of the conceptual architecture is limited to roughly a page of text and a large figure with caption. This is followed by a section (2.4.2), "Cognitive capacities", which lists a set of requirements for a cognitive robot, such as the incorporation of learning into every process of the robot, and a discussion of some of the rationale for the conceptual architecture. Unfortunately, much of the text here is devoted to hypothesizing capabilities that few, if any, of the partners plan on implementing for the key experiments, such as generalized goals for achieving human satisfaction.

The goal of this workpackage is to provide "a framework where the various contributions in Cogniron will fit together" (DIP, page 97). We take this to mean that the architecture is intended to aid integration and guide component development. While the conceptual architecture has value for thinking about some of the necessary ingredients of robot architectures in general, it is unclear how it would help different partners to either better design their software components in advance to support integration, or aid them in integration after the fact. Indeed it might even be a small hindrance since it adds the additional administrative burden of documenting how each group's architecture maps to the conceptual architecture. If this workpackage is to truly support integration, it must develop a software systems architecture at a much more detailed level. One might have expected work on this to have begun in the second year of the project. Failing such a detailed software architecture, the conceptual architecture should at least focus on those architectural capabilities that the partners actually intend to implement.

WP6.2

The main output of this WP reported is the development of a theoretical framework based on the joint intentions model of Cohen and Levesque. There is some discussion in sections 3 and 4 of an implementation framework but it is far from clear how this relates to the generic architecture of 3.1 Is it an expansion of one of the boxes on that diagram? Of several boxes? The report also doesn't make clear what specific functionality is supported by this implementation of Cohen and Levesque that wasn't already supported by previous systems like Tambe's STEAM or Mateas' ABeL, or how it will be used in the Cogniron project as a whole.

A definite issue is that the discussion of section 4 has not yet been mapped onto the interaction requirements: both those of cooperative dialogue from RA1 and those of collaborative movement from RA3. For example, constantly asking a checking question of the human to establish commitment could become extremely irritating and impact the human's perception of the robot's social competence. The links with multi-modal dialogue are mentioned as 'future work' – it seems a little surprising given the allocation of effort to RA6 and the importance of the topic that it has not yet been addressed.

More generally, it is unclear how this work relates to the overall integrated project. While it is presumably intended as an implementation of the CF-ROT ("reasoning about tasks, about its own abilities") Cogniron function, it is not available as a Cogniron service, it is not used by any other partner, and there are no joint tasks in any of the key experiments, either in their present forms or in the future planned forms described in the updated DIP. The report states that the joint-intentions executive is used for the KE2 demonstration. However, it appears only to be used for one of the three KE2 robots (RACKHAM) and it's unclear what its actual role in the key experiment is, since the KE2-2 scenario doesn't involve any joint plans. The only use of the joint plan functionality described in the report is for a tour-giving demo apparently unrelated to KE2. KE1, which does involve giving tours, doesn't use the joint-intentions executive. Although we would not go so far as to say that every piece of technology developed within an EU Integrated Project needs to be used in the final integrated demo(s), it is problematic for an entire work package never to be used for any of the KEs, or indeed ever to leave the lab in which it was developed. Joint intentions is a 15 year old theory and it has been used in many systems. If a new implementation of it is to be built, then it must be used in one of the KEs and some argument must be made for why a new implementation was needed in order to implement the KE. As it stands, the executive has only been used to solve a task that other partners solved using much simpler technology.

Finally, one might have expected some discussion in this section of whether goal or plan recognition capabilities would be included to allow a robot to be proactive in assisting a human. If it were included, where would this figure in the conceptual architecture? Or if it were not included, what are the implications for social competence?

Section 5 discusses a set of different implementations by different partners, without making it at all clear whether they are consistent with each other and with the development of the architecture for a cognitive robot of WP6.1. This was covered to some extent in the presentation of RA6 but the deficiencies in the conceptual architecture referred to above make the mappings supplied there somewhat vacuous.

In spite of the feedback we offered for the first year review on the need for integration, one still has the overwhelming impression from this section of partners ploughing their own furrows starting from the diverse architectures each brought to the project. Entirely lacking is any discussion of metrics for success or even adequacy in these architectures, and the vagueness of the conceptual architecture means that it does not provide a framework within which sensible comparisons and evaluations can be made.

WP6.3
The main work reported concerns the way in which humans attribute intentionality to robots. The analysis of a user study from 2004, carried out in the first part of 2005, is reported. Interesting and useful design implications have been extracted from this data. There does seem to be some overlap between some of the work reported – on human comfort-levels – and the work of RA3 but insofar as this indicates inputs into RA6 from RA3 it is to be welcomed.

2.1.3.7 RA7: System-level integration and evaluation

Work in RA7 involved the integration of technologies developed by the different partners under the different research activities, their demonstration in a set of Key Experiments, and the development of evaluation techniques for the integrated systems and the results of the demonstrations. Under the updated implementation plan, this work is divided into four workpackages:

WP7.1. Evaluation of functions and key experiments
WP7.2. Full implementation of the "Robot Home Tour"

WP7.3. Full implementation of the "Curious Robot"
WP7.4. Full implementation of "Learning Skills and Tasks"

In the first performance period, this work centered on developing a preliminary specification document for the key experiments and defining the Cogniron Functions (CFs) involved in them. In the second performance period, the scripts were revised and initial implementations were performed. This involved a significant effort in software integration. A major component of this work was the definition of a set of software services that could be shared between the partners and between the hardware platforms for the different key experiments, along with a set of data structure definitions, and a minimal API for set-up, take-down and periodic scheduling of services. These definitions were shared using a wiki to support distributed collaboration. This is an important start and the researchers should be praised for this.

WP7.1 involves the development of test metrics, test cases, measuring failure rates, and identifying the necessary environmental conditions in which the Cogniron functions work properly, and usability evaluation. The lead contractor is IPA and the PAR lists the workpackage as having used 60.4 person-months. Very little was said about this workpackage in the PAR. It is referred to only at the end of p.38, which states that it contributed to appendices C and E of the D7.2006 report. However, appendix C consists of the results of a questionnaire of all the partners as to their contributions to the different Cogniron functions and the second part is a table giving the current implementation status and execution time of the different CF implementations. The latter table includes a "test methods" column, but most of the entries are sketchy at best. Many are blank.

Appendix E gives the scripts for the different key experiments. These represent a significant improvement over last year's plans. However, for KE2 and KE3, these are different than the scripts actually being used in the workpackages. KE2 uses three different scripts, all of which are different from the KE2 script in appendix E and KE3 uses two different scripts, but both are similar to the script in appendix E. Also, much of the scripts, as presently written, could be implemented using the current state of the art, and so do not fully demonstrate the capabilities of the technologies being developed in the project. If we understand the reports properly, this is because the scripts represent only what has been integrated so far and that the plan is to extend the scripts as the partners integrate more capability into the demonstrations. However, this makes them more of documentation of what has been implemented than a planning tool to guide future implementation. It is important that the partners define goals for the key experiments that can guide their research, that represent a clear improvement over the current state of the art, and that properly demonstrate the technologies being developed in the other workpackages.

The biggest weakness of the RA7 deliverables is the lack of a substantive discussion of evaluation. The *Updated Detailed Implementation Plan* states that the goals for RA7 in the second phase were to (1) move further in the implementation of the key experiments and (2) to:

> … refine the evaluation process of the *Cogniron Functions*. This is still ongoing, and was accomplished mainly through the underlying *Services level*. *Test* cases, test databases, and target performance (e.g. frequency for tracking algorithms or some recognition rate) were specified. (p.48, italics original)

However this work does not appear to have made it into D7. As stated above, appendix C.2 does include some information about testing, but very little; the only performance information appears to be current performance, not target performance. As it stands, the only evaluation metric that seems to have been specified for the key experiments are binary (did it/didn't do it). Again, this is disappointing for a workpackage involving 60 person-months of effort.

That said, there has been definite progress on the individual key experiments. WP7.2, which involves the implementation and testing of the Robot Home Tour, is proceeding well, with laboratory tests given in the reports, and experiments in an unmodified apartment environment having been given in the review presentations. These are to be commended.

WP7.3, which involves the implementation of the curious robot KE, involved 60.4 person-months in this project period, under the direction of LAAS. The experiment has been split into three sub-experiments to allow parallel implementation and ease integration difficulties. While the desire to allow parallel implementation is understandable, this gives the workpackages more of a feel of being a set of unit tests of modules than a true integrated project. Worse, *none of these experiments appear to have any significant involvement with the other partners*. They appear to involve almost exclusively components that were developed by LAAS, only one of which appears to be used elsewhere in the project (NHP); the rest are not even mentioned on the wiki. Finally, appendix E, produced by IPA as a component of WP7.1 gives only a single scenario for KE2, and it is different from all the scenarios mentioned for the individual sub-experiments. Although they have very definite overlaps, it is unclear what the relationship between these different scenarios is.

KE2-1 involves a robot driving to a table with a single isolated object on it, and grasping it. The scenario states that "it can then be given to a person or placed elsewhere", but it is unclear whether that would be part of the key experiment. This appears to involve the integration of only two cogniron services, Mobj and GraspPlan, both of which appear to have been developed at LAAS. Thus, while mention is made of comparing it to work done by other partners, the current implementation appears to be the sole work of one partner. It should also be mentioned that neither service appears to be mentioned in appendix C or D, nor do they appear anywhere on the wiki. This seems to run contrary to the spirit of an IP.

KE2-2 involves navigating while safely maintaining a "comfort zone" around any humans who may enter the area. Again, this appears to involve the integration of only two Cogniron services, in this case ICU and NHP, again, both apparently developed at LAAS. The NHP service is documented on the wiki, but ICU is not. It is unclear how ICU corresponds to the services and modules listed in appendices C and D.

KE2-3 involves acquiring knowledge from sensory-motor information using a neural-network architecture. It involves only one Cogniron function (OR) and no Cogniron services. It is not mentioned on the wiki (at least the string "neural" does not appear on the wiki) and it is unclear how its architecture could be integrated with any of the architectures used in the rest of the project or even how it could be mapped to the conceptual architecture which is supposed to be guiding the design process. *It is very important that it be made clear exactly how this component of the project relates to the rest of the project because as it stands it appears to have no interactions with the rest of the project*.

WP7.4 involves the development of KE3, the demonstration in which the robot is taught pick-and-place tasks. This package involved 60.4 person-months of effort in phase 2 and was led by UniKarl. Their results are encouraging. Although much more integrated with the other partners than KE2, this experiment does feel in some ways like a unit test of the learning algorithms.

The last section of report D7 involves software integration. While in phase 1 the partners defined the Cogniron functions, in this phase, they have begun mapping them to concrete software, expressed as "Cogniron services". The service framework requires that all services provide a main routine, and optionally init, end, and interrupt routines. The main routine appears to be intended to run either one-shot (i.e. the service is really a subroutine), or periodically (i.e. the service is a coroutine). The documented services on the wiki seem to depart somewhat from this framework (for example, none appear to have interrupt routines, but some have an additional “control” routine not sanctioned by the architecture). The wiki also documents the data structures used by the routines. Both of these are good first starts. However, the current documentation is currently limited to the provider(s) and seeker(s), implementation status and language, and the calling sequences for the standard procedures (main, etc.), and much of this is left empty. There is no documentation of OS or library requirements, whether the software needs to make its own thread or process, or what its CPU requirements are. Nor does there appear to be any attempt to centralize source-code control. A shared CVS, Sub/version, or SourceForge repository would presumably facilitate code sharing considerably, as would the adoption of a bug-tracking tool like Bugzilla.

It should also be mentioned that nearly every Cogniron service implementation mentioned in appendix C lists itself as using 100% of the CPU, and many still list themselves as being unable to operate in real time. Although implementations will doubtless be optimized before the final demo, this is still cause for serious concern. If the components are going to be properly integrated on a real running platform, then realistic CPU budgets need to be worked out now and *partners need to commit to delivering implementations that will run within their CPU budgets*. If not, more and/or bigger computers need to be ordered for the demonstration platforms. If more machines cannot be placed on the platform, then different algorithms that can be realistically implemented on the available hardware need to be developed.

2.2 Work planned for the next 18-month period:

As noted in the first-year review report, there is clearly interaction and synergy between the partners: the descriptions of the RAs bear this out, and some 20 documented intersite visits. However, this has still only resulted in 4 (out of 75) joint papers. We recognize that it takes time to write up and publish work, so we hope to see a much larger number of joint papers in future years. Although we believe there is an improvement in the amount of genuine collaboration, we are still concerned that overall the level does not reflect the amount that should be present in an *Integrated Project*. In particular, we repeat our comments with regard to the Key Experiments (see the discussion above in RA7). The revised structure of WA5 has resulted in much reduced worry about parallel development and there is some evidence of genuine collaboration (including one of the project’s joint papers) but research is still predominantly focused in individual labs. Last year we hoped that “the development of the Cogniron functions and their use in the experiments is a sensible attempt to overcome these potential problems, and indeed the functions have been chosen it seems specifically in relation to partner interaction rather than for mainly functional or architectural reasons”; however, it appears (from table C1, D7) that most CFs are single partner (only LOC, NAV, OR, RG and LIF having multi-partner input, and even in these cases the joint work is “parallel” rather than integrated).

Despite these comments, we believe there is genuine value in the collaborations across the project, with the consortium indeed being *“more than the sum of its individual members”*; however, much more could be achieved and we look forward to seeing the fruition of the integration of the (in many cases excellent) pieces of work conducted at individual labs in the final two years of the project.

With regard to the commitment and performance of individual partners, we did not detect any unsatisfactory performance, and in many cases it is clear that the partners have worked very hard to achieve the results reported.

The proposed implementation plan for the period of months 24 to 48 is designed to provide greater integration, allowing modules produced by different partners to operate together, and taking into account operational constraints. The consortium proposes to enhance robustness and generality for the systems. The updated DIP is a clear improvement over last year's DIP. However we are all concerned by the amount of work that is left until the last year in the updated DIP. This might simply be a matter of casual phrasing and/or a tendency to quantize milestones to units of a year. But we are concerned that a number of the work packages, if taken at their word, aren't delivering finished component technologies until the end of the project, leaving it unclear what version of those components will be integrated into the KEs, and in many cases, whether they will be integrated at all. Nearly all WPs need to be revised to list separate milestones for completion of components, integration into KE robots, and evaluation of the components in the integrated system. While the last of these can sensibly be placed in month 48, the others must necessarily be much earlier.

2.2.1 RA1 Multi-modal dialogues.

Research area 1 proposes to adapt dialog strategies and to increase situation awareness by using contextual information on the environment and system state based on multi-modal cues and the use of user models. The work is composed of 3 workpackages:

WP1.1 Adaptive multi-modal dialogue
WP1.2 Representation and integration of knowledge for an embodied multi-modal dialogue system
WP1.3 Evaluation methods

As requested by the reviewers at the first year's review, a joint deliverable report for the entire Research area is proposed for each year. A total of 68 man-months are proposed for this effort, with 36 mm proposed by Univ Bielefeld, 19 mm proposed by KTH, 10 mm proposed by LAAS and 3 by University of Hertfordshire. These appear to be reasonable estimates for the work proposed.

The plan for WP1.1 is good, but as mentioned above, we are concerned that the adaptation strategies aren't scheduled to be delivered until month 48. We presume this means "implemented, integrated, and evaluated in KE1", but if not, this needs to be made explicit and some explanation given for doing the work without integrating it.

WP1.2 will focus on integrating task state from a planner and social cues from humans. This is good, but the statement of work needs to clarify its relationship with RA6 and RA7. It appears to state that it will use the executive developed in WP6.2 and that the output of WP1.2 will be used in KE2. However, neither WP6.2 nor WP7.3 seem to indicate this. Nor do any of the work packages suggest that the output of WP6.2 will be used in KE1. Moreover, while the milestones mention integration of multimodal cues to assess situation (which we take to mean social cues), integration of task state from planners is not mentioned in the milestones. WP1.2 also has the problem of not appearing to deliver its component until it will be too late to integrate it.

WP1.3, again, proposes a good plan for research but as written is not scheduled to deliver its output (a proposal for evaluation guidelines) until the completion of the overall project when, strictly speaking, it would be too late to do the evaluation.

2.2.2 RA2 - Detection and understanding of human activity

Research area 1 proposes to devote efforts toward developing techniques for interpretation of human activities and integration of activity recognition with body tracking and object tracking. The work is organized in 3 work packages.

WP2.1 Detection and perception of body parts based on sensor features
WP2.2 Human body model: Integration and fusion

WP2.3 Context based interpretation and classification of activities

WP2.1 addresses detection, identification and tracking approaches. Objectives are to combine 2D image-based features with acoustic tracking methods, as well as to develop 3-D methods for human detection, model initialization and human motion capture.

WP2.2 proposes work on integration of features into a global description of the human body movement. Objectives include fusion of 2-D and 3-D tracking methods for body pose estimation, evaluation and comparison of complementary fusion approaches, and interactive human model parameter acquisition.

We were concerned, as discussed in section 2.1.3.2, that the relevance of the 3D articulated tracking work in WP2.1 and WP2.2 to the overall integrated project has not been made clear. It is important that this work be integrated into one of the RA7 systems and tied to actual functionality in one of the KE scenarios. However, for this to work, the milestones must be revised so that the tracking systems are be delivered well before month 48 and the milestones must commit to delivering implementations that can run in real-time on whatever computation budget will be available for KE1.

WP2.3 seeks to interpret and classify human motion in order to gain information on current activities and the focus of attention. Work includes classification of human activities for the context of the KEs, recognition of simple manipulative action within a dialogue context, and segmentation of dynamic gestures. While we assume that the output of WP2.3 will indeed be used in RA7, it must be pointed out again, that the current KE scenarios for RA7 do not actually require this work. Again, the relationship between the output of this workpackage and the overall integrated project needs to be clarified, the delivery milestone needs to be moved earlier, and integration and evaluation milestones need to be added.

A total of 70 man-months are proposed for this effort, spread relatively evenly over 6 partners. A total of 15mm are requested for LAAS, 14 mm are requested for UKa, 13 mm are requested for IPA, 12 mm are requested for UniBi, 10 mm are requested for UvA and 6 mm are requested for UH. With so many partners involved, one worries whether there will be genuine collaboration. Joint deliverable reports are proposed for each year. It will be a challenge for the partners to maintain a coherent integration of the contributions from all 6 partners in this area.

2.2.3 RA3 - Social Behaviour and Embodied Interaction

Research area 3 will study issues of personal spaces and postures and provide guidelines for implementing human-robot interaction schemes and for interpreting human activities. It will also include work on development of motion planners in presence of humans.

Efforts in this area are organized in 4 workpackages:

WP3.1 Personal Spaces and Social Rules in Human-Robot Interaction
WP3.2 Posture and Positioning
WP3.3 Models and Algorithms for Motion in Presence and in the Vicinity of Humans
WP3.4 Requirements for Contextual Interpretation of Body Postures and Human Activity

Plans for WP3.1-3.3 are all reasonable and should be commended for delivering their outputs in month 42, when they can reasonably be integrated.

Workpackage 3.1 will develop default social rules for implementing socially acceptable, non-verbal behaviour in robots with respect to social distances in human-robot and robot-human approach scenarios, including handing over an object. The experimental scenarios will focus on naturalistic settings and long-term human-robot interactions.

Workpackage 3.2 addresses the role of posture and positioning in task-oriented robot-human interaction, based on user studies that will provide data on how robot and human coordinate their actions in a shared space, including a multiple room scenario, where spatial coordination is necessary for natural communication between human and robot.

Workpackage 3.3 will develop models and algorithms for motion in presence and in the vicinity of humans. A planning module will be implemented to allow the robot to determine the feasibility of motion and manipulation tasks in the presence of humans.

Workpackage 3.4 will investigate gestures and body postures/movements that occur naturally in situations such as described in the KEs. This work will use annotated databases of image sequences of human activities and social behaviour relevant to the KEs to consolidate and expand taxonomies developed at UH and KTH in year 2. We are somewhat concerned that WP3.4, which delivers a taxonomy of body postures and human activities in month 42, might still be too late to facilitate the development of code that could implement this taxonomy.

A total of 102.5 man months are proposed for this effort in years 3 and 4 with the largest share for UH (38), KTH (29 mm), and LAAS (28.5 mm). UKa and UvA will have minor roles with 4 mm and 3 mm each respectively for this 2 year period. As with the other research areas, consolidated joint deliverables are proposed for each year.

2.2.4 RA4 - Learning Skills and Tasks

Research area 4 will devote 3 work packages to imitation and social learning. For learning more complex tasks, the consortium will investigate interactive learning in which the human helps to refine the robot’s knowledge and to correct its mistakes using various forms of deixis, verbal and non-verbal.

WP4.1 Inferring the goals, representation and metric of the task
WP4.2 Social Context of Robot Social Learning
WP4.3 Incremental Knowledge Acquisition of Complex Tasks

Workpackage 4.1 will focus on methods to infer the goals, representations and metrics of a task, in order for the robot to determine what to imitate. It will investigate ways in which the human could more directly and more precisely guide the robot’s learning.

Workpackage 4.2 will investigate methods for the robot to autonomously determine *how*, *what* and *who* to imitate, that take into account the social context and the relationship between the robot learner and a teacher.

Workpackage 4.3 will build a meta-representation of the overall task knowledge the robot companion has learned and to apply more complex reasoning methods on the underlying task knowledge data base. This will transform the task knowledge from action sequences to more expressive real programs that incorporate control flow, explications of repetitiveness and selection of alternatives at runtime.

As with RA3, we are concerned about the relationship between the DIP and the overall integrated project. While p.26 of the updated DIP states that the results of each workpackage will be used in KE3, the current scenarios for KE3 do not seem to involve these technologies. As with other RAs, the milestones listed do not mention the key experiments or any integration activities, so it is unclear whether the delivery of a method in month 48 means first results, or finished evaluation within the context of a key experiment. Again, we recommend that development of methods, integration in the key experiments, and evaluation be unpacked as separate milestones.

A total of 70 man-months are proposed for years 3 and 4 in this research area, with UH providing 36 mm, EPFL providing 20 and UniKarl providing 13. LAAS will follow the work with 1 mm over 2 years.

2.2.5 RA5 - Spatial Cognition and multi-modal situation awareness

Research area 5 addresses the understanding of how an embodied system can obtain a conceptualisation of sensory and sensory-motor data, generate plans and actions to navigate and manipulate in typical home settings. Work will include space modeling, in conjunction with the key experiments scenarios, to come to representations that can be intelligible to humans. The consortium will also address dynamic situations. They will further focus on learning object categories by manipulating and showing.

The work is organized in 2 Workpackages.

WP5.1 Models of space
WP5.2 Models of objects

Workpackage 5.1 will explore four research questions: How can the robot come to a hierarchical representation of space? How does the robot achieve consistency between 'local' representations and the higher level (categorical, semantic) representations of space? Are the formed categories meaningful to humans and how can human supervision aid the categorization? Can the robot develop methods to represent space by detecting objects? With 5 sub-areas identified within workpackage 5.1, this almost looks likes a research area in itself.

Workpackage 5.2 will address autonomous and interactive learning methods, fusion of appearance-based methods and geometrical models, and open-ended learning algorithms for object models. This workpackage is also organised as 3 sub-workpackages, making it also appear to be a mini-research area.

The impression that Research Area 5 is, in fact, a composition of 2 research areas, is reinforced by the projected use of 194 man months during year 3 and year 4. Just about every partner is involved with UvA proposing 48 mm, LAAS proposing 36 mm, 32 mm for EPFL, 24 mm for KTH, 18 for UniBi and 12 for UnivKarl.

As with the other RAs, we are concerned that the relationship between specific subprojects and the overall integrated project isn't always clear. Steps 5.1.2.2 and 5.1.2.3 specifically mention implementation and testing on robots in KE1. However, there is virtually no mention of the key experiments, the Cogniron functions, Cogniron services, or indeed integration of any kind in the rest of the DIP for RA5. And again, it is worrisome that much of the technology being developed here isn't required by the current designs of the key experiments and that most of the systems aren't scheduled to be implemented until month 48. To be an *implementation plan*, and not just a research proposal (compelling as it may be), the DIP needs to make much more clear what components of the research are expected to be used in the KEs, when the basic functionality will be operational by, and when it will be integrated into the KEs. It should also include some kind of indication of the expected resource requirements of the systems so that one can determine whether it is realistic to expect to implement it in real-time on the available hardware platforms.

2.2.6 RA6 - Intentionality and Initiative

The conceptual architecture will be further studied and refined to account for all the sought cognitive capabilities, taking into account lessons learned from the partial implementations and progress of the state of the art. Collaborative problem solving and task achievement supervision will also be investigated in more details and further implemented. Guidelines from intentionality attribution studies will be used for the key experiments.

WP6.1 seeks to develop “a generic architecture for a cognitive robot”. This WP is somewhat confusing in that different reports speak variously about the “generic architecture”, the “conceptual architecture”, and the “architecture for a cognitive robot”, but it isn’t clear whether these are all different names for the same architecture or different architectures at different levels of detail (e.g. is the generic architecture a specific instance of the conceptual architecture?). The problem here is that the conceptual architecture is so abstract that it’s hard to know what the proposed work would even mean in the context of it. For example, how does one investigate “the links between dialog, interaction, and deliberation”, in an architecture that doesn’t contain dialog or interaction as components?. Or if the neural net architecture in KE2-3 really is an instance of the conceptual architecture, what does it mean to “investigate control structures” when one of the instantiations doesn’t seem to have control structures in the traditional sense of the term? If the goal of this WP is to revise the existing conceptual architecture, then it is unclear why 3 additional person-years of effort are required to do it.

The only milestone listed for this work package on p.45 of the updated DIP is “an architecture for a cognitive robot” in M36. Again, it’s unclear whether this means an updated conceptual architecture, a software architecture, or a more theoretical cognitive architecture. The description on page 97 lists an additional milestone, an updated version of the architecture, which suggests that the M36 architecture isn’t an update, and so might be different from the conceptual architecture. In any case, this needs to be made clear. Also, the role of the milestone in M48 needs to be made clear. Is there really an additional milestone? And if so, what is the purpose of releasing a new reference architecture for the different KEs after the KEs are complete?

WP6.2 seeks to develop a software framework that will allow human-robot decisional interaction. As with WP6.1, there is a mismatch between the milestones listed on p.45 and on p. 98 (the second milestone is missing form p.45). Here the phrasing of the milestones is unclear. M36 promises to deliver the software framework. However, M48 promises to deliver the “implemented and demonstrated” software framework, seemingly implying that the M36 version won’t be implemented. As with other RAs, it’s unclear whether the implemented and demonstrated version means implemented and demonstrated as part of the KEs or just tested on its own. The DIP states that “The effective implementation will be performed in coherence with interactive decisional situations that may arise in the key experiments”. This appears to say that the implementation will not be used in the KEs but that the researchers will try to test it in scenarios that might be relevant to the key experiments, but again, this should be made clearer.

WP6.3 seeks to produce guidelines for robot appearance and behavior when attracting the attention of humans. The DIP states that these will be used in KE1 and KE2 during phase 4. Again, we would be more comfortable if the current designs for KE1 and KE2 actually involved the robot attracting humans' attention, but we believe that the partners intend to alter the KE designs accordingly. Again, there is an inconsistency between the milestones for this workpackage as listed on pages 46 and 99. Page 46 states the guidelines will be available in M42, whereas p.99 says M48. We assume M42 is correct, since M48 would be too late to incorporate the results into the key experiments.

2.2.7 RA7 - Systems level integration and evaluation.

Work on RA7 is intended to focus on (1) implementation of the key experiments, and (2) testing and evaluation.

WP7.1 involves the definition of the key experiments and their evaluation criteria, and the actual usability testing of the KE robots. The current updated DIP for this WP is unacceptable. The version on p.49 lists only a one-sentence definition of the objectives and three sentences of description of the work to be done. It also lists two milestones, but they are different from the milestones on p.100. Page 100, however, also lists a different objective and description from the version of page 49. The page 100 version is a modified version of the objectives and description for WP7.2, found on p.101. We would suggest that the actual milestones be the union of the milestones given on p.49 and p.100: that in M36, preliminary tests be done **and** the final definition of the tests be done, and that in M48 the final testing be complete.

WP7.2 involves the implementation of KE1. Again, there are some mismatches between the versions described on p.50 and p.101. We would suggest adopting the milestones on p.101, and the incorporation of the additions to the scenario and script into the version on p.101.

WP7.3 and WP7.4, involve the implementation of KE2 and KE3, respectively. These have the same mismatches between sections 1.2 and 1.5 as in WP7.2. Again, we would suggest adopting the milestones as listed in section 1.5, while copying the additions to the key experiments discussed in section 1.2 into the versions in section 1.5

WP7.2-7.4 should also consider CPU budgets for the different components they intend to incorporate. As mentioned above, nearly all the Cogniron services as specified are using 100% of the CPU. Unless the partners expect to be able to add large numbers of CPU upgrades to the demonstration platforms, advance planning now will be necessary to prevent failure in phase 4 due to CPU overload.

3 CONSORTIUM PARTNERSHIP

While there has been a great deal of excellent technical work, it is not yet truly integrated. The project is intended to perform research in a number of related areas that combine synergistically in a capstone demonstration of a cognitive companion robot. The demonstration is supposed to act as the forcing function for integration. It is understandable, and perhaps even wise, to choose to split the final demonstration into multiple KEs on different hardware platforms. But at present, it has been split into 6 experiments on 4 hardware platforms. Half of these experiments are being implemented entirely by one partner (LAAS), apparently without the use of any Cogniron functionality from the other partners. Few of the WP implementation plans in the updated DIP mention integration, much less the key experiments. And WP6.1 and 7.1, which ought to be driving the integration, don't seem to be well integrated with the work being done on the ground.

That said, there are very definite signs of progress. The partners are indeed beginning to cooperate on group demonstrations, especially for KE1. The definition of Cogniron services and code sharing through the wiki is also an important step; we encourage the partners to continue with this. And the updated DIP for RA7 does list further pieces of functionality that the partners intend to integrate in the KEs. If this is followed up by aggressive technology transfer in phase 3, along with a willingness to adapt one research program to fit the needs of integration, then we believe this could be a very successful project. However, time is running short; these steps can't be left to phase 4.

It should also be added that the consortium appears to be continuing to be playing an active role in the EURON NoE. There continues to be potential for collaboration with IPs and STREPs in the Cognitive Systems area (one partner is already joint with one of these) and Cogniron should try to play a full part, as appropriate in events and activities organized by those projects, including the new coordinated action, Eucognition (www.eucognition.org), and their annual meeting (e.g. http://www.socsci.ru.nl/CogSys2/index.html).

4 PROJECT MANAGEMENT AND CO-ORDINATION

While there has been a great deal of excellent technical work, it may be that much of this was not explicitly managed. We are concerned that the issues we have raised elsewhere in this document, particularly with respect to WA6 and WA7 reflect on the technical management. Integration typically does not happen without a great deal of careful watching, coordination, encouragement, monitoring and even coercion! We expect the technical management to be more active in future. We recommend the adoption of more detailed milestones than are presently included in the implementation plan, whether these are placed in the revised DIP or not. One milestone per year is insufficient to flag when a part of the project is getting behind. We also recommend the adoption of standard software project management tools and methodologies, such as a source code control, bug and feature request trackers, etc.

The administrative and financial management generally seem adequate. We have found some minor discrepancies between different parts of the deliverables, but this is not entirely surprising given the timescales involved and the quantity of documentation.

The consortium did report the delay in the EU payment due, and we hope that the next payment can be made in a timely fashion.

5 USE AND DISSEMINATION OF KNOWLEDGE

As noted last year, this is primarily a basic research project, so considerations of commercial exploitation are not a key consideration. The project did consider demonstrating at CeBIT 06 or Automatica, but the costs were rather high, and we agree that this did not represent value for money, at least at this stage of the project. The project has now formulated revised plans for three different avenues of dissemination to industry which seem entirely appropriate given the current stage of the research. Longer term, we expect the research may have a significant impact on the commercial robotics industry.

With regard to academic dissemination, the main vehicle has been publication of papers describing the research conducted. These have been published in a variety of outlets, including some in top quality outlets (such as IJCAI), and one paper won a best paper prize. The publications are all listed on www.cogniron.org, but the lack of pdfs here makes it hard for interested readers to directly access the results. One is told that "Most papers are available directly on each partners' publications webpage", but direct links to these pages are not given. Presumably pdfs are not given here for reasons of copyright and/or problems in maintenance, but we note that the pdfs are not even available on the internal wiki. Moreover, in our experience publishers never object to modification of copyright agreements to allow publication on any appropriate web pages. We recommend that it be made easier to directly access the pdfs.

Cogniron partners have been very active in workshop organisation which we commend (though none of the events appear to have had explicit Cogniron "branding", unlike the IJCAI'05 tutorial organised by the related EU COSY project).

6 CONCLUSION and SUMMARY of RECOMMENDATIONS

☐ Unsatisfactory project (The project has failed to achieve critical objectives and/or is not at all on schedule)
☒ Acceptable project (The project has achieved most of its objectives and technical goals for the period with relatively minor deviations)
☐ Good to excellent project (The project has fully achieved its objectives and technical goals for the period and has even exceeded expectations)

Recommendation
☐ the project should continue without modifications
☒ the project should continue with the following modifications (technical or administrative)
☐ the project should be terminated (list main reasons)

At the first year review, two formal recommendations were given for the second year:

1) Reinforce work on both cognitive and systems architectures.
2) Consolidate the deliverables.

Concerning the first recommendation, the consortium has increased the manpower in this area. They have increased the number of meetings and workshops devoted to this problem. However, as noted above, much of this effort has been focussed on defining a "Conceptual architecture" relating perception/interpretation, learning, decision-making and interaction. What is needed is additional work on software architectures, integration tools and development methods, particularly for the KEs. In addition, a cognitive architecture should be defined to describe the core cognitive capabilities of the system, and how these fit together and interact to provide capabilities for tasks and social interaction.

Concerning the consolidation of deliverables, the consortium has satisfactorily implemented this recommendation, and the results are a set of deliverables that, while still large, allow more resources to be devoted to fundamental issues rather than superficial formalities.

For the third year of the project we make the following formal and informal recommendations.

Formal Recommendations:

1) Reinforce work on both cognitive and systems architectures.

2) Clarify the scripts, tests and metrics to be used for evaluation in the KE2 and KE3 experiments. Don't limit the KE designs to what you have running; design the final versions of the KEs *now* and use them to guide implementation
3) Nearly all WPs need to be revised to list separate milestones for completion of components, integration into KE robots, and evaluation of the components in the integrated system. While the last of these can sensibly be placed in month 48, the others must necessarily be much earlier. Intermediate (e.g. 6 month) milestones should be specified for all tasks..
4) Make clearer how the research in WPs will serve the integrated system.

Informal Recommendations:

1. The consortium should build on its strengths, and seek to become a visible success by redefining the state of the art in cognitive robotics.
2. Publish the XML schema used for module integration, and promote these as the basis for industry standard protocols.
3. The consortium should meet the challenges of a cognitive robot by more fully exploiting social interaction for learning in the key experiments.
4. Integrate other consortium partners into the KE2 experiments. For example, these tasks inherently involve social interaction with contributions that can come from KTH and/or UH.
5. It is recommended that the consortium use KE1 as a model for KE2 and KE3. The reviewers expect to see the same degree of integration, allowing the system to be used for experiments in robot behaviour and learning rather than simply the focus for systems building.
6. Be very explicit about the problems for which the consortium seeks to advance the state of the art. Use these planned innovations to define the experiments and the requirements for technology development. Use the requirements to define the technological work in the remaining areas.
7. Adopt professional methods for integration. Proper project management requires clear specification for components. Reliable integration and proper performance-evaluation experiments require component interoperability. This can only be achieved if components are assembled according to clearly defined architectural principles, using well defined component interfaces.
8. We recommend the adoption of standard software project management tools and methodologies, such as a source code control, bug and feature request trackers, etc.
9. We recommend revisiting the timing of the milestones in WP6 – all the milestones are very late in the project at M36 or M42, so that it is very hard to see how results can actually be fed back into the project. Given that RA6 offers theoretical integration just as RA7 covers practical integration, the late milestones are particularly worrisome. We recommend that this be revisited and a milestone be set perhaps in six months time (month 30) for a first version of the architecture for a cognitive robot.
10. We recommend that it be made easier to directly access the pdfs of papers published by the consortium that are listed on www.cogniron.org.
11. We recommend that consideration be given as to whether 3D articulated motion capture is necessary for Cogniron, and if so, how it will be used.
12. We recommend that in order to ensure that the components of the service functions will be properly integrated on a real running platform, that realistic CPU budgets be worked out now and partners should commit to delivering implementations that will run within their CPU budgets. If not, more and/or bigger computers need to be ordered for the demonstration platforms. If more machines cannot be placed on the platform, then different algorithms that can be realistically implemented on the available hardware need to be developed.
13. We recommend partners devote more effort to writing joint papers – this will naturally help drive genuine collaboration to achieve the project's goals.

Name(s) of the reviewer(s):

Date:

Signature(s):

7 APPENDICES

1. Status of project reports and deliverables (submitted/delayed, accepted/rejected/to be modified)
2. List of participants
3. Agenda

Appendix 1
Status of project reports and deliverables.

Deliverable number	Title	Status (submitted/ delayed)	Accepted/ Rejected/ To be modified	Deadline for re-submissions	Remarks
D1.4.1	Joint RA1 deliverable	Submitted	accepted		
D2.4.1	Joint RA2 deliverable	Submitted	accepted		
D3.6.1	Joint RA3 deliverable	Submitted	accepted		
D4.4.1	Joint RA4 deliverable	Submitted	accepted		
D5.6.1	Joint RA5 deliverable	Submitted	accepted		
D6.1.1	Progress on a definition of the architecture of Cogniron Robots	Submitted within joint RA6 deliverable	Clarification required		Unclear relation between conceptual architecture and generic architecture; unclear what role the conceptual architecture plays in the integrated project
D6.4.1	Joint RA6 deliverable	Submitted	To be modified		Relationship between WP6.2 and KEs should be clarified; WP descriptions should be reconciled with DIP.
D7.1.2	Refined specification document	Submitted within joint RA7 deliverable	To be modified		Appendix E appears inconsistent with scripts being used in KEs
D7.2.2	Report and video clips of KE1	Submitted within joint RA7 deliverable	accepted		Very impressive
D7.3.2	Report and video clips of KE2	Submitted within joint RA7 deliverable	accepted		
D7.4.2	Report and video clips of KE3	Submitted within joint RA7 deliverable	accepted		
D8.1.4	Electronic copy of information available within project repository	verified	accepted		
D8.2.1	Confidential list of filed patents	Submitted as part of joint RA8 deliverable	accepted		

Deliverable number	Title	Status (submitted/ delayed)	Accepted/ Rejected/ To be modified	Deadline for re-submissions	Remarks
D8.3.3	Expert symposia/workshops	Submitted as part of joint RA8 deliverable	accepted		
D8.4.1	Proposal for industrial dissemination meeting	Submitted as part of joint RA8 deliverable	accepted		
D9.1.2	Summer school website with repository with available teaching material	Not available	To be provided next year.		
D10.1.3	Executive summary	Submitted	accepted		
D10.2.1	Periodic reporting reports and documents	Submitted	accepted		

Appendix 2
List of participants at the review.

Appendix 3
Agenda of the meeting

COGNIRON 2nd Review meeting
Brussels, February 14, 2006
Agenda

9:00 - Project status:
Recall of project objectives - project coordinator (10')
Main achievements so far - project coordinator (20')
Other activities year2 (innovation, training, management) - project manager (10')
Highlights of work plan next phase - project coordinator (10')
Discussion 10'

10:00 - Integration and key experiments (RA7)
RA7 leader 10'+ KE leaders (3 x 10') + 10' discussion

10:50 – Break

11:05 - RAs 1-6 :
Past work and plans for next phase – RA leaders (20' presentation+10' discussion)
11:05 – RA1
11:35 – RA2
12:05 – RA3

12:35 – Lunch

14:00 RAs 1-6 cont'd
14:00 – RA4
14:30 – RA5
15:00 – RA6
15:30 – RA7 – integration tools

15:50 – Break

16:05 – Summary and General discussion (25')
16:30 – Reviewers private meeting with Project Officer
17:45 – Feedback from Project Officer and Reviewers
18:00 - End