EU FP7 CogX
ICT-215181
May 1 2008 (52months)

DR 7.3:
Analysis of a Robot that Explains Surprise

Jeremy L. Wyatt[1], Marc Hanheide[2], Andrzej Pronobis[3], Kristoffer Sjöö[3], Alper Aydemir[3], Patric Jensfelt[3], Moritz Göbelbecker[4], Charles Gretton[5], Graham S. Horn[1], Richard Dearden[1], Miroslav Janicek[6], Hendrik Zender[6], and Geert-Jan Kruijff[6]

[1]*BHAM, Birmingham* [2]*Lincoln University (formerly BHAM)*
[3]*KTH, Stockholm* [4] *ALU-FREI, Freiburg*
[5]*NICTA, formerly BHAM* [6] *DFKI GmbH, Saarbrücken*

⟨jlw@cs.bham.ac.uk⟩

Due date of deliverable:	May 30 2012
Actual submission date:	May 30 2012
Lead partner:	BHAM
Revision:	final
Dissemination level:	PU

In this report we describe our experimental analysis of the the Year 3 Dora system, a robot that could plan to achieve a variety of tasks in an environment with undiscovered objects, rooms etc. In addition this robot was able to explain surprising planning failures. In the second attachment to this deliverable we give more technical detail on the method used to achieve this.

Executive Summary

This report presents an analysis of the period 3-4 Dora system. This system was able to plan information gathering activities necessary to achieve a task given by a human. The robot can reason about a variety of entities, including making assumptions about the world necessary to form a plan, modelling limited open worlds, modelling the epistemic effects of actions, replanning on the fly, and switching between decision theoretic and non-decision theoretic planning as necessary. In addition the system is able to explain surprises that result in planning failures, such as when an expected outcome essential to a plan does not occur. It seeks to explain these using the framework of assumptions (and additional background knowledge), just as it uses assumptions to produce plans under incomplete knowledge in the first place. We present a case based experimental analysis of the main system in the first attachment, and present the use of assumptions to produce explanations for surprising planning failures in the second attachment.

Role of explanations and surprises in CogX

In our plans for CogX the ability to explain surprises is one of the final stages of the scheme for self-extension. By explaining surprising results in terms of assumptions and additional background knowledge the robot is able to form hypotheses about the existence of additional objects, or additional relations between objects, or both. These hypotheses can then be tested using additional plans, and then added to the robot's knowledge thus creating a second route to self-extension (the first being exploration based on the initial knowledge and tasks as demonstrated in several of the systems presented in CogX (Dora 1, Dora2, George 1-3, Dexter 2).

Contribution to the CogX scenarios and prototypes

The results presented in this report are an analysis of the Dora system, and are thus related to the scenario on task driven information gathering and self-extension.

1 Tasks, objectives, results

1.1 Planned work

The task to which this work contributes is Task 7.6 (*Experimental study of explanation with limited extension*). The aim was to analyse a robot that can extend its representations in a limited way (e.g. extending its map). The objective addressed in the work is Objective 11 (*A robotic implementation of our theory able to complete a task involving mobility, interaction and manipulation. In the face of novelty, uncertainty, partial task specification and incomplete knowledge*). In this deliverable we have focussed on analysing a robot implementation concerned with mobility and interaction with a human, where the knowledge of the robot is incomplete at the start of the task.

1.2 Actual work performed

The attachments to the deliverable describe the results of the work performed during year 3 and 4 on the Dora system to enable it to explore and fill knowledge gaps in a task driven manner, and to explain surprises that result in planning failures when they occur. The Task 7.6 has been addressed here by a case based experimental analysis of the robot's ability to perform a variety of tasks in an environment under incomplete knowledge, by the ability of the robot to extend its representations, by the ability of the robot to plan to do so, and also by the ability of the robot to explain surprising planning failures. The objective O11 has been addressed by the development of this same system, and by the integration in the robot of various elements of our theory.

1.3 Relation to the state-of-the-art

The major difference between Dora and other previous robot systems that perform similar tasks (object search, autonomous mapping, room categorisation) is that i) Dora is re-taskable across a range of tasks whereas previous systems typically perform just one task, ii) Dora reasons about open worldness, and iii) employs a switching planner that enables both satisficing and optimising style planning.

With regard to retaskability, there are many instances of systems that perform active SLAM [4, 6] by using path planners in continuous or quantised state spaces that explicitly plan information gain over many steps, but only for that specific task. Similarly, in object search there are a number of approaches that plan within an information space, expressing the value of particular viewing locations by modelling both sensor behaviour and prior belief about object location [8]. In these planning is often greedy one step lookahead for view selection [1, 5], although [7] reasons about information over multiple steps.

All these approaches, however, path plan rather than perform task level planning, and do so within an essentially closed world using probabilistic representations of state uncertainty. Most other probabilistic planners or path planners for robots employ unstructured representations of state ([7] is an exception) that make path planning or task specific planning easy, but which do not easily lend themselves either to re-taskability or to planning in open worlds. It is a difficult problem to extend probabilistic approaches to reasoning about open worlds (i.e. where new objects, rooms etc may appear). We have developed instead developed extensions to our continual planning approach that allow planning in limited open worlds, and thus enable the robot to reason about the benefit of activities such as searching for a room of the type that is likely to contain the object searched for [2]. In addition we employ a switching planner, that swaps between a classical and a decision theoretic representation of the planning domain. This has allowed us to produce plans that reason about trade-offs quantitatively when this is computationally feasible, and the remainder of the time produce satisficing plans. Other approaches (of which we are aware) to task planning for robots that are retaskable only produce satisficing plans e.g. [3].

References

[1] A. Andreopoulos, S. Hasler, H. Wersing, H. Janssen, J.K. Tsotsos, and E. Korner. Active 3d object localization using a humanoid robot. *Robotics, IEEE Transactions on*, (99):1–18, 2011.

[2] A. Aydemir, M. Göbelbecker, A. Pronobis, K. Sjöö, and P. Jensfelt. Plan-based object search and exploration using semantic spatial knowledge in the real world. *Proc. of the European Conference on Mobile Robotics (ECMR 2011), Orebro, Sweden*, 2011.

[3] C. Galindo, J.A. Fernández-Madrigal, J. González, and A. Saffiotti. Robot task planning using semantic maps. *Robotics and Autonomous Systems*, 56(11):955–966, 2008.

[4] R. Martinez-Cantin, N. de Freitas, E. Brochu, J. Castellanos, and A. Doucet. A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, 27(2):93–103, 2009.

[5] K. Shubina and J.K. Tsotsos. Visual search for an object in a 3d environment using a mobile robot. *Computer Vision and Image Understanding*, 114(5):535–547, 2010.

[6] R. Sim and N. Roy. Global a-optimal robot exploration in slam. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 661–666. IEEE, 2005.

[7] J. Velez, G. Hemann, A. Huang, I. Posner, and N. Roy. Planning to perceive: Exploiting mobility for robust object detection. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Freiburg, Germany, 2011.

[8] Y. Ye and J.K. Tsotsos. Sensor planning for 3d object search. *Computer Vision and Image Understanding*, 73(2):145–168, 1999.

2 Annexes

2.1 Reasoning about Epistemic Actions and Uncertainty for Autonomous Knowledge Gathering

Bibliography M. Hanheide, A. Pronobis, K. Sjöö, A. Aydemir, P. Jensfelt, M. Göbelbecker, C. Gretton, G.S. Horn, R. Dearden, M.Janicek, H.Zender, G.J. Kruijff, J.L. Wyatt"Reasoning about Epistemic Actions and Uncertainty for Autonomous Knowledge Gathering". To be submitted.

Abstract In any real world task a robot tries to accomplish, it faces two significant challenges that it needs to deal with: (i) its knowledge about the world is incomplete, so substantial knowledge required to successfully achieve a given goal is missing; and (ii) the knowledge it might have about the world is uncertain, due to noise in sensing and/or unreliability of accessible knowledge sources. Nevertheless, we expect our robot to behave intelligently to achieve the given goal robustly and efficiently. In this paper we present a robot system and its architectural underpinning that addresses these challenges. It can accomplish a variety of different epistemic goals in order to extend its knowledge about the world and it features a probabilistic approach to representation, reasoning and planning that allows it to generate goal-driven behaviour in a world full of uncertainties. The robot features a novel *switching planner* that allows the system to schedule actions implemented by a number of *competences* that gather knowledge from various knowledge sources, such as: interactively by asking humans; from sensing the world through the robot's own sensors; or by exploiting common-sense background knowledge gathered from web resources. We demonstrate and analyse the behaviour of our system in real world runs with three different goals given to the robot: To autonomously explore an unknown map; to learn about the category of rooms in this map (e.g. kitchen, corridor, etc.); and to autonomously determine the location of a specific object. We show that the robot autonomously invokes competences that yield the intended information gain in order to accomplish each task.

Relation to WP WP7 is about the demonstration of the ideas in CogX as a complete systems level theory that works in real robot systems. The work presented in this deliverable is a description of the Dora system that brings together numerous contributions from across CogX. WP7 also concerns the analysis of the robot systems we develop. This deliverable also presents an case based analysis of the abilities of the Dora system and a detailed presentation of the methods used to performance explanations of surprising planning failures.

2.2 Explaining Execution Failures in Continual Planning

Bibliography M. Göbelbecker "Explaining Execution Failures in Continual Planning". Technical Report.

Abstract Continual planning is an effective approach to decision making in uncertain dynamic worlds. It involves creating plans based on assumptions about the real world and replanning if those plans fail. We discuss methods for making these assumptions explicit and providing explanations why a continual planning task may have failed or produced unexpected outcomes.

Relation to WP In this WP we developed the Dora system, which is able to explain surprising planning failures. This attachment explains how that explanation mechanism is implemented in terms of assumptions.