Dora: A Robot that Plans and Acts Under Uncertainty *

M. Göbelbecker[1], M. Hanheide[2], C. Gretton[3], N. Hawes[4], A. Pronobis, A. Aydemir, K. Sjöö[5], and H. Zender[6]

[1] Albert-Ludwigs-Universität Freiburg, Germany
[2] University of Lincoln, UK
[3] NICTA, Australia
[4] University of Birminham, UK
[5] Royal Institute of Technology (KTH), Stockholm, Sweden
[6] DFKI GmbH, Saarbrücken, Germany

Abstract. Dealing with uncertainty is one of the major challenges when constructing autonomous mobile robots. The CogX project addressed key aspects of that by developing and implementing mechanisms for self-understanding and self-extension – i.e. awareness of gaps in knowledge, and the ability to reason and act to fill those gaps.
We discuss our robot called Dora, a showcase outcome of that project. Dora is able to perform a variety of search tasks in unexplored environments.
One of the results of the project is the Dora robot, that can perform a variety of search tasks in unexplored environments by exploiting probabilistic knowledge representations while retaining efficiency by using a fast planning system.

1 Introduction

The mission statement of the CogX project is

> [...] to develop a unified theory of self-understanding and self-extension with a convincing instantiation and implementation of this theory in a robot. By self-understanding we mean that the robot has representations of gaps in its knowledge or uncertainty in its beliefs. By self-extension we mean the ability of the robot to extend its own abilities or knowledge by planning learning activities and carrying them out.[7]

The Dora robot is one instantiation of this theory, demonstrating self-extension in a task-driven way, i.e. ultimately it is driven to satisfy goals given by a user. Dora can perform a variety of object search and exploration tasks in a dynamic real-world environments. To achieve this, we have contributed to three areas of research:

* The research reported here was performed in the EU FP7 IP "CogX: Cognitive Systems that Self-Understand and Self-Extend" (ICT-215181).

[7] `www.cogx.eu`

S. Wölfl (Ed.): Poster and Demo Track of the 35th German Conference on Artificial Intelligence (KI-2012), pp. 12-16, 2012.

- Reliability in the presence of an uncertain environment requires a probabilistic representation of the world and taking these probabilities into account for decision making. Our system uses *probabilistic background knowledge* to perform tasks faster in the common case while still being able to perform in unlikely situations.
- Flexibility of tasks is provided by using a domain-independent planner. This has a number of important benefits in our setting. First, it means Dora can be given a variety of goals – e.g., search for one or more specific objects, classify a space, explore a space, etc. Second, as dialogue and sensing capabilities are added to the robot, the planner is able to plan for those without any modification to its architecture.
- To operate in an *open world*, the planner needs to be able to reason about possible gaps in the robot's knowledge and on how to fill those gaps.

In the following section we will give a short overview on related work and then present the two central components of the system. A more comprehensive description of many aspects of Dora is provided in an earlier paper by Hanheide et al. [1].

2 Related Work

A number of autonomous robot systems have been developed that use domain independent planning as their primary decision making component [2, 3], some of them also operating in open worlds. We are, however, not aware of any robot system that uses probabilistic models and domain independent planning for high-level decision making.

A number of systems have been developed specifically for the task of object search [4], but they usually treat the problems as purely geometric, not leveraging semantic information or additional sources of information, such as dialogue with a human. The system that comes closest to ours [5] uses decision theoretic planning to locate objects in large environments but requires a pre-build and annotated map.

3 The Dora System

The Dora system architecture is based on PECAS [6], consisting of various components which are grouped into subarchitectures. For this abstract, we would like to focus on two of those that are central to decision making: The *Conceptual Map*, that represents and updates the probabilistic knowledge of the system and the *Planner*, which uses this information to decide which actions to perform.

Some other subarchitectures, on which we will not expand, are the *spatial* subarchitecture which is responsible for navigation and mapping, *vision* for object detection and *dialogue* for interaction with humans. All these subarchitectures can provide state information directly to the planner, or to the conceptual map to integrate with prior knowledge.

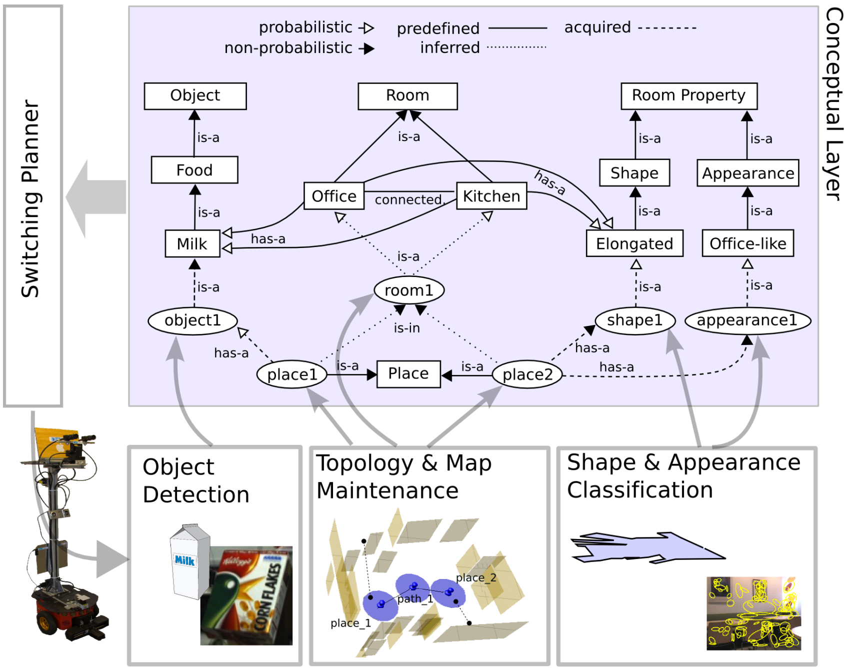

Fig. 1. A visualisation of the entities contained in the conceptual map. Sensing processes (at the bottom) transform modal sensor input into instances (shown as ellipses) and acquired relations with either other instances or known concepts (shown as boxes). Using acquired and predefined relations, the conceptual map can then infer further relations which is then used by the planning system.

3.1 Conceptual Map

The conceptual map [7], illustrated in Figure 1, uses probabilistic common sense knowledge to make connections between concepts, such as room categories and the existence of certain objects. For example, the knowledge that cereal boxes are likely to be found in kitchens can (a) help in classifying the category (e.g. kitchen) of a room in which cereals are seen, and (b) help the robot prioritise a search for cereal if a room is likely to be a kitchen. These relations were initially obtained from the *Open Mind Indoor Common Sense* database[8] and quantified by queries to an online image search engine [8, 9].

3.2 Planning

Our system uses a domain independent planner that combines decision theoretic reasoning with fast classical continual planning [10]. For high level decision making, which we call *sequential sessions*, it operates according to the continual

[8] `http://openmind.hri-us.com/`

planning principle: It does not create a full plan for all contingencies but computes one serial plan and monitors its execution, replanning if the robot ends up in a state that makes the rest of the plan invalid. While the classical planner used in this step is inherently deterministic, is can take probabilities into account by making *assumptions* (such as "room 0 may be a kitchen"). Less likely assumptions lead to higher costs of the resulting plan, thus leading to plans that tend to rely on more likely facts.

The replanning method employed in the sequential session makes it easy to support *open worlds*, in which entities (such as rooms or objects) can appear and disappear from the environment. To reason about finding new objects in the planner, we add a number of virtual objects to the initial planning state which the planner can instantiate and use later on [11]. For example, if the only room known to the robot has a low likelihood of containing a cereal box, it may decide to explore unknown space, hoping to eventually find a room where finding cereals is more likely (e.g. a kitchen).

For problems that involve uncertainty and (possibly noisy) sensing, it is usually more appropriate to model it as a partially observable Markov decision process (POMDP). POMDPs allow accurate modelling noisy sensing, such as the effect of false positive and false negative rates in object detectors. As solving large-scale POMDPs is infeasible, our planner tries to identify the parts of the sequential plan in which taking these effects into account is important, and uses a separate, *decision theoretic planner* to solve these subproblems.

4 Conclusion

We presented a robot system that approaches the problem of acting and reasoning in dynamic, open and uncertain worlds by integrating two approaches: the conceptual map that integrates uncertain observations with probabilistic conceptual knowledge and a fast continual planning system that can exploit these representations to efficiently find plans to solve the given tasks. The resulting system can operate autonomously in unknown environments while still being able to solve tasks quickly.

References

1. Hanheide, M., Gretton, C., Dearden, R., Hawes, N., Wyatt, J., Pronobis, A., Aydemir, A., Goebelbecke, M., Zender, H.: Exploiting probabilistic knowledge under uncertain sensing for efficient robot behaviour. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011). (2011)
2. Talamadupula, K., Benton, J., Kambhampati, S., Schermerhorn, P., Scheutz, M.: Planning for human-robot teaming in open worlds. ACM Trans. Intell. Syst. Technol. **1** (December 2010) 14:1–14:24
3. Kraft, D., Başeski, E., Popović, M., Batog, A.M., Kjær-Nielsen, A., Krüger, N., Petrick, R., Geib, C., Pugeault, N., Steedman, M., Asfour, T., Dillmann, R., Kalkan,

S., Wörgötter, F., Hommel, B., Detry, R., Piater, J.: Exploration and planning in a three-level cognitive architecture. In: CogSys. (2008)

4. Shubina, K., Tsotsos, J.: Visual search for an object in a 3D environment using a mobile robot. Computer Vision and Image Understanding **114**(5) (2010) 535–547
5. Kunze, L., Beetz, M., Saito, M., Azuma, H., Okada, K., Inaba, M.: Searching objects in large-scale indoor environments: A decision-thereotic approach. In: IEEE International Conference on Robotics and Automation (ICRA), St. Paul, MN, USA (May 14–18 2012)
6. Hawes, N., Brenner, M., Sjöö, K.: Planning as an architectural control mechanism. In: HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction, New York, NY, USA, ACM (2009) 229–230
7. Pronobis, A., Mozos, O.M., Caputo, B., Jensfelt, P.: Multi-modal semantic place classification. Int. J. Robot. Res. **29**(2-3) (February 2010) 298–320
8. Zhou, K., Zillich, M., Zender, H., Vincze, M.: Web mining driven object locality knowledge acquisition for efficient robot behavior. In: Proc. Int. Conf. Intelligent Robots and Systems (IROS). (to appear)
9. Zender, H.: Multi-layered conceptual spatial mapping – representing spatial knowledge for situated action and human-robot interaction. In Amirat, Y., Chibani, A., Zarri, G.P., eds.: Bridges Between the Methodological and Practical Work of the Robotics and Cognitive Systems Communities – From Sensors to Concepts. Intelligent Systems Reference Library. Springer Verlag, Berlin/Heidelberg, Germany (to appear)
10. Göbelbecker, M., Gretton, C., Dearden, R.: A switching planner for combined task and observation planning. In: Twenty-Fifth Conference on Artificial Intelligence (AAAI-11). (August 2011)
11. Aydemir, A., Göbelbecker, M., Pronobis, A., Sjöö, K., Jensfelt, P.: Plan-based object search and exploration using semantic spatial knowledge in the real world. In: Proc. of the European Conference on Mobile Robotics (ECMR 2011), Örebro, Sweden (sep 2011)