# Dora, a Robot Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Object Search*

**Marc Hanheide, Charles Gretton**
University of Birmingham
Birmingham, UK

**Moritz Göbelbecker**
Albert-Ludwigs-Universität
Freiburg, Germany

**Abstract**

Dora, the robot, is trying to find object in its environment. Instead of just exhaustively searching everywhere, Dora is equipped with probabilistic reasoning, representations, and planning to exploit uncertain common-sense knowledge, such as that cornflakes are usually found in kitchens, while also accounting for the uncertainty of sensing in the real-world. Dora demonstrates how to combine task and observation planning in the presence of uncertainty by autonomously switching between contingent and sequential planning sessions. The demonstration emphasises the benefit of employing a robot with common-sense knowledge and the benefit of the switching planner.

## Introduction

With Dora, we are presenting results of our efforts to build a robot capable of performing tasks on demand in dynamic real-world environments. With this paper and demonstration we explicitly address the challenge to perform task and observation planning under uncertainty in pursuit of current robot goals by presenting a new planning approach to reason with new representations of space. For Dora we integrate probabilistic models of background conceptual knowledge, and the visual appearance of objects and of room categories, to solve an object search task. These models are used to create and maintain a probability distribution over possible states with respect to the spatial structure, the categories of objects and rooms, and their relations to each other. Dora, as presented in this paper, is a successor of a previous system (Hawes et al. 2011) that did not make use of probabilistic representations and featured only a classical, sequential planner (Helmert 2006) to achieve exploration and categorisation of rooms.

## Related Work

Probabilistic representations are employed for many localised functions in robots operating in the real world. For example, Thrun et al. (2000) use such representations in most of their system's individual components, but their robot behaviour is generated using a reactive controller rather than a domain-independent planner as here.

A number of recent integrated robotic systems incorporate a high-level *continual planning* and execution monitoring subsystem (Talamadupula et al. 2010; Kraft et al. 2008). For the purpose of planning, sensing is modelled deterministically, and beliefs about the underlying state are modelled qualitatively. We are not aware of any robot system that features both a unifying probabilistic representation, and a domain-independent planner which is able to reason quickly over that unified decision-theoretic model to generate behaviour.

Object search with mobile robots has been studied for almost 20 years (Shubina and Tsotsos 2010), yet no previous system reasons with uncertain conceptual knowledge about room and object categories. Instead, most dedicated systems treat the problem as a geometric one. Closest to our approach is the work by Sjöö et al. (2010) who used common-sense knowledge encoded into a rule-based ontology to inform a deterministic planner which previously categorised room to search for a particular object. Bouguerra, Karlsson, and Saffiotti (2007) extended this approach to treat some of the conceptual knowledge as uncertain, although restricted to the number of occurrences of object types in rooms. Vasudevan and Siegwart (2008) went beyond this to perform room categorisation through Bayesian reasoning about the presence of objects, but did not (as none of these did) include observation models in their reasoning (thus perception was still considered to be deterministic).

## The Dora System

In order to perform its object search task Dora is equipped with a camera on a pan-tilt unit, a laser scanner, and one laptop accommodating all the processes. The system architecture itself is an extension of PECAS (Hawes, Brenner, and Sjöö 2009) composed of many components functionally structured into subarchitectures. In general, Dora features speech understanding and dialogue components to receive commands from humans, a goal management subsystem (Hanheide et al. 2010) translating commands into goals for the planning subsystem, and many other utility components whose description goes beyond the scope of this extended abstract. At the core of the system is the *switching planner*. Its role is to deliberate Dora's behaviour to ef-

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

ficient object search, exploiting common-sense, relational, and conceptual knowledge. It operates on the probabilistic belief state defined by the conceptual layer of the spatial representation. This probabilistic representation utilises a chain graph model (Lauritzen and Richardson 2002) for inference and integrates conceptual and instance knowledge as detailed by Hanheide et al. (2011), with the latter continuously being maintained by perceptual processes.

Among these are processes that maintain metric and topological maps (Pronobis et al. 2009). Following an approach described by Hawes et al. (2011) the map is discretised into places and rooms, yielding discrete instances to plan with. Also, we employ a continuously running process recognising properties which are evident of the category of rooms following work by Pronobis et al. (2010). Properties being recognised here are the shape of rooms and their visual appearance. Also, we employ an object detector (Mörwald et al. 2010), pre-trained for a set of 19 objects of interest. The planning system can invoke this object detector as part of a sensing action to get evidence about the existence of an object in the current view of the robot.

The probabilistic relations in the conceptual layer have to be quantified appropriately. For probabilistic relations between instances and concepts these are derived from the sensing processes. For the relations of common-sense and conceptual knowledge we either derive them from training sets or from harvesting information from the web. As will be demonstrated, Dora is capable of exploiting the probabilistic knowledge about the co-occurrence of objects and rooms. This relation was quantified employing a combination of qualitative bootstrapping from the *Open Mind Indoor Common Sense* database[1] and queries to an online image search engine. This offline acquisition process yields conditional probabilities, such as $P(room = \text{kitchen}|object = \text{cereal\_box}) = 0.33$.

## The Switching Planner

To generate flexible goal-oriented behaviour our system employs a domain-independent planner. The object search scenario poses several challenges to the planning system: On the one hand, planning and execution monitoring must be lightweight, robust, timely, and should span the lifetime of the robot. Those processes must seamlessly accommodate exogenous events, changing objectives, and the underlying *unpredictability* of the environment. On the other hand, in order to act intelligently the agent must perform computationally expensive reasoning about *contingencies*, and possible revisions of subjective belief according to quantitatively modelled uncertainty in acting and sensing.

In our work we take a concrete step towards addressing the challenges we outlined. We have developed a *switching* domain-independent planning system that operates according to the continual planning paradigm. It uses first-order declarative problem and domain representations, expressed in a novel extension of PPDDL (Younes et al. 2005) called *Decision-Theoretic (DT)PPDDL*, for modelling stochastic decision problems that feature partial observability. The system *switches*, in the sense that the underlying planning procedure changes depending on our robot's subjective degrees of belief, and progress in plan execution. When the underlying planner is a deterministic sequential planner, i.e., a *classical* planner, we say planning is in a *sequential* session, and otherwise it is in a *contingent* session. Finally, planning is continual in the usual sense that, whatever the session, plans are adapted and rebuilt online in reaction to changes to the planning model (e.g. when objectives are modified, or when our robot's path is obstructed by a door being closed). By autonomously mixing these two types of sessions our robot is able to be robust and responsive to changes in its environment *and* make appropriate decisions in the face of uncertainty. We will give a brief overview of the approach, a more detailed description can be found in the literature (Göbelbecker, Gretton, and Dearden 2011).

[1]http://openmind.hri-us.com/

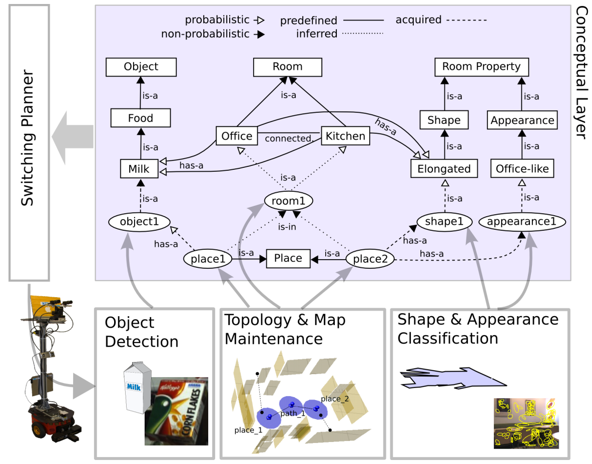


Figure 1: An abstract view of the processes and representations of the system. Sensing processes (at the bottom) discretise and categorise sensor input into instances (shown as ellipses) and acquired relations in conceptual layer. This layer also comprises knowledge about concepts (rectangles) of which only an excerpt in shown. The switching planner reasons upon the state distribution given by the conceptual map.

### Sequential Planning

During a sequential session, a rewarding *trace* of a possible execution is computed using a modified version of the cost-optimising satisficing planner *Fast Downward* (Helmert 2006) which trades action costs, goal rewards, and determinacy.

The planning model we use for specifying the sequential planning problems is an extended $\text{SAS}^+$formalism (Bäckström and Nebel 1995). In contrast to probabilities in more expressive models like MDPs, actions do not have multiple possible outcomes, they just can succeed with probability $\mathrm{p}(a)$ or transition into a sink state with probability of $1 - \mathrm{p}(a)$. "Real" probabilistic actions can be approximated by creating a separate action for every possible outcome (Yoon, Fern, and Givan 2007). The planner plans according to a cost function $\mathrm{c}$ that weights the cost of a plan against its probability. There are several possible choices for how to combine costs and probabilities, we chose a function that resembles the expected reward adjusted to our restricted planning model. With $R$ being a reward constant, we minimise the formula $c(\pi) = \sum_{a\in\pi} c(a) + R\left(1 - \prod_{a\in\pi} p(a)\right)$. For small values of $R$ the planner will prefer cheaper but more unlikely plans, for larger values more expensive plans will be considered.

**Assumptions** To model uncertain initial states (which are an essential feature of exploration problems), we introduce the concept of *assumptive actions*. The initial state of the planning problem is the set of necessarily true propositions. Assumptive actions are then used to add other, uncertain, propositions (assumptions) to the state. Provided that the plan is optimal, only assumptions that help achieving the goal will be included, preferring ones that are more likely.

If, for example, the initial state contains uncertainty about the category of a room, with $P(\text{cat}(\text{room}) = \text{kitchen}) = 0.5$, $P(\text{cat}(\text{room}) = \text{office}) = 0.3$ $P(\text{cat}(\text{room}) = \text{corridor}) = 0.2$. We would then add the assumptions:

$$
\begin{aligned}
\text{pre}(a_1) =& \text{pre}(a_2) = \text{pre}(a_3) = \{\text{def}_{cat(\text{room})} = \bot\} \\
\text{eff}(a_1) =& \{\text{cat}(\text{room}) = \text{kitchen}, \text{def}_{\text{cat}(\text{room})} = \top\} \\
\text{p}(a_1) =& 0.5 \qquad \text{c}(a_1) = 0 \\
\text{eff}(a_2) =& \{\text{cat}(\text{room}) = \text{office}, \text{def}_{\text{cat}(\text{room})} = \top\} \\
\text{p}(a_2) =& 0.3 \qquad \text{c}(a_2) = 0 \\
\text{eff}(a_3) =& \{\text{cat}(\text{room}) = \text{corridor}, \text{def}_{\text{cat}(\text{room})} = \top\} \\
\text{p}(a_3) =& 0.2 \qquad \text{c}(a_3) = 0
\end{aligned}
$$

The def-variable makes sure that we cannot make more than one assumption about the same variable.

To utilise background conceptual knowledge, e.g. the probability of finding an object in a certain type of room, we use operators that model the conditional dependencies[2]:

```
(:action object-in-room
   :parameters (?cl - class ?r - room
                ?c - category)
   :probability (P-obj-given-category ?cl ?c)
   :precondition (= (cat ?r) ?c)
   :effect (obj-exists ?cl in ?r))
```

where `(P-obj-given-category ?cl ?c)` are fluents containing the probabilities. Using these operators, we do not have to construct the entire initial state description of the problem explicitely (as we did in the original description of the switching planner).

The system always begins with a sequential session, and once *Fast Downward* produces a trace, plan execution proceeds by applying actions from that trace in sequence until the applicability of the next scheduled action is too uncertain according to a threshold parameter (here, set at 95%). A contingent session then begins which tailors sensory processing to determine whether the assumptions made in the trace hold, or which otherwise acts to achieve the overall objectives.

[2] def-conditions and effects are omitted for clarity

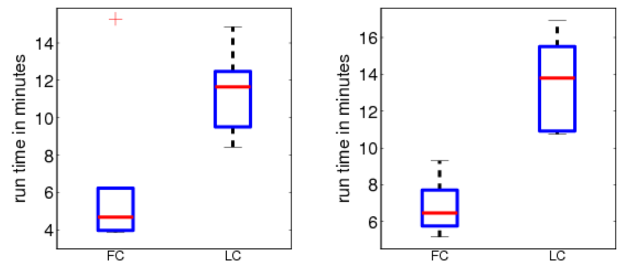


(a) BHAM: 19 runs. (b) KTH: 10 runs.

Figure 2: Box and whisker diagrams of total runtime to achieve the given task in two environments comparing the 'full' system (FC) to the 'lesioned' case (LC).

## Contingent Planning

Because decision-theoretic planning in large problems is too slow for our purpose (we seek response times in seconds), contingent sessions plan in an abstract decision process determined by the current trace and underlying belief-state. This abstraction is constructed by first excluding all propositions that are not true of any state in the trace, then adding them back, using as a heuristic the entropy of the trace assumptions conditional on a candidate proposition. Propositions are added, one at a time, until the number of states in the initial belief-state reaches a given threshold (here, 150 states). To the resulting abstract model we also add *disconfirm* and *confirm* actions that the contingent session can schedule in order to judge an atomic assumption in the trace. In the abstract model these actions yield a small reward if the corresponding judgement is true (or small penalty otherwise). Once a judgement action is scheduled for execution the contingent session is terminated, and control is returned to a sequential session.

# Experimental Evaluation

In order to test the effectiveness of (i) exploiting default probabilistic knowledge in a conceptual layer of our representation, (ii) the switching planner, and (iii) our implementation of the overall system, we ran two configurations ('full' and 'lesioned') of the system in two natural world environments; a residential house in Birmingham (BHAM) and a floor of offices and a kitchen at KTH Stockholm (KTH).

Our evaluation compares the full system with a lesioned system in which the categorisation of visual appearance and shape properties has been disabled, emulating the limited reasoning capabilities available in our previous system (Hawes et al. 2011), where no such evidence was available. The task in all these runs was to find a box of cornflakes. The starting position of the robot was either the living room (in BHAM) or an office (in KTH), i.e. rooms that according to the acquired common-sense knowledge are quite unlikely to contain objects of type cornflakes. This was chosen to showcase the benefit of the probabilistic representation and planning.

Fig. 2 shows the overall runtime to complete the object search task in the lesioned (denoted as 'LC' in the figure) and the full system ('FC') in both environments. What can clearly be seen from the figure is that the full system which can exploit the evidence about the categories of rooms achieves the task significantly faster (Mann-Whitney test $p < 0.01$ for both environments) on average. It benefits from the probabilistic common-sense knowledge that it is quite unlikely to find cornflakes in the room the robot was in and made it decide to first drive to the kitchen to start the search there. On the contrary, in the lesioned case the robot had less information and had to conduct a full exhaustive search. So it started its search in the living room or office, respectively, because the object is as likely to be in this room than in any other. Further details and a more exhaustive analysis of the results are given in (Hanheide et al. 2011).

## References


Bäckström, C., and Nebel, B. 1995. Complexity results for SAS$^+$ planning. *Comp. Intell.* 11(4):625–655.

Bouguerra, A.; Karlsson, L.; and Saffiotti, A. 2007. Handling uncertainty in semantic-knowledge based execution monitoring. In *Proc. Int. Conf. Intelligent Robots and Systems (IROS)*, 437–443.

Göbelbecker, M.; Gretton, C.; and Dearden, R. 2011. A switching planner for combined task and observation planning. In *Twenty-Fifth Conference on Artificial Intelligence (AAAI-11)*.

Hanheide, M.; Hawes, N.; Wyatt, J. L.; Göbelbecker, M.; Brenner, M.; Sjöö, K.; Aydemir, A.; Jensfelt, P.; Zender, H.; and Kruijff, G.-J. M. 2010. A Framework for Goal Generation and Management. In *Proceedings of the AAAI Workshop on Goal-Directed Autonomy*.

Hanheide, M.; Hawes, N.; Gretton, C.; Zender, H.; Pronobis, A.; Wyatt, J.; Göbelbecker, M.; and Aydemir, A. 2011. Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI)*. to appear.

Hawes, N.; Hanheide, M.; Hargreaves, J.; Page, B.; Zender, H.; and Jensfelt, P. 2011. Home Alone : Autonomous Extension and Correction of Spatial Representations. In *Proc. Int. Conf. on Robotics and Automation*.

Hawes, N.; Brenner, M.; and Sjöö, K. 2009. Planning as an architectural control mechanism. In *HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, 229–230. New York, NY, USA: ACM.

Helmert, M. 2006. The fast downward planning system. *J. Artif. Intell. Research* 26:191–246.

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.; and Piater, J. 2008. Exploration and planning in a three-level cognitive architecture. In *CogSys*.

Lauritzen, S. L., and Richardson, T. S. 2002. Chain graph models and their causal interpretations. *J. Roy. Statistical Society, Series B* 64(3):321–348.

Mörwald, T.; Prankl, J.; Richtsfeld, A.; Zillich, M.; and Vincze, M. 2010. BLORT – The Blocks World Robotic Vision Toolbox. In *Proc. ICRA Workshop Best Practice in 3D Perception and Modeling for Mobile Manipulation*.

Pronobis, A.; Sjöö, K.; Aydemir, A.; Bishop, A. N.; and Jensfelt, P. 2009. A framework for robust cognitive spatial mapping. In *Proceedings of the 14th International Conference on Advanced Robotics (ICAR09)*.

Pronobis, A.; Mozos, O. M.; Caputo, B.; and Jensfelt, P. 2010. Multi-modal semantic place classification. *Int. J. Robot. Res.* 29(2-3):298–320.

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

Sjöö, K.; Zender, H.; Jensfelt, P.; Kruijff, G.-J. M.; Pronobis, A.; Hawes, N.; and Brenner, M. 2010. The Explorer system. In Christensen, H. I.; Kruijff, G.-J. M.; and Wyatt, J. L., eds., *Cognitive Systems*. Springer. 395–421.

Talamadupula, K.; Benton, J.; Kambhampati, S.; Schermerhorn, P.; and Scheutz, M. 2010. Planning for human-robot teaming in open worlds. *ACM Trans. Intell. Syst. Technol.* 1:14:1–14:24.

Thrun, S.; Beetz, M.; Bennewitz, M.; Burgard, W.; Cremers, A.; Dellaert, F.; Fox, D.; Hähnel, D.; Rosenberg, C.; Roy, N.; Schulte, J.; and Schulz, D. 2000. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *International Journal of Robotics Research* 19(11):972–999.

Vasudevan, S., and Siegwart, R. 2008. Bayesian space conceptualization and place classification for semantic maps in mobile robotics. *Robot. Auton. Syst.* 56:522–537.

Yoon, S.-W.; Fern, A.; and Givan, R. 2007. FF-replan: A baseline for probabilistic planning. In *Proc. Int. Conf. Automated Planning and Scheduling (ICAPS)*.

Younes, H. L. S.; Littman, M. L.; Weissman, D.; and Asmuth, J. 2005. The first probabilistic track of the international planning competition. *J. Artif. Intell. Res.* 24:851–887.