

Dora The Explorer: A Motivated Robot

Nick Hawes^{*}, Marc
Hanheide
Intelligent Robotics Lab
School of Computer Science
University of Birmingham, UK

Kristoffer Sjöö, Alper
Aydemir, Patric Jensfelt
Centre for Autonomous
Systems
Royal Institute of Technology
Stockholm, Sweden

Moritz Göbelbecker,
Michael Brenner
Institut für Informatik
Albert-Ludwigs-Universität
Freiburg, Germany

Hendrik Zender, Pierre
Lison, Ivana
Kruijff-Korbayová,
Geert-Jan Kruijff
Language Technology Lab
German Research Center for
Artificial Intelligence (DFKI)
Saarbrücken, Germany

Michael Zillich
Vision for Robotics (V4R)
Automation and Control
Institute
Vienna University of
Technology

ABSTRACT

Dora the Explorer is a mobile robot with a sense of curiosity and a drive to explore its world. Given an incomplete tour of an indoor environment, Dora is driven by internal motivations to probe the gaps in her spatial knowledge. She actively explores regions of space which she hasn't previously visited but which she expects will lead her to further unexplored space. She will also attempt to determine the categories of rooms through active visual search for functionally important objects, and through ontology-driven inference on the results of this search.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Robotics

General Terms

Algorithms, Experimentation

Keywords

cognitive robotics, motivation, exploration, mapping, reasoning

1. INTRODUCTION

It has been a long standing aim of the robotics community to develop a robot capable of being a useful assistant in the home or workplace. There are a great many barriers facing such a development. One such barrier is that

^{*}Contact author: n.a.hawes@cs.bham.ac.uk

Cite as: Dora The Explorer: A Motivated Robot, Hawes et al., *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX. Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

current systems require a lot of knowledge about an area before they can perform tasks in it. If you were to ask your interactive robot assistant “bring me the milk from the kitchen”, you would only be likely to get the milk if the robot knew the complete layout of the building, how the humans working there describe the rooms, where objects are typically found, and many other things. This information could be programmed *a priori*, or could be provided by a human during a tour when the robot was first received. These approaches have two problems. First, they are rather demanding on the time of humans; the more information the robot requires, the more work a human has to do to provide it. This will become increasingly true as mobile vision and manipulation improves. Second, the world will continually change throughout the robot's lifetime. This will render the initial information useless, and require additional programming or human-led training.

Our solution to this problem is to allow the robot to gather knowledge autonomously. We do this by allowing it to explicitly model *gaps in its own knowledge*, which it can then proactively attempt to fill by performing knowledge gathering actions such as sensing and reasoning. This paper summarises a demo which instantiates this approach in Dora the Explorer, a mobile robot intended to perform human-specified tasks (such as the one described above) in an office environment. Dora is able to model two different types of knowledge gaps: gaps in her spatial knowledge and gaps in her knowledge about the functional categories of rooms. Spatial knowledge gaps represent areas in space which Dora knows about but hasn't visited yet. They are derived from laser scan readings combined with a metric map (built at run-time). These gaps are filled by Dora driving into the previously unvisited space. Categorical knowledge gaps represent rooms which Dora knows about, but which haven't been assigned categories. Categorical gaps are generated by ontology-based reasoning over a topological map built on top of the metric map. These gaps are filled by searching for objects in the current room and using the results to infer

its function. For example, if a stove was found in a room, Dora might hypothesise that the room is a kitchen. The following section summarises the techniques used in the system to support such behaviour.

2. ARCHITECTURE

Dora's knowledge gathering is performed by following plans generated at run-time. Embedding planning into a heterogeneous robot system which itself is embodied in a dynamic, unpredictable world, requires a supporting architecture. Our architectural approach is an extension of PECAS [1]. The whole system is divided into function-based *subarchitectures*, each of which contain processing components sharing information via a working memory (WM). Modal (i.e. sensor-based) subarchitectures (e.g. mapping, vision, language) each store local representations on their WM. These modal representations are then fused into a single amodal representation by a *binding* subarchitecture, which reasons about connections between modalities. Binding provides a single view of the system's knowledge which can be used to generate planning states. The representation used by the system at this level is comparable to propositional logic. Because PECAS is intended for systems operating in multi-agent, dynamic worlds, it uses *continual planning* and *execution monitoring* to cope with partial observability and remain responsive to change.

In addition to this existing core, the Dora system incorporates a number of innovations driven by the demands of autonomous knowledge gathering: goal generation and management; planned exploration of unknown space in a new spatial model; and active visual search leading to ontology-based room categorisation. These developments, and the role they play in the demonstration, are described in the following paragraphs.

Although the process of planning has been widely researched, a comparatively small amount of attention has been directed towards where the goals for planning processes come from. In Dora we have been exploring an architecture for goal generation and management based on the work of Wright et al. [3]. This architecture is composed of reactive *goal generators* which create new goals from modal and amodal WM content; a collection of *filters* which do a first pass selection of goals to be considered for activation; and *management mechanisms* which determine which of the remaining goals should be *activated* (i.e. planned for). The architecture allows multiple new goals to be generated asynchronously by the system (e.g. when a new area of space is sensed, or when a command is given), whilst also determining which collection of goals should currently be pursued by the system (e.g. which bit of space should be explored, or which class of goal should be pursued).

A representation of space is an essential part of any mobile robot. Most current techniques provide the ability to map an area and localise within this map, but do not lend themselves to the generation of symbols for planning or other higher-level reasoning tasks. In Dora we use a new place-based representation developed with this purpose in mind [2]. In particular, Dora has been used to investigate how unexplored space can be represented in such a model. Areas where Dora's laser detects free space which is not already part of an existing place is noted as a *frontier*. Frontiers are aggregated into *placeholders* which indicate the potential for generating a new place (and thus a new spatial symbol). The

presence of a placeholder triggers a goal generator to create a goal to fill the corresponding area of space by exploration. This goal is only selected if it passes through the filter and management mechanisms.

In Dora we make the assumption that the presence of particular objects determines the functional category of a room. To this end we have given Dora a decision logic-based reasoner populated with facts from the Open Mind Indoor Commonsense database describing relationships between object presence and room type (e.g. if you see a printer you might be in an office or computer room). When Dora detects a room without a category label, a goal generator creates a goal to categorise it. If this goal is activated, the plan produced causes Dora to travel to the room in question and perform a visual search for known objects. This is done by generating a view plan of regions of the room which might contain objects, then running an object recogniser from these views. When an object is found, a representation is stored on WM where the reasoner accesses it and adds it to its database. These additions, coupled with the aforementioned rules, allow Dora to infer the category of the room being searched (satisfying the planning goal).

3. DEMO

In the demo, Dora is given a short tour of an indoor area. After the tour, her goal filters are switched to allow previously generated goals to compete for activation. The user can manually adjust the filters to set priorities for classes of goals. Depending on these filters, and a cost/benefit analysis of the individual goals, Dora will select the goal or goals to pursue next, creating and executing plans to fill knowledge gaps. As she explores the world, new goals are created which enter the management architecture and influence behaviour. An example of this behaviour is that Dora can pass an open door leading to unexplored space, choose to change direction to pass through the door, then decide to explore and categorise the room beyond it. After this she can choose to readopt the goal that led her past the door originally, or choose something else that appears more rewarding. A video of the demo can be seen at <http://cogx.eu/results/dora/>.

4. ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement No. 215181, CogX.

5. REFERENCES

- [1] N. Hawes, M. Brenner, and K. Sjöö. Planning as an architectural control mechanism. In *HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 229–230, New York, NY, USA, 2009. ACM.
- [2] A. Pronobis, K. Sjöö, A. Aydemir, A. N. Bishop, and P. Jensfelt. A framework for robust cognitive spatial mapping. In *Proceedings of the 14th International Conference on Advanced Robotics (ICAR09)*, Munich, Germany, June 2009.
- [3] I. Wright, A. Sloman, and L. Beaudoin. Towards a design-based analysis of emotional episodes. *Philosophy Psychiatry and Psychology*, 3(2):101–126, 1996.