KI – Künstliche Intelligenz manuscript No.
(will be inserted by the editor)

OPEN-EASE – A Cloud-Based Knowledge Service for Autonomous Learning

Moritz Tenorth · Jan Winkler · Daniel Beßler · Michael Beetz

the date of receipt and acceptance should be inserted later

Abstract We present OPEN-EASE, a cloud-based knowledge base of robot experience data that can serve as episodic memory, providing a robot with comprehensive information for autonomously learning manipulation tasks. OPEN-EASE combines both robot and human activity data in a common, semantically annotated knowledge base, including robot poses, object information, environment models, the robot's intentions and beliefs, as well as information about the actions that have been performed. A powerful query language and inference tools support reasoning about the data and retrieving information based on semantic queries. In this paper, we focus on applications of OPEN-EASE in the context of autonomous learning.

1 Introduction

Autonomous learning for robotic agents has recently gained substantial interest in the research community. The basic idea is that the learning methods do not have to depend on programmers who provide suitable training examples for learning. Rather, autonomous learning investigates how "autonomous systems can efficiently learn from the interaction with the environment, especially by having an integrated approach to decision making and learning, allowing systems to autonomously decide on actions, representations, hyperparameters and model structures for the purpose of efficient learning" [1]. Indeed, research on the co-development of learning algorithms and representations suitable for learning has achieved substantial progress, with deep learning possibly being the most prominent example.

Yet, two aspects of autonomous learning seem to be barriers for even larger impact on autonomous robot control. First, autonomous learning so far typically works within a carefully defined mathematical framework without the agent being able to look beyond the structural, representational limits of said framework. Characteristic examples are robotic agents that learn motion skills. Such robots can learn motion skills that optimize some given objective functions, but they are in most cases unable to reason about the consequences of their motions. Consider a robot that learns to flip pancakes. Such a robot could learn to perform dexterous motions of the spatula, but not to understand how the motion is to be adapted if the pancake sticks to the pan or if a knife is used instead of a spatula. Handling

Institute for Artificial Intelligence, Universität Bremen, Am Fallturm 1, 28359 Bremen, Germany

such situations requires common sense (e.g. regarding physics). Second, the data sets that are used for learning are often too narrow in their scope. If a robot is to learn flipping pancakes it is not sufficient to learn from the position, velocity, acceleration, and force data of the motion. The learning data does also have to include the intentions, beliefs, perceived information, predicted effects, etc. Humans have developed an episodic memory that records experiences in a very comprehensive and detailed fashion in order to provide the informational basis for autonomous learning. We propose a publicly accessible knowledge service, called OPEN-EASE, to overcome these two barriers for autonomous robot learning. OPEN-EASE provides an infrastructure for enabling robots to collect episodic memories that do not only provide the pose and sensor data streams but also a symbolic structure on top of this data that enables the robots to reason about what they did, how, what happened when they did it, why they did it, what they believed and what decisions they made. This data is available in the form of a knowledge representation that is based on description logic and Prolog-based reasoning.

If we aim at robots that are able to autonomously improve their course of action based on data they record during complex tasks (which forms a kind of "episodic memory"), we have to account for this complexity in the learning setup. A memory system for providing data for autonomously learning models of robot manipulation tasks therefore has to fulfill the following requirements:

(○) *Comprehensive data* that covers all relevant aspects a robot needs to consider, ranging from the geometry of the robot and the scene, the robot's movements over time and its sensory percepts up to the types, parameters, durations, and results of the performed actions

(○) An *expressive representation* that sets the individual pieces of information into relation and assigns meaning to them, e.g. that a pose denotes the position of the robot's gripper

(○) A *query language* operating on the data that allows the robot to select the information that it needs for a learning problem at hand, for example all trajectories of the right gripper during reaching motions, or the pose of its camera and all known objects in the surroundings at the times of failed perception tasks

With the OPEN-EASE project [2], we aim at establishing an episodic memory system that provides autonomous robots with unprecedented memorization and reasoning capabilities. It consists of a large semantically annotated database of comprehensive log data of robot manipulation tasks which is combined with a representation and query language based on time interval logics. Examples of such annotations are the hierarchical relation between tasks, their start and end times, types, which robots operated on which objects in which contexts, and what an action's input parameterization and output was. The symbolic annotations are closely linked to subsymbolic information such as current joint states of a robot, absolute viewing direction of a camera, and characteristics of detected objects (color, shape, type, etc.). The system further includes tools for recording this data during task execution without slowing down the robot, for exploring and visualizing the data using a web-based frontend, and for sharing the data with other researchers via a cloud-based platform. In this article, we focus on the applications of the system and its contributions in the context of autonomously learning models of robot manipulation tasks. For details on the implementation, the system components and the query language, the interested reader is referred to the original conference paper [2].

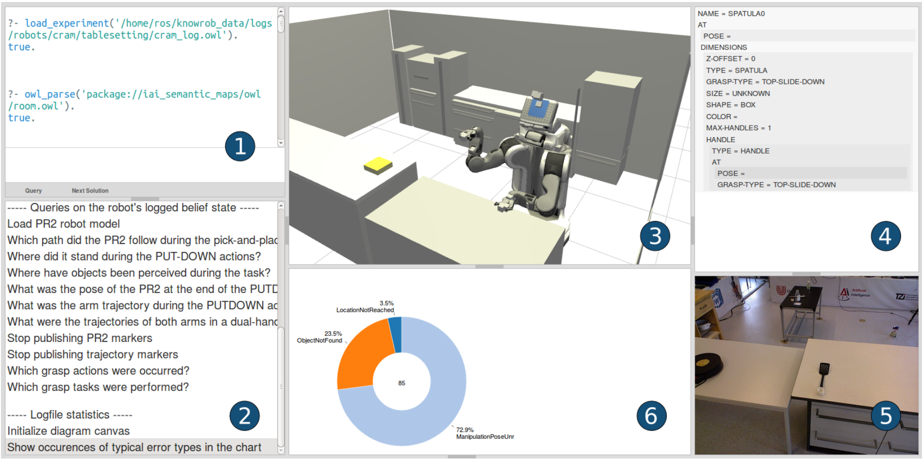

Fig. 1: Web interface of OPEN-EASE.

2 System Overview

Figure 1 shows the web-based frontend of the OPEN-EASE system. The *Prolog interaction pane* (1) allows the user to type Prolog queries and commands and to see the answers to these queries. Prepared queries with English translation are provided in the *query list pane* (2). The *3D display pane* (3) visualizes query results such as robot poses, its environment, trajectories and object poses. The *belief pane* (4) enables the user to inspect the internal data structures of the robot's control program including descriptions of objects, actions and locations. The *image pane* (5) can display images captured by the robot's camera, and the *visual analytics pane* (6) can visualize statistical data as bar charts and pie charts. Examples of statistical data to display are failure- and task type distributions for getting a quick overview of what happened in the analyzed dataset.

The OPEN-EASE system consists of two main components, one for recording data during manipulation episodes and one for analyzing this data using the web-based frontend. The former has been described in detail in a previous article [3], the latter has been implemented as a cloud-based version of the KNOWROB robot knowledge base [4]. The data recording system logs comprehensive data during robot manipulation tasks and stores it in the form of an episodic memory. High-volume, continuous sensor data is stored in an efficient, schema-less MongoDB database. Symbolic plan events, such as the hierarchy of actions and sub-actions that are performed, their parameters, results and durations, are stored in a knowledge base. Both the symbolic and continuous log data are represented with respect to a common ontology, which allows standardized semantic access to all information in the system. The query interface can either be used by humans via the aforementioned web-based frontend, or by robots via a WebSocket connection. It offers a library of query predicates for reasoning about the logs that extract symbolic knowledge from the subsymbolic log data at query time as needed. We call this concept a "virtual knowledge base" that is created on top of the semi-structured and often high-volume log data. The conceptual connection between these components is shown in Figure 2: A "virtual knowledge base" provides a symbolic view on the continuous, subsymbolic sensor data stored in the database, which can then easily be combined with symbolic plan events.

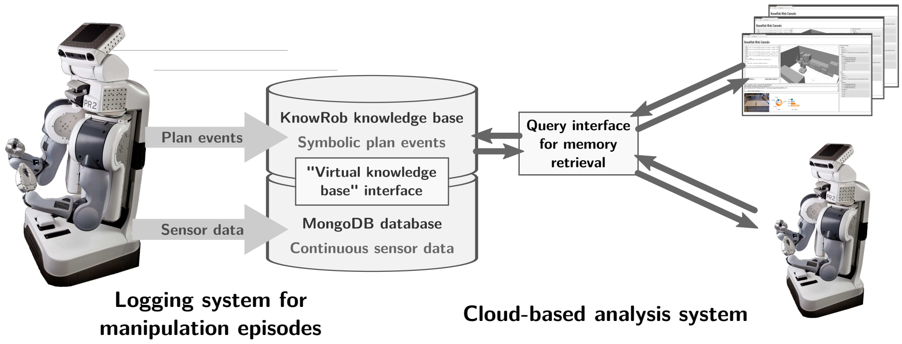

Fig. 2: Overview of the OPEN-EASE system. The logging components record detailed log data of manipulation episodes into a knowledge base. A library of query predicates is provided for reasoning about the data and can be called by human or robot clients via the Internet.

3 Projected Applications

OPEN-EASE is designed to facilitate a wide range of applications. At the moment, we are already using a prototype of the OPEN-EASE system for some of the following applications, and are working on extending it towards the others.

3.1 Analysis, Benchmarking and Debugging of Robot Behavior

The detailed logs recorded by OPEN-EASE allow to analyze robot tasks in great detail, both at the level of single sub-actions and at the aggregate level of complete tasks. For example, the user can can ask questions such as: How long did actions take? Which actions were successful? Under which circumstances did which sub-actions fail? Could these failures be handled automatically? How did the robot move while performing its tasks? The logs further allow to reconstruct the world state as the robot believed it to be as well as the internal state of the robot's control program at a given moment in time. This can be very helpful for debugging complex robot tasks, for identifying subtle anomalies in robot behavior, and for analyzing problems *after* a task has been performed.

3.2 Autonomous Learning from Episodic Memories

The log data includes both the motion commands sent by the robot and the observed outcomes of these actions, i.e. the performed movements and their success or failure. This allows to learn models correlating the plan outcomes and the initial plan parameters, turning the supervised learning problem of identifying suitable plan parameters into a unsupervised one. Examples of learning tasks are the selection of good locations for performing manipulation actions, expectations about the duration of tasks, or the 'compilation' of optimized motion primitives for common tasks. Figure 3 shows the episodic memory of a PR2 robot grasping a cup. Trajectories can be semantically distinguished by the task during which they were performed.

Based on a declarative description of a learning task, that specifies which features are to be used for predicting which output variables, a robot can autonomously query the knowledge base to retrieve the required training data from its log files. As the data in the knowledge

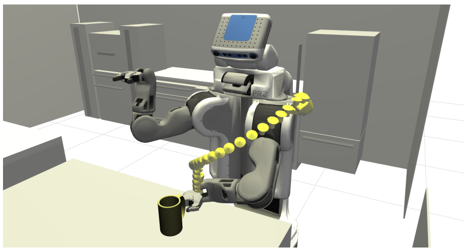

Fig. 3: Memorized grasping trajectory for picking up a cup

base is semantically annotated, semantic meaning can also be assigned to the learned models to help a robot decide what to use these models for.

3.3 Platform for Publishing Benchmark Datasets

Due to its cloud-based nature, it is very easy to share datasets in OPEN-EASE with other researchers. The standardized common ontology and query predicates map the individual data sets, which may be structured in different ways, into a common format, thereby facilitating the re-use of analyses and algorithms on other data sets. A common data format is also the prerequisite for comparing the performance of algorithms on different benchmark data sets. Such benchmarks exist for other domains such as SLAM [5] or human activity recognition [6–8], though these benchmarking problems are much more narrow and better defined than robot manipulation tasks.

Having comprehensive and easy-to-use benchmark datasets available may also attract researchers in machine learning and AI that do not have the robot equipment to perform experiments themselves. While robotics is often used as example application of new algorithms in these fields, the danger of developing algorithms without access to real data is that assumptions are made that are not justified in real applications. By providing datasets recorded during experiments of real robots, we hope to alleviate this problem.

3.4 Making Robot Experiments Reproducible

Reproducing robot experiments is often difficult, despite all progress made towards common open-source components and middleware, as it usually requires access to the same robot hardware and a similar environment setup. In addition, expertise in robot technology and software components is often needed to run the experiments. With OPEN-EASE, researchers can make a complete dataset of their experiments public and let reviewers and peer researchers explore the data to better understand the experiment setup, the exact actions that have been performed, their timing, and all other aspects that characterize the robot's performance and the difficulty of the tasks it had to perform.

3.5 Teaching

We are starting to use OPEN-EASE for teaching students about AI-enabled robotics. The web-based interface provides an easy method for giving them access to a robot's knowledge base and data from real experiments without requiring any software installation. Using these tools, students can for example learn which data is available in which format, develop new inference mechanisms, and test them in the recorded situations.

4 Discussion

In this paper, we gave an overview of the OPEN-EASE system and its applications in the context of autonomous robot learning. OPEN-EASE provides an open platform for recording, visualizing, analyzing and sharing robot data. Its reasoning mechanisms allow robots to extract the data they need for a specific learning problem from automatically recorded, semi-structured log data.

To explore the existing data sets and try out the analysis system, users can simply create a free account at *http://www.open-ease.org*. Users that would like to record their own data sets can use the provided open-source logging infrastructure and either start their own, local analysis server or upload the data to the public OPEN-EASE server. We have tried to keep the toolkit as modular as possible to allow users to selectively use those components that fit their setup and use case. Our robots perform their tasks under the supervision of the CRAM executive [9] that automatically records comprehensive log data as described in [3]. This setup is well integrated and currently provides the most comprehensive reasoning support. However, if users would like to use their own robot executive, they can follow the documentation to adapt the logging components and to implement the query predicates based on their data structures.

Acknowledgements This work was supported in part by the DFG Project BayCogRob within the DFG Priority Programme 1527 for Autonomous Learning and the EU FP7 Projects *RoboHow* (Grant Agreement Number 288533) and *SAPHARI* (Grant Agreement Number 287513).

References

1. Autonomous learning: Summer school 2014. Accessed: 2015-02-15. [Online]. Available: http://www.mis.mpg.de/calendar/conferences/2014/al.html
2. M. Beetz, M. Tenorth, and J. Winkler, "Open-EASE – a knowledge processing service for robots and robotics/ai researchers," in *IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, Washington, USA, 2015, accepted for publication.
3. J. Winkler, M. Tenorth, A. K. Bozcuoglu, and M. Beetz, "CRAMm – memories for robots performing everyday manipulation activities," *Advances in Cognitive Systems*, vol. 3, pp. 47–66, 2014.
4. M. Tenorth and M. Beetz, "KnowRob – A Knowledge Processing Infrastructure for Cognition-enabled Robots," *International Journal of Robotics Research (IJRR)*, vol. 32, no. 5, pp. 566 – 590, April 2013.
5. S. Ceriani, G. Fontana, A. Giusti, D. Marzorati, M. Matteucci, D. Migliore, D. Rizzi, D. G. Sorrenti, and P. Taddei, "Rawseeds ground truth collection systems for indoor self-localization and mapping," *Autonomous Robots*, vol. 27, no. 4, pp. 353–371, 2009.
6. M. Tenorth, J. Bandouch, and M. Beetz, "The TUM Kitchen Data Set of Everyday Manipulation Activities for Motion Tracking and Action Recognition," in *IEEE International Workshop on Tracking Humans for the Evaluation of their Motion in Image Sequences (THEMIS), in conjunction with ICCV2009*, 2009.
7. M. Rohrbach, S. Amin, M. Andriluka, and B. Schiele, "A database for fine grained activity detection of cooking activities," in *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Providence, United States, June 2012.
8. F. De la Torre, J. Hodgins, J. Montano, S. Valcarcel, and J. Macey, "Guide to the Carnegie Mellon University Multimodal Activity (CMU-MMAC) Database," CMU-RI-TR-08-22, Robotics Institute, Carnegie Mellon University, Tech. Rep., 2009.
9. M. Beetz, L. Mösenlechner, and M. Tenorth, "CRAM – A Cognitive Robot Abstract Machine for Everyday Manipulation in Human Environments," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Taipei, Taiwan, October 18-22 2010, pp. 1012–1017.

About the Authors

Moritz Tenorth received his doctoral degree in Computer Science in 2011 from TU München, after studying electrical engineering at RWTH Aachen and ENSTA-ParisTech and obtaining his Diplom degree (equiv. M.Eng.) in 2007 from RWTH Aachen, Germany. His research interests include grounded knowledge representations which integrate information from web sources, observed sensor data and data mining techniques, and their applications to knowledge-based action interpretation and robot control.

Jan Winkler studied mechanical engineering and control theory at University of Stuttgart and Toyohashi University of Technology and obtained his Diplom degree (equiv. M.Eng.) in Engineering Cybernetics in 2012 from University of Stuttgart, Germany. His main research field is high-level plan design and situationally aware execution, as well as plan abstraction primarily focused on autonomous mobile manipulation robots, and generation of episodic robot memories for autonomous experience collection.

Daniel Beßler studied computer science at University of Bremen and obtained his Diplom degree (equiv. M.Eng.) in 2014 from University of Bremen, Germany. His research interest consists of knowledge processing and representation for robots, knowledge-enabled perception and knowledge-enabled robot control programs.

Michael Beetz is a professor for computer science at the Faculty for Mathematics & Informatics of the University Bremen and head of the Institute for Artificial Intelligence (IAI). He also was the vice coordinator of the German cluster of excellence CoTeSys. He received his MSc, MPhil and PhD degrees from Yale University in 1993, 1994 and 1996 and his Venia Legendi from the University of Bonn in 2000. Michael Beetz was a member of the steering committee of the European network of excellence in AI planning (PLANET) and coordinated the research area robot planning. He was associate editor of the Artificial Intelligence journal.