

Learning Probabilistic Models for Mobile Manipulation Robots *

Jürgen Sturm and Wolfram Burgard

University of Freiburg

Germany

{sturm,burgard}@informatik.uni-freiburg.de


Abstract

Mobile manipulation robots are envisioned to provide many useful services both in domestic environments as well as in the industrial context. In this paper, we present novel approaches to allow mobile manipulation systems to autonomously adapt to new or changing situations. The approaches developed in this paper cover the following four topics: (1) learning the robot's kinematic structure and properties using actuation and visual feedback, (2) learning about articulated objects in the environment in which the robot is operating, (3) using tactile feedback to augment visual perception, and (4) learning novel manipulation tasks from human demonstrations.

1 Introduction

The development of flexible mobile manipulation robots is widely envisioned as a large breakthrough in technology and is expected to have a significant impact on our economy and society in the future. Mobile manipulation robots that are equipped with one or more gripper arms could fulfill various useful services in private homes including cleaning, tidying up as well as fetch and carry tasks. Robust solutions to all of these tasks obviously would mean a significant time benefit to their owners. For example, by supporting elderly and mobility-impaired people in the activities of daily living, appropriate service robots can reduce the dependency on external caregivers and support such people to live a self-determined and autonomous life. In addition, small and medium-sized enterprises would profit enormously from robotic co-workers that they can easily reconfigure to new production tasks. This technology would significantly lower the production costs of smaller companies and thus provide them with a significant competitive advantage. The goal of this work is to provide novel approaches that enable mobile manipulation robots to be flexibly used in everyday life. The

*The dissertation on which this extended abstract is based was the recipient of the ECCAI Artificial Intelligence Dissertation Award 2011 and was shortlisted for the EURON Georges Girault Award 2012 [Sturm, 2011].



(a) body schema learning



(b) articulated objects



(c) tactile sensing



(d) imitation learning

Figure 1: Illustration of the four research questions addressed in this thesis. (a) Body schema learning using visual self-observation. (b) Learning to operate articulated objects, here: a fridge. (c) Using tactile sensing to estimate the state of a container. (d) Imitation learning to acquire novel manipulation skills.

challenge in these applications is that robots operating in unstructured environments have to cope with less prior knowledge about themselves and their surroundings. Therefore, they need to be able to autonomously learn suitable probabilistic models from their own sensor data to robustly fulfill their tasks.

For decades, stationary manipulation robots have successfully been used in industrial mass production. In these applications strong assumptions about the physical setup and

a controlled environment allow the creation of efficient but highly engineered approaches. These solutions are custom-tailored to specific applications which makes them difficult to adapt: typically, changes in the application require the manual adaptation of the robot's control code, a new layout of its work cell, and possibly the reconfiguration of its hardware. For this reason, industrial manipulators require the supervision of experts on a regular basis, and are therefore only cost-effective for the mass production. In contrast, the environment of mobile manipulators used for domestic service tasks or in small series production is largely unstructured, i.e., it can neither be exactly specified nor easily controlled. To deal with these uncertainties, mobile manipulation robots need to be considerably more flexible, robust, and adaptive than their stationary predecessors.

This paper provides an overview on the PhD thesis of Jürgen Sturm [Sturm, 2011], which is available online at

http://vision.in.tum.de/members/sturmju/phd_thesis

and published as a hardcover by Springer [Sturm, 2013]. Next to the thesis itself, this website contains additional material such as videos, research papers, and freely available datasets.

2 Challenges

To illustrate the relevance of the topics presented in this thesis, we motivate our work using a typical example task of a domestic service robot. We assume that the robot is given the task to deliver a drink, which requires the robot to open the fridge, pick up the right bottle, and pour its content into a glass. To be able to accurately use its manipulator, the robot first needs to verify its body schema using visual self-observation (Figure 1a) This enables the robot to compensate for mechanical inaccuracies and detect potential hardware failures. Once the robot established its body schema, it navigates to the fridge to retrieve a drink (Figure 1b). To open the fridge, the robot identifies the fridge door and generates a suitable trajectory for opening it. This, in turn, requires a kinematic model of the fridge. Being able to learn kinematic models is fundamental for versatile service robots, as there are too many different cabinet doors and drawers in domestic environments to exclusively rely on predefined models. After the robot has successfully opened the fridge, it picks up a bottle. By using its tactile sensors (Figure 1c), the robot verifies that it has grasped the correct object and that this object is in the expected state. The next step of the delivery task is to pour the drink into a glass (Figure 1d). This skill, however, might not be part of the robot's current programming. In this case, the user can teach the robot this novel manipulation skill by demonstrating it to the robot. From this demonstration, the robot learns and generalizes a description of the task that it can subsequently use to reliably reproduce it. Such an intuitive programming interface is an essential prerequisite for the usability of service robots in everyday life.

This motivating example leads us to the four research questions that we tackle in this thesis:

- How can a manipulation robot learn to accurately position its arm?
- How can a manipulation robot robustly operate doors and drawers?

- How can a manipulation robot infer the state of the objects it manipulates?
- How can a user intuitively teach novel manipulation tasks to a robot?

A robot that operates in unstructured environments with no or minimal human supervision needs to be able to perceive the world through its own sensors, and subsequently, build from this data an internal, up-to-date representation of the world. As sensor data is always noisy and potentially incomplete, a robot requires robust techniques to interpret and integrate it intelligently into its own models of the world. A robot can then use these models to estimate the state of objects in the world, simulate the consequences of its actions, generate plans, and, finally, verify the success of its actions.

Our work is based on state-of-the-art Bayesian learning techniques such as graphical models, Gaussian processes and robust estimation methods. The probabilistic formulation of our approaches allows a robot to deal with uncertainties in the sensor observations and action execution and to consider them adequately during action planning. In an exhaustive set of experiments on real robots and in simulation we demonstrate that our approaches significantly reduce the dependency of manipulation robots on hand-crafted models and structured environments. In sum, this thesis provides novel probabilistic learning techniques that enable a manipulation robot

- to learn the body schema of its arm from scratch using self-observation, and to monitor and adapt this model over extended periods of time,
- to learn kinematic models of articulated objects from observation or interaction to reliably operate doors and drawers,
- to learn tactile object models to estimate the identity and state of the objects being manipulated, and
- to learn novel manipulation tasks from human demonstrations, and to reproduce them robustly in similar situations.


3 Thesis Outline

The thesis is organized as follows. In Chapter 2, we provide the technical background in machine learning and probabilistic modeling that we require for the remainder of the thesis. In Chapter 3, we present a novel approach that enables a robot to learn the body schema of its manipulator from scratch using visual self-observation (see Figure 2a). In contrast to previous approaches, we estimate both the kinematic structure and the kinematic properties of the robot arm [Sturm *et al.*, 2008b]. We model the observations of each link of the arm as a Gaussian process and learn a Bayesian network that describes the kinematics of the whole system. An example of such a learned body schema is visualized in Figure 4. The explicit representation of the kinematic structure allows the robot to detect and localize deviations between the model and the real arm to specific components of the network [Sturm *et al.*, 2008a]. Our approach provides a flexible, probabilistic representation of robot kinematics and, furthermore, enables



(a) online adaptation of the body schema (b) learned kinematic model of a fridge door

Figure 2: Examples of our solutions to kinematic model learning. (a) With our approach, robots can autonomously adapt their body schema in case of hardware failures and tool changes. (b) Our framework also applies to passively-actuated articulated objects and enables robots to reliably operate typical household objects such as cabinets, fridges and dishwashers.




(a) tactile sensing (b) imitation learning


Figure 3: More examples of model learning using our approaches. (a) This service robot learns to haptically discriminate empty from full bottles and uses this knowledge to tidy up a table. (b) This robot learns its instructions (here: cleaning the white board) from human demonstrations.

a manipulation robot to position its end effector accurately even in the presence of hardware failures. Our results show that robots using our approach require less maintenance and can be used over longer periods of time without human intervention [Sturm *et al.*, 2009a].

A central task of service robots is to interact with articulated objects, for example, to open doors in order to navigate between rooms or to pick up objects from cabinets or drawers. In Chapter 4, we show how our approach on body schema learning can be generalized to such articulated objects. We extend our approach by additional parametric models and use Bayesian model comparison to choose between the alternatives [Sturm *et al.*, 2009b; 2010a]. This increases the robustness and efficiency of our approach while we keep the high flexibility of the Gaussian process models. In contrast to previous work, our approach applies to a significantly larger class of articulated objects and provides more accurate kinematic models. Furthermore, we can estimate the degrees of freedom of an articulated object and dis-



(a) Gaussian process regression



(b) Bayesian network

Figure 4: The central concept in our work is to learn compact probabilistic models from sensor data. This enables service robots to robustly deal with substantial changes in their environment. For example for body schema learning, we use (a) Gaussian processes and (b) Bayesian networks to represent the kinematic model of a robot.

cover kinematic loops [Sturm *et al.*, 2011]. Complimentary to this, we demonstrate in Chapter 5 how a manipulation robot can recognize cabinet doors and drawers on dense depth images without requiring visual markers [Sturm *et al.*, 2010b; Rühr *et al.*, 2012].

In addition to articulated objects, service robots also need to manipulate many other objects such as bottles, silverware, or dishes. If a robot has tactile sensors in its gripper, it can use them to obtain additional information about the objects it is interacting with. In Chapter 6 and Chapter 7, we present two novel approaches that manipulation robots can use to learn tactile object models. The first approach is based on the bag-of-features method and enables a robot to verify whether it has grasped the correct object [Schneider *et al.*, 2009]. In our second approach, we analyze the dynamics of the tactile signal to recognize the internal state of liquid containers [Chitta *et al.*, 2010; 2011]. This ability is, for example, important for a domestic service robot that tidies up a table and needs to decide whether a juice bottle is full or empty and should be stored in the fridge or disposed in the trash can (as illustrated in Figure 3a). Our results indicate that tactile sensing is a useful source of information for a robot to augment its perceptions during object manipulation.

Another prerequisite for successful service robotics applications is that normal users can quickly and intuitively instruct the robot to perform novel tasks. Inspired by work on




Figure 5: Three state-of-the-art mobile manipulation robots that we used for developing and testing our approaches.

imitation learning, we develop in Chapter 8 a novel approach to learn manipulation tasks by observing a human instructor demonstrating a certain manipulation task [Eppner *et al.*, 2009]. From these demonstrations, the robot extracts invariances in the execution of the task and infers from them a generalized task model (see Figure 3b). In contrast to existing approaches, the factorized representation of the manipulation task as a dynamic Bayesian network allows us to dynamically add new constraints, for example, to avoid obstacles during reproduction, or to prefer a particular body posture. Our approach allows normal users to provide novel task descriptions to a manipulation robot in an intuitive way, which we consider an important prerequisite for the daily use of manipulation robots. Finally, we conclude the thesis with a summary of our results in Chapter 9 and give an outlook to future work.

To develop and test our approaches, we used three different state-of-the-art mobile manipulators as depicted in Figure 5. By evaluating our approaches successfully on different experimental platforms, we ensure that our approaches also generalize to other mobile manipulation robots.

All of our approaches are based on state-of-the-art Bayesian learning techniques such as Gaussian processes, sample consensus methods, and graphical models. The probabilistic formulation of our approaches allows a robot to deal with uncertainties in the sensor observations and action execution and to consider them adequately during action planning. Furthermore, we show that our approaches substantially increase the flexibility, adaptability and robustness of manipulation robots.

4 Software

We released parts of our software as open-source to offer other researchers the opportunity to verify our results, evaluate our approaches on different data, and use our software in their research. In particular, we provide free software implementations of our body schema learning approach and the complete framework for kinematic model learning of articulated objects.

- The ZORA framework¹ implements our approach on body schema learning as described in Chapter 3. It is freely available under the GPL license. Furthermore, a

¹<http://www.informatik.uni-freiburg.de/~sturm/zora.html>

detailed tutorial explains how to reproduce our results on various simulated manipulators.

- The ARTICULATION stack² provides several software libraries for learning kinematic models of articulated objects as described in Chapter 4 and Chapter 5. We released the software stack under the BSD license. Further, we provide several tutorials that explain in detail how kinematic models of articulated objects can be learned from observed trajectories and how the framework can be used with Python and C++.

Our open-source software for operating articulated objects with mobile manipulators is currently being used in the demonstrators of several renowned research labs across Europe (U Freiburg, TU Munich, TU Eindhoven, ETH Zurich) and the United States (Bosch Research, Georgia Tech) and thus both in academia and in industry. Several ongoing research projects are currently using or extending our approach on learning kinematic models of articulated objects. The RoboEarth project³ aims at the creation of a worldwide object database and annotates articulated objects with the models learned using our approach. The goal of the SFB/TR 8⁴ is to investigate the cognitive foundations for human-centered spatial assistance systems, and plans in project A8 to extend our approach to learn 3D models of the rigid parts of articulated objects. The First-MM project⁵ aims to enable robots to acquire new manipulation skills which also involve grasping and operating articulated objects using our approach. The goal of the TidyUpRobot project⁶ is to use the PR2 robot in various tidying-up tasks.

5 Conclusion

We think that the field of mobile manipulation bears a large market potential in the near future. In this work, we presented several innovative approaches to relevant problems that arise when mobile manipulators are applied in unstructured environments and changing situations. We hope that our work increases the dependability, flexibility, and ease of use of manipulation robots and thereby contributes to the development of truly useful robotic assistants for industry and society.

References

[Chitta *et al.*, 2010] S. Chitta, M. Piccoli, and J. Sturm. Tactile object class and internal state recognition for mobile manipulation. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Anchorage, AK, USA, 2010.

²<http://www.ros.org/wiki/articulation>

³RoboEarth is a research project funded by the European Union Seventh Framework Programme FP7/248942 (2009–2013).

⁴The Transregional Collaborative Research Center Spatial Cognition: Reasoning, Action, Interaction has been established by the German Research Foundation (DFG) (2003–2014).

⁵First-MM is another research project founded under the European Union Seventh Framework Programme FP7/248258 (2010–2014).

⁶The TidyUpRobot project is part of the PR2 beta program sponsored by Willow Garage (2010–2012).

- [Chitta *et al.*, 2011] S. Chitta, J. Sturm, M. Piccoli, and W. Burgard. Tactile sensing for mobile manipulation. *IEEE Transactions on Robotics (T-RO)*, 27:558–568, June 2011.
- [Eppner *et al.*, 2009] C. Eppner, J. Sturm, M. Bennewitz, C. Stachniss, and W. Burgard. Imitation learning with generalized task descriptions. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Kobe, Japan, May 2009.
- [Rühr *et al.*, 2012] T. Rühr, J. Sturm, D. Pangercic, M. Beetz, and D. Cremers. A generalized framework for opening doors and drawers in kitchen environments. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, St. Paul, MA, USA, May 2012.
- [Schneider *et al.*, 2009] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard. Object identification with tactile sensors using bag-of-features. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, St. Louis, MO, USA, Oct. 2009.
- [Sturm *et al.*, 2008a] J. Sturm, C. Plagemann, and W. Burgard. Adaptive body scheme models for robust robotic manipulation. In *Proc. of Robotics: Science and Systems (RSS)*, Zurich, Switzerland, June 2008.
- [Sturm *et al.*, 2008b] J. Sturm, C. Plagemann, and W. Burgard. Unsupervised body scheme learning through self-perception. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Pasadena, CA, USA, May 2008.
- [Sturm *et al.*, 2009a] J. Sturm, C. Plagemann, and W. Burgard. Body schema learning for robotic manipulators from visual self-perception. *Journal of Physiology-Paris*, 103(3-5):220–231, Sept. 2009. Neurorobotics.
- [Sturm *et al.*, 2009b] J. Sturm, V. Pradeep, C. Stachniss, C. Plagemann, K. Konolige, and W. Burgard. Learning kinematic models for articulated objects. In *Proc. of the Intl. Conf. on Artificial Intelligence (IJCAI)*, Pasadena, CA, USA, July 2009.
- [Sturm *et al.*, 2010a] J. Sturm, A. Jain, C. Stachniss, C. Kemp, and W. Burgard. Operating articulated objects based on experience. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct. 2010.
- [Sturm *et al.*, 2010b] J. Sturm, K. Konolige, C. Stachniss, and W. Burgard. Vision-based detection for learning articulation models of cabinet doors and drawers in household environments. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Anchorage, AK, USA, 2010.
- [Sturm *et al.*, 2011] J. Sturm, C. Stachniss, and W. Burgard. A probabilistic framework for learning kinematic models of articulated objects. *Journal on Artificial Intelligence Research (JAIR)*, 41:477–626, Aug. 2011.
- [Sturm, 2011] J. Sturm. *Approaches to Probabilistic Model Learning for Mobile Manipulation Robots*. PhD thesis, University of Freiburg, Germany, May 2011.
- [Sturm, 2013] J. Sturm. *Approaches to Probabilistic Model Learning for Mobile Manipulation Robots*. Springer

Tracts in Advanced Robotics (STAR). Springer, June 2013.