

Operating Articulated Objects Based on Experience

Jürgen Sturm¹

Advait Jain²

Cyrill Stachniss¹

Charles C. Kemp²

Wolfram Burgard¹

Abstract—Many tasks that would be of benefit to users in domestic environments require that robots manipulate articulated objects such as doors and drawers. In this paper, we present a novel approach that simultaneously estimates the kinematic model of an articulated object based on the trajectory described by the robot’s end effector, and uses this model to predict the future trajectory of the end effector. One advantage of our approach is that the robot can directly use these predictions to generate an equilibrium point control path for operating the mechanism. Additionally, our approach can improve these predictions based on previously learned articulation models. We have implemented and tested our approach on a real mobile manipulator. Through 40 trials, we show that the robot can reliably open various household objects, including cabinet doors, sliding doors, office drawers, and a dishwasher. Furthermore, we demonstrate that using the information from previous runs as a prior significantly improves the prediction accuracy.

I. INTRODUCTION

Applications in domestic environments are envisioned as a promising area for mobile service robots. For many beneficial tasks, robots operating in such environments would have to deal with articulated objects such as doors or drawers. Typically, a large variety of doors, drawers, and similar mechanisms are present in human environments. This makes it hard to pre-program a robot to operate the objects that it would encounter while performing its job.

In this paper, we consider the problem of modeling and operating articulated objects with a mobile manipulator. Articulated objects are objects with movable parts such as a cabinet with doors or drawers. Our approach enables a real robot to learn articulation models, which describe the kinematics of the corresponding parts of the object. Furthermore, we utilize these models to control the motion of the manipulator. In the past, several researchers have addressed the problem of estimating and handling doors and drawers [5, 9, 12, 20]. Most of these approaches, however, are either entirely model-free or assume substantial knowledge about the model and its parameters. Whereas model-free approaches release designers from providing any a-priori model information, the knowledge about objects and their articulation properties supports state estimation, motion prediction, and planning.

¹ J. Sturm, C. Stachniss, and W. Burgard are with the Autonomous Intelligent Systems Lab, Institute of Computer Science, University of Freiburg, Germany.

² A. Jain, C. Kemp are with the Healthcare Robotics Lab, Georgia Tech, USA.

This work has been partly supported by the EC within the 7th Framework Programme under grant agreement no FP7-248258-First-MM, by the German Research Foundation (DFG) under contract number SFB/TR-8, by the National Science Foundation (NSF) under grants CBET-0932592 and IIS-0705130, and by Willow Garage.



Fig. 1. Our service robot successfully opens various articulated objects (here: a dishwasher) by generating equilibrium control points from the kinematic model that it learns online from its end effector trajectory.

The contribution of this paper is two-fold: First, we present an approach to learn the kinematic model of an articulated object while the robot is operating it with its end effector. For the manipulation of the object, the predictions of the model are directly used to generate an equilibrium point control path for the manipulator. Our model not only supports rotational and prismatic joints, but also more complex mechanisms which the robot estimates through non-parametric learning. Figure 1 shows an illustrative example in which our robot opens the door of a dishwasher using this approach. As a second contribution, we enable the robot to improve its model estimation capabilities over time. The learned model is used as prior information in future manipulation tasks. As we will demonstrate in the experimental section, this information can help a robot make better predictions about the movements of an object given that the robot has seen a similar object before. We regard this as a reasonable assumption, since similar mechanisms can be found both within a single home, which may contain multiple instances of a kinematic design, and between homes, which can share specific kinematic designs as well as more general characteristics, such as rotational axes that are typically parallel or normal to gravity.

II. RELATED WORK

Several researchers have addressed the problem of operating articulated objects with a manipulator. A large number of these techniques have focused on handling doors and drawers [1, 3, 5, 9, 10, 12, 20]. The majority of approaches, however, have assumed an implicit kinematic model of the articulated object. Meeussen *et al.* [12] describe an integrated

navigation system for mobile robots including vision- and laser-based detection of doors and door handles that enables the robot to successfully open doors using a compliant arm. Diankov *et al.* [3] formulate door and drawer operation as a kinematically constrained planning problem and propose to use caging grasps to enlarge the configuration space. Wieland *et al.* [20] combine force and visual feedback to reduce the interaction forces when opening kitchen cabinets and drawers. In contrast to our work, these approaches do not learn a kinematic model while manipulating an object and thus cannot re-use information learned during previous interactions. Moreover, as we show in our experiments, our method can successfully operate a variety of mechanisms, even when encountering them for the first time.

Estimating kinematic models and the kinematic structure from observation has been studied previously [7, 19, 23] without subsequently using the models for manipulation. Katz and Brock [8] have enabled a robot to first interact with a planar kinematic object on a table top in order to visually learn a kinematic model, and then manipulate the object using this model to achieve a goal state. In contrast, our approach learns from the 3D kinematic trajectory of the robot’s end effector. It also simultaneously learns a model and manipulates the object to achieve its goal.

The articulated objects found within typical household settings tend to exhibit common kinematic structure. Consequently, learning models of articulated objects offers the possibility to increase the performance of a mobile manipulation system over time. Whenever a robot that learns models for articulated objects operates a previously unseen object, it needs to decide whether to merge the novel information into an existing model or to learn a new model from scratch. This can be formulated as a model selection problem [11]. The Bayesian Information Criterion (BIC) [16] has often been used in clustering applications. For example, Pelleg *et al.* use the BIC in the X -means algorithm [15] to efficiently estimate the right number of clusters for a set of data points in Euclidean space. In this paper, we use the BIC to estimate the number of models for the motion trajectories of articulated objects.

To actually operate an articulated object, we use equilibrium point control (EPC) [6] which is a form of impedance control inspired by the equilibrium point hypothesis. Using EPC, the motion of the robot’s arm is commanded by adjusting the position of a Cartesian-space equilibrium point (CEP) that denotes where the robot’s end effector would settle in the absence of externally applied forces other than gravity. Previous robotics research has looked at similar robotic control strategies in simulation [4], in free-space motions [21], legged locomotion [13], rhythmic manipulation from a fixed base [22], and in the design and control of compliant actuators [14, 2].

In previous work [6], we demonstrated how a robot can use equilibrium point control and low mechanical impedance to open cabinet doors, whose axes of rotation are parallel to gravity, and open cabinet drawers. We furthermore presented a technique [19] for learning kinematic models



Fig. 2. Left: The mobile manipulator used in this work. Right: Examples of a human using his hand as a hook and corresponding orientations for the robot’s hook end effector.

of objects where we assumed that a person manipulates the object and that the robot only observes these activities. In contrast to this, the approach presented here enables a robot to learn kinematic models, to predict the motion of the articulated object, and feed this information into the controller of the manipulator. To the best of our knowledge, none of the previous approaches learns a kinematic model in 3D of the manipulated object and simultaneously uses the corresponding predictions in the controller. In addition, we show how a robot can exploit its own experience from previous manipulations to improve the prediction accuracy of the articulation models.

III. OUR APPROACH TO OPERATING ARTICULATED OBJECTS USING PREVIOUS EXPERIENCE

The robot we use for this research is a statically stable mobile manipulator equipped with a hook as the end effector (see Fig. 2). Whereas the left image of Fig. 3 shows an overall block diagram for the complete system, the right image of Fig. 3 illustrates the idea with an example. The robot observes the pose of its end effector in Cartesian space, denoted by $\mathbf{x} \in \mathbb{R}^3$. While operating the mechanism, the robot records the trajectory $\mathbf{x}_{1:t}$ over time as a sequence of poses. In Section III-A, we describe how we fit a model \mathcal{M} to such a trajectory online. The robot then uses this model to predict the continuation of the trajectory including the Jacobian \hat{J}_t at the estimated current configuration of the mechanism. In Section III-B, we explain how we include prior information in the model fitting step and how a joint model is learned from previously observed trajectories. Finally, in Section III-C, we describe how our implementation of the equilibrium point control employs the prediction of the kinematics of the mechanism to generate a set of control points $\mathbf{x}_{1:t}^{CEP}$.

A. Learning Models of Articulated Objects

Since we do not know what kind of kinematic mechanism exists between the handle and the cabinet, we fit several candidate models representing different kinds of links. This set includes a rotational model for doors \mathcal{M}^{rot} , a prismatic model for drawers \mathcal{M}^{prism} , a rigid model for non-articulated objects \mathcal{M}^{rigid} , and a more flexible model based on Gaussian processes $\mathcal{M}^{LLE/GP}$ that can describe any other 1-DOF trajectory, for example 2-bar-links in garage doors.

For the moment, we consider a single trajectory obtained from a mechanism and denote this observed trajectory as

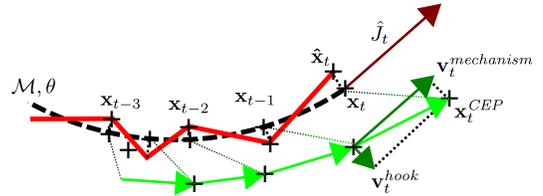
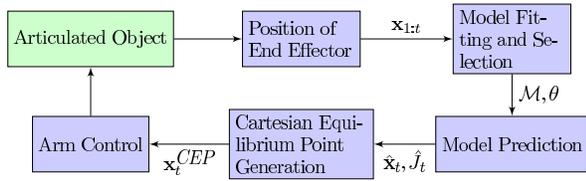


Fig. 3. Left: Overall control structure. The robot observes the trajectory $\mathbf{x}_{1:t}$ of its end effector. From that, it estimates the model \mathcal{M}, θ and predicts the current pose $\hat{\mathbf{x}}_t$, and the Jacobian \hat{J}_t . Using these predictions, it generates the next Cartesian equilibrium point \mathbf{x}_t^{CEP} . Right: Example trajectory illustrated for the right cabinet door (top view). The generated CEP trajectory smoothly pulls the door open, while keeping the hook robustly on the handle.

$\mathcal{D} = \mathbf{x}_{1:t}$. Later, we will explain how the model fitting can be improved by additionally using data from previous trials for the estimation.

1) *Model Fitting*: For estimating the parameters of any of the above-mentioned models, we need to find a parameter vector $\theta \in \mathbb{R}^k$ that maximizes the data likelihood given this model, i.e.,

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} p(\mathcal{D} | \mathcal{M}, \theta) \quad (1)$$

This estimation procedure (and the number of parameters k) is different for each of the models. For example, the prismatic model is parametrized by a 3D line ($k = 5$). We estimate one point on the line as the mean of the trajectory and the direction vector as the principal component using SVD. The rotational model is parametrized by the rotation center, the rotation axis, and the radius. We estimate the $k = 7$ parameters of this model using two RANSAC steps: first, we recover the rotational plane, then we fit a circle on the inliers on the plane. All mechanisms that we consider here (except for the rigid link) have a latent variable $q \in \mathbb{R}$ that describes the configuration of the mechanism. In the case of a door, for example, the variable q describes the opening angle.

2) *Model evaluation*: To evaluate how well a single pose observation \mathbf{x} can be explained by a model, we have to evaluate

$$p(\mathbf{x} | \mathcal{M}, \theta) = \int p(\mathbf{x} | q, \mathcal{M}, \theta) p(q | \mathcal{M}, \theta) dq. \quad (2)$$

Under the assumption that no configuration state q is more likely than another one, this simplifies to

$$p(\mathbf{x} | \mathcal{M}, \theta) = \int p(\mathbf{x} | q, \mathcal{M}, \theta) dq. \quad (3)$$

If we assume that $p(\mathbf{x} | q, \mathcal{M}, \theta)$ is a uni-modal distribution, an approximation of the integral is to evaluate it only at the expected configuration \hat{q} given the observed pose \mathbf{x} , i.e.,

$$\hat{q} = \mathbb{E}_{\mathcal{M}, \theta} [q | \mathbf{x}]. \quad (4)$$

For this configuration, we can compute the expected pose $\hat{\mathbf{x}}$,

$$\hat{\mathbf{x}} = \mathbb{E}_{\mathcal{M}, \theta} [\mathbf{x} | \hat{q}]. \quad (5)$$

If we now assume that the error between the observed pose \mathbf{x} and the expected pose $\hat{\mathbf{x}}$ given the model \mathcal{M} and its parameter θ is normally distributed, the observation likelihood from Eq. 2 can be approximated as

$$p(\mathbf{x} | \mathcal{M}, \theta) \propto \exp(-\|\hat{\mathbf{x}} - \mathbf{x}\|^2 / \sigma^2). \quad (6)$$

Note that the expectations of Eq. 4 and Eq. 5 for the models in our candidate set have a closed form solution [19]. In the case of the prismatic model, for example, this corresponds to projecting the noisy observation \mathbf{x} onto the line model; the difference $\|\hat{\mathbf{x}} - \mathbf{x}\|$ then corresponds to the observation error that we assume to be normally distributed with $\|\hat{\mathbf{x}} - \mathbf{x}\| \sim \mathcal{N}(0, \sigma)$. Given this, we can efficiently compute Eq. 6. Finally, the marginal data likelihood over the whole trajectory becomes

$$p(\mathcal{D} | \mathcal{M}, \theta) = \prod_{\mathbf{x} \in \mathcal{D}} p(\mathbf{x} | \mathcal{M}, \theta). \quad (7)$$

3) *Model Selection*: For Bayesian model selection, we need to compare the posterior probability of the models given the data

$$p(\mathcal{M}, \theta | \mathcal{D}) = \frac{p(\mathcal{D} | \mathcal{M}, \theta) p(\theta | \mathcal{M}) p(\mathcal{M})}{p(\mathcal{D})}. \quad (8)$$

While the direct evaluation of the posterior generally is difficult, it can be approximated efficiently based on the Bayesian Information Criterion (BIC). Let k be the number of parameters and n the number of observations ($n = |\mathcal{D}|$). Then, the BIC is defined as

$$\text{BIC}(\mathcal{M}) = -2 \log p(\mathcal{D} | \mathcal{M}, \hat{\theta}) + k \log n, \quad (9)$$

where $\hat{\theta}$ is the maximum-likelihood (ML) parameter vector. Model selection now reduces to selecting the model that has the lowest BIC, i.e.,

$$\hat{\mathcal{M}} = \underset{\mathcal{M}}{\operatorname{argmin}} \text{BIC}(\mathcal{M}). \quad (10)$$

Given a partial trajectory $\mathbf{x}_{1:t}$ of the end effector, the robot continuously estimates the kinematic model of the articulated object. By deriving the expected pose $\hat{\mathbf{x}}_t$ given in Eq. 5 w.r.t. the configuration of the object, the robot can compute the Jacobian from the model at its current configuration as

$$\hat{J}_t = \frac{\partial \hat{\mathbf{x}}(\hat{q})}{\partial \hat{q}}. \quad (11)$$

B. Using Prior Information to Improve Model Prediction

Using the approach described above, a robot always starts from scratch when observing movements of a new articulated object. From a learning perspective, this may be seen as unsatisfactory since most articulated objects encountered in man-made environments belong to few different classes. For example, in a specific office or kitchen, many cabinet doors will tend to be the same size. Furthermore, in some countries



Fig. 4. Images showing the robot after it has operated the five mechanisms used in the experiments. The objects are (from left to right): a cabinet door that opens to the right, a cabinet door that opens to the left, a dishwasher, a drawer, and a sliding cabinet door.

the size of such furniture is standardized. Thus, a robot operating in such environments can significantly boost its performance by learning priors over the space of possible articulated object models.

This section describes our approach to learning priors for articulated objects and a means for exploiting them during the manipulation of an unknown object. The key idea is to identify a small set of representative models for the articulated objects and utilize this as prior information to increase the prediction accuracy when handling new objects.

To keep the notation simple, consider the case that the robot has interacted with two articulated objects. Their observed motion is given by two data trajectories \mathcal{D}_1 and \mathcal{D}_2 . The robot has to decide if both trajectories should be described by two distinct models \mathcal{M}_1 and \mathcal{M}_2 or by a joint model \mathcal{M}_{1+2} , each with corresponding maximum-likelihood parameter vectors. In the first case, this results in having $n = n_1 + n_2$ data points, $k = k_1 + k_2$ parameters, and a marginal data likelihood of $L = p(\mathcal{D}_1 | \mathcal{M}_1, \hat{\theta}_1)p(\mathcal{D}_2 | \mathcal{M}_2, \hat{\theta}_2)$.

In the latter case both trajectories are explained by a single, joint model \mathcal{M}_{1+2} with a parameter vector $\hat{\theta}_{1+2}$, that is estimated from the joint data $\mathcal{D}_1 \cup \mathcal{D}_2$ as described in Section III-A. This model also is learned from $n = n_1 + n_2$ data points, using only a single set of parameters $k' = k_1 = k_2$. The marginal data likelihood for this hypothesis is $L' = p(\mathcal{D}_1, \mathcal{D}_2 | \mathcal{M}_{1+2}, \hat{\theta}_{1+2})$. We can evaluate whether describing the data using a joint model instead of two separate models is beneficial by evaluating

$$BIC(\mathcal{M}_1, \mathcal{M}_2) > BIC(\mathcal{M}_{1+2}) \quad (12)$$

$$-2 \log L + k \log n > -2 \log L' + k' \log n. \quad (13)$$

If more than two trajectories are considered, one has to evaluate all possible assignments of these trajectories to models and select the assignment with the lowest BIC. However, this quickly becomes intractable. Thus, we use an approximation and consider the trajectories sequentially, in the order the robot observes them. We check whether merging the new trajectory with one of the existing models is more beneficial in terms of the BIC compared to adding a new model for that trajectory to the set of previously encountered models.

After having identified a set of models as prior information, we can exploit this knowledge for making better predictions when observing a so far unseen articulated object. Consider the situation in which a partial trajectory of a new object has been observed. To exploit the prior information,

TABLE I
PERFORMANCE OF THE ROBOT ON THE 5 MECHANISMS.

Mechanism	Angle/Distance Pulled	Estimated Radius (True Radius)	Max Interaction forces	Success Rate
Right Door	$96.5^\circ \pm 6.4^\circ$	$0.38\text{m} \pm 0.01\text{m}$ (0.39m)	$20.4\text{N} \pm 3.0\text{N}$	7/8
Left Door	$90.5^\circ \pm 7.6^\circ$	$0.35\text{m} \pm 0.02\text{m}$ (0.34m)	$36.3\text{N} \pm 2.3\text{N}$	7/8
Dishwasher	$72.0^\circ \pm 4.0^\circ$	$0.66\text{m} \pm 0.03\text{m}$ (0.65m)	$22.5\text{N} \pm 10.4\text{N}$	7/8
Drawer	$0.44\text{m} \pm 0.00\text{m}$		$38.4\text{N} \pm 2.6\text{N}$	8/8
Sliding Cabinet	$0.54\text{m} \pm 0.09\text{m}$		$26.4\text{N} \pm 7.5\text{N}$	8/8

we proceed exactly as before. We compute and compare the BIC according to Eq. 12 treating the newly observed data points as a new model or merge them into one of the m previously identified models by

$$BIC(\mathcal{M}_{new}, \mathcal{M}_1, \dots, \mathcal{M}_m) > \underset{j=1, \dots, m}{\operatorname{argmin}} BIC(\mathcal{M}_1, \dots, \mathcal{M}_{j+new}, \dots, \mathcal{M}_m). \quad (14)$$

If the newly observed data is merged with an existing model, the parameter vector is estimated from a much larger dataset ($\mathcal{D}_j \cup \mathcal{D}_{new}$ instead of \mathcal{D}_{new}) which leads to a better estimation. Note that this step is carried out after each observation of the new sequence. Thus, if the currently manipulated object ceases to be explained by the known models, the method instantaneously creates a new model. After successful object manipulation, this model serves as additional prior information for the future.

C. Equilibrium Point Control using Learned Articulation Models

The equilibrium point controller uses these estimates to generate Cartesian equilibrium point (CEP) trajectories in a fixed world frame, attached to the initial location of the handle. At each time step t , the controller computes a new equilibrium point \mathbf{x}_t^{CEP} as

$$\mathbf{x}_t^{CEP} = \mathbf{x}_{t-1}^{CEP} + \mathbf{v}_t^{mechanism} + \mathbf{v}_t^{hook}, \quad (15)$$

where $\mathbf{v}_t^{mechanism}$ is a vector intended to operate the mechanism, and \mathbf{v}_t^{hook} is a vector intended to keep the hook from slipping off of the handle. \mathbf{v}_t^{hook} is computed using the force between the hook and the handle measured by the wrist force-torque sensor of the robot. The controller computes

$$\mathbf{v}_t^{mechanism} = d^{mechanism} \frac{\hat{\mathbf{J}}_t^w}{\|\hat{\mathbf{J}}_t^w\|} \quad (16)$$

as a vector of length $d^{mechanism} = 0.01\text{m}$ along the estimated Jacobian of the mechanism.

Finally, we transform \mathbf{x}^{CEP} from the world coordinate frame into the local coordinate frame of the robot. We refer the reader to our previous work [6] for details about our implementation of equilibrium point control, including how it can be used to coordinate the motion of a mobile base and a compliant arm.

IV. EXPERIMENTS

We evaluated the performance of our approach on five different mechanisms: a cabinet door that opens to the right, a cabinet door that opens to the left, a dishwasher, a drawer, and a sliding cabinet door. We performed 8 trials for each mechanism. The robot started approximately 1m from the location of the handle. The task for the robot was to navigate to the mechanism and operate it. We deemed a trial to be successful if the robot navigated to the mechanism and opened it through an angle greater than 60° for rotary mechanisms or 0.3m for linear mechanisms. The input to the robot was the approximate 3D location of the handle, an orientation for the hook, and an initial direction for the robot’s end effector to move. We did not provide any further information about the object.

A. Task Performance

Fig. 4 shows the robot after it has pulled open each of the five mechanisms in one of the trials and Table I summarizes the performance of the robot. The robot successfully opened the 3 rotary mechanisms in 21 out of 24 trials and the 2 linear mechanisms in all 16 trials. Overall the robot was successful in 37 out of 40 trials (92.5%).

All three failures were due to the robot failing to hook onto the handle prior to operating the mechanism, most likely due to odometry errors and errors in the provided location of the handle. The learning and estimation approach we present in this paper did not cause any errors.

B. Evaluation of Model Fitting and Selection

In our previous work, we used different estimators to find the model parameters. In [6], we used a gradient-descent minimizer with 3 degrees of freedom. This estimator was only suitable for planar motions and required two initial poses for left and right doors, while a drawer was assumed when the estimated radius exceeded 5m. In [18], we use a RANSAC-based estimator for rotational models that determines all degrees of freedom including the plane of the rotation, the center of the rotation, the radius, and the object orientation.

A quantitative evaluation of the accuracy of the estimators is shown in Fig. 5. We compared the average prediction error of the in-plane minimizer versus the in-plane and full RANSAC estimator. We found that the RANSAC-based estimators outperform the minimizer, probably due to the outlier rejection. Using our RANSAC-based technique, estimating the plane of rotation comes at no additional cost and even allows us to improve the accuracy of our models

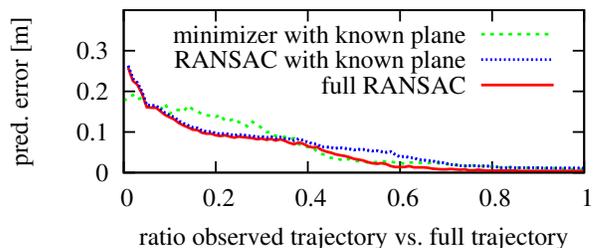


Fig. 5. Comparison of the performance of the model estimators in terms of the average prediction error, evaluated on the full trajectory. For this experiment, we excluded the dishwasher, because its door does not move in the horizontal plane.

– presumably, because the externally provided estimate of the horizontal plane was not perfect. The runtime of our estimators is linear in the number of data points. At the moment, the whole model fitting and selection step takes 5ms for a trajectory of 100 data points on a single 2.4 GHz Intel Core.

Further, we found that our approach selected the correct model candidate for all 37 trajectories after fully observing them, i.e., a rotational model for the left door, the right door, and the dishwasher, and a prismatic model for the drawer and the sliding cabinet door.

C. Improving Model Estimation based on Experience

By using the approach described in Section III-B on data from the robot operating the mechanisms in the 37 trials of Section IV-A, we search for the set of models for all trajectories with the lowest overall BIC. Figure 6 shows the result of this experiment. The colors indicate the cluster to which the trajectories have been assigned to. Our approach correctly recognized that the robot had operated 5 different mechanisms and assigned the 37 different trajectories correctly to the corresponding models.

In theory, our incremental clustering method may lead to different models depending on the order in which the data arrives. In order to evaluate this, we randomized the order of the previously recorded dataset and analyzed the resulting assignment of trajectories to models. We found that the trajectories from different mechanisms were never assigned to the same model. However, in 26 of 100 trials, the trajectories of one of the five objects were clustered into two models instead of one. After carefully inspecting the data, we found that this happens when the first model for an object is created from a trajectory where the robot was slightly de-localized.

We measured the average prediction error with and without learning prior models (see Fig. 7), using leave-one-out cross-validation and a randomized ordering of the trajectories. We found that the prior models reduce the prediction error considerably especially if the new trajectory is observed only partially. When 30% to 70% of the new trajectory have been observed, the prediction error is reduced by a factor of three (and more). As a result, the robot comes up with a substantially more accurate model early and can utilize this knowledge for better controlling its manipulator.

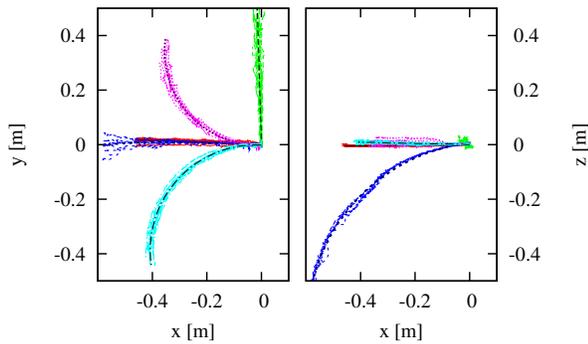


Fig. 6. These plots show the observed trajectories and the 5 recovered models when minimizing the overall BIC using our approach. Trajectories assigned to the same model are depicted in the same color.

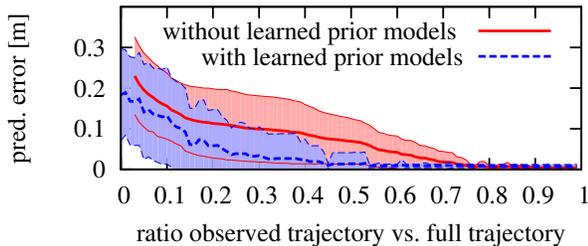


Fig. 7. This graph shows the average prediction error with and without prior model learning.

Throughout all experiments described so far, we used a fixed global noise term of $\sigma = 0.05\text{m}$. This accounts for inaccuracies in the observation of the end effector position, due to variations in the hooking position, and small errors in the kinematic forward model and robot base localization. We found in repeated experiments that in the range between $0.02\text{m} \leq \sigma \leq 0.20\text{m}$, the results are similar to our previous results obtained with $\sigma = 0.05\text{m}$. Only for significantly smaller values, each trajectory generates its own model. For much larger values, observations from different mechanisms are clustered into a joint model. Thus, our results are insensitive to (small) variations in σ .

V. CONCLUSION

In this paper, we presented a novel approach that combines kinematic model learning with equilibrium point control of articulated objects. We implemented and tested our proposed system on a real mobile manipulator. We showed that with our approach the robot can reliably learn and operate various mechanisms typically found in household environments. In addition to that, we showed that a robot can gain experience during manipulation and exploit information from previously manipulated mechanisms to significantly improve the prediction accuracy of its model.

Despite these encouraging results, there is space for improvements. In our current implementation, we re-estimate its parameters from all data points when adding a new trajectory to an existing model. This may become intractable for a robot that operates over an extended period of time. One possible solution here is to estimate both the model parameters and the model likelihood only from a sampled

subset of the data of a fixed size, to allow the system to be executed online. Such ideas have been used in related contexts, for example by Snelson and Ghahramani [17].

REFERENCES

- [1] D. Anguelov, D. Koller, E. Parker, and S. Thrun. Detecting and modeling doors with mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3777–3784, 2004.
- [2] D.J. Clapa, E.A. Croft, and A.J. Hodgson. Equilibrium point control of a 2-DOF manipulator. *Journal of Dynamic Systems, Measurement, and Control*, 128:134, 2006.
- [3] R. Diankov, S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning with caging grasps. In *Proc. of IEEE-RAS Intl. Conf. on Humanoid Robots (Humanoids)*, 2008.
- [4] X. Gu and D. Ballard. An equilibrium point based model unifying movement control in humanoids. In *Proc. of Robotics: Science and Systems (RSS)*, 2006.
- [5] A. Jain and C.C. Kemp. Behavior-based door opening with equilibrium point control. In *Proc. of the RSS Workshop on Mobile Manipulation in Human Environments*, 2009.
- [6] A. Jain and C.C. Kemp. Pulling open doors and drawers: Coordinating an omni-directional base and a compliant arm with equilibrium point control. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [7] D. Katz and O. Brock. Extracting planar kinematic models using interactive perception. In *RSS Workshop on Robot Manipulation*, 2008.
- [8] D. Katz and O. Brock. Manipulating articulated objects with interactive perception. In *Proc. of Robotics: Science and Systems (RSS)*, pages 272–277, Pasadena, USA, May 19–23 2008.
- [9] E. Klingbeil, A. Saxena, and A.Y. Ng. Learning to open new doors. In *Proc. of the RSS Workshop on Robot Manipulation*, 2009.
- [10] D. Kragic, L. Petersson, and H.I. Christensen. Visually guided manipulation tasks. *Robotics and Autonomous Systems*, 40(2-3):193 – 203, 2002.
- [11] D.J.C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [12] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, Patrick M., E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote, J. Hsu, R. Rusu, B. Marthi, G. Bradski, K. Konolige, B. Gerkey, and E. Berger. Autonomous door opening and plugging in with a personal robot. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [13] S.A. Migliore. *The Role of Passive Joint Stiffness and Active Knee Control in Robotic Leg Swinging: Applications to Dynamic Walking*. PhD thesis, Georgia Institute of Technology, 2009.
- [14] Y. Mukaibo, S. Park, and T. Maeno. Equilibrium Point Control of a Robot Arm with a Double Actuator Joint. *International Symposium on Robotics and Automation*, 2004.
- [15] D. Pelleg and A. Moore. X-means: Extending k-means with efficient estimation of the number of clusters. In *Proc. of the Int. Conf. on Machine Learning (ICML)*. Morgan Kaufmann, 2000.
- [16] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 1978.
- [17] E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems 18*, pages 1259–1266, 2006.
- [18] 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 Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [19] J. Sturm, V. Pradeep, C. Stachniss, C. Plagemann, K. Konolige, and W. Burgard. Learning kinematic models for articulated objects. In *Proc. of the Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2009.
- [20] S. Wieland, D. Gonzalez-Aguirre, N. Vahrenkamp, T. Asfour, and R. Dillmann. Combining force and visual feedback for physical interaction tasks in humanoid robots. In *Proc. of IEEE-RAS Intl. Conf. on Humanoid Robots (Humanoids)*, 2009.
- [21] M.M. Williamson. Postural primitives: Interactive behavior for a humanoid robot arm. In *Proc. of the Int. Conf. on Simulation of Adaptive Behavior*, 1996.
- [22] M.M. Williamson. *Robot arm control exploiting natural dynamics*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [23] J. Yan and M. Pollefeys. Automatic kinematic chain building from feature trajectories of articulated objects. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2006.