

Interactive Perception for Learning the Dynamics of Articulated Objects

Felix Endres, Jeff Trinkle, Wolfram Burgard

Abstract—Manipulating articulated objects is an important skill for robots operating in human environments. We propose to learn a physical model of the dynamics of articulated objects to accurately predict the motion of the object. Being aware of the dynamic effects of its actions, the robot no longer needs to maintain a firm grasp of the handle over the full course of the manipulation, which allows for one-point-contact manipulation or early release. This ability reduces the degrees of freedom required of the manipulator and allows for high speed execution. We present an approach to learn the objects’ dynamics from sensor observations of the moving door. The observations can incorporate information from force sensing or depth measurements, as obtained from a laser range scanner. Our method allows the robot to interactively learn the door dynamics, updating the learned model from observations gathered during manipulation. We devise an algorithm to predict the dynamic behavior of doors within the first manipulation, which allows the robot to bootstrap the model itself.

Current approaches to robotic manipulation of articulated objects do not make use of explicit knowledge about the dynamics of the object. Most approaches assume the manipulation execution to be quasi-static, i.e. slow enough that the inertial forces are negligible. This substantially reduces the execution speed for systems that do not maintain a closure grasp of the handle over the course of the manipulation [1]. Approaches in which the robot maintains such a grasp are robust to small inertial forces but require specialized controllers [2] or high dimensional motion planning [3] to avoid lateral forces between manipulator and handle. In all cases, the object needs to be released at rest, which means the end effector needs to be in contact until the desired position is reached. In general, this is a challenging problem, particularly for robots with limited reachability or low number of joints.


We propose to take advantage of the dynamics of articulated objects. We learn the model of the object’s dynamics, i.e. mass or moment of inertia, and the deceleration of the object by friction and air drag with respect to velocity and position. The observations can be acquired during manipulation using force sensing capabilities of the manipulator or using a laser range scanner or depth camera. For a door opening task, we extract the state θ from the mean \bar{p} and principle component axes m, n of the depth measurements in a small bounding polygon placed at the initial guess for the door location. The angle is computed from the door normal. We adaptively grow the boundary by tracking the door (see Figure 1b). The hinge location h is estimated from the sequence $[\bar{p}_t, n_t]$ using a least squares approximation.

We use locally weighted regression to robustly fit a second order polynomial to the time series $[\theta_t]$. The velocity ω_t and deceleration α_t can be obtained by the first and second derivative of the local fit. We use a


This work has partly been supported by the European Commission under the contract number FP7-ICT-248258-First-MM, the National Science Foundation under grant CCF-1208468, FRIAS, University of Freiburg and the DARPA under W15P7T-12-1-0002



(a) Being able to predict the dynamic behavior of the door, the robot can swing it open precisely to a desired state, here 60° . The manipulation time from touch to release of the door is approximately one second.



(b) Computation of door state



(c) Regression of the deceleration and kinematics from depth observations (blue and green dots). forces using a Gaussian process for predicting the object’s trajectory.

Gaussian process to learn the function $f : (\theta, \omega) \rightarrow \alpha$, which allows us to generalize the observations to unseen states (θ', ω') and make accurate predictions of the future trajectory, including the stopping point at any time during the manipulation. A learned model and the corresponding input data is shown in Figure 1c.

For the robot to safely bootstrap the model, we let it move the object at constant velocity for a short distance (10 cm) and use the force sensors to acquire an initial estimate of the object’s braking force (mainly friction and air drag). For a door, this requires to transfer the end-effector’s force and position measurements to torque and angle measurements. Therefore, we learn the kinematic model in parallel (see video attachment) and use the integrals of the measurements at the end of the motion. We compute the moment of inertia during the following acceleration of the door by integrating the difference between applied torque and the estimate of the braking torque. The robot then releases the door as soon as the kinetic energy is sufficient to reach the desired position.

To validate our approach experimentally, we let a robot swing doors open in 45 trials with target states between 45° and 90° . The desired state was achieved with a standard deviation of only 2.6° . The described bootstrapping strategy converged in three trials to within 5° accuracy.

REFERENCES

- [1] J. Sturm, A. Jain, C. Stachniss, C. C. Kemp, and W. Burgard, “Operating articulated objects based on experience,” in *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Taipei, Taiwan, Oct. 2010.
- [2] Y. Karayiannidis, C. Smith, F. Vina, P. Ögren, and D. Kragic, ““open sesame!” - adaptive force/velocity control for opening unknown doors,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 4040–4047.
- [3] S. Chitta, B. Cohen, and M. Likhachev, “Planning for autonomous door opening with a mobile manipulator,” in *Robotics and Automation, 2010. Proceedings. ICRA '10. 2010 IEEE International Conference on*, Anchorage, Alaska, 2010.