

Learning the Dynamics of Doors for Robotic Manipulation

Felix Endres¹, Jeff Trinkle², Wolfram Burgard¹

Abstract—Opening doors is a fundamental skill for mobile robots operating in human environments. In this paper we present an approach to learn a dynamic model of a door from sensor observations and utilize it for effectively swinging the door open to a desired angle. The learned model enables the realization of dynamic door-opening strategies and reduces the complexity of the door opening task. For example, the robot does not need to maintain a grasp of the handle, which would form a closed kinematic chain. Accordingly, it reduces the degrees of freedom required of the manipulator and facilitates motion planning. Additionally, execution is faster, because the robot merely needs to push the door long enough to achieve the right combination of position and speed such that the door stops at the desired state. Our approach applies Gaussian process regression to learn the deceleration of the door with respect to position and velocity of the door. This model of the dynamics can be easily learned from observing a human teacher or by interactive experimentation.

I. INTRODUCTION

Opening doors is fundamental for mobile robots. There is a variety of ways to open a door, which we humans unconsciously choose and execute depending on the situation. For latched doors, we first need to turn the knob or handle. In the following opening motion, though, we generally do not maintain a firm grasp on the handle, as that would restrict our motion when passing through the doorway. As we perceive the dynamics of the door in the course of opening, we can predict its trajectory when released. This ability makes us highly flexible during the execution. For instance, we can release the moving door as soon as its kinetic energy suffices to reach the desired state. Further, it allows us to dynamically switch the contact points and even the manipulating hand without interruption.

Current approaches to robotic door opening do not make explicit use of knowledge of door dynamics. Most approaches assume quasi-static motion, i.e. slow enough that inertial forces are negligible. This substantially reduces the execution speed. Approaches in which the robot maintains such a grasp are robust to small inertial forces but require specialized controllers [1], [2] or high dimensional motion planning [3], [4] to avoid large forces between manipulator and handle. In all cases, the door needs to be released at rest, which means the end effector needs to be in contact




Fig. 1. Top: The robot shoved the door. Bottom: Visualization of the stopping angles of the doors in our experiments with three doors. For each door the robot has been commanded to open to 45, 60 and 90 degrees. The colored circle segments show the range from minimum to maximum, the solid black lines designate the mean, the dashed lines the respective commands. The dynamics model was learned from three observations for each door.

until the desired door state is reached. In general, quasi-static door opening is a challenging problem, particularly for robots with limited reachability or a low number of joints.

We present an approach to learn a model of the door’s dynamics from sensor observations in Section IV. This allows us to make accurate predictions of the door’s behavior. The learned model can therefore be used during the execution to predict the motion of the object at any time. Because the model captures the physical properties of the door, it generalizes over different starting conditions and desired stopping positions. We further propose an approach to let the robot bootstrap the model during the first opening of an unknown door.

We experimentally evaluate our approach in several door opening tasks with a real robot and show that it can be used to open a door quickly with accurate results. The approach is applicable on robots with few degrees of freedom and limited reachability and does not require extensive computational resources.

¹Felix Endres and Wolfram Burgard are with the Department of Computer Science, University of Freiburg, Germany. {endres, burgard}@informatik.uni-freiburg.de

²Jeff Trinkle is with the Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY, USA. trink@cs.rpi.edu

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

II. RELATED WORK

Robotic manipulation of articulated objects has been intensively researched for over a decade.

1) *Detecting Doors*: Several methods for detecting doors [5], [6] and handles [7], [8], [9], [10] have been proposed in the literature. We assume the door location to be roughly known, such that we can identify the door plane.

2) *Unlatching the Door*: Many approaches to door opening focus on unlatching the door and consider the door to be open after unlatching. Reliable results have been experimentally demonstrated [3], [11]. Contrarily, we do not consider unlatching of the door in this work but assume the door to be unlatched to focus on quickly and accurately moving the door to a desired state.

3) *Control Approaches*: Several research groups proposed compliant control algorithms that aim at applying force only along the allowed trajectory, while keeping lateral forces to a minimum [12]. Some approaches learn the kinematic model during the manipulation [13], [1], [2] and use position control along the path, while force control minimizes the forces perpendicular to the allowed motion. Reachability issues can be handled by moving the base [12], [1].

4) *Motion Planning Approaches*: In contrast, motion planning approaches compute a sequence of actions that leads to the desired goal state [3], [4]. Unfortunately the planning space is high dimensional and highly constrained. Finding a valid plan in acceptable time is a topic of intensive research [14], [15], [16], [17].

5) *Estimating the Kinematic Structure*: There has been intensive research on accurately estimating the kinematic model of articulated objects [18], [19]. However, the proposed methods all require tracking of fixed points on the mechanism, e.g., the firmly grasped handle via forward kinematics or visual features or markers. We therefore propose an approach to learn the kinematics using a 2.5D depth sensor such as a laser range scanner in Section IV-A.

III. ANALYTICAL DYNAMICS OF DOORS

A. Dynamics of Hinged Doors

To predict the dynamic behavior of a door from its physical properties, it is sufficient to describe the door by its *moment of inertia* I and the *kinetic friction* τ_f , which describes the resisting torque due to friction within the door hinges and air resistance. According to the law of conservation of energy, assuming the friction torque is constant over a time interval $t-t'$ and neglecting other effects than friction, the reduction in rotational kinetic energy needs to be equal to the work of friction, i.e.,

$$\frac{1}{2}I(\omega_t^2 - \omega_{t'}^2) = \tau_f(\theta_t - \theta_{t'}), \quad (1)$$

where ω_t is the angular velocity at time t and θ_t is the corresponding opening angle. Note that the direction of the friction torque is always opposite to the current velocity ω . Given the angular velocity ω , we can compute the stopping

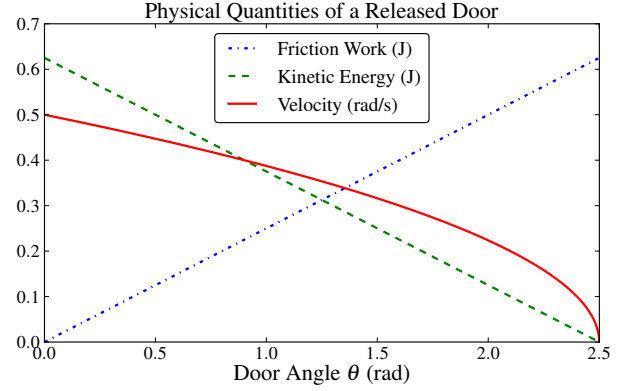


Fig. 2. Door trajectory example for constant friction. See also the trajectory estimates from real data in Figure 6. Note that the velocity has a different physical unit than friction work and kinetic energy and in general need not be in the same scalar range.

distance $\Delta\theta = \theta_t - \theta_{t'}$ by setting $\omega_{t'}$ to zero,

$$\Delta\theta = \frac{I\omega_t^2}{2\tau_f} \quad (2)$$

This lets us determine the position to release the door (i.e. stop accelerating it, if we maintain no grasp) during manipulation such that the door stops at the desired angle.

Conversely, we determine the required velocity at a given release angle θ_0 in order for the door to stop at θ_T as

$$\omega = \sqrt{\frac{2}{I}\tau_f(\theta_T - \theta_0)} \quad (3)$$

Note that I and τ_f only appear in form of the fraction $\alpha = \tau_f/I$ in Equations 2 and 3, we can also use the acceleration α instead.

Given the current position θ_0 , velocity ω_0 and the (constant) deceleration from friction we can compute the trajectory of a door over time as

$$\theta(t) = \frac{1}{2}\alpha t^2 + \omega_0 t + \theta_0. \quad (4)$$

Figure 2 shows an example trajectory computed with the described physical model.

Note that the deceleration of the door is not only influenced by friction and air drag. For example gravity affects non-vertical doors. In this work, all decelerating torques are included in the friction term τ_f .

B. Dynamics of Sliding Doors

Modelling the dynamics of sliding doors is conceptually the same as for hinged doors. However, given an environment with unknown doors of both types, we assume the door type to be identified [18].

The equality of kinetic energy and friction work given in Equation 1 for linear motion is

$$\frac{1}{2}m(v_t^2 - v_{t'}^2) = F_f(x_t - x_{t'}), \quad (5)$$

The stopping distance Δx for a given linear velocity v is thus computed by

$$\Delta x = \frac{mv^2}{2F_f}. \quad (6)$$

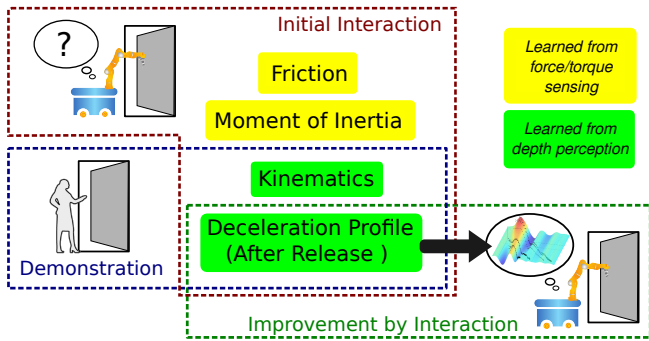


Fig. 3. The proposed methods for learning the various aspects of dynamics model of a door.

The required velocity at a position x_0 to reach $x_0 + \Delta x$ is

$$v = \sqrt{\frac{2F_f \Delta x}{m}}. \quad (7)$$

The kinetic friction of a sliding object can be found during constant-velocity motion as $F_f = F_{ee}^T \hat{x}$ the dot product of the measured force vector and a unit vector along the sliding direction. During experimentation, instead of the moment of inertia (Equation 28), we would estimate the mass m during acceleration from

$$m = \frac{\int_0^T (F_{ee} - F_f)^T \mathbf{x} dt}{v_t - v_0} \quad (8)$$

The approach for learning the friction profile of the door presented in IV-B directly applies to the linear case, by translating measurements and predictions of angle, angular velocity and acceleration to position, linear velocity and acceleration respectively.

IV. LEARNING THE DYNAMICS OF DOORS

To predict the effects of the robot's actions on an articulated object, we need to learn a model of the dynamic behavior of the manipulated object. In this section we present the building blocks that allow the robot to learn the door dynamics interactively or from demonstrations. An overview of the presented methods is shown in Figure 3.

A. Learning the Door Geometry from Observation

To convert between the positions, velocities and forces at the end effector of the robot and the angular positions, velocities and torques of the door, we need to extract a precise geometric description of the door, i.e., the hinge location and the opening angle, from sensor observations. Sturm et al. [18] and others proposed methods to estimate the kinematics of articulated objects, by tracking a fitted rectangle, a marker on the door, or the end effector pose. Here we propose an approach that requires only an arbitrary point and normal on the surface of the door. This allows us to compute the geometry using only a planar depth sensor such as a laser range scanner or a single scan line from a depth camera. Without loss of generality, we assume the depth sensor to provide Cartesian point measurements $\{\mathbf{p}_i \in \mathbb{R}^2\}_0^N$ as obtained by reprojecting the measurements of a horizontal laser range scan.




Fig. 4. Online estimation of the door state and its kinematics from laser range scans (top view). Laser measurements are shown as blue (outliers) and green (inliers) dots. The bounding polygon is initialized as a box at the initial estimate of the door location an expanded online.

Assuming the robot is positioned in front of the door, we filter the measurements with a bounding box to make sure we observe the door only. The dimensions of the box depend on the uncertainty in the estimate of the robot's relative position to the door. We then apply a statistical filter to the remaining measurements to reject points far from their neighbors, e.g., as possibly obtained at the edge of the door.

To estimate the angle of the door, we subtract the mean $\bar{\mathbf{p}}$ from the remaining measurements and determine the door normal of the data set using the principal components analysis (PCA) [20]. In 2D, the eigenvector with the smaller eigenvalue, \mathbf{n} , is the normal. We obtain the current angle θ of the door from \mathbf{n} as

$$\theta = \text{atan2}(n_y, n_x) \quad (9)$$

By repeating this procedure while the door moves, we get several measurements for the mean point $\bar{\mathbf{p}}$ and the door normal \mathbf{n}_j . We track the direction of the door normal, to make sure that it always points in the same direction. After at least three measurements, we can apply a least squares optimization to compute the location of the hinge in closed form. We define the vector $\mathbf{h} = [h_x, h_y, r_h]^T$, where $[h_x, h_y]$ is the hinge location and r_h is the distance of the hinge to the line at the surface of the door. This explicitly allows the hinge to be non-collinear to the door surface. Each measurement must satisfy

$$[\mathbf{n}_j^T \mathbf{1}] \mathbf{h} = \mathbf{n}_j^T \bar{\mathbf{p}}_j. \quad (10)$$

Defining the vector $\mathbf{g} = [\dots \mathbf{n}_j^T \bar{\mathbf{p}}_j \dots]^T$ and the matrix $\mathbf{M} = [\dots [\mathbf{n}_j^T \mathbf{1}]^T \dots]^T$, we can directly compute the hinge parameters by solving the linear system

$$\mathbf{h} = \mathbf{M}^+ \mathbf{g},$$

where $\mathbf{M}^+ = \mathbf{M}^T (\mathbf{M} \mathbf{M}^T)^{-1}$ is the pseudo inverse.

This approach accurately determines the actual center of rotation, but requires observations of a variety of door states before converging. Observing a 90° door swing results in a highly precise estimate. However, for only few degrees of motion, the estimate is far less accurate and stable than for the estimation algorithm of a collinear hinge. We therefore use the measurement $\mathbf{n}_0 = [0, 0]$ (and $\mathbf{n}_0^T \bar{\mathbf{p}}_0 = 0$) as a prior,

which sets the radius to zero while giving no information about the hinge location. This makes the initial estimates equivalent to the collinear estimation case but still converges to the correct location without noticeable delay.

To ensure tracking of the door over the full range of angles, we use the available state estimates to extrude the bounding box into a ring segment. See Figure 4 for an illustration of the results after observing a full opening motion.

B. Learning the Deceleration from Observation

To learn the dynamic behavior of the door, we can use the door state estimation method from Section IV-A to determine the deceleration of the door caused by friction. This is applicable, e.g., when learning from human demonstration or from experimentation.

The deceleration is the second derivative of the angle with respect to time. However, direct numerical differentiation of noisy data amplifies the noise. Assuming the hinge friction to be constant over time, we can describe the trajectory by a second order polynomial. Given N subsequent measurements θ_i at times t_i , we therefore desire the coefficients of the polynomial

$$\theta_i = c_1 t_i^2 + c_2 t_i^1 + c_3 t_i^0 = \mathbf{c}^T \mathbf{t}_i \quad (11)$$

that best fits to the measured data. To determine the least squares fit, we define the squared error to minimize as

$$E(\mathbf{c}) = \sum_i w_i \frac{1}{2} (\theta_i - \mathbf{c}^T \mathbf{t}_i)^2. \quad (12)$$

Here w_i are weights, which we assume to be all equal to one for now. The parameter vector \mathbf{c} that minimizes the weighted squared error can be found by setting the derivative with respect to \mathbf{c} to zero, i.e.

$$\frac{dE(\mathbf{c})}{d\mathbf{c}} = \sum_i [-w_i \theta_i \mathbf{t}_i + w_i \mathbf{t}_i \mathbf{t}_i^T \mathbf{c}] = 0 \quad (13)$$

$$\underbrace{\sum_i w_i \theta_i \mathbf{t}_i}_{3 \times 1} = \underbrace{\sum_i w_i \mathbf{t}_i \mathbf{t}_i^T}_{3 \times 3} \mathbf{c} \quad (14)$$

$$\mathbf{b} = \mathbf{A} \mathbf{c}. \quad (15)$$

After solving the linear system for the polynomial coefficients \mathbf{c} , the deceleration is obtained by the second derivative of the polynomial, i.e., $2c_1$. It can then be used for the prediction of the door behavior as described in Section III-A.

Care has to be taken, to segment the recorded data, as only the part from release to stop follows Equation 4. When the robot learns by experimentation, the release time is known. In case of human demonstrations, we need to determine the release time otherwise. In our experiments, this was achieved by using a wireless mouse to push the door. The demonstrator holds the mouse, such that the button makes the contact with the door during pushing. The release time is then obtained by the button release event. The stopping time is determined by searching for the maximum angle of the trajectory within some seconds of the release time. With this procedure demonstrations can be captured within seconds without requiring specific software or hardware.




Fig. 5. Locally weighted regression (thin black curve) of measured angles (magenta crosses) of a door opening demonstration. The dashed green and red curves correspond to the first and second derivative of the blue curve, and thus constitute the estimations of angular velocity and acceleration.

In practice, the deceleration of doors through friction and other effects may significantly change throughout the trajectory of the door, e.g., when the hinges are slightly non-collinear, non-upright or the door makes contact with the floor. For the quasi-static case, a detailed investigation of *friction profiles* of articulated household objects has been conducted by Jain et al. [21]. We want the predictions to generalize with respect to the starting position and velocity of the door. We thus need to take this variation into account as the total friction work over a distance changes depending on the point of release – and therefore the required kinetic energy for the desired stopping distance. To ensure generalization, we need to learn a profile of the frictional deceleration $\alpha_f(\theta)$ of the door that depends on the opening angle.

To accommodate for varying friction, we need to drop the assumption of a constant deceleration of the door. We therefore adapt the regression approach described above to do locally weighted regression [22]. Instead of solving Equation 15 for the whole trajectory, we compute a local solution at time t_c by setting the weights w_i such that only a small (time-)neighborhood influences the solution. The weights are computed using a weighting function centered at t_c . A common choice is the tricube kernel [22]

$$\Delta_t = \frac{|t_c - t_i|}{l} \quad (16)$$

$$w_i = (1 - \Delta_t^3)^3, \quad \text{for } |\Delta_t| < 1 \quad (17)$$

$$w_i = 0, \quad \text{for } |\Delta_t| > 1. \quad (18)$$

Here, l is a length scale, that sets the trade-off between locality and insensitivity to noise. It should therefore be set depending on the update frequency and accuracy of the state estimation.

Figure 5 shows a demonstrated trajectory, the fitted polynomial and its derivatives. With the ability to compute the acceleration $\alpha_f(\theta)$ at any (observed) position, we can predict the door trajectory given the current position θ_0 and velocity

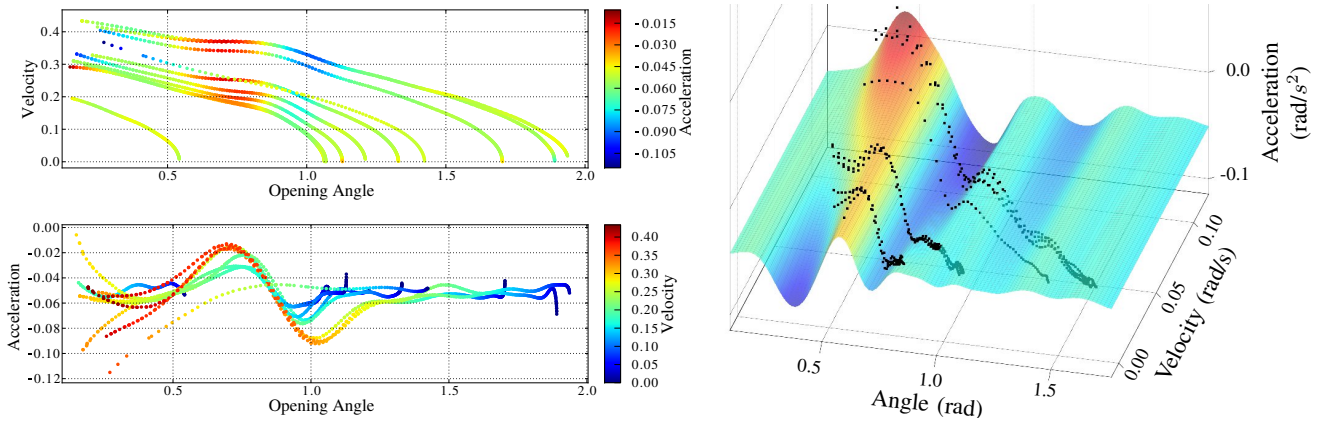


Fig. 6. The estimates of the angular velocity (upper left) and deceleration (lower left) show a clear pattern with respect to the door opening angle. The resulting acceleration profile as modeled by the learned Gaussian process (right) reflects the variations in deceleration over the position and the velocity.

ω_0 . By numerically integrating $\alpha_f(\theta)$, we can use Equation 3 to compute the future velocity at angle θ as

$$\omega^2 = \omega_0^2 + 2 \sum_{i=1}^N \alpha_f(\theta_0 + (i + \frac{1}{2})\theta_s) \theta_s, \quad (19)$$

where $\theta_s = \frac{1}{N}(\theta - \theta_0)$ is the discretization step size. The stopping distance can be computed by summing until the right hand side becomes zero. Figure 6 shows trajectories extracted from several observations.

C. Prediction from Multiple Observations

Using more than one observation to predict the door behavior allows us to substantially improve our prediction. Further, we have found in our experiments that the velocity-dependent deceleration forces, i.e., viscous friction and air drag, are not negligible. Integrating several measurements allows us to learn this dependency. Unfortunately, the dependency of the deceleration with respect to the door angle is not known in advance. In addition, we found that the dependency on the velocity can not easily be described using a parametric model.

We therefore use a Gaussian process (GP) [23], a non-parametric regression method, to model $\alpha_f(\theta, \omega)$, the acceleration w.r.t. angle and velocity. We chose to use a GP because of its beneficial extrapolation behavior for generalization w.r.t. unseen velocities. A GP can be seen as a infinite-dimensional Gaussian distribution. A dimension of the Gaussian is defined by the input location $\mathbf{x} = [\theta, \omega]^T$, the target value $y = \alpha$ constitutes an observation of the Gaussian in the respective dimension. The covariance function defines how target values y and y' covariate based on the input locations \mathbf{x} and \mathbf{x}' . We chose the commonly used squared exponential covariance function (cf. [23], Eq. 2.31)

$$k_{SE}(\Delta \mathbf{x}) = \sigma_f^2 \exp(-\Delta \mathbf{x}^T \Lambda^{-1} \Delta \mathbf{x}) \quad (20)$$

where $\Delta \mathbf{x} = \mathbf{x} - \mathbf{x}'$ and $\Lambda = \text{diag}(l_\theta^2, l_\omega^2)$ contains the different length scales l for each input dimension of \mathbf{x} . Since the GP represents a Gaussian distribution, it is completely defined by the mean vector μ and the covariance matrix K . We use the mean of the training dataset, i.e., $\mu = \frac{1}{N} \sum_{i=1}^N \alpha_i$

and then subtract μ from the target values to obtain the vector $\mathbf{y} = [\dots (\alpha_i - \mu) \dots]^T$. The entries of the covariance matrix K are computed from the input dimensions of the training data as

$$K_{ij} = k_{SE}(\mathbf{x}_i - \mathbf{x}_j) + \delta_{ij} \sigma_n^2. \quad (21)$$

The Kronecker delta δ_{ij} is used to add the measurement variance σ_n^2 to the covariance on the main diagonal.

To predict the acceleration given an angle and a velocity \mathbf{x}_* we infer the mean μ_* of the conditional distribution $p(y_* | \mathbf{y})$

$$\mu_* = k(\mathbf{x}_*, \mathbf{X})(K + \sigma_n^2 I)^{-1} \mathbf{y}, \quad (22)$$

where $k(\mathbf{x}_*, \mathbf{X}) \in \mathbb{R}^N$ is a row vector of the covariances of the new location \mathbf{x}_* with the training dataset $\mathbf{X} = [\dots \mathbf{x} \dots]$. The vector $\mathbf{y}_k = (K + \sigma_n^2 I)^{-1} \mathbf{y}$ does not change during the opening motion and can therefore be precomputed, reducing the online computations to N -fold evaluation of $k(\mathbf{x}_*, \mathbf{X})$ and the dot product with \mathbf{y}_k . To maintain real-time performance for big training datasets, the trajectory data can be downsampled in accordance with the length scale of the weighting function in Equation 16.

Since the deceleration can vary quickly over a few degrees, we set the length scale l_θ to 5° . The velocity dependent variations were found to be nearly linear, with varying slope over θ . We thus enforced approximate linearity with $l_\omega = 1$, which is longer than the range of encountered velocities. The learned GP is used to numerically integrate the estimates of the deceleration analogous to Equation 19 with added dependency on the velocity.

D. Learning the Dynamics from Experimentation

To bootstrap the dynamics model for unknown doors without demonstration more safely, the robot can estimate the friction τ_f and the door's moment of inertia I using force or torque sensing during the contact phase of the first opening action. These estimates allow it to determine when to release the door using Equation 2. The friction in the hinge can be estimated by moving the door at constant angular velocity. For a constant velocity motion of the door, the torque applied

to it by the end effector τ_{ee} is equal to the friction torque τ_f . We can use the linear force F_{ee} at the end effector to compute the applied torque

$$\tau_f = \tau_{ee} = F_{ee} \times r, \quad (23)$$

where the vector r is perpendicular to the axis of rotation and connects said axis with the contact point of the door and the end effector of the robot. Since r is computed from the relative position of the hinge and the end effector, we require an estimate of the hinge axis. This can be estimated online by the method described in Section IV-A. In our experiments, the estimate obtained before releasing the door was always within few centimeters of the final estimate. Since our sensor data is noisy, we average over the measurements from an interval of about 5° .

To determine the moment of inertia of the door we need to estimate ω , the angular velocity, of the door. Analogous to the torque, we can compute the estimate from the linear velocity v_{ee} of the end effector,

$$\omega = v_{ee} \times r. \quad (24)$$

Given the velocity ω and the friction τ_f we can compute the moment of inertia I by accelerating the door. When accelerating the door, we can compute I , from the relation

$$\alpha = \frac{\tau_{ee} - \tau_f}{I}. \quad (25)$$

To be robust to sensor noise we integrate the measurements over time. Without further information, we assume the friction to be constant over the trajectory of the door and independent of the velocity.

$$\int_0^T \alpha(t) dt = \int_0^T \frac{\tau_{ee}(t) - \tau_f}{I} dt \quad (26)$$

$$\omega(T) - \omega(0) = \frac{1}{I} \int_0^T \tau_{ee}(t) - \tau_f dt \quad (27)$$

$$I = \frac{\int_0^T \tau_{ee}(t) dt - \tau_f T}{\omega(T) - \omega(0)}. \quad (28)$$

Given force and position measurements at the end effector at times $\{t_0, \dots, t_i, \dots, t_T\}$, we can compute the corresponding effective torques $\tau_i = |F_i \times r_i - \tau_f|$. The door's moment of inertia can then be estimated by

$$I = \frac{\sum_{i=1}^T \tau_i (t_i - t_{i-1})}{\omega(T) - \omega(0)}. \quad (29)$$

Inserting the estimates for I and τ_f into Equation 2, we can predict the final opening angle during the manipulation. However, the accuracy of this method depends greatly on how well the assumption of constant friction is met. While we have no exact ground truth for the moment of inertia, the estimates obtained in our experiments approximately match our expectations given the mass of the door and assuming a uniform mass distribution.

TABLE I
AVERAGE STOPPING ANGLE AND THE (UNBIASED) STANDARD DEVIATION OF 45 DOOR OPENINGS, FIVE FOR EACH DOOR AND EACH OF THREE TARGET STATES.

Command	90.0°	60.0°	45.0°
Steel Door	91.5° ± 4.7°	61.3° ± 0.4°	46.4° ± 0.6°
Wooden Door A	86.9° ± 0.6°	62.4° ± 1.2°	47.4° ± 0.6°
Wooden Door B	93.6° ± 1.6°	62.8° ± 0.7°	47.2° ± 0.8°
Summary	90.6° ± 4.0°	62.2° ± 1.0°	47.0° ± 0.8°

V. EXPERIMENTS

To evaluate the presented approach we investigate the ability of the robot to learn and apply the model of the door dynamics to swing closed doors to a desired opening angle.

A. Experimental Setup

1) *Robot*: We use the DLR Light Weight Robot (LWR), a seven degree of freedom manipulator, with a passive end effector. It is mounted on a KUKA omniRob mobile base. We use a Hokuyo UTM-30LX laser range scanner to observe the door at 40 Hz. We let the robot push the door using a linear position-controlled motion, such that the door achieves a velocity sufficient to reach the goal state.

2) *Doors*: We demonstrate our approach on substantially different door types. The first door is a metal door attached to a concrete wall of the building (see Figure 1). The door is comparably heavy but smooth-running. The second door is made of a veneer on a wooden framework and is attached to a freestanding wooden frame. It weighs 27.8 kg. The friction is very low in the beginning, but the hinges are slightly misaligned, which increases the friction substantially towards 90° . To increase the variety, we conducted a further set of experiments with a brush seal attached to the wooden door to increase the friction and an additional weight of 1.6 kg firmly attached to the handle. We will refer to these two configurations of the wooden door as “A” and “B”.

B. Door Dynamics Estimation

1) *Human Demonstration*: We first evaluate the accuracy and precision of the opening task using models generated from observing a human demonstrator opening the doors. We demonstrated three pushes with varying stopping angle for each door. The robot was then commanded to open each door to 45° , 60° and 90° . To evaluate the performance, each opening was repeated five times, resulting in a total of 45 executions. The model has not been updated from the robot's own actions. The results are given in Table I and in Figure 1. The time the robot was in contact with the door in our experiments ranges from 0.6 s to 2.0 s with an average of only 1.0 s. The release angle was between 6 and 16 degrees.

2) *Interactive Experimentation*: While demonstrations with the method described in Section IV-B are quickly done, the robot should be able to learn from its own actions. We evaluated the opening performance in a task sequence, where the robot generates and updates the model from its

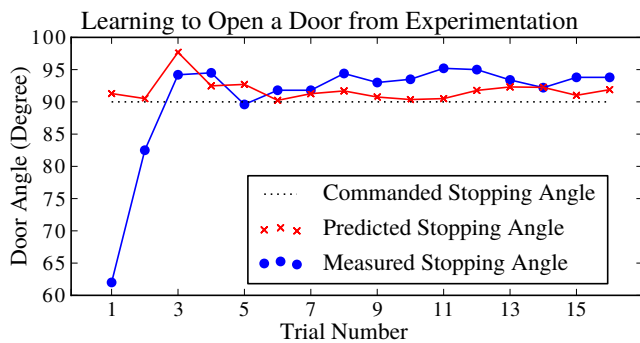


Fig. 7. We let the robot experiment with the door. Initially it has no prior knowledge of the door, therefore it estimates the friction and moment of inertia as described in Section IV-D. From the second trial on the robot uses the observations from previous trials.

own experimentation. For the first opening, we apply the procedure described in Section IV-D. To make the task challenging, we chose the door with the increasing friction, such that the initial estimate is guaranteed to fail. The robot therefore falls almost 30° short of opening the door to 90° . However, after only two observations, the result is within 5° of the desired opening angle. Figure 7 shows the whole sequence of 16 trials. The robot releases the door as soon as it predicts the target state to be reached. The prediction system may notice this to late, because of the time required to get the end-effector pose updates from the robot’s operating system. In this case the robot is aware that the door is overshooting and the estimate of how much is given by the red dots in the figure.

VI. CONCLUSION

In this paper we investigated the manipulation of doors with explicit consideration of dynamic effects. Our approach allows the robot to learn the dynamic behavior of a door by interactive experimentation or observation of demonstrations by a human. The learned model of the door includes the kinematic model, moment of inertia and the deceleration of the door through friction. Using a Gaussian process for non-parametric regression, we modeled the deceleration with respect to angle and velocity of the door. We use the model to predict the trajectory of the door throughout the manipulation. In experiments with a real robot, we demonstrated that the dynamic model allows to reduce the complexity of a door opening task, as a point contact is then sufficient, for accurately swinging the door quickly to a desired angle.

REFERENCES

[1] L. Peterson, D. Austin, and D. Kragic, “High-level control of a mobile manipulator for door opening,” in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2000, pp. 2333–2338.

[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 *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 1799–1806.

[4] F. Burget, A. Hornung, and M. Bénéwitz, “Whole-body motion planning for manipulation of articulated objects,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2013.

[5] D. Anguelov, D. Koller, E. Parker, and S. Thrun, “Detecting and modeling doors with mobile robots,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2004, pp. 3777–3784.

[6] R. B. Rusu, W. Meeussen, S. Chitta, and M. Beetz, “Laser-based Perception for Door and Handle Identification,” in *Proc. of the Intl. Conf. on Advanced Robotics (ICAR)*, 2009.

[7] D. Kragic, L. Petersson, and H. I. Christensen, “Visually guided manipulation tasks,” in *Robotics and Autonomous Systems*, 40(2-3):193–203, 2002.

[8] E. Klingbeil, A. Saxena, and A. Ng, “Learning to open new doors,” in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2010, pp. 2751–2757.

[9] M. Quigley, S. Batra, S. Gould, E. Klingbeil, Q. Le, A. Wellman, and A. Ng, “High-accuracy 3d sensing for mobile manipulation: Improving object detection and door opening,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2009, pp. 2816–2822.

[10] A. Andreopoulos and J. K. Tsotsos, “Active vision for door localization and door opening using playbot: A computer controlled wheelchair for people with mobility impairments,” in *CRV*, 2008, pp. 3–10.

[11] A. Jain and C. C. Kemp, “Behavior-Based Door Opening with Equilibrium Point Control,” in *RSS Workshop: Mobile Manipulation in Human Environments*, 2009.

[12] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Erubimov, T. Foote, J. Hsu, R. B. Rusu, B. Marthi, G. Bradski, K. Konolige, B. P. Gerkey, and E. Berger, “Autonomous door opening and plugging in with a personal robot,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 729–736.

[13] J. Sturm, A. Jain, C. Stachniss, C. 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)*, 2010, pp. 2739–2744.

[14] J. Kuffner and S. M. LaValle, “RRT-connect: An efficient approach to single-query path planning,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2000, pp. 995–1001.

[15] Y. Yang and O. Brock, “Elastic roadmaps - motion generation for autonomous mobile manipulation,” *Autonomous Robots*, vol. 28, no. 1, pp. 113–130, 2010.

[16] J. Trinkle and R. Milgram, “Complete path planning for closed kinematic chains with spherical joints,” *International Journal of Robotics Research*, vol. 21, no. 9, pp. 773–789, 2002.

[17] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner, “Manipulation planning with caging grasps,” in *Humanoids*, 2008, pp. 285–292.

[18] J. Sturm, C. Stachniss, and W. Burgard, “A probabilistic framework for learning kinematic models of articulated objects,” *Journal on Artificial Intelligence Research (JAIR)*, vol. 41, pp. 477–526, 2011.

[19] D. Katz and O. Brock, “Extracting planar kinematic models using interactive perception,” in *Unifying Perspectives in Computational and Robot Vision*. Springer US, 2008, vol. 8, pp. 11–23.

[20] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[21] A. Jain, H. Nguyen, M. Rath, J. Okerman, and C. C. Kemp, “The Complex Structure of Simple Devices: A Survey of Trajectories and Forces that Open Doors and Drawers,” in *Proc. of the IEEE RAS/EMBS*

Intl Conf. on Biomedical Robotics and Biomechatronics (BIOROB), 2010, pp. 729–736.

- [22] W. S. Cleveland, “Robust locally weighted regression and smoothing scatterplots,” *Journal of the American Statistical Association*, vol. 74, no. 368, pp. 829–836, 1979.
- [23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2005.