Behavior-Based Door Opening with Equilibrium Point Control

Advait Jain and Charles C. Kemp
advait@cc.gatech.edu, charlie.kemp@bme.gatech.edu

***Abstract*— Within this paper we present a set of behaviors that enable a mobile manipulator to reliably open a variety of doors. After a user designates a location within 20cm of the door handle, the robot autonomously locates the door handle using a tilting laser range finder, approaches the handle using its omni-directional base, reaches out to haptically find the door, makes contact with the handle, twists it, and pushes open the door.**

The robot uses equilibrium point control for all arm motions. Our implementation uses a 7 DoF anthropomorphic arm with series elastic actuators (SEAs). For our control scheme, each SEA applies a gravity compensating torque plus a torque from a simulated, torsional, viscoelastic spring. Each virtual spring has constant stiffness and damping, and a variable equilibrium point. The behaviors use inverse kinematics to generate trajectories for these joint-space equilibrium points that correspond with Cartesian equilibrium point trajectories for the end effector.

With 43 trials and 8 different doors, we show that these compliant trajectories enable the robot to robustly reach out to make contact with doors (100%), operate door handles (96.9%), and push doors open (100%). The complete system including perception and navigation succeeded with unlocked doors in 28 out of 32 trials (87.5%) and locked doors in 8 out of 8 trials (100%). Through 157 trials with a single door, we empirically show that our method for door handle twisting reduces interaction forces and is robust to variations in arm stiffness, the end effector trajectory, and the friction between the end effector and the handle.

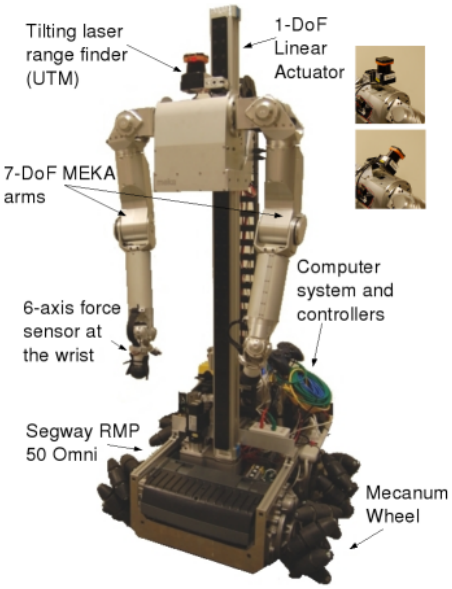

Fig. 1. The mobile manipulator used in this paper and the servo tilting the laser range finder to generate a 3D point cloud.

I. Introduction

Fully autonomous service robots will often need to open doors to freely operate within human environments. And, assistive robots that open doors on command would be valuable for physically-impaired users. In spite of these opportunities, autonomous manipulation of doors remains a challenging problem after more than a decade of research [1, 2]. In addition to societal utility, door opening is a worthwhile challenge problem for autonomous mobile manipulation due to the clear measures of success, ubiquity of examples for testing, large real-world variation, and relatively simple mechanics.

Within this paper we make two main contributions. First, we present a full system that can robustly approach and push open a novel, handle-operated door designated by a user. Unlike previous work, our system does not require a map [3, 4], training [5, 6], an explicit door model [7, 8, 9], nor estimates of the handle and door kinematics [10, 2, 11].

Much of the prior work on door opening has focussed on only one or two aspects of the complete task [12, 13, 14, 15, 11]. Instead, we present a system that integrates perception, navigation and manipulation to servo up to the door handle and open the door. Moreover, the performance of our system in the overall task, including perception, navigation and manipulation, is competitive with the best results published to date [5], and integrates haptic feedback to detect common circumstances such as a locked door or blocked door.

In contrast to our previous work [16], our new robot operates more efficiently. It opens a door in approximately 2 minutes total (including navigation over 1.5 meters) and twists the handle in around 4 seconds. Our new system also uses distinct control algorithms and behaviors.

Our second contribution is that we present a straightforward implementation of equilibrium point control and empirically demonstrate that it is successful and robust under real-world conditions. Even though both the handle and the door have rotary joints, we found that compliant trajectories that command the end effector to follow lines in Cartesian space enable the robot to succeed at unlatching a door and pushing it open by approximately 30cm. Moreover, the behaviors are robust to variations in the control parameters and the environment.

Long-standing debates continue about whether or not the Equilibrium Point Hypothesis is true for human motor control [17, 18]. Fortunately, we only need to determine if equilibrium point control is useful for robots, not whether it is used by humans. Previous robotics research has looked at similar robotic

control strategies in simulation [19], in freespace motions [20], in legged locomotion [21], in rhythmic manipulation from a fixed based [22], and in the design and control of compliant actuators [23, 24]. However, few, if any, studies have empirically evaluated this form of control in the context of task-oriented mobile manipulation.

Equilibrium point control offers a promising alternative to other forms of control. As we demonstrate, it can be effectively used for both freespace reaching trajectories and mechanically constrained manipulation tasks under real-world conditions in unstructured environments. In contrast to position control methods, our implementation does not require high-fidelity estimates of the kinematics of the handle and door [6]. Unlike some approaches to force control, we do not use inverse dynamics [25].

Like other forms of impedance control, equilibrium point control naturally handles unexpected perturbations, such as collisions. This active compliance in conjunction with the passive, mechanical compliance of the series elastic actuators reduces the chances of damage to the robot, the environment, and nearby people. The end effector stiffness of the manipulator in the plane of the door in our system is relatively low. For example, it is lower by around a factor of five compared to work on door opening with Cartesian space impedance control using the DLR-Lightweight-Robot-II [26]. Since the freespace motions used by the robot are designed to move until contact at an unknown distance along the trajectory, this compliance is advantageous.

In general, one can think of equilibrium point control as generating compliant trajectories with a stiffness that varies as a function of both the stiffness of the springs and the arm's posture. While similar compliant trajectories could be implemented using other forms of impedance control, such as Cartesian space impedance control [26], we believe equilibrium point control is less complex in practice and effective for real-world tasks in unstructured environments.

II. The Robot

The robot is a new, as yet unnamed, statically stable mobile manipulator that our lab, the Healthcare Robotics Lab, assembled in early 2009 (see Figure 1). It consists of arms from MEKA Robotics (MEKA A1), an omni-directional mobile base from Segway (RMP 50 Omni), and a 1-DoF linear actuator from Festo that can lift the manipulator and sensors from ground level to 1.2m above the ground. For this work, a rubber ball serves as the end effector. Distinctive features of this new robot include the use of series elastic actuators [27] in all 14 DoF of the two arms (7 DoF each) and four Mecanum wheels for the base.

A. The Software and the Sensors

A Mac Mini running Ubuntu GNU/Linux performs all of the computation for sensing and high-level control. There is also a Dell Studio Hybrid that runs Ubuntu GNU/Linux with a kernel patched with RTAI for real-time operation. It performs computations for the MEKA arms. We have written all our software in Python and make use of a variety of open source packages including SciPy, KDL, ROBOOP, OpenCV and ROS (Robot Operating System).

For this work, the robot uses two distinct types of sensors. First, the robot uses a laser range finder (Hokuyo UTM-30LX) mounted on a servo motor (Robotis Dynamixel RX-28) on top of the torso. The servo tilts the laser range finder about the horizontal axis (Figure 1). We use this tilting laser range finder to obtain 3D point clouds of the environment. The laser range finder has a resolution of 0.25° and we obtain planar scans at 20Hz. The servo encoders have a resolution of 0.3°. This sensing configuration has been inspired by the Personal Robot 2 from Willow Garage.

Second, the robot senses forces and torques using a wrist-mounted 6-axis force/torque sensor (ATI Mini40 from ATI Industrial Automation). The arm's joints also sense torque, but the current behaviors only use this sensing implicitly in the context of virtual spring control.

III. System Overview

We now briefly describe the different steps that the robot goes through to navigate to a door handle and push open the door as shown in Figure 2. The input to the door opening behaviors presented in this paper is a rough 3D location for the door handle in the robot's coordinate frame. We have previously presented interfaces (laser pointer interface and a touch screen) with which users can select a 3D location in the world [28]. More generally, this 3D location could be provided by an autonomous perception system. In this work, we display a scan from the laser range finder and allow the user to click on any point to designate the location of the handle. We also fix the initial estimate of the height of the door handle to 1m.

In our experiments, the robot starts around 1.5m from the door handle and the user selects a 3D location in the neighborhood of the door handle. The robot then makes a 3D scan with the tilting laser range finder, segments the door handle (Section IV-C) and lowers the laser range finder so that it scans at the height of the door handle (Section IV-D). The robot then approaches the door handle (Section V) and takes another 3D scan. After segmenting the door handle, the robot haptically finds the surface of the door and makes contact with the door handle near the appropriate tip (Section VI-B). Finally, the robot twists the door handle and pushes the door open (Sections VI-C and VI-D).

IV. Perceiving the Door

The robot uses its laser range finder to make estimates about different task-relevant features related to the door. The robot's perception starts with raw sensor data (a point cloud consisting of $\sim$ 40,000 3D points or a planar laser scan of around 600 points) and reduces it to low-dimensional, task-relevant features: the estimated location of the door handle, the two tips of the handle (left or right), the estimated height of the handle above the ground, and the estimated orientation of the door. Our methods for segmenting and detecting the door and door handle have similarities with [29] and contemporaneous

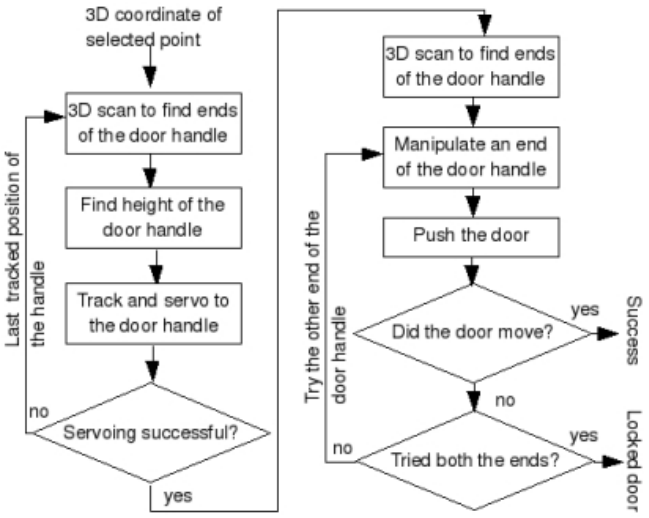

Fig. 2. The different behaviors that the robot executes to open a door.

work on perceiving doors [30], although we only perform these perceptual tasks within a small volume around a user-selected 3D location.

A. Estimating Door Orientation and Distance

The robot first estimates the orientation of the door (with respect to an axis of rotation parallel to gravity) and the distance of the door from the robot using a non-tilting scan from the laser range finder. In addition to this scan, the algorithm requires that a location near the door be provided. The orientation and distance estimates are used when approaching the door and when detecting and segmenting the handle.

The algorithm assumes that the door is wider than 0.3m, so it first finds line segments of length greater than 0.3m in the laser scan using a Hough transform from OpenCV (see Figure 5). For each line segment, it then finds the shortest distance between the provided location and the line segment. The algorithm rejects segments with a distance greater than 0.5m. It also assumes that the robot is within 45° of facing the door and throws out any line segments that violate this assumption. Finally, it assumes that the door will either be recessed or flush with the walls, so the algorithm returns the distance and orientation of the line segment whose perpendicular distance from the robot is the greatest.

B. Detecting the Door

The door detection algorithm takes as input a 3D point cloud and a location that is believed to be less than 20cm from the door handle. This algorithm and the handle segmentation algorithm that follows use a coordinate frame whose origin is at the robot with the floor parallel to the X-Y plane and the door parallel to the Y-Z plane, as shown in Figure 3. The robot maintains this coordinate frame by estimating the orientation of the door in the X-Y plane as described in the previous section. Both algorithms only analyze points from the 3D point cloud that fall within a volume of interest (VOI). This VOI is an axis aligned box of 40cm x 60cm x 40cm along the X, Y, and Z axes and is centered on the provided location near the handle.

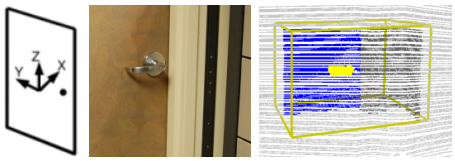

Fig. 3. **Left:** Coordinate axes parallel to the axes of the frame with the origin at the robot that is used by the door and door handle detection algorithms. **Middle:** Image of a door and handle. **Right:** The output of the door handle segmentation algorithm. The volume of interest (VOI) is the yellow box, the door is blue, the door handle is yellow, and the points inside the VOI are dark grey. The point cloud outside the VOI is light grey.

The detection algorithm assumes that the door is planar, parallel to gravity within the VOI, and is represented by a significant number of points in the VOI. Since the robot already has an estimate for the door's orientation in the X-Y plane, these assumptions imply that it only needs to estimate the distance of the door from its origin along the X-axis. To do this, the algorithm uses a histogram in which each bin covers a mutually exclusive 5mm range of depths along the X-axis in the VOI. It then estimates the depth of the door as the depth corresponding to the bin with the largest number of entries.

C. Segmenting the Door Handle

Once the robot has found the plane associated with the door, it finds the handle protruding from its surface. To segment the door handle within the VOI, the handle segmentation algorithm removes all points on and behind the plane of the door, and any point less than 3.5cm in front of the detected plane of the door. It then converts the remaining point cloud into a 3D occupancy grid. The resolution of the grid is 2cm along the X and Z axes and 0.25cm along the Y-axis. The algorithm clusters these grid cells into objects by performing a 3D connected components labeling with 26-connectivity.

The algorithm assumes that the door handle will not be shorter than 4cm (dimension of the cluster along the Y-axis) and will not be wider than 10cm (dimension of the cluster along the Z-axis), so it throws out all clusters whose dimensions violate these assumptions. It then selects the cluster closest to the robot's initial estimate of the door handle location as the door handle. Figure 3 shows the output of the door and door handle segmentation algorithms.

D. Estimating the Door Handle's Height

In this implementation, the laser range finder needs to scan parallel to the ground at the height of the door handle when servoing to the door handle. To estimate the height of the handle with high precision, the robot scans parallel to the ground while the linear actuator moves the robot's torso down. Points on the scan line that are to the right or left of the segmented handle are removed from the scan. While descending, the robot stores the average distance of the remaining points from the estimated surface of the door. The height with the maximum average distance is estimated to be the door handle's height.

Fig. 4. These images show the robot servoing to the door handle.

Fig. 5. **Left:** Position of the robot (red circle), the points from the laser scan (dark blue), the estimate of the door (green line) and the tracked position of the door handle (black circle) during the servoing behavior. **Right:** (clockwise from the left) Laser scan rotated and converted into an image, the template of the door handle and the result of correlation of the template with the image.

V. Approaching the Door

The robot first performs a 3D scan of the area surrounding the user-selected location and segments the door handle. This segmentation gives the robot the approximate position of the door handle and the positions of the tips of the door handle in its frame of reference. With reliable odometry, this would likely be sufficient for the robot to approach the handle. In contrast to our previous work, however, we do not have reliable odometry estimates due to slipping by the Mecanum wheels. Consequently, for this implementation we servo the robot to the handle using the laser range finder.

Prior to servoing, the robot lowers its laser range finder to look directly at the handle using the method described in the previous section (see Figure 4). It then uses its omni-directional base to simultaneously and independently servo its orientation and its position based on the estimated orientation of the door and the estimated position of the handle. The robot continues to servo to the door handle until it is within the workspace of the door opening controller for the arm.

To estimate the position of the handle while servoing, the robot tracks the handle. The handle tracker converts the laser scan into an image and uses a template and 2D image correlation to estimate the handle's location. Prior to this conversion the tracker rotates the laser scan into a canonical orientation based on the orientation estimate for the door. The tracker acquires its initial template from the estimated location of the door handle tips from the handle segmentation (see Figure 5). As it tracks the handle, it updates the template. If the tracker detects that it has lost the handle, due to a jump in the handle's estimated pose (change exceeds 0.1m), the robot rescans the environment and tries again. For our experiments the robot is only allowed to try again once.

VI. Manipulating the Door

Once the servoing has stopped, the robot assumes that it is facing the door, and the door handle is within the workspace of the door opening controller for its right arm.

A. Equilibrium Point Control

The robot uses a form of equilibrium point control for all arm motions, except the wrist. We control the wrist using a high-stiffness position controller. The wrist is held such that it will be normal to the plane of the door if the robot is facing the door.

On the Dell Studio Hybrid, a control loop computes desired torques for the joints, τ, as the sum of two torque vectors at a rate of 1kHz.

$$\tau = -g(q) + (-K_p\tilde{q} - K_d\dot{q}) \quad (1)$$

The first torque vector, $g(q)$, is the torque due to gravity as a function of the current joint angles q. Subtracting it provides gravity compensation. The second torque vector, $-K_p\tilde{q} - K_d\dot{q}$, simulates a torsional, viscoelastic spring with constant stiffness and damping at each joint. $\tilde{q} = q - q_{eq}$ where q_{eq} is the equilibrium point in the joint space, and K_p and K_d are diagonal stiffness and damping matrices.

We wrote the outer control loop which runs on a Mac Mini running Ubuntu GNU/Linux. This sends equilibrium points to all of the joints at a rate of approximately 10Hz. At each time step, the control loop uses the inverse kinematics solver from KDL[1] to compute joint angles that would lead to the end effector following the commanded trajectory, if the arm had no compliance. These joint angles are used as the joint-space equilibrium point, q_{eq}, which corresponds to a Cartesian equilibrium point for the end effector that moves along the commanded trajectory. For most of the motions in this work, the trajectory moves the Cartesian equilibrium point along a line in Cartesian space at a constant rate, until sensing indicates that the arm should stop.

B. Making Contact with the Handle

The robot reaches out towards a position on the door that is above the desired contact location on the handle. It first uses the linear actuator to raise the arm up towards this location. It then positions the arm laterally so that it can reach out using a trajectory normal to the door that passes through this target location. Finally, it reaches out toward the position and stops when it detects contact with the door. It then moves the end effector away from the door by 2cm and uses the linear actuator to move down until the end effector makes contact with the handle near the tip. During this downward movement, the arm is held in a fixed compliant posture. This is shown in Figure 6.

C. Twisting the Handle

Figure 6 shows the robot twisting a door handle in the counter-clockwise direction. The twisting behavior moves the Cartesian equilibrium point of the manipulator (Section VI-A) along a line in the plane parallel to the surface of the door. The twisting behavior updates the equilibrium point by 0.5cm along the line at a rate of 10Hz. For most of our experiments, the line has an angle of 20° or -20° with the

[1]Kinematics and Dynamics Library (http://www.orocos.org/kdl)

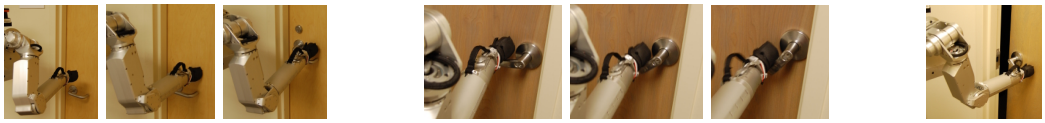

Fig. 6. **Left:** The robot reaches out to haptically detect the surface of the door and then moves down until it makes contact with the handle near the tip. **Middle:** The robot twists the handle. **Right:** The robot pushes open the door.

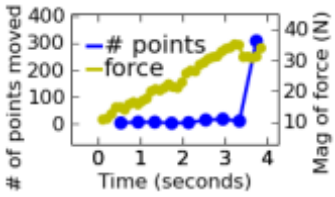

Fig. 7. Motion of the door as observed by looking for points that moved in the laser scans (blue) and the magnitude of the contact force at the end effector (yellow) while the robot twists the door handle. The X-axis is the time in seconds (twisting the door handle took around 4 seconds).

vertical, depending on which tip of the handle the robot is manipulating (left or right).

To decide which tip of the door handle to manipulate first, we use a heuristic based on the estimated line segment of the door (Section IV-A). If the handle is closer to the right end of the line segment, the robot manipulates the left end of the handle and vice-versa. For a flush door, the line segment may not be informative. To compensate for errors in estimating the extent of the door, the robot tries both the ends of the door handle.

In addition to the line parallel to the plane of the door, the twisting behavior also changes the equilibrium point in the direction normal to the door with the goal of maintaining a pushing force of 10N. It uses a bang-bang controller that at each time step moves the equilibrium point out of the door by 0.5cm if the pushing force exceeds 10N, and into the door by 0.5cm if the force falls below 10N.

While it twists the door handle, the robot monitors its progress using the laser range finder, which scans parallel to the ground, and the wrist-mounted force/torque sensor. The robot looks for motion of the door by comparing the current laser scan with a scan taken just prior to twisting. It subtracts these scans and counts the number of points that have moved by more than 0.5cm. The pushing force of 10N that the twisting controller tries to maintain results in a sharp increase in this count when the handle is twisted enough, as shown in Figure 7. The twisting controller continues until either the robot detects significant door motion, the magnitude of the total resultant force exceeds 45N (indicating that the robot has twisted the handle to the end stop or that the door is locked), or the magnitude of the force falls below 2N (indicating that the end effector has slipped off the handle).

D. Pushing the Door Open

Without moving the mobile base, the robot attempts to push the door open with a straight line equilibrium point trajectory. It continues along this trajectory until the equilibrium point has moved forward by 30cm or it senses an opposing force of 30N along the X-axis. If this force limit is exceeded, the robot assumes that either the twisting behavior was unsuccessful or an obstruction is blocking the door.

After completing this push, the robot once again uses its laser range finder to see if the door has moved. If the door did not move, the twisting behavior failed on that tip of the handle and the robot tries to manipulate the other tip. If twisting both the tips of the handle is unsuccessful, the robot reports that the door is locked.

VII. Evaluation

In Section VII-A, we report on 157 trials with a single door, and empirically show that our method of twisting door handles is robust with respect to the orientation of the trajectory, the stiffness of the arm, and the friction between the end effector and the handle. Then in Section VII-B, we present the performance of all the behaviors and the overall system on 8 doors in 43 trials. We carried out all trials in mid-March, 2009.

A. Twisting Door Handles

To provide intuition about the mechanics of door handle twisting, Figure 8 shows a simple quasistatic model of the interaction forces between the end effector of the manipulator and a door handle. If we model the door handle as a pin joint with a torsional spring of stiffness K and ignore the mass of the door handle, then

$$N = K\theta/l \quad (2)$$

$$f \leq \mu_s N \quad \text{(no slip)} \quad (3)$$

$$f = \mu_k N \quad \text{(slip)} \quad (4)$$

where N is the normal force, θ is the angle of the door handle with the horizontal, l is the distance between the point of contact and the axis of the handle, f is the friction force and μ_s and μ_k are the coefficients of static and kinetic friction. The normal force increases as the door handle is twisted, and the direction of friction force can change depending on the direction of impending (or actual) slip.

Friction is important for our handle twisting behavior as it prevents the end effector from slipping off (or slipping up) the handle. If a robot rigidly grasps the door handle, it can be modeled as a high or infinite coefficient of friction. If the end effector motion corresponds exactly to the geometry of the door handle, there will be no impending slip and the interaction force will be as small as possible for a fixed contact location on the handle and equal to the normal force. With zero friction, the end effector will only sense the normal force. Our

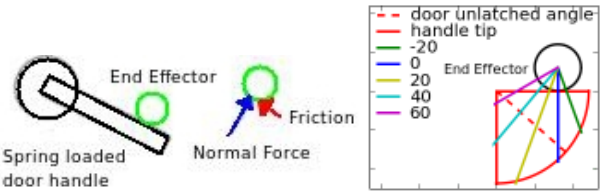

Fig. 8. **Left:** Model of the end effector twisting a door handle. **Middle:** In plane forces on the end effector from the handle. **Right:** Figure showing how the handle is constrained to move (red), the different lines along which we varied the equilibrium point for the end effector, and the initial position of the end effector when the twisting behavior starts (black). The length of the handle is 12cm and the end effector made contact approximately 4cm from the tip of the handle. The angle of the door handle at which the door can be pushed open is shown as a dashed red line.

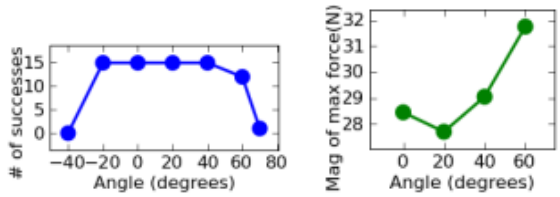

Fig. 9. **Left:** Number of successes (out of 15) for different angles for the equilibrium point line of the twisting behavior. **Right:** Average (over 15 trials) of the magnitude of the maximum force measured by the wrist force-torque sensor for different angles of the equilibrium point line.

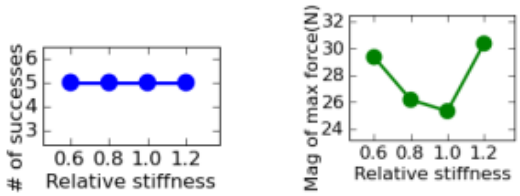

Fig. 10. **Left:** Number of successes (out of 5) for different stiffness settings. **Right:** Average (over 5 trials) of the magnitude of the maximum force measured by the wrist force-torque sensor while twisting the handle for different stiffness settings of the virtual joint springs.

previous work relied on quasistatic motion and low friction to estimate the normal force and push against it while twisting the handle [16].

We now describe experiments on one door handle to evaluate the performance of the twisting behavior (Section VI-C) when its parameters are varied.

1) Angle of the Equilibrium Point Trajectory: While twisting the handle, the trajectory of the equilibrium point is the sum of a line parallel to the surface of the door and small motions normal to the surface (see Section VI-C). In this section we empirically show that the performance of the system is robust with respect to the orientation of the line.

To simulate variation due to navigation, we put the robot in three different orientations in front of the door (-10°, 0° and 10°) with the handle in the workspace of the manipulator. For each orientation, we varied the angle of the line in increments of 20° from -40° to 70° as shown in Figure 8 and performed five trials for each angle. We kept all the other control parameters constant.

The robot successfully twisted the handle and pushed the door in 15 out of 15 trials when the angle of the equilibrium point line was between -20° and 40°, failed three times out of 15 when the angle was 60°, and was unsuccessful in all trials for -40° and 70° (see Figure 9). For -40° the end effector slipped off the door handle and for 70° it slipped up the handle towards the pin joint. This suggests that there is a range of orientations of the equilibrium point line for which our control strategy will be successful.

Figure 9 also shows the magnitude of the maximum interaction force between the end effector and the handle and door for the different trajectories. This magnitude includes the component of the interaction force normal to the surface of the door. From the model of the door handle and Figures 8 and 9, we see that as the equilibrium point line deviates from the trajectory of a point on the handle, the maximum force measured at the end effector increases. Given our model, we expect that this results from increased frictional forces and a reduced moment arm.

Minimizing the interaction forces during manipulation reduces the chance of damage. Based on the results shown in Figure 9, 20° appears to reduce the interaction forces and sits near the middle of the successful range of angles. Unless otherwise noted, we use 20° for the trajectory angle of twisting in the remaining experiments.

2) Stiffness of the Manipulator: In this section we show how the maximum interaction force between the handle and the manipulator varies as the stiffness of the manipulator changes. For our method of twisting handles, the interaction forces increase as the stiffness of the manipulator increases.

The manipulator that we have used in this work is compliant due to the SEAs and virtual springs. The joint stiffness settings that we used for all the experiments result in a stiffness at the end effector of approximately $1800N/m$ normal to the surface of the door, and $250N/m$ and $200N/m$ in the horizontal and vertical directions parallel to the plane of the door. The stiffness is almost equal in the plane of the door and much higher in the direction normal to the surface of the door.

To scale the stiffness of the arm, we multiply the output torque of the virtual springs by a scalar, α, which gives

$$\tau = -g(q) + \alpha(-K_p\tilde{q} - K_d\dot{q}). \quad (5)$$

This is equivalent to scaling the stiffness and damping at each joint by α, which we set equal to 1.0 for all experiments other than this one. For this experiment, we vary α from 0.6 to 1.2 to assess the performance of handle twisting as a function of the arm's overall stiffness and damping (see Figure 10).

The robot successfully twisted the handle and pushed the door open in five out of five trials for each value of α. At the lower values, 0.6 and 0.8, the end effector had a tendency to slip on the handle, but still managed to succeed at the task. When we increased α to 1.2, the magnitude of the maximum interaction force increased. As the stiffness increases, the arm should behave more like a position controlled arm, and we would expect greater interaction forces.

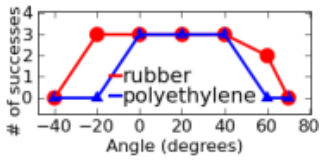

Coefficient of friction

	Static	Kinetic
rubber-steel	0.6-0.9	0.3-0.6
polyethylene-steel	0.2	

Fig. 11. **Top:** Number of times (out of three) the robot successfully twisted the handle and pushed open the door with different materials on the end effector (rubber and polyethylene) as a function of the angle of the equilibrium point line. **Bottom:** Coefficient of static and kinetic friction for rubber-steel and polyethylene-steel contact [31].

3) Friction: In all the other experiments, we covered the end effector with a sheet of rubber to increase the friction between the end effector and the door handle. We now show how the performance of the handle twisting behavior changes with variation in the friction. In this experiment we covered the end effector with polyethylene which has a lower coefficient of friction on steel than rubber (see Figure 11) [31]. In practice, we would also expect a change in the contact properties due to variations across door handles.

Figure 11 shows the results of repeating the experiments of Section VII-A.1 with rubber and polyethylene on the end effector. When polyethylene was in contact with the handle, the end effector slipped more easily. The range of angles for the equilibrium point line for which door handle twisting was successful was smaller, but our control strategy was successful in pushing open the door with both types of contact.

B. The Complete System

We carried out a total of 43 door opening trials on 8 different doors, as shown in Figure 12. In each trial, the robot started out approximately 1.5m perpendicular to the door and less than 0.5m to the right or left of the door handle in the direction parallel to the door. The starting orientation of the robot was either approximately perpendicular to the door or facing towards the door handle (a variation of around 30°). We clicked on a point in the laser scan to select a location in the world (as explained in Section III). The robot's task was to navigate to the door handle, twist it and push open the door.

Table I shows the performance of the individual behaviors on 8 different doors. We tested each component behavior four times and the preceding behaviors were successfully executed prior to each test. For these experiments, 'Servo to handle' was deemed successful if the robot servoed to the door such that the door handle was within the workspace of the manipulator. 'Segment and touch handle' required the robot to segment the door handle and make contact with it. 'Twist and push' was successful if the robot successfully twisted the door handle and pushed open the door. 'Determine if locked' required the robot to correctly report whether the door was locked or not. Failure to open an unlocked door also results in a failure of 'Determine if locked' because the robot incorrectly reports the door to be locked.

Fig. 12. The door handles of the 8 doors used in the experiments. For the results we report, the handles are numbered 1 to 8 from left to right.

TABLE I

PERFORMANCE OF THE COMPONENT BEHAVIORS ON 8 DIFFERENT DOORS.

Door #	Door state	Servo to handle	Segment and touch handle	Twist and push	Determine if locked
1	Unlocked	4/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
2	Unlocked	3/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
3	Unlocked	4/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
4	Unlocked	4/4	4/4	3/4	3/4
	Locked	1/1	1/1		1/1
5	Unlocked	4/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
6	Unlocked	4/4	3/4	4/4	4/4
	Locked	1/1	1/1		1/1
7	Unlocked	3/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
8	Unlocked	4/4	4/4	4/4	4/4
	Locked	1/1	1/1		1/1
Success Rate		38/40	39/40	31/32	39/40
Success %		95.0%	97.5%	96.9%	97.5%

Our implementation of manipulation with compliant trajectories performed well in these tests. The robot successfully twisted the door handle in 31 out of 32 trials (96.9%). The complete system including perception and navigation succeeded with unlocked doors in 28 out of 32 trials (87.5%) and locked doors in 8 out of 8 trials (100%). There were two failures in 'Servo to handle', one failure in segmenting the handle, and one failure in twisting the handle.

Our method of servoing to the door handle relies on detecting the handle in every planar scan from the laser range finder. Thin door handles, small errors in the estimated height of the handle, and motion of the sensor while the robot moves can result in the robot losing track of the door handle. We believe that more sophisticated methods for servoing and reliable odometry would reduce the chance of the robot failing to navigate to the door handle.

The segmentation failure was due to incorrect estimation of the orientation of the door from the Hough transform (Section IV-A). The robot failed to twist the door handle once when the end effector slipped towards the axis of rotation of the handle, and the twisting force exceeded its threshold even though the door was unlocked.

VIII. DISCUSSION AND CONCLUSION

Within this paper we have presented a behavior-based system for door opening with empirical success over many doors. Currently, this system operates door handles and pushes open doors. For many applications the robot would also need to pull open doors, operate door knobs, traverse doorways, and operate doors with door closers. We believe that many aspects of our approach will generalize to these tasks, including the structure of the behavior system, the segmentation-based

perceptual methods, and the use of compliant trajectories for manipulation.

We have also shown that simple compliant trajectories using equilibrium point control can robustly operate a door handle, push open a door, and reach out to find a door. This suggests that compliant trajectories may be a generally useful approach to manipulation. Whether programming a robot by hand or having a robot learn on its own, our results indicate that searching over a space of simple compliant trajectories can be profitable. In the case of door opening, linear compliant trajectories appear to be quite capable and robust, even though both the handle and the door have rotary joints. For door opening, the space of compliant trajectories that will achieve success appears to be large and the consequences of a failed trajectory can be mitigated. This naturally leads to the question of how to select a specific compliant trajectory from the many successful options. Our results suggest that searching for successful trajectories and stiffness settings that minimize the maximum interaction force would be advantageous. We hope to further explore these possibilities in future work.

IX. Acknowledgments

We thank Cressel Anderson for his work building the robot. We thank Aaron Edsinger and Jeff Weber for their prompt assistance. This research was supported in part by NSF grant IIS-0705130 and Willow Garage.

References

[1] K. Nagatani and S. Yuta, "An experiment on opening-door-behavior by an autonomous mobile robot with a manipulator," *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pp. 45–50, 1995.

[2] G. Niemeyer and J. Slotine, "A simple strategy for opening an unknown door," in *1997 IEEE International Conference on Robotics and Automation, 1997. Proceedings.*, vol. 2, 1997.

[3] A. Petrovskaya and A. Ng, "Probabilistic Mobile Manipulation in Dynamic Environments, with Application to Opening Doors," *International joint conference on artificial intelligence (IJCAI07), Hyderabad*, 2007.

[4] C. Rhee, W. Chung, M. Kim, Y. Shim, and H. Lee, "Door opening control using the multi-fingered robotic hand for the indoor service robot," *Proceedings. ICRA'04. 2004 IEEE International Conference on Robotics and Automation*, 2004.

[5] E. Klingbeil, A. Saxena, and A. Y. Ng, "Learning to open new doors," in *RSS Workshop on Robot Manipulation: Intelligence in Human Environments*, 2008.

[6] M. Quigley, S. Batra, S. Gould, E. Klingbeil, Q. Le, A. Wellman, and A. Y. Ng, "High-accuracy 3d sensing for mobile manipulation: Improving object detection and door opening," in *Proceedings of IEEE International Conference on Robotics and Automation*, 2009.

[7] A. Schmid, N. Gorges, D. Göger, and H. Wörn, "Opening a door with a humanoid robot using multi-sensory tactile feedback," *International Conference on Robotics and Automation*, 2008.

[8] A. Petrovskya, O. Khatib, S. Thrun, and A. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," *Proc. of ICRA*, 2006.

[9] R. Diankov, S. Srinivasa, D. Ferguson, and J. Kuffner, "Manipulation Planning with Caging Grasps," in *IEEE International Conference on Humanoid Robots*, 2008.

[10] L. Petersson, D. Austin, and D. Kragic, "High-level control of a mobile manipulator for door opening," *International Conference on Intelligent Robots and Systems*, 2000.

[11] M. Prats, S. Wieland, T. Asfour, A. del Pobil, and R. Dillmann, "Compliant interaction in household environments by the Armar-III humanoid robot," in *8th IEEE-RAS International Conference on Humanoid Robots, 2008. Humanoids 2008*, 2008, pp. 475–480.

[12] R. Brooks, L. Aryananda, A. Edsinger, P. Fitzpatrick, C. Kemp, U. O'Reilly, E. Torres-Jara, P. Varshavskaya, and J. Weber, "Sensing and Manipulating Built-for-Human Environments," *International Journal of Humanoid Robotics*, vol. 1, no. 1, pp. 1–28, 2004.

[13] D. Kragic and H. Christensen, "A Framework for Visual Servoing Tasks," *Intelligent Autonomous Systems*, 2000.

[14] C. Eberst, M. Andersson, and H. Christensen, "Vision-based door-traversal for autonomous mobile robots," in *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2000.

[15] M. Prats, P. Sanz, A. del Pobil, E. Martínez, and R. Marín, "Towards multipurpose autonomous manipulation with the UJI service robot," *Robotica*, vol. 25, no. 02, pp. 245–256, 2007.

[16] A. Jain and C. C. Kemp, "Behaviors for robust door opening and doorway traversal with a force-sensing mobile manipulator," in *RSS Workshop on Robot Manipulation: Intelligence in Human Environments*, 2008.

[17] E. Bizzi, N. Hogan, F. Mussa-Ivaldi, and S. Giszter, "Does the nervous system use equilibrium-point control to guide single and multiple joint movements?" *Behavioral and brain sciences(Print)*, vol. 15, no. 4, pp. 603–613, 1992.

[18] M. Hinder and T. Milner, "The case for an internal dynamics model versus equilibrium point control in human movement," *The Journal of Physiology*, vol. 549, no. 3, pp. 953–963, 2003.

[19] X. Gu and D. Ballard, "An equilibrium point based model unifying movement control in humanoids," in *Robotics: Science and Systems*, 2006.

[20] M. Williamson, "Postural primitives: Interactive behavior for a humanoid robot arm," in *From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. MIT Press, 1996, p. 124.

[21] S. Migliore, "The Role of Passive Joint Stiffness and Active Knee Control in Robotic Leg Swinging: Applications to Dynamic Walking," Ph.D. dissertation, Georgia Institute of Technology, 2009.

[22] M. Williamson, "Robot arm control exploiting natural dynamics," Ph.D. dissertation, Massachusetts Institute of Technology, 1999.

[23] Y. Mukaibo, S. Park, and T. Maeno, "Equilibrium Point Control of a Robot Arm with a Double Actuator Joint," *International Simposium on Robotics and Automation*, 2004.

[24] D. Clapa, E. Croft, and A. Hodgson, "Equilibrium point control of a 2-DOF manipulator," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, p. 134, 2006.

[25] L. Sentis and O. Khatib, "Synthesis of whole-body behaviors through hierarchical control of behavioral primitives," *International Journal of Humanoid Robotics*, 2005.

[26] C. Ott, B. Baeuml, C. Borst, and G. Hirzinger, "Autonomous opening of a door with a mobile manipulator: A case study," *IFAC Symposium on Intelligent Autonomous Vehicles*, 2007.

[27] G. A. Pratt and M. M. Williamson, "Series elastic actuators," vol. 1, jul 1995, pp. 399–406.

[28] Y. S. Choi, C. D. Anderson, J. D. Glass, and C. C. Kemp, "Laser pointers and a touch screen: Intuitive interfaces to an autonomous mobile robot for the motor impaired," in *ACM SIGACCESS Conference on Computers and Accessibility*, 2008.

[29] R. Rusu, Z. Marton, N. Blodow, M. Dolha, and M. Beetz, "Towards 3D Point cloud based object maps for household environments," *Robotics and Autonomous Systems*, vol. 56, no. 11, pp. 927–941, 2008.

[30] R. B. Rusu, W. Meeussen, S. Chitta, and M. Beetz, "Laser-based perception for door and handle identification," in *Proceedings of International Conference on Advanced Robotics*, 2009.

[31] . [Online]. Available: http://www.diracdelta.co.uk/science/source/f/r/friction/source.htmlhttp://www.easycalculation.com/physics/classical-physics/friction-table.php