Experiments in Impulsive Manipulation

Wesley H. Huang
Institute for Complex Engineered Systems
Carnegie Mellon University

Matthew T. Mason
Robotics Institute
Carnegie Mellon University

Abstract

In this paper, we present the results of our experimental effort in one form of impulsive manipulation: tapping. Our previous work studied the mechanics of tapping a planar object which then slides on a support surface, coming to rest due to friction. This work addresses the practical issues in creating a system which uses this mode of manipulation. We begin with the design of tapping devices—end effectors designed to deliver an impulse to an object, and report some of the issues we have found to be important in their design. Our next step was to perform single-tap experiments in order to fit and evaluate the models of impact and sliding. These experiments have shown that objects rotate less than predicted; we have found that the addition of a scaling factor for the torque due to friction enables the models to predict object motion reasonably well. In order to do positioning experiments, we developed a number of planning methods (or feedback control strategies) to compensate for errors in modeling, parameters, and actuation. These planning methods were successfully used to demonstrate a positioning task. We also have experimentally demonstrated that tapping can be used to position an object more precisely than the manipulator can position the tapping device. We offer some sensitivity analysis in support of this result.

1 Introduction

Impulsive manipulation is the use of impulsive forces for manipulation; it consists of two phases: a strike to an object, followed by free motion of that object, subject to forces and constraints in its environment. Such manipulation can be very fast, used for precision positioning, and requires relatively simple manipulators. By studying this mode of manipulation, we can increase the repertoire of robotic manipulators.

In our work, we have focused on one particular form of impulsive manipulation: tapping a planar object which then slides on a support surface, coming to rest due to frictional forces. There are many ways that tapping could be used to position objects on a support surface, e.g. using a robotic arm to position a tapping device relative to the object or by using a "positioning cell" with a number of tapping devices at fixed positions around its perimeter. In this work, we have explored the first approach.

Our previous work (Huang *et al.* [4]) analyzed the mechanics of tapping, showing how to strike an object in order to achieve a given displacement (translation and rotation). This analysis made a number of assumptions about the world: objects are planar and have a uniform support distribution; all relevant object properties are known; only Coulomb friction acts on the object after impact; and the impact follows the classical model of two dimensional impact with friction. It was our intuition that a real experimental setup would be close enough to these models to allow feedback control to compensate for any modeling error and for errors in object parameters. Huang *et al.* [4] assumed that objects were axisymmetric (rotationally symmetric); in these experiments, we have used circular objects and objects that, while not axisymmetric, are nearly so.

Our experimental effort began with the design and construction of a device to deliver a controlled impact. We present the design of that device and some of the issues relevant to its design.

Our first experiments were single-tap experiments which were designed to test how well the impact and friction models predicted object motion. A high speed measurement system was used in order to determine the velocity profiles of the object. These experiments showed that while the linear translation matched the models well, the rotation did not. Objects rotated less than predicted, so we introduced a scaling factor for the torque due to friction. With this addition, the models can predict object motion reasonably well.

Once we had a reasonable model of the mechanics of tapping an object, we started to consider positioning tasks. Tapping, although a discrete rather than continuous process, is nonholonomic in the sense that a tap can only produce transitions in state space in a certain range of directions. We must in general plan a sequence of taps to accomplish a positioning task. Huang [3] showed that any configuration can be reached within two taps. However, for practical purposes, the minimum-tap plan may not be the best plan. Errors in modeling, object parameters, and actuation will produce variations in the resulting displacement for a tap; we have formulated a number of planning methods which essentially serve as feedback control schemes to deal with these errors.

The final result of this paper is the experimental confirmation that tapping can be used to position an object more precisely than the manipulator can position a tapping device, i.e. more precisely than the manipulator could position the object using pick-and-place manipulation.

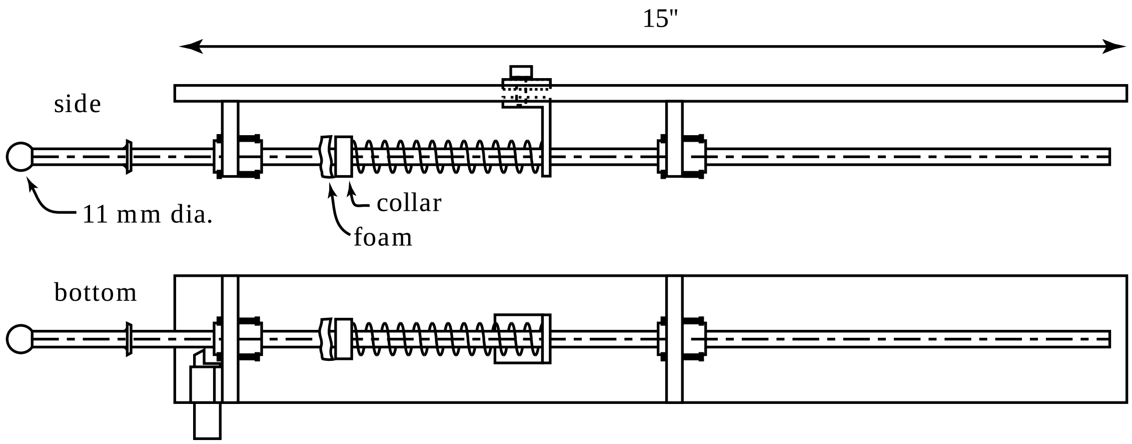

Figure 1: The tapping device used for these experiments features linear ball bearings, a manually adjustable backstop, and a ball bearing at the tapper tip.

1.1 Related work

The first instance of impulsive manipulation in the robotics literature was work done by Higuchi [2] who used an electromagnetic coil to deliver an impulsive force for linear micropositioning. More recently, Yamagata and Higuchi [7] have developed a piezoelectric device to deliver impulsive forces for micropositioning of optical sensors.

Zhu *et al.* [8] have performed experiments with sliding objects in conjunction with what they call “releasing manipulation.” Instead of striking objects to produce initial velocities, the objects are smoothly accelerated with a manipulator and then released. Their experimental results, like ours, show significant variation in the angle rotated and distance translated.

Our work builds upon that of Voyenli and Eriksen [6] who analyzed the physics of sliding rotating disks and rings and upon that of Goyal *et al.* [1] who have studied the dynamic behavior of planar sliders. Our previous work (Huang *et al.* [4]) presented a solution to the inverse mechanics for a sliding rotating axisymmetric object; this analysis is the basis for the planning and analysis in this paper.

This work draws from a number of other areas, including friction, impact, nonprehensile manipulation, minimalism, parts feeding, and robotic juggling.

The results presented in this paper appear in greater detail in Huang [3].

2 Design of a tapping device

A tapping device designed for manipulation should be able to deliver an adjustable amount of impulse to the object (both in the normal and tangential directions), should be as consistent and repeatable as possible, and should produce a clean impact (i.e. should not push the object for a distance). The approach we have chosen is to use a striker which collides with the object.

Figure 1 shows the design of our most recent tapping device, the one used in the experiments described in this paper. This device is designed to be mounted as an end effector on a robotic manipulator. The striker is a steel shaft which is drawn back to compress a spring against an adjustable backstop. A catch engages a ring around the front end of the striker in order to hold and release the striker. A large ball bearing at the tip allows the angle of incidence of the striker to be changed without affecting the contact normal and contact point.

Through designing this device and its predecessors, we have found a number of important issues in the design of such devices. These include the following:

- In order for the striker not to push the object for some distance, there should be no forces acting on the striker at the time of impact. This device is calibrated so that the spring stops acting upon the striker before the point of impact, leaving the striker free to rebound. This sort of design, however, makes designing a completely automatic device more involved.

- Release mechanisms can cause significant variations in a tapping device. These mechanisms generally serve to remove some kinematic constraint to release stored mechanical potential energy. Ideally, the release should be quick and affect the state of the striker as little as possible. Some approaches to this problem are to use sharp edges sliding against each other, to minimize friction between parts, and to minimize clearances in the bearings or bushings of the striker and of the release mechanism.

- A tapping device should be designed to be easily calibrated. Careful construction will allow easier kinematic calibration (determining the transform from the wrist to the striker tip). Calibration between the striker parameters (the amount of draw in this case) and the impulse delivered by the tapping device is difficult unless the physical phenomena driving the device is well

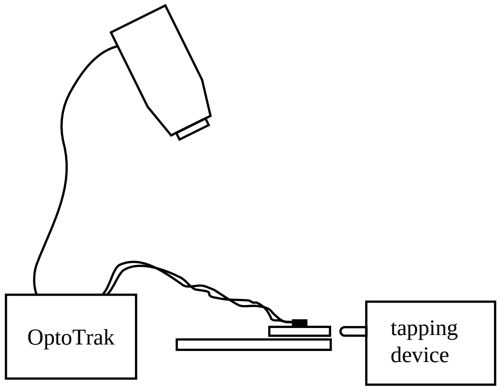

Figure 2: Setup for the single-tap experiments. A fixed tapping device is used to tap an object which is instrumented for measurement by a high speed tracking system.

known (a linear spring or gravitational potential energy) or unless careful measurements of the device can be made.

3 Single-tap experiments

The single-tap experiments consisted of many runs, each of which used different tapper settings and angles of incidence. The object was instrumented with two infrared LEDs so that it could be tracked by an OptoTrak system which provided us with position and orientation measurements at 500 Hz.

We performed a number of experiments with different objects and surfaces as well as tapping devices. In this paper, we describe one of these experiments in detail and summarize the others.

3.1 Aluminum disk experiment

The object used in this experiment was an aluminum disk with a radius of 48 mm, a height of 12 mm and a weight of 245.5 g. This object slides on a formica surface; the coefficient of friction between the object and the surface was determined to be 0.17. By placing one of the LEDs on the striker (to measure both striker and object velocities), the coefficient of restitution was determined to be 0.86.

There were four runs in this experiment, each of which consisted of 10 trials. The displacements and initial velocities for all trials are shown in Figure 3. Since the position data can be used to calculate velocity profiles and initial object velocities, the impact model can be evaluated independently from the sliding model.

3.1.1 Evaluating the impact model

The impact model can be evaluated by comparing the measured initial velocities with the calculated initial velocities (calculated from the measured nominal striker velocity and the angle of incidence). The translational velocities matched well, but the measured rotational velocities were much larger than predicted. By fitting two of the constants in the impact model (which consist of various mass and geometric terms), we obtained a usable model of impact which predicted average initial translation velocities to within 5 to 15 percent, initial rotational velocities to within 30 percent. There is considerable variation within each run — the standard deviation is as much as 15 percent in initial rotational velocity and 5 percent in initial translational velocity.

3.1.2 Evaluating the sliding model

Ideally, we would evaluate the sliding model by how well the predicted and measured velocity profiles matched, but the measured velocity profiles are noisy enough that it would be difficult to do a meaningful comparison. Instead, we have turned to a comparison of the measured and predicted initial velocities and displacements. The measured initial velocities are compared to the initial velocities calculated from the measured displacement, and the measured displacements are compared to the displacements calculated from the measured initial velocities.

We found that the translations and initial translational velocities matched well, but the measured rotations are much less than calculated, and the initial rotational velocities were much greater than calculated. This problem was addressed by using a scaling factor to increase the torque due to friction. The scaling factor was determined to be 3.64 by minimizing the sum of the squares of the average difference between measured and calculated initial velocities over each run.

With this scaling factor, calculated translations were within 1.2% to 10.8%, rotations to within 3.2% to 12.3% (on average). Figure 4 shows an example of measured and predicted velocity profiles for one trial.

3.2 Other experiments

Over the course of this work, we performed a number of other experiments. One of these experiments used a plexiglas disk sliding on an aluminum surface; this experiment did not require any scaling factor for the frictional torque. Another used an aluminum square sliding on an aluminum surface. This object was modeled well by the axisymmetric mechanics using a torque scaling factor of 2.17.

3.3 Discussion

The most important result of these experiments is that real objects can be reasonably modeled with the axisymmetric mechanics despite deviations from our model assumptions — primarily that objects have a uniform support distributions and do not have three dimensional extent. There experiments also showed that there are noncircular objects that are also reasonably modeled by the axisymmetric mechanics.

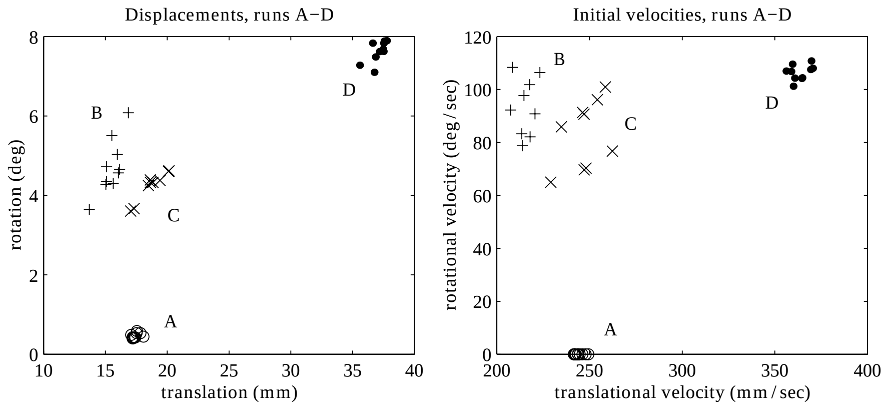

Run	A	B
Angle of incidence (deg)	0	14.5
Striker velocity (mm/sec)	419	419
Average distance (mm)	17.4	15.5
Average rotation (deg)	0.5	4.7

Run	C	D
Angle of incidence (deg)	7	7
Striker velocity (mm/sec)	431	632
Average distance (mm)	18.7	37.1
Average rotation (deg)	4.2	7.6

Figure 3: Displacements and initial velocities for the aluminum disk experiment.

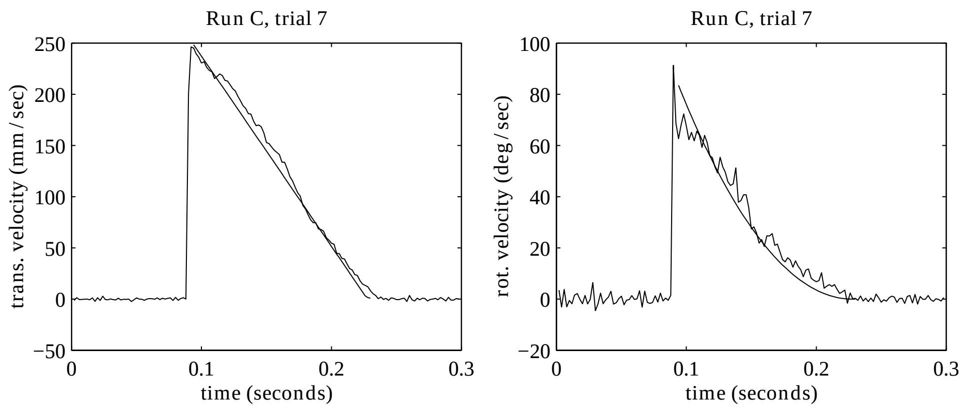

Figure 4: Measured and predicted velocity profiles for one run of the aluminum disk experiment. Predicted profiles are based on measured displacement and use the torque scaling factor.

The necessity of a scaling factor for the frictional torque is likely due, in part, to deviations from the model assumptions. However, the magnitude of the scaling factors leads us to believe that there may be some other mechanism that results in increased frictional torque.

These single-tap experiments were done with reasonable care. While they could be done with more precisely machined surfaces, better tapping devices, and so on, it is not clear that we would want to do this. In order for tapping to be easily used for real world applications, we must deal with typical situations — which generally involve three dimensional parts, nonideal materials, low precision tapping devices, etc. We deal with these deviations with the planning methods (feedback control) described in the next section.

4 Planning methods for positioning tasks

A single tap cannot position an object arbitrarily in the plane. It turns out that two suffice, but because of the large variations in single taps, two-tap plans seldom succeed. In this section, we describe planning methods that deal with the error of a tap in different ways. These planning methods serve as feedback control laws because we replan after every tap from the object's current configuration. We are able to do this because of the discrete nature of impulsive manipulation.

We consider positioning tasks using the axisymmetric mechanics. Under these mechanics, objects will always slide in a straight line, so the task of positioning an object in the plane is then a two dimensional problem: one in translation (along the line between the start and goal configuration) and one in rotation.

Analysis of the mechanics of tapping reveals that a tap can position an axisymmetric object at any state within a cone in state space. This cone is determined by the object mass, moment of inertia, geometry, and material properties. For circular axisymmetric objects, this cone is symmetric about the x (translation) axis. It is fairly easy to see that such an object can reach any point in state space in at most two taps (in opposite directions) — the minkowski sum of two cones pointing in opposite directions covers the entire configuration space.

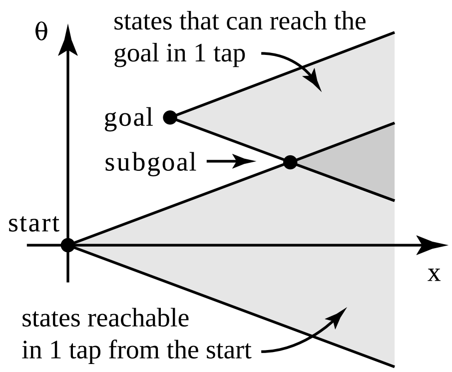

Figure 5: The exact two-tap planning method chooses the closest feasible subgoal.

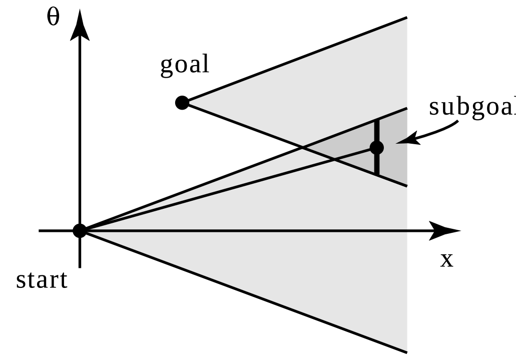

Figure 6: The conservative two-tap planning method chooses a subgoal such that the error ellipse fits entirely within the space of feasible subgoals.

There is a two dimensional set of points that could be used as an intermediate state to reach the goal. By choosing different intermediate points, the planning methods make different tradeoffs between the expected number of taps and its standard deviation.

We have assumed that the size of the error ellipse for a tap is linearly related to the size of the tap. We have found that the error in rotation is generally larger (and more important) than the error in translation, so for simplicity, we let our error ellipses degenerate into lines.

In these experiments, we use the same aluminum disk and formica surface described in Section 3.1. The positioning task we consider here is to move the object 10 mm and rotate it 10 degrees; this is about three times as much rotation as can be achieved with a single tap. The acceptable positioning accuracy is 1 mm and 1 degree, a limit primarily due to the tapping device — it cannot consistently produce translations less than 1.7 mm.

For these experiments, a SCARA robot is used to position the tapping device with respect to the object. An overhead camera measures the position and orientation of the object to an approximate accuracy of 0.7 mm and 0.8 degrees.

4.1 Exact "two-tap" plans

The exact two-tap planning method chooses the closest feasible intermediate state as illustrated in Figure 5. The problem with this planning method is that errors may cause the object, after the first tap, to be in a state unable to reach the goal with a second tap. This planning method is implemented as follows: if the goal can be reached in one tap, then we execute that tap, otherwise, we execute the first of an exact two-tap plan.

Of all the planning methods, this was the most robust, consistently taking 5 or fewer taps to position the object.

4.2 Conservative "two-tap" plans

The conservative two-tap planning method chooses an intermediate state such that the error ellipse from the first tap fits within the space of feasible subgoals, as illustrated in Figure 6. Although the object is more likely to be in a position to reach the goal after the first tap, the effect of this strategy is to cause larger excursions as compared with the exact two-tap plans. The second tap will therefore be larger and have a larger error ellipse about the goal. Like for the exact two-tap plans, we use the following rule: if the object can reach the goal in a single tap, then we execute that tap; otherwise, we execute the first of a conservative two-tap plan to reach the goal.

This planning method did slightly worse than the exact two-tap plans; most runs took about 5 taps.

4.3 Multi-tap plans

The simplest way to guarantee positioning the object within a given tolerance of the goal is to limit the size of the last tap. For a given positioning accuracy, there is a small "cone" of states that can reach the goal with a sufficiently small tap. There is a larger cone that can reach this first cone such that the error ellipse from that tap lies entirely within the first cone. Through this back-chaining, we can construct a sequence of cones extending away from the goal. If the object starts at a configuration not inside one of these cones, then a tap must be planned so that the error ellipse from that tap lies entirely within a cone. We execute this planning method by first checking whether the object is within the sequence of cones leading to the goal. If so, we execute a tap to take it to the next cone. If not, we execute a tap to take the object into the nearest cone.

For the given task of these experiments, the nominal plan consisted of 5 taps, and in almost every run, it succeeded in 5 taps.

4.4 Discussion

Although the exact two tap plans generally require replanning along the way, they have the potential to take exactly two taps to reach the goal. The multi-tap plans, on the other hand, do not attempt to take less than five taps to reach the goal for the given positioning task and there is greater certainty that this plan will succeed in the planned number of taps.

5 High precision positioning

Early in the course of this work we conjectured that by tapping, an object could be positioned more precisely than the robot could position the tapping device. To test this idea, we performed experiments in which the position and the orientation of the tapping device were rounded with respect to the global frame. The experimental setup and the positioning task were the same as in the previous section (moving the same disk 10 mm and 10 degrees to an accuracy of 1 mm and 1 degree). First, the positions were rounded to the nearest 10 mm and orientations to the nearest 5 degrees. Then, positions were rounded to the nearest 15 mm and the nearest 10 degrees. In all trials, the goal was successfully reached; a typical trial took 5–6 taps.

5.1 Sensitivity analysis

In order to explain this phenomenon, we examined the Jacobian relating changes in the striker position and orientation to changes in the resulting object configuration. This Jacobian combines the effect of the manipulator configuration on the striker contact point with the mechanics of impact and of a sliding object. Since there is no analytic form for the mechanics of a sliding rotating object, that part of the Jacobian must be calculated numerically. For a 10 mm, 3 degree tap, we have:

$$\begin{bmatrix} \partial x \\ \partial y \\ \partial \theta \end{bmatrix} = \begin{bmatrix} -0.0689 & 0.135 & 0.126 \\ -4.7 \cdot 10^{-6} & -0.000026 & 0.176 \\ -0.0088 & 0.0251 & 0.0234 \end{bmatrix} \begin{bmatrix} \partial i \\ \partial j \\ \partial \beta \end{bmatrix}$$

where (i, j) and β are the error in striker position and orientation. Multiplying this Jacobian by the striker error vector $[7.5\ 7.5\ 5]^T$ (which is the average error in the 15 mm, 10 degree rounding experiment) yields $[1.1\ 0.9\ 0.2]^T$, a smaller error than the striker error. For the object and materials of this experiment, there is a broad range of taps for which the norm of the Jacobian is less than one (implying that the error in the resulting object configuration will always be less than the error in striker configuration). This range covers the entire range of displacements possible with the current tapping device.

6 Conclusions

We have demonstrated a practical implementation of manipulation via tapping and have found several interesting results.

First, we have developed a device to deliver a controlled, repeatable impact and have identified a number of issues that are important in designing such devices. Secondly, we have found as a result of the single-tap experiments that sliding objects do not rotate as much as our models predict. While this is undoubtedly in part due to deviations from our model assumptions, there may be some other mechanism giving rise to a larger frictional torque. We have demonstrated several planning methods (or feedback control laws) to compensate for errors in modeling, parameter estimation, and actuation. Finally, we have shown, both experimentally and analytically, that it is possible to position objects more precisely by tapping than the robot can position the tapping device.

Acknowledgements

We would like to thank Jason Powell for building and helping design the tapping devices and the Medical Robotics and Computer Assisted Surgery lab at Carnegie Mellon for use of their OptoTrak system.

This work was supported in part by NASA through the Graduate Student Researchers Program and through grant NAGW 1175 and by the NSF under grants IRI–9318496 and IRI-9114208.

References

[1] S. Goyal, A. Ruina, and J. Papadopoulos. Planar sliding with dry friction. Part 2. Dynamics of motion. *Wear*, 143:331–352, 1991.

[2] T. Higuchi. Application of electromagnetic impulsive force to precise positioning tools in robot systems. In O. Faugeras and G. Giralt, editors, *International Symposium on Robotics Research: The Third International Symposium*. MIT Press, 1986.

[3] W. H. Huang. *Impulsive Manipulation*. PhD thesis, Carnegie Mellon University, Aug. 1997. Available as Carnegie Mellon Robotics Institute technical report CMU–RI–TR–97–29.

[4] W. H. Huang, E. P. Krotkov, and M. T. Mason. Impulsive manipulation. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 120–5, 1995.

[5] W. H. Huang and M. T. Mason. Mechanics for vibratory manipulation. In *IEEE International Conference on Robotics and Automation*, 1997.

[6] K. Voyenli and E. Eriksen. On the motion of an ice hockey puck. *American Journal of Physics*, 53(12):1149–53, December 1985.

[7] Y. Yamagata and T. Higuchi. A micropositioning device for precision automatic assembly using impact force of piezoelectric elements. In *IEEE International Conference on Robotics and Automation*, 1995.

[8] C. Zhu, Y. Aiyama, T. Chawanya, and T. Arai. Releasing manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, page 1, Osaka, Japan, Nov. 1996.