解 説

Physics-Based Manipulation in Human Environments

Mehmet R. Dogar* Siddhartha S. Srinivasa*
*The Robotics Institute, School of Computer Science, Carnegie Mellon University

1. Introduction

There are striking differences between the way humans and current robots manipulate objects. One difference is in the variety of actions used. The list of actions that we humans use to push, pull, throw, tumble, and play with the objects around us is nearly endless. Robots, however, manipulate objects almost exclusively through pick-and-place actions. As a consequence, robots are also limited in the variety of tasks that they can perform.

Robots are limited to pick-and-place actions because they use motion planners which are agnostic to physics. Pick-and-place actions do not require physics models to predict how the manipulated object moves: it is rigidly attached to the hand. However, complex manipulation skills require complex physics-based models to predict how the world behaves. For example, to push a heavy piece of furniture out of the way, a robot needs a physics model that predicts how the furniture will move.

At the Personal Robotics Lab at Carnegie Mellon University we investigate methods to use realistic physics models in manipulation planning. We develop planners that enable robots to physically interact with the environment in order to perform useful tasks. Some of these tasks are impossible to perform using only pick-and-place actions.

In this paper we present our first steps in building a physics-based manipulation planner. We use a quasi-static analysis to predict how objects move when they are pushed. We integrate these predictions into a manipulation planner which produces pushing actions as well as pick-and-place actions. We demonstrate the effectiveness of our approach in three domains:

Reconfiguring clutter In cluttered environments robots need to move obstacles out of the way in order to reach other objects. These obstacles may not be movable by pick-and-place actions if they are large or heavy. Our planner uses physics-based pushing actions to move such objects.

Manipulation through clutter When we humans reach into a cluttered fridge shelf to grasp a milk jug or into a gym bag to pull a towel out, we frequently contact multiple objects other than the one we want to grasp. Robots are the opposite. Motion planners avoid contacting any object other than the goal at all costs, since the contacted objects' motion cannot be predicted. We show that physics-based manipulation planners are free of this constraint. Our planner enables a robot to reach for and grasp the target while simultaneously contacting and moving aside obstacles in order to clear a desired path.

Manipulation under uncertainty Object pose uncertainty is a major source of failure during robotic manipulation. Our planner can use the mechanics of pushing to reduce uncertainty, resulting in robust manipulation even under high uncertainty.

Clutter and uncertainty are two main problems for robotic manipulation in human environments. We are excited to see that physics-based planning has the potential to improve robot manipulation capabilities in the face of both issues.

A future goal for us is to develop physics-based planners which extend beyond pushing and can accommodate actions such as rolling, throwing, and toppling. We need to address two major issues:

Planning time Physics simulations are slow. Manipulation planners need to consider alternative cases, running simulations many times, which results in long planning times. We propose a solution to this problem based on pre-computing and caching of the physical interactions between the robot manipulator and an object.

Robustness Physics simulations are inaccurate.

原稿受付 2013 年 2 月 8 日
キーワード：Physics-based Manipulation, Non-prehensile Actions, Pushing
*Pittsburgh, PA 15260, USA

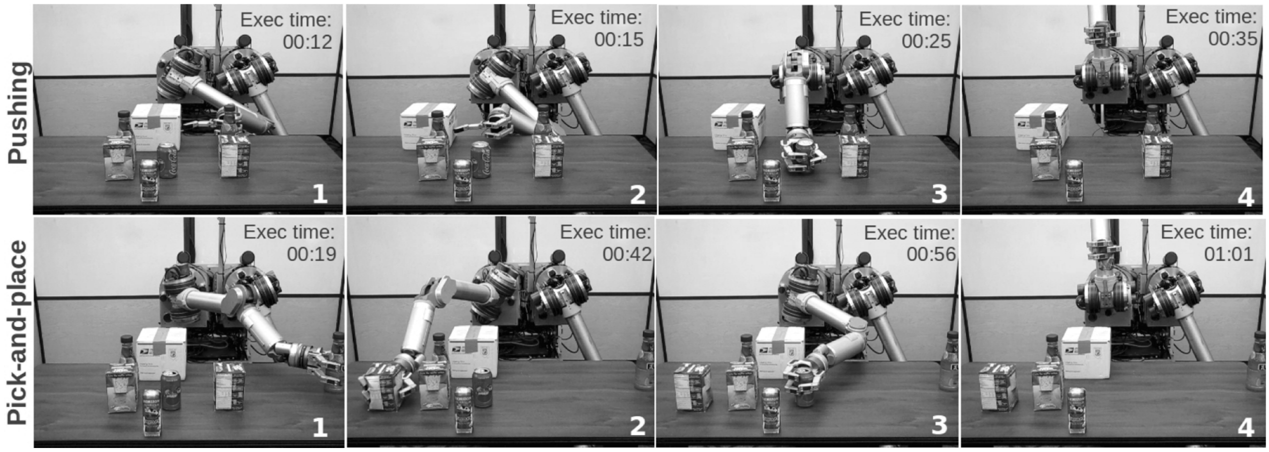

Fig. 1 Two reconfiguration plans where the goal is to reach to the red can hidden behind a large box that cannot be grasped. (Top) Reconfiguration planning with pushing actions. (Bottom) Reconfiguration planning with pick-and-place actions. Image taken from [1]

Plans based on inaccurate predictions run the risk of failure. We propose two solutions and provide examples in the pushing domain: (i) conservative planning, and (ii) using sensor feedback.

Robots can perform remarkable manipulation tasks in human environments using physics-based planning. The pushing based approach that we take in this paper is an important step in this direction.

2. Manipulation with Pushing Actions

We present a framework which produces manipulation actions including pushing, as well as pick-and-place. In order to verify a plan, the framework predicts how a certain action will move an object.

2.1 Predicting object motion

Our planner uses a simulation to predict how objects move when they are pushed. We developed this simulation with the theoretical background provided by a number of studies which analyze the *quasi-static* interactions between a pusher and a pushed object. Mason [2] presents such an analysis and proposes a method, the *voting theorem*, to find the sense of rotation of a pushed object. Other studies build on this and analyze the controllability, planning, and uncertainty-reducing properties of pushing [3]~[7].

Our simulation uses the *limit surface* [8] to relate the generalized forces applied on an object to the resulting generalized velocity. We approximate the limit surface with a three-dimensional ellipsoid. Howe and Cutkosky [9] present the conditions under which this approximation can be done. Our planner uses three-dimensional models of objects and this simulator to predict object motion.

2.2 Reconfiguring Clutter

A robot can reconfigure clutter efficiently using pushing actions. We illustrate this with an example in **Fig. 1**. The robot's goal is to reach the red can hidden behind a box which is too large to be grasped. Our planner pushes the large box to the side and reaches the goal object (Fig. 1-Top). However, if the robot plans for the same task using only pick-and-place actions, it needs to pick up two other objects and avoid the large box (Fig. 1-Bottom). This results in longer execution and planning times. One can also construct scenes where an ungraspable object *must* be moved to reach the goal, in which case the pick-and-place reconfiguration approach will fail completely.

A reconfiguration planner identifies the objects to move, the order to move them, and where to move them. The general problem is NP-Hard [10]. Planners in the literature produce feasible results by performing a search over the order of objects [11] [12]. Stilman et al. [13] [14] perform this search using the back-projection of robot actions, starting from the final action of reaching to the goal. Our planner [1] uses back-projections of pushing actions as well as pick-and-place actions.

Our planner searches the space of different pushing actions by discretizing the different directions to push an object. Along each direction the hand can also be offset laterally with respect to the object pose, and different hand preshapes can be used.

2.3 Manipulation through Clutter

In cluttered spaces, we humans can simultaneously contact multiple objects while manipulating a goal object. For example, when we reach towards the back of a fridge shelf our hands may contact to the objects at the front edge; or, when we reach into a gym bag to pull a towel out our hands contact many other objects as well as the towel.

Our physics-based planner can afford such simultaneous interactions with cluttering objects [15]. This is different from the reconfiguration actions, since in this case the planner does *not* plan a separate action for each contacted object. Instead, given a single manipulator trajectory to reach the goal object, the planner predicts how each contacted object will move (**Fig. 2**-Left). Then the planner verifies that these objects' motions will not cause any problems, e.g. objects falling off the edge of the table, or the wrong object being grasped.

We present an example plan in Fig. 2-Right. In this figure, as the robot is reaching a can, the hand contacts and simultaneously pushes two other objects blocking the way. This is possible because the physical predictions verify the motions of the blocking objects.

Our robot can execute many manipulator trajectories which would be labeled as infeasible by a traditional motion planner based on collision checking. As a result, our robot is more successful in planning and executing manipulation tasks in cluttered environments.

2.4 Manipulation under Uncertainty

Robotic manipulation systems suffer from *uncertainty* in human environments. Consider the task of grasping an object. In such a task, the robot detects the object and estimates its pose. If there is significant uncertainty in the estimated pose of the object, the robot hand can miss it, or worse, collide with it in an uncontrolled way.

Fig. 2 Manipulation of objects through clutter. (Left) Given a manipulator trajectory our planner predicts how each contacted object will move, and verifies that the goal will be reached successfully. (Right) An example execution, where the robot pushes two blocking objects out of the way simultaneously as it is reaching a can. Images taken from [15]

Physics-based manipulation can address this problem by employing *uncertainty reducing actions* [17]. Our planner harnesses the mechanics of pushing to *funnel* an object into a stable grasp despite high uncertainty [16].

We present an example in **Fig. 3**-Left. In the figure, there is high initial uncertainty about the pose of the bottle. As the hand pushed forward, the uncertainty funnels into the hand, and the object can be grasped. We call this action *push-grasping*.

We define a push-grasp as a straight motion of the hand parallel to the pushing surface along a certain direction. The *pushing distance*, indicating the translation along the pushing direction, and the *aperture*, indicating the distance between the symmetrically shaped fingertips, are important parameters of the push-grasp. The larger these values, the larger the uncertainty that the hand can funnel in. However, such push-grasps may be more difficult to execute in cluttered environments. Our planner searches over different pushing directions, different hand apertures, and different pushing distances to find a successful push-grasp.

We present an example push-grasp in Fig. 3-Right. The robot sweeps a region over the table during which the box rolls into its hand, before closing the fingers. The large swept area ensures that the box is grasped even if its position is estimated with some error.

3. Challenges in Physics-Based Manipulation

As we extend our framework to use a wider variety of actions, there are challenges that we need to address. In this section we present two major challenges and discuss how we addressed them in the context of pushing.

3.1 Planning Time

The planner must predict the consequences of each action it considers. One way to do this is running a

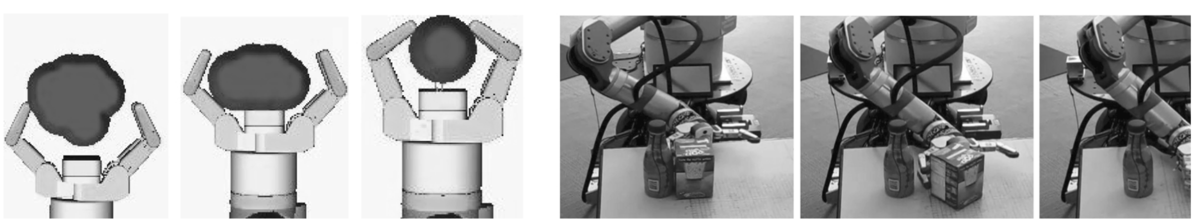

Fig. 3 Push-grasping. (Left) Using pushing to reduce object pose uncertainty. (Right) Execution of a push-grasp. Images taken from [1] and [16]

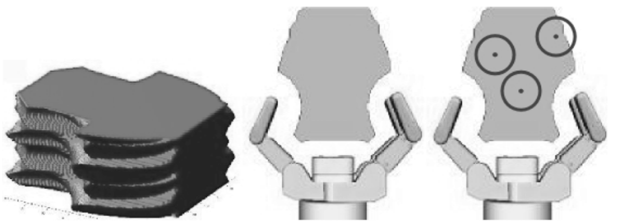

Fig. 4 The capture region. (Left) The capture region of a box shaped object. The horizontal dimension is the rotation of the object. The other two dimensions are the translation of the object. (Center) The capture region of a circularly symmetric object. (Right) If the object's center falls into the capture region, it will be grasped. The figure shows three such poses. Images taken from [16]

physics simulation during planning. Physics simulations are slow, however, and this approach may result in long planning times.

We address this issue by pre-computing and caching the physical interactions between the robot manipulator and objects. For example, given a push-grasp, an object's geometry, and its physical properties, we can compute the set of all object poses that result in a stable push-grasp. We call this set the *capture region*. We present two capture regions in **Fig. 4**-Left and Fig. 4-Center. We compute capture regions assuming that the object is located on a planar surface. Hence, the capture region is a set of points $(x, y, \theta) \in SE(2)$ (Fig. 4-Left). If the object is circularly symmetric, we can drop θ and represent the capture region in two dimensions (Fig. 4-Center).

During planning, we use the capture region to determine whether a push-grasp will succeed, instead of running a simulation: If the object is in the capture region of the push-grasp, then the push-grasp will succeed (Fig. 4-Right). If there is uncertainty associated with the object pose, our planner checks to see if all the uncertainty is included in the capture region. We can perform this check very fast. As a result our planner can find a successful push-grasp in a few seconds, even when there is significant object pose uncertainty and clutter.

We also pre-compute the actual trajectories objects follow when they are pushed in a certain way. We then use these trajectories during planning instead of running simulations.

This approach also has limitations. It is not possible to enumerate all possible cluttered scenes. Therefore, we limit our pre-computations to the interactions between the robot manipulator and an object. In a given scene we verify that these pre-computed structures are still valid, e.g. no object-object interaction occurs.

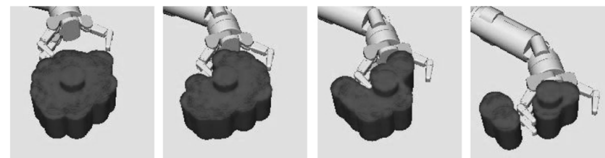

Fig. 5 Particle filtering while pushing a can. The actual location of the can is shown raised

3. 2 Robustness

The accuracy of the physics-based predictions determine the robustness of our manipulation plans: If an object moves in an unpredicted manner, the execution may fail. These inaccuracies can be due to two different reasons:

- Uncertainties in parameters. For example, the pressure distribution of an object affects the way it moves during pushing. If the robot does not have a good estimate of this parameter, its predictions will be inaccurate.
- Inaccuracies in the physics model. For example, the predictions of our quasi-static model will get inaccurate if significant dynamic forces arise during pushing.

We are investigating different methods to address the problem of inaccurate predictions. One approach we take is making conservative plans, i.e. plans that will work in spite of the uncertainties. We identify the different parameters that affect our pushing predictions. These include the pressure distribution of the object, and the coefficient of friction between the manipulator and the object. We then run our pushing simulations with a range of values of these parameters, where the range is determined by our uncertainty about that parameter. We accept an action if it achieves the goal for all of these simulations.

While the conservative planning approach addresses the uncertainties in parameters, it does not address the second source of inaccuracy, namely, the inaccuracies in our physics models. Therefore, we are currently investigating methods of using sensor feedback that can account for all kinds of inaccuracies.

We are modeling manipulation as a stochastic process where the forward models are provided by noisy physics-based predictions. We then use sensor feedback to filter the uncertainty induced to the system by the physics-based actions. We present a simulated example in **Fig. 5** where we use particle filtering [18] to track the pose of an object as it is pushed. We use a sensor that can detect a contact, but not the location of the contact, between the hand and an object. The filter succeeds in reducing the uncertainty but ends up with a bimodal distribution since the sensor can not differentiate between the contacts on the outer and inner surfaces of the hand. In our ongoing study we are investigating different probabilistic methods and different sensor models that can be used during physics-based manipulation.

4. Conclusion

Physics-based planning will enable robots to perform a wide variety of manipulation tasks in human environments. We developed a planner which can use physics-based pushing actions. This planner performs remarkable tasks in cluttered and uncertain human environments. Many of these tasks are impossible to perform with a planner using only pick-and-place actions. We identify the challenges in building physics-based planners and present how we tackle these challenges in the pushing domain.

Acknowledgements This material is based upon work partially supported by the National Science Foundation under Grant No. EEC-0540865. Part of this work was conducted when Mehmet Dogar was a research intern at Willow Garage, Inc. Special thanks to Kaijen Hsiao, Matei Ciocarlie, Matt Mason, Chris Atkeson, and the members of the Personal Robotics Lab at CMU.

References

[1] M. Dogar and S. Srinivasa: "A planning framework for non-prehensile manipulation under clutter and uncertainty," Autonomous Robots, vol.33, no.3, pp.217–236, 2012.
[2] M.T. Mason: "Mechanics and Planning of Manipulator Pushing Operations," International Journal of Robotics Research, vol.5, no.3, pp.53–71, 1986.
[3] R.C. Brost: "Automatic grasp planning in the presence of uncertainty," International Journal of Robotics Research, vol.7, no.1, 1988.
[4] M. Peshkin and A. Sanderson: "The motion of a pushed, sliding workpiece," IEEE Journal on Robotics and Automation, vol.4, no.1, pp.569–598, 1988.
[5] K.M. Lynch and M.T. Mason: "Stable Pushing: Mechanics, Controllability, and Planning," IJRR, vol.15, no.6, pp.533–556, 1996.
[6] S. Akella and M.T. Mason: "Posing polygonal objects in the plane by pushing," International Journal of Robotics Research, vol.17, no.1, pp.70–88, 1998.
[7] K.M. Lynch: "Locally controllable manipulation by stable pushing," IEEE Transactions on Robotics and Automation, vol.15, no.2, pp.318–327, 1999.
[8] S. Goyal, A. Ruina and J. Papadopoulos: "Planar sliding with dry friction. Part 1. Limit surface and moment function," Wear, no.143, pp.307–330, 1991.
[9] R.D. Howe and M.R. Cutkosky: "Practical Force-Motion Models for Sliding Manipulation," International Journal of Robotics Research, vol.15, no.6, pp.557–572, 1996.
[10] G. Wilfong: "Motion planning in the presence of movable obstacles," Proceedings of the Fourth Annual Symposium on Computational Geometry, pp.279–288, 1988.
[11] O. Ben-Shahar and E. Rivlin: "Practical pushing planning for rearrangement tasks," IEEE Transactions on Robotics and Automation, vol.14, pp.549–565, 1998.
[12] M.H. Overmars, D. Nieuwenhuisen, D. Nieuwenhuisen, A. Frank and H. Overmars: "An effective framework for path planning amidst movable obstacles," In International Workshop on the Algorithmic Foundations of Robotics, 2006.
[13] M. Stilman and J.J. Kuffner: "Planning among movable obstacles with artificial constraints," In International Workshop on the Algorithmic Foundations of Robotics, pp.1–20, 2006.
[14] M. Stilman, J.-U. Schamburek, J. Kuffner and T. Asfour: "Manipulation planning among movable obstacles," IEEE International Conference on Robotics and Automation, 2007.
[15] M. Dogar, K. Hsiao, M. Ciocarlie and S. Srinivasa: "Physics-based grasp planning through clutter," Robotics: Science and Systems VIII, 2012.
[16] M. Dogar and S. Srinivasa: "Push-Grasping with Dexterous Hands," IEEE/RSJ International Conference on Intelligent Robots and Systems, 2010.
[17] M. Erdmann and M.T. Mason: "An exploration of sensorless manipulation," IEEE Journal of Robotics and Automation, vol.4, no.4, pp.369–379, 1988.
[18] S. Thrun, W. Burgard and D. Fox: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press, 2005.

Mehmet R. Dogar

Mehmet Dogar is a Ph.D. student at the Robotics Institute at Carnegie Mellon University. He is a member of the Personal Robotics Lab. His interests include physics-based non-prehensile manipulation, manipulation planning, and grasping. His work on push-grasping was a Best Paper Award Finalist at IROS 2010. Mehmet received his B.S. and M.S. degrees from Middle East Technical University, Turkey.

Siddhartha S. Srinivasa

Prof. Siddhartha Srinivasa is an Associate Professor at the Robotics Institute at Carnegie Mellon University. He founded the Personal Robotics Lab (http://personalrobotics.ri.cmu.edu) and co-directs the Manipulation Lab at CMU. Prof. Srinivasa's research focus is on developing manipulation, perception, planning, and learning algorithms that enable robots to accomplish useful manipulation tasks in dynamic and cluttered indoor environments. Prof. Srinivasa's work has won numerous awards including at RO-MAN 2012, HRI 2010, ICRA 2009 and 2010, IROS 2010, and RSS 2012. His work was also recently featured on the September 2011 issue of the National Geographic Magazine.