



UNIVERSITY OF LEEDS

This is a repository copy of *Pushing Fast and Slow: Task-Adaptive Planning for Non-prehensile Manipulation Under Uncertainty*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/141371/>

Version: Accepted Version

Proceedings Paper:

Agboh, WC orcid.org/0000-0002-0242-0215 and Dogar, MR orcid.org/0000-0002-6896-5461 (2020) Pushing Fast and Slow: Task-Adaptive Planning for Non-prehensile Manipulation Under Uncertainty. In: Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics, Springer Proceedings in Advanced Robotics (SPAR). 13th International Workshop on the Algorithmic Foundations of Robotics (WAFR 2018), 09-11 Dec 2018, Mérida, Yucatán, México. Springer . ISBN 978-3-030-44051-0

https://doi.org/10.1007/978-3-030-44051-0_10

© Springer Nature Switzerland AG 2020. This is an author produced version of a paper published in Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics, Springer Proceedings in Advanced Robotics (SPAR). Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Pushing Fast and Slow: Task-Adaptive Planning for Non-prehensile Manipulation Under Uncertainty

Wisdom C. Agboh and Mehmet R. Dogar

School of Computing, University of Leeds
Leeds, LS29JT, United Kingdom,
{scwca,m.r.dogar}@leeds.ac.uk

Abstract. We propose a planning and control approach to physics-based manipulation. The key feature of the algorithm is that it can adapt to the accuracy requirements of a task, by slowing down and generating “careful” motion when the task requires high accuracy, and by speeding up and moving fast when the task tolerates inaccuracy. We formulate the problem as an MDP with action-dependent stochasticity and propose an approximate online solution to it. We use a trajectory optimizer with a deterministic model to suggest promising actions to the MDP, to reduce computation time spent on evaluating different actions. We conducted experiments in simulation and on a real robotic system. Our results show that with a task-adaptive planning and control approach, a robot can choose fast or slow actions depending on the task accuracy and uncertainty level. The robot makes these decisions online and is able to maintain high success rates while completing manipulation tasks as fast as possible.

Keywords: Manipulation & Grasping, Motion and Path Planning

1 INTRODUCTION

We propose a planning and control algorithm for non-prehensile manipulation. The key feature of our algorithm is *task-adaptivity*: the planner can adapt to the accuracy requirements of a task, performing fast or slow pushes. For example in Fig. 1 (top), the robot is pushing an object on a narrow strip. The task requires high-accuracy during pushing — otherwise the object can fall down. The controller therefore generates slow pushing actions that make small but careful progress to the goal pose of the object. In Fig. 1 (bottom), however, the object is on a wide table and the goal region for the object is large (circle drawn on the table). In this case, the controller generates only a small number of fast pushing actions to reach the goal quickly — even if this creates more uncertainty about the object’s pose after each action, the task can still be completed successfully. We present a controller that can adapt to tasks with different accuracy requirements, such as in these examples. There has been significant recent interest in non-prehensile pushing-based manipulation. Most existing work use motion planning and *open-loop* execution to address the problem of generating a sequence of actions to complete a non-prehensile manipulation task [16, 17, 7, 23, 11]. Others developed closed-loop approaches. Ruiz-Ugalde et al. [26] proposed a compact predictive model for sliding objects and a feedback controller to push an object to a desired location. Hogan and Rodriguez [9] proposed a model predictive controller that uses integer

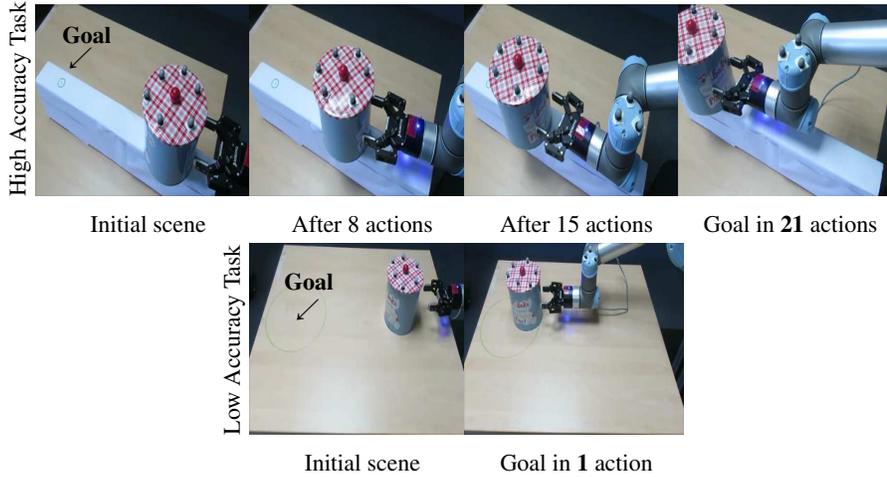


Fig. 1: Task-adaptive pushing with 21 slow actions for a high accuracy task (top) and a single fast action for a low accuracy task (bottom).

programming to handle the different contact modes associated with the pusher-slider system. Arruda et al. [2] considered the task of pushing an object to a goal region while avoiding regions of high predictive uncertainty through model predictive path integral control. We take a similar model predictive control (MPC) approach and propose a closed-loop planning and control algorithm for pushing tasks.

A common feature of existing work is the reliance on the quasi-static model of pushing [21, 10]. While one reason of the popularity of the quasi-static model may be the simpler analytic equations of motion it enables one to derive, another reason is the slow nature of quasi-static interactions, which keeps the uncertainty during pushing tightly bounded and therefore easier to control accurately.

However, accuracy is not the main criterion for every task, as we illustrate in Fig. 1. Fast motions, even if inaccurate, may be desired during some tasks. We humans also adapt our actions to the task (Fitts’s law [8]). Imagine reaching into a fridge shelf that is crowded with fragile objects, such as glass jars and containers. You move slowly and carefully. However, if the fridge shelf is almost empty, only with a few plastic containers that are difficult-to-break, you move faster with less care.

The major requirements to build a task-adaptive planner/controller are:

1. The planner must **consider a variety of actions (different pushing speeds)**: The robot should not be limited to moving at quasi-static speeds. It must consider dynamic actions wherever possible to complete a given task as fast as possible.
2. The planner must **consider action-dependent uncertainty**: Different actions can induce different amounts of uncertainty into the system. For example, pushing an object for a longer distance (or equivalently pushing faster for a fixed amount of time) would induce more uncertainty than pushing a short distance (or equivalently pushing slower for a fixed amount of time) [31].

One way to build such a controller is to model the problem as a Markov Decision Process (MDP) with stochastic dynamics, where the stochasticity is action-dependent.

Then, if this MDP is solved for an optimal policy under a cost that includes time to reach the goal, the resulting policy will use fast actions when it can, and fall back to slow actions for tasks that require higher accuracy.

In this paper, we model the problem as an MDP. However, we do not search for a globally optimal policy as this would be prohibitively computationally expensive. Instead, we solve the MDP online [24, 15] with an approximate solution. Even in this online setting, evaluating the value of all possible actions (including actions of a wide variety of speeds), proves computationally expensive, since the cost of physics-based predictions is high. Therefore, instead of evaluating all possible actions, at any given state, we first use a fast trajectory optimizer to suggest a reduced set of promising actions, i.e. actions that are known to drive the system to the goal under the deterministic setting. We then evaluate these actions under uncertainty to pick the best one.

Our specific contributions include a task-adaptive online solution to the stochastic MDP for pushing-based manipulation and a trajectory optimizer to generate actions for evaluation under the MDP setting. Additionally, we compare our task-adaptive planner with a standard model predictive control approach using only slow actions for high and low accuracy tasks under different levels of uncertainty. We show that our approach achieves higher success rates in significantly smaller amounts of time, especially for tasks that do not require high accuracy. Finally, we implement our approach on a real robotic system for tasks requiring different accuracy levels and compare it with standard MPC. Results can be found in the video at <https://youtu.be/A77CACvow5g>.

2 PROBLEM FORMULATION

We consider the problem where a robot must plan a sequence of non-prehensile actions to take an environment from an initial configuration to a desired goal configuration. We consider two task categories: In the *pushing* task the goal is to push a target object into a goal region; and in the *grasping in clutter* task the goal is to bring a target object, among other objects, into the robot’s hand. Our scenes contain D dynamic objects. \mathbf{q}^i refers to the full pose of each dynamic object, for $i = 1, \dots, D$. We assume a flat working surface and the robot is not allowed to drop objects off the edges.

Our robot is planar with a 1-DOF gripper. The robot’s configuration is defined by a vector of joint values $\mathbf{q}^R = \{\theta_x, \theta_y, \theta_{rotation}, \theta_{gripper}\}$. We represent the complete state of our system as \mathbf{x}_t at time t . This includes the configuration and velocity of the robot and all dynamic objects; $\mathbf{x}_t = \{\mathbf{q}^R, \mathbf{q}^1, \dots, \mathbf{q}^D, \dot{\mathbf{q}}^R, \dot{\mathbf{q}}^1, \dots, \dot{\mathbf{q}}^D\}$. Our control inputs are velocities: $\mathbf{u}_t = \dot{\mathbf{q}}^R$ applied to the robot’s joints. We then define the stochastic discrete time dynamics of our system as:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \zeta(\mathbf{u}_t) \quad (1)$$

where f is a deterministic function that describes the evolution of state \mathbf{x}_t given the action \mathbf{u}_t . We induce stochasticity in the system dynamics through $\zeta(\mathbf{u}_t) \propto \|\mathbf{u}_t\|$, which is proportional to the magnitude of action \mathbf{u}_t . When we push an object over a long distance, there is a large number of interactions/contacts especially in cluttered environments. This implies that the uncertainty in the resulting state at the end of a long push should be larger than that for a shorter push.

We assume an initial state of the system \mathbf{x}_0 . Our goal is to generate a sequence of actions for the robot such that the desired final goal configuration of the environment is reached as quickly as possible without dropping objects off the edge of our working surface. In the foregoing paragraphs, $\mathbf{U} = \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{n-1}\}$ denotes a sequence of control signals of fixed duration Δ_t applied in n time steps and we use brackets to refer to the control at a certain time step, i.e. $\mathbf{U}[t] = \mathbf{u}_t$. Similarly, \mathbf{X} is a sequence of states.

To build a task-adaptive controller, we formulate our problem as an MDP, and we provide an approximate solution to it. An MDP is defined by a tuple $\langle S, A, P, L' \rangle$, where S is the set of states, A is the set of actions, P is the probabilistic transition function, and L' defines the costs. In our problem S is given by all possible values of \mathbf{x}_t . Similarly, A is given by all possible values of \mathbf{u}_t , and P can be computed using the stochastic transition function in Eq. 1. The optimal policy for an MDP is given by:

$$\pi^*(\mathbf{x}_t) = \arg \min_{\mathbf{u}_t \in A} \left[L'(\mathbf{x}_t, \mathbf{u}_t) + \gamma \cdot \int_S P(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t) \cdot V^*(\mathbf{x}_{t+1}) \cdot dx_{t+1} \right] \quad (2)$$

where $0 < \gamma < 1$ is the discount factor, and V^* is the optimal value function. An on-line one-step lookahead approximate solution to the MDP problem can be found by sampling and evaluating the average value over samples as in [24, 15]:

$$\tilde{\pi}(\mathbf{x}_t) = \arg \min_{\mathbf{u}_t \in A} \left[L'(\mathbf{x}_t, \mathbf{u}_t) + \frac{1}{Q} \cdot \sum_{\mathbf{x}_{t+1} \in S(\mathbf{x}_t, \mathbf{u}_t, Q)} \tilde{V}(\mathbf{x}_{t+1}) \right] \quad (3)$$

where $S(\mathbf{x}_t, \mathbf{u}_t, Q)$ is the set of Q samples found by stochastically propagating $(\mathbf{x}_t, \mathbf{u}_t)$, and \tilde{V} is an approximation of the value function. For our problem, to compute the cost $L'(\mathbf{x}_t, \mathbf{u}_t)$, we use a cost function L which also takes into account the next state \mathbf{x}_{t+1} :

$$L(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{u}_t) = \sum_i^D \{w_e \cdot e^{k \cdot d_p^i} + w_s \cdot (\mathbf{x}_{t+1}^i - \mathbf{x}_t^i)^2\} + k_{act} \quad (4)$$

The first term in the cost which we call the *edge cost* penalizes pushing an object close to the table's boundaries or static obstacles. We show the edge cost in Fig. 4 where we define a safe zone smaller than the table's boundaries. If an object is pushed out of this safe region as a result of an action between t and $t + 1$, we compute the pushed distance d_p . Also note that k is a constant term and no edge costs are computed for objects in the safe zone. The second term is the *environment disturbance cost* which penalizes moving dynamic objects away from their current states. The third term, k_{act} is a constant cost incurred for each action taken by the robot. We use w_e and w_s to represent weights for the edge and environment disturbance costs respectively. Then, we compute $L'(\mathbf{x}_t, \mathbf{u}_t)$ using the same set of Q samples:

$$L'(\mathbf{x}_t, \mathbf{u}_t) = \frac{1}{Q} \cdot \sum_{\mathbf{x}_{t+1} \in S(\mathbf{x}_t, \mathbf{u}_t, Q)} L(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{u}_t)$$

This solution requires propagating Q samples for every possible action \mathbf{u}_t , to find the one with the minimum total cost. Performing this for all actions $\mathbf{u}_t \in A$ is not feasible for our purposes for a variety of reasons: First, we are interested in actions that span a wide range of speed profiles (i.e. fast and slow), which make our potential action

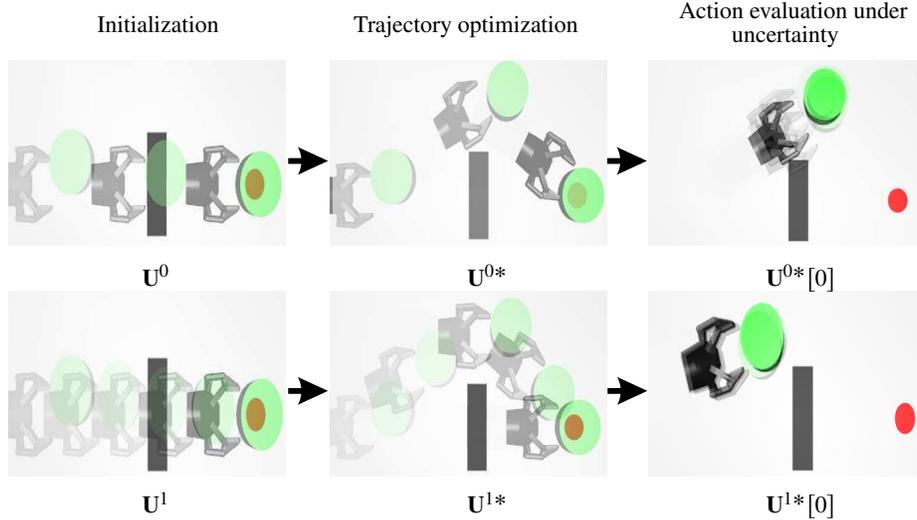


Fig. 2: First column: Initialization of our task-adaptive approach with control sequences including fast (top) and slow (bottom) actions. Second column: Stochastic trajectory optimization of the initial candidate control sequences. Last column: Action evaluation under uncertainty through sampling.

set large; second, each propagation in our domain is a physics simulation which is computationally expensive; and third, our goal is closed-loop pushing behaviour close to real-time speeds.

3 PROPOSED APPROACH

Instead of considering a large action set in our online MDP solution (Eq. 3), we propose to use a small set of promising actions including both fast and slow actions. We identify such a set of actions using a trajectory optimizer based on the deterministic dynamics function f of the system.

Note that our approach does not discretize the action space or the state space *a priori*. We adaptively sample the action space using a trajectory optimizer to find high value actions to consider at a given current state. We also get stochastic next state samples by applying these actions through a physics simulator.

In Alg. 1, we present our online approximate solution to the MDP.

Consider the scene in Fig. 2, where we have a planar gripper and an object. Our task is to push the object to a desired goal location (the red spot) while avoiding the rectangular black obstacle. We begin by generating N candidate action sequences $\{U^0, \dots, U^{N-1}\}$ to the goal by using the *GetActionSequences* procedure (line 1). The number of actions in each sequence varies between n_{min} and n_{max} , and each action is of fixed duration Δ_t . In the example task (Fig. 2), we show the candidate action sequences in the first column where $N = 2$, $n_{min} = 2$, and $n_{max} = 4$. Since each action is of fixed duration, the set of action sequences contain both fast (top) and slow (bottom) actions. Details of how we generate these candidate control sequences are explained in Sec 4.

Algorithm 1: Online MDP solver

Input : \mathbf{x}_0 : Initial state
Parameters : Q : Number of stochastic samples
 N : Number of initial candidate control sequences
 n_{min} : Minimum number of control actions in a control sequence
 n_{max} : Maximum number of control actions in a control sequence

- 1 $\{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\} \leftarrow \text{GetActionSequences}(\mathbf{x}_0, n_{min}, n_{max}, N)$
- 2 **while** *task not complete* **do**
- 3 $\{\mathbf{U}^{0*}, \dots, \mathbf{U}^{N-1*}\}, \{\tilde{\mathbf{V}}^0, \dots, \tilde{\mathbf{V}}^{N-1}\} \leftarrow \text{GetOptActionSequences}(\mathbf{x}_0, \{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\})$
- 4 **for** $i \leftarrow 0$ **to** $N - 1$, **do**
- 5 $V^i = 0$
- 6 **for** *each sample in* Q , **do**
- 7 $\mathbf{x}_1 = \text{StochasticExecution}(\mathbf{x}_0, \mathbf{U}^{i*}[0])$
- 8 $V^i = V^i + L(\mathbf{x}_0, \mathbf{x}_1, \mathbf{U}^{i*}[0]) + \tilde{\mathbf{V}}^i[0]$
- 9 $V^i = V^i / Q$
- 10 $i_{min} = \arg \min_{i \in N} V^i$
- 11 $\mathbf{x}_1 \leftarrow \text{execute } \mathbf{U}^{i_{min}}[0]$
- 12 check task completion
- 13 $\mathbf{x}_0 \leftarrow \mathbf{x}_1$
- 14 $\{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\} \leftarrow \{\text{GetActionSequences}(n_{min}, n_{max}, N - 1), \mathbf{U}^{i_{min}}[1 : n - 1]\}$

Using a trajectory optimizer (Sec. 5), the procedure *GetOptActionSequences* returns N optimized control sequences $\{\mathbf{U}^{0*}, \dots, \mathbf{U}^{N-1*}\}$, and an approximation of the value function $\{\tilde{\mathbf{V}}^0, \dots, \tilde{\mathbf{V}}^{N-1}\}$ along the optimal trajectories. We visualize the optimized trajectories for the example task in the second column.

We seek a one-step lookahead solution to the *MDP*, hence, we get the first actions $\mathbf{U}^{i*}[0]$ from each of the optimized control sequences. Our task is now to select the best amongst N actions even under uncertainty. We apply each action Q times to the stochastic state transition function in Eq. 1 to yield Q next state samples (line 7). More specifically, using our system dynamics model, we apply the controls (velocities) to the robot for the control duration Δ_t and thereafter we wait for a fixed extra time t_{rest} for objects to come to rest before returning the next state and computing the cost. We can see the Q samples in the third column (Fig. 2) for the example pushing task. Thereafter, on line 8, for each sample we add the immediate cost L to the approximate value $\tilde{\mathbf{V}}^i[0]$ for the resulting state. We use the same approximate value for all Q samples. We then compute an average value for each action (line 9), select the best one and execute it (line 11). If the task is not yet complete, we repeat the whole process but also re-use the remaining portion of the best control sequence $\mathbf{U}^{i_{min}}[1 : n - 1]$ from the current iteration. This algorithm chooses slow, low velocity actions for tasks that require high accuracy but faster actions for tasks that allow inaccuracies.

In the example pushing task, Fig. 2 (third column), we see a wider distribution in the resulting state for the fast action (top) compared to the slower action (bottom) as a result of our uncertainty model. Such a wide distribution in the resulting state increases the probability of undesired events happening especially in high accuracy tasks. For example, we see that some samples for the fast action result in collisions between the

Algorithm 2: GetActionSequences ($\mathbf{x}_0, n_{min}, n_{max}, N$)

Output : $\{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\}$: A set of candidate action sequences
Parameters : Δ_t : Control duration for each action in an action sequence

- 1 **for** $k \leftarrow 0$ **to** $N - 1$ **do**
- 2 $n^k = \lceil \frac{n_{max} - n_{min}}{N - 1} k \rceil + n_{min}$
- 3 $\mathbf{U}^k[0 : n^k - 1] = \left\{ \frac{\text{Distance to goal}}{n^k \Delta_t} \right\} \mathbf{1}$
- 4 **return** $\{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\}$

Algorithm 3: GetOptActionSequences($\mathbf{x}_0, \{\mathbf{U}^0, \dots, \mathbf{U}^{N-1}\}$)

Output : $\{\mathbf{U}^{0*}, \dots, \mathbf{U}^{N-1*}\}$: Optimized set of action sequences
 $\{\tilde{\mathbf{V}}^0, \dots, \tilde{\mathbf{V}}^{N-1}\} \triangleright$ Approx. value function along N optimal trajectories

- 1 **for** $i \leftarrow 0$ **to** $N - 1$ **do**
- 2 $\mathbf{U}^{i*}, \tilde{\mathbf{V}}^i \leftarrow \text{TrajectoryOptimization}(\mathbf{x}_0, \mathbf{U}^i)$
- 3 **return** $\{\mathbf{U}^{0*}, \dots, \mathbf{U}^{N-1*}\}, \{\tilde{\mathbf{V}}^0, \dots, \tilde{\mathbf{V}}^{N-1}\}$

robot and the obstacle. This implies high costs with respect to the slower action, hence our planner chooses the slow action in this case, executes it and starts the whole process again from the resulting state.

4 GENERATING A VARIETY OF ACTIONS

At each iteration of our online MDP solver, we provide N actions that are evaluated under uncertainty. First, using the *GetActionSequences* procedure in Alg. 2, we generate N candidate action sequences where the number of actions in an action sequence increases linearly from n_{min} to n_{max} (line 2).

On line 3, each of the action sequences is set to a straight line constant velocity profile to the goal. This is a simple approach, other more complicated velocity profiles could also be used here. Furthermore, in the *GetOptActionSequences* procedure, Alg. 3, we use these candidate action sequences to initialize a stochastic trajectory optimization algorithm (Sec. 5). The algorithm quickly finds and returns a locally optimal solution for each of the candidate control sequences. It also returns an approximation of the value function along N optimal trajectories. We use this approximate value while evaluating actions in our online MDP solver.

5 TRAJECTORY OPTIMIZATION

Trajectory optimization involves finding an optimal control sequence \mathbf{U}^* for a planning horizon n , given an initial state \mathbf{x}_0 , an initial candidate control sequence \mathbf{U} , and an objective J which can be written as:

$$J(\mathbf{X}, \mathbf{U}) = \sum_{t=0}^{n-1} L(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{u}_t) + w_f L_f(\mathbf{x}_n) \quad (5)$$

J is obtained by applying the control sequence \mathbf{U} starting from a given initial state and includes the sum of running costs L and a final cost L_f . We use the constant w_f to

Algorithm 4: Stochastic Trajectory Optimization

Input : \mathbf{x}_0 : Initial state
: \mathbf{U} : Candidate control sequence containing n actions

Output : \mathbf{U}^* : Optimal control sequence

Parameters : K : Number of noisy trajectory rollouts
: v : Sampling variance vector
: C_{thresh} : Success definition in terms of cost
: I_{max} : Maximum number of iterations

- 1 $\mathbf{X}, \mathbf{C} \leftarrow \text{TrajectoryRollout}(\mathbf{x}_0, \mathbf{U})$
- 2 **while** I_{max} not reached **and** $\text{Sum}(\mathbf{C}) > C_{thresh}$ **do**
- 3 **for** $k \leftarrow 0$ **to** $K - 1$ **do**
- 4 $\delta \mathbf{U}^k \leftarrow N(\mathbf{0}, v)$ \triangleright Random control sequence variation
- 5 $\mathbf{U}^k = \mathbf{U} + \delta \mathbf{U}^k$
- 6 $\mathbf{X}^k, \mathbf{C}^k \leftarrow \text{TrajectoryRollout}(\mathbf{x}_0, \mathbf{U}^k)$
- 7 $\mathbf{U}^* \leftarrow \text{UpdateTrajectory}(\mathbf{U}, \{\delta \mathbf{U}^0, \dots, \delta \mathbf{U}^{K-1}\}, \{\mathbf{C}^0, \dots, \mathbf{C}^{K-1}\})$
- 8 $\mathbf{X}^*, \mathbf{C}^* \leftarrow \text{TrajectoryRollout}(\mathbf{x}_0, \mathbf{U}^*)$
- 9 **if** $\text{Sum}(\mathbf{C}^*) < \text{Sum}(\mathbf{C})$ **then**
- 10 $\mathbf{U} \leftarrow \mathbf{U}^*, \quad \mathbf{X} \leftarrow \mathbf{X}^*, \quad \mathbf{C} \leftarrow \mathbf{C}^*$
- 11 **for** $j \leftarrow 0$ **to** $n - 1$ **do**
- 12 $\tilde{\mathbf{V}}[j] = \sum_{h=j}^{n-1} \mathbf{C}^*[h]$ \triangleright Approx. value function for state \mathbf{x}_j^* along the optimal trajectory
- 13 **return** $\mathbf{U}^*, \tilde{\mathbf{V}}$

weight the terminal cost with respect to the running cost. We consider a deterministic environment defined by the state transition function $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$. This is a constraint that must be satisfied at all times. Then the output of trajectory optimization is the minimizing control sequence:

$$\mathbf{U}^* = \arg \min_{\mathbf{U}} J(\mathbf{X}, \mathbf{U}) \quad (6)$$

In this work, we consider stochastic trajectory optimization methods such as STOMP [13] and MPPI [30]. With parallel rollouts on multiple cores of a PC, these methods show impressive speed. They also easily accept arbitrary cost functions that do not have to be differentiable. In contrast with sampling-based methods such as RRTs and PRMs [16, 23], optimization approaches are able to produce lower cost trajectories within a time limit even if they do not take the system to the desired goal state. These benefits make stochastic trajectory optimization very attractive. In this work, we propose Alg. 4 which adapts the STOMP algorithm [13] for non-prehensile object manipulation.

We begin with an initial candidate control sequence \mathbf{U} and iteratively seek lower cost trajectories (lines 2-10) until the cost reaches a threshold or until the maximum number of iterations is reached (line 2). We add random control sequence variations $\delta \mathbf{U}^k$ on the candidate control sequence to generate K new control sequences at each iteration (line 5). Thereafter on line 6, we do a trajectory rollout for each sample control sequence using the *TrajectoryRollout* procedure in Alg. 5. It returns the corresponding state sequence \mathbf{X} and costs \mathbf{C} calculated for each state along the resulting trajectory. I.e $C[t]$ is the cost of applying action u_t in state x_t . After generating K sample control sequences and their corresponding costs, the next step is to update the candidate control sequence

Algorithm 5: TrajectoryRollout

Input : \mathbf{x}_0 : Initial state
: \mathbf{U} : Control sequence with n actions

Output : \mathbf{X} : State sequence
: \mathbf{C} : Costs along the trajectory

- 1 **for** $t \leftarrow 0$ **to** $n - 1$ **do**
- 2 $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$
- 3 $\mathbf{C}[t] = L(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{u}_t)$ ▷ Calculate cost using Eq. 4
- 4 **if** $t == n - 1$ **then**
- 5 $\mathbf{C}[t] = \mathbf{C}[t] + L_f(\mathbf{x}_{t+1})$ ▷ Add final cost
- 6 **return** \mathbf{X}, \mathbf{C}

using the *UpdateTrajectory* procedure. One way to do this is a straightforward greedy approach where the minimum cost trajectory is selected as the update:

$$k^* = \arg \min_k \sum_{t=0}^{n-1} \mathbf{C}^k[t] \quad , \quad \mathbf{U}^* = \mathbf{U} + \delta \mathbf{U}^{k^*} \quad (7)$$

Another approach is a cost-weighted convex combination similar to [30]:

$$\mathbf{U}^*[t] = \mathbf{U}[t] + \frac{\sum_{k=0}^{K-1} [\exp(-(\frac{1}{\lambda})\mathbf{C}^k[t])] \delta \mathbf{U}^k[t]}{\sum_{k=0}^{K-1} \exp(-(\frac{1}{\lambda})\mathbf{C}^k[t])} \quad (8)$$

Where λ is a parameter to regulate the exponentiated cost's sensitivity. In our experiments, for a small number of noisy trajectory rollouts K (e.g. $K = 8$), a greedy update performs better. Hence, we use the greedy update in all our experiments.

Once the trajectory update step is complete, we update the candidate control sequence only if the new sequence has a lower cost (line 9, Alg. 4). The trajectory optimization algorithm then returns the locally optimal control sequence and an approximation of the value function for each state along the trajectory, where the value function is approximated using the sum of costs starting from that state.

The cost terms for the state-action sequences in this algorithm are equal to the running costs in Eq. 4, with the addition of a terminal cost on the final state depending on the task. The terminal cost for the pushing task is given by:

$$L_f = \begin{cases} 0 & \text{if } R_o - R_g < 0 \\ (R_o - R_g)^2 & \text{if } R_o - R_g > 0 \end{cases}$$

where R_o is the distance between the pushed object and the center of a circular goal region of radius R_g . The terminal cost term for the task of grasping in clutter is given by: $L_f = d_T^2 + w_\phi \cdot \phi_T^2$. We show how the distance d_T and the angle ϕ_T are computed in Fig. 3. First, create a vector from a fixed point in the gripper to the target object where d_T is the length of this vector and ϕ_T is the angle between the forward direction of the gripper and the vector. We use w_ϕ to weight angles relative to distances.

6 BASELINE APPROACH

We implement a standard model predictive control algorithm (*MPC*) as a baseline approach. It involves repeatedly solving a finite horizon optimal control problem using the

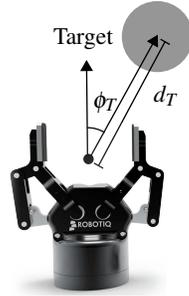


Fig. 3: Goal cost terms

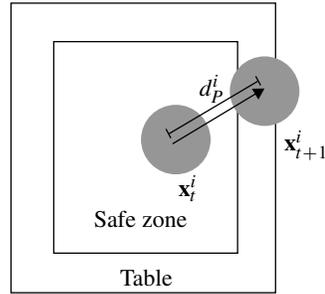


Fig. 4: Edge cost terms

stochastic trajectory optimizer presented in Alg. 4 and proceeds as follows: optimize a finite horizon control sequence, execute the first action, get the resulting state of the environment and then re-optimize to find a new control sequence. When re-optimizing, we warm-start the optimization with the remaining portion of the control sequence from the previous iteration such that optimization now becomes faster. We initialize the trajectory optimizer for standard MPC with actions at the quasi-static speed. In addition, we propose another baseline approach to compare against in this work: uncertainty aware model predictive control (UAMPC). This is a version of our online MDP solver where only low speed actions are considered. Specifically, all the candidate control sequences are generated with the maximum number of actions i.e. $n_{min} \leftarrow n_{max}$.

7 EXPERIMENTS

We call our planning approach task-adaptive model predictive control (*TAMPC*). We verify how well our approach is able to handle uncertainty and adapt to varying tasks. First, we compare the performance of *TAMPC* with a standard model predictive control (*MPC*) approach. Here we hypothesize that *TAMPC* will complete a given pushing task within a significantly shorter period of time and will be able to adapt to different tasks, maintaining a high success rate under varying levels of uncertainty.

Next, we compare the performance of our approach with uncertainty aware MPC (*UAMPC*). Here we hypothesize that: *UAMPC* will have a similar success rate and will take a longer amount of time to complete the task in comparison with *TAMPC*.

We conduct experiments in simulation and on a real robotic system. We consider the tasks of pushing an object to a goal region and grasping an object in clutter. Given an environment for planning, we create two instantiations:

- **Planning environment:** The robot generates plans in the simulated planning environment. The trajectory optimizer (Alg. 4) uses deterministic physics during planning and our online MDP solver uses stochastic physics to evaluate actions.
- **Execution environment:** Here, the robot executes actions and observes the state evolution. It is the physical world for real robot experiments but it is simulated for simulation experiments. The execution environment is stochastic.

For the planning environment and the execution environment when it is simulated, we use a physics engine, Mujoco[28], to model the deterministic state transition function f in Eq. 1. We model stochasticity in the physics engine by adding Gaussian noise

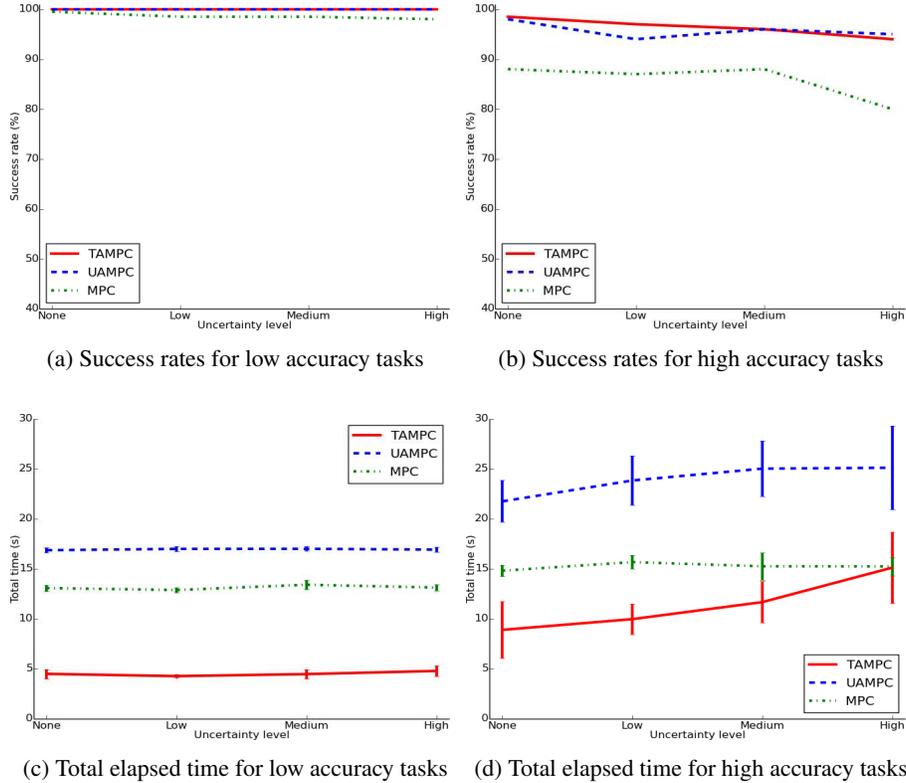


Fig. 5: Success rate and total elapsed time versus uncertainty level for low and high accuracy tasks.

on the velocities $\{\dot{\mathbf{q}}^R, \dot{\mathbf{q}}^1, \dots, \dot{\mathbf{q}}^D\}$ of the robot and objects at every simulation time step:

$$\{\tilde{\mathbf{q}}^R, \tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^D\} = \{\dot{\mathbf{q}}^R, \dot{\mathbf{q}}^1, \dots, \dot{\mathbf{q}}^D\} + \mu, \quad \mu \sim \mathcal{N}(0, \beta(\mathbf{u}_t)) \quad (9)$$

where \mathcal{N} is the Gaussian distribution and β is an action dependent variance function. We create a linear model for the variance function as:

$$\beta(\mathbf{u}_t) = b \|\mathbf{u}_t\| \quad (10)$$

Where b is a constant. In our simulation experiments, *uncertainty level* refers to the degree of stochasticity dictated by the slope b of the variance function β used to generate the Gaussian noise μ injected at every simulation time step in Eq. 9.

7.1 Push planning simulation experiments

We present a high accuracy task in Fig. 1 (top). It is made up of a thin strip and a small goal region. We also define a low accuracy task in Fig. 1 (bottom) which is a much larger table with a wider goal region. We create 200 such planning environments for each of the high and low accuracy tasks. For each environment:

1. We randomly select the shape (box or cylinder) of the pushed object.

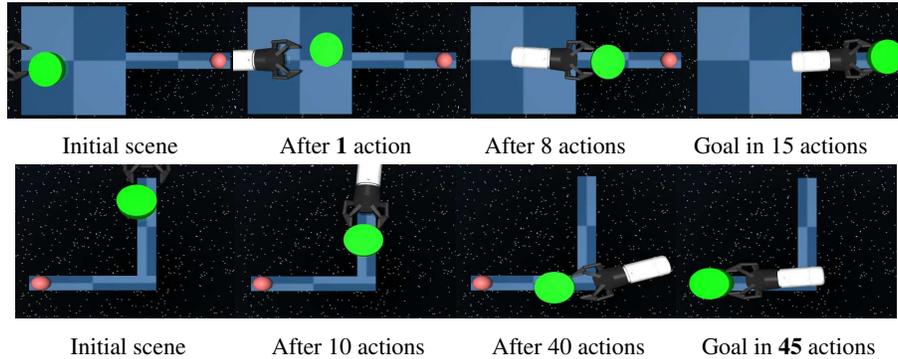


Fig. 6: Push planning in a changing environment (top) using a single fast push initially and then slow pushes later on due to the narrow strip. For the L-shaped environment (bottom), the robot executes many actions to successfully navigate the edge.

2. For each object, we randomly¹ select shape dimensions (radius and height for the cylinder, extents for the boxes), mass, and coefficient of friction.
3. We randomly² select a position on the working surface for the pushed object.

We create four uncertainty levels: no uncertainty, low uncertainty, medium uncertainty and high uncertainty. For the no uncertainty case, no extra noise was added to the physics engine. For low, medium and high levels of uncertainty, $b = \{0.05, 0.075, 0.1\}$ respectively. We test the different planning and control approaches and specify a time-out of 3 minutes including all planning, re-planning and execution.

Success rates: We declare success when the robot is able to push an object to the target region without dropping it off the edge of the table within the specified time limit. We plot the results in Fig. 5a and Fig. 5b. For the low accuracy level push planning task (Fig. 5a), *TAMPC* and *UAMPC* were able to maintain a 100 % success rate while *MPC* showed a slight decrease in success rates as uncertainty grew. For the high accuracy pushing task (Fig. 5b), *TAMPC* and *UAMPC* were also able to maintain a good average success rate. *MPC* on the other hand maintains a poor success rate. The major reason for this is uncertainty.

Total time: The total time in our experiments includes all planning and execution time. Fig 5c and Fig 5d show the average of 200 scenes with 95 % confidence interval of the mean. For the low accuracy level task, our *TAMPC* planner is able to achieve the goal in under 5s (Fig 5c), while *UAMPC* and *MPC* took significantly more time to complete the task. This clearly shows that our method is able to generate successful fast actions while maintaining a high success rate. For the high accuracy level task (Fig 5d), our planner is able to generate as many small actions as needed as the uncertainty grew.

¹ The uniform range used for each parameter is given here. Box x-y extents: $[0.05m, 0.075m]$; box height: $[0.036m, 0.05m]$; cylinder radius: $[0.04m, 0.07m]$; cylinder height: $[0.04m, 0.05m]$; mass: $[0.2kg, 0.8kg]$; coef. fric.: $[0.2, 0.6]$.

² The random position for the pushed object is sampled from a Gaussian with a mean at the lower end of the table (0.1m from the edge of a 0.6m long table along the center axis and a variance of 0.01m.)

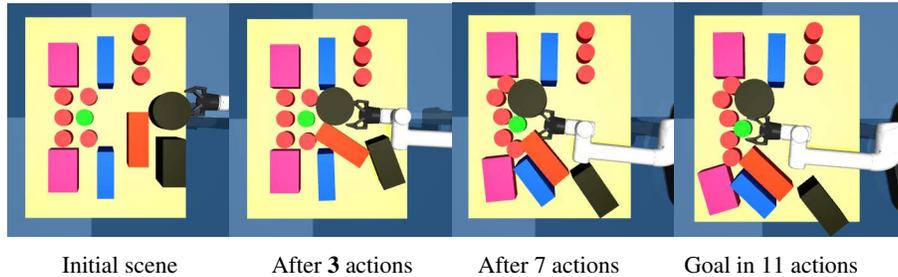


Fig. 7: Grasping in clutter: The robot uses fast actions initially but chooses slower actions as it gets closer to the goal object near the edge of the table.

Hence it was able to maintain a high success rate and still complete the task within a very small amount of time in comparison with the baseline approach. Furthermore, we also test the adaptive behavior of our approach for the environments in Fig. 6. In the changing environment (Fig. 6, top), the robot begins with a fast push due to a large initial area. Thereafter, it naturally switches to slow pushes on the thin strip to complete the task. For pushing in the L-shaped environment (Fig. 6, bottom), the robot generally pushes slow. However, it spends a lot of time to navigate the corner.

7.2 Grasping in clutter simulation experiments

We conducted simulation experiments for grasping in clutter in scenes similar to Fig. 7. Our scenes are randomly generated containing boxes and cylinders. In addition, our robot now has four control inputs (including the gripper). We tested our task adaptive planner in clutter to observe how the planner adapts given different environment configurations. We see that the robot manipulates clutter and is able to grasp the target object. An example scene is shown in Fig. 7 where the aim is to grasp the target object in green without pushing any other objects off the edge of the table. The robot initially begins with fast actions to push obstacles out of the way. However, as the robot gets closer to the target object, it chooses slower actions due to a higher probability of task failure in that region.

7.3 Real robot experiments

In the real robot experiments we use a UR5 robot with a Robotiq 2 finger gripper. We restrict the motion of the gripper to a plane parallel to the working surface such that we have a planar robot. We use OpenRave[6] to find kinematics solutions at every time step. For the push planning experiments, the gripper is completely open such that the robot receives three control inputs $\mathbf{u}_t = (\theta_x, \theta_y, \theta_{rotation})$ at every time step. We use a medium uncertainty level to model the real world stochasticity. We place markers on the pushed object and track its full pose with a motion capture system (OptiTrack). We manually replicated three execution worlds for each task accuracy level from the randomly generated environments we created during push planning simulation experiments. We tested our planners in these environments. We show snapshots from our real robot experiments. In Fig. 8 (top), we have a low task accuracy environment where the standard MPC approach is successful after 20 actions. However, by using a single dynamic push in Fig. 1 (bottom), our task-adaptive control approach is able to complete the push planning task in under 2 seconds.

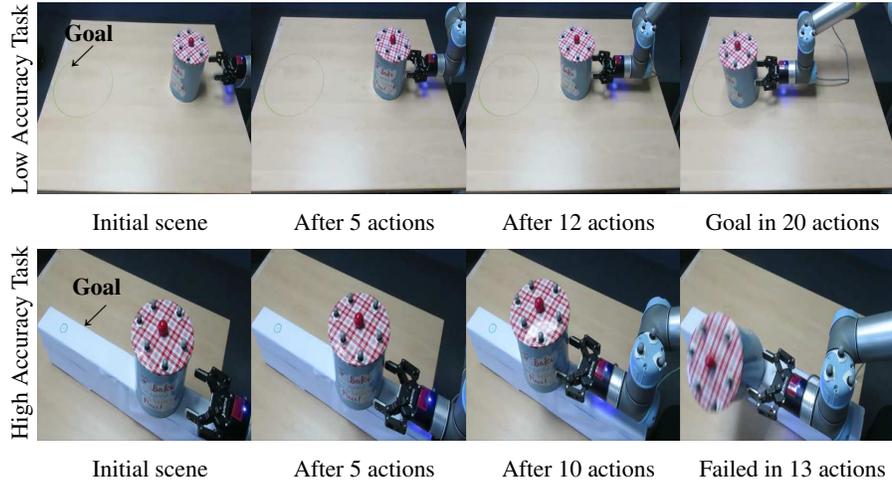


Fig. 8: *MPC* using a large number of actions to complete a low accuracy level task (top), and causing the pushed object to fall off for a high accuracy level task (bottom).

Moreover, for the high task accuracy problem, *MPC* was unable to push the target object to the desired goal location (Fig. 8 (bottom)). It executes actions without reasoning about uncertainty and pushed the goal object off the edge. Our task-adaptive controller was able to consider uncertainty while generating small pushes (Fig. 1 (top)) to complete the task. These results can be found in the accompanying video at <https://youtu.be/A77CACvow5g>.

8 RELATED WORK

Robots are becoming increasingly capable of performing complex manipulation tasks [33] by leveraging the environment’s physics [29]. A traditional approach to complete manipulation tasks is motion planning [18, 14], followed by open-loop execution. However, these plans are likely to fail due to uncertainty. Others [27, 32, 11, 19] have proposed using conservative open-loop motion plans that are robust to uncertainty. Recently, *MPC* algorithms for manipulation [2, 9, 4, 1] have been developed. They use feedback and can handle bounded uncertainty [22]. We take a similar closed-loop approach to non-prehensile manipulation in this paper.

There are recent works that develop learning-based uncertainty-aware controllers in robotics mainly for navigation and collision avoidance. Richter and Roy [25] proposed safe visual navigation with deep learning. Given that current deep learning methods produce unsafe predictions when faced with out-of-distribution inputs, they detect these novel inputs using the reconstruction loss of an autoencoder and switch to a rule-based safe policy. Similarly, Choi et al. [5] switch between a learned policy and a rule-based policy for autonomous driving. However, they estimate uncertainty using a single mixture density network without Monte Carlo sampling. We focus on non-prehensile manipulation where designing a rule-based policy for each new task is not feasible.

Furthermore, Kahn et al. [12] proposed an uncertainty-aware controller for autonomous navigation with collision avoidance. First, they estimate uncertainty from multiple bootstrapped neural networks using dropout. Thereafter, they consider a very

large number of fixed action sequences at a given current state. They evaluate these action sequences under an uncertainty-dependent cost function within an MPC framework. The resulting policy chooses low speed actions in unfamiliar environments and naturally chooses high speed actions in familiar environments. While this approach relieves the burden of designing a rule-based policy, it requires the evaluation of a large number of *a priori* fixed action sequences.

Our approach of adaptively sampling the action space given a current state is similar to prior work [3, 20] which propose sampling-based planning in continuous action MDPs. We use a trajectory optimizer to guide our action sampling process.

9 DISCUSSION AND FUTURE WORK

We presented a closed-loop planning and control algorithm capable of adapting to the accuracy requirements of a task by generating both fast and slow actions. This is an exciting first step toward realizing task-adaptive manipulation planners. In this work, we use a stochastic trajectory optimizer that outputs locally optimal control sequences. Thus, the resulting policy can get stuck in a local optima. Moreover, the trajectory optimizer may not return a good control sequence that reaches the goal for a given task if some design parameters (e.g. n_{max}) are chosen poorly. Furthermore, our uncertainty model is simple and is only an approximation of the real world stochastic phenomena. We will work toward finding uncertainty models that better describe the real world physics stochasticity especially for manipulation in clutter. Finally, we will investigate the generalization of our task-adaptive approach to other manipulation primitives.

10 ACKNOWLEDGEMENTS

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No. 746143, and from the UK Engineering and Physical Sciences Research Council under grants EP/P019560/1 and EP/R031193/1.

REFERENCES

- [1] Wisdom C. Agboh and Mehmet R. Dogar. Real-time online re-planning for grasping under clutter and uncertainty. In *IEEE-RAS Humanoids*, 2018.
- [2] E. Arruda, M. J. Mathew, M. Kopicki, M. Mistry, M. Azad, and J. L. Wyatt. Uncertainty averse pushing with model predictive path integral control. In *Humanoids*, 2017.
- [3] Sébastien Bubeck, Gilles Stoltz, Csaba Szepesvári, and Rémi Munos. Online optimization in x-armed bandits. In *NIPS*. 2009.
- [4] B. Calli and A. M. Dollar. Vision-based model predictive control for within-hand precision manipulation with underactuated grippers. In *ICRA*, 2017.
- [5] S. Choi, K. Lee, S. Lim, and S. Oh. Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling. In *ICRA*, 2018.
- [6] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning with caging grasps. In *Humanoids*, 2008.
- [7] M. R Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa. Physics-based grasp planning through clutter. In *Robotics: Science and Systems*, 2012.
- [8] Paul M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Experimental Psychology*, 1954.

- [9] F. R. Hogan and A. Rodriguez. Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. *WAFR*, 2016.
- [10] Robert D Howe and Mark R Cutkosky. Practical force-motion models for sliding manipulation. *IJRR*, 15(6):557–572, 1996.
- [11] A. M. Johnson, J. King, and S. Srinivasa. Convergent planning. *IEEE RA-L*, 2016.
- [12] Gregory Kahn, Adam Villaflor, Vitchyr Pong, Pieter Abbeel, and Sergey Levine. Uncertainty-aware reinforcement learning for collision avoidance. *CoRR*, 2017.
- [13] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. Stomp: Stochastic trajectory optimization for motion planning. In *ICRA*, 2011.
- [14] L. Kavraki and J. C. Latombe. Randomized preprocessing of configuration for fast path planning. In *ICRA*, 1994.
- [15] Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. A sparse sampling algorithm for near-optimal planning in large markov decision processes. In *IJCAI*, 1999.
- [16] J. E. King, J. A. Haustein, S. Srinivasa, and T. Asfour. Nonprehensile whole arm rearrangement planning on physics manifolds. In *ICRA*, 2015.
- [17] N. Kitaev, I. Mordatch, S. Patil, and P. Abbeel. Physics-based trajectory optimization for grasping in cluttered environments. In *ICRA*, 2015.
- [18] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO*, 2004.
- [19] B. Luders, M. Kothari, and J.P. How. Chance constrained rrt for probabilistic robustness to environmental uncertainty. In *AIAA Guidance, Navigation, and Control Conference*, 2010.
- [20] C. Mansley, A. Weinstein, and M. L. Littman. Sample-based planning for continuous action markov decision processes. In *ICAPS*, 2011.
- [21] Matthew T. Mason. Mechanics and planning of manipulator pushing operations. *The International Journal of Robotics Research*, 5(3):53–71, 1986.
- [22] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789 – 814, 2000.
- [23] Muhayyuddin, M. Moll, L. Kavraki, and J. Rosell. Randomized physics-based motion planning for grasping in cluttered and uncertain environments. *IEEE RA-L*, 3(2):712–719, 2018.
- [24] Laurent Péret and Frédéric Garcia. On-line search for solving markov decision processes via heuristic sampling. In *ECAI*. IOS Press, 2004.
- [25] Charles Richter and Nicholas Roy. Safe visual navigation via deep learning and novelty detection. In *RSS*, 2017.
- [26] F. Ruiz-Ugalde, G. Cheng, and M. Beetz. Fast adaptation for effect-aware pushing. In *Humanoids*, 2011.
- [27] A. Sieverling, C. Eppner, F. Wolff, and O. Brock. Interleaving motion in contact and in free space for planning under uncertainty. In *IROS*, 2017.
- [28] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IROS*, 2012.
- [29] M. Toussaint, K. Allen, K. Smith, and J. Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. In *RSS*, 2018.
- [30] G. Williams, A. Aldrich, and E. Theodorou. Model predictive path integral control using covariance variable importance sampling. *CoRR*, 2015.
- [31] K. T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez. More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing. In *IROS*, 2016.
- [32] J. Zhou, R. Paolini, A. M. Johnson, J. A. Bagnell, and M. T. Mason. A probabilistic planning framework for planar grasping under uncertainty. *IEEE RA-L*, 2(4):2111–2118, 2017.
- [33] Y. Zhu, Z. Wang, J. Merel, A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, Nando de Freitas, and N. Heess. Reinforcement and imitation learning for diverse visuomotor skills. In *RSS*, 2018.