



R × R: Rapid eXploration for Reinforcement learning via sampling-based reset distributions and imitation pre-training

Gagan Khandate¹ · Tristan L. Saidi¹ · Siqi Shang¹ · Eric T. Chang² · Yang Liu² · Seth Dennis² · Johnson Adams² · Matei Ciocarlie²

Received: 16 January 2024 / Accepted: 17 July 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

We present a method for enabling Reinforcement Learning of motor control policies for complex skills such as dexterous manipulation. We posit that a key difficulty for training such policies is the difficulty of exploring the problem state space, as the accessible and useful regions of this space form a complex structure along manifolds of the original high-dimensional state space. This work presents a method to enable and support exploration with Sampling-based Planning. We use a generally applicable non-holonomic Rapidly-exploring Random Trees algorithm and present multiple methods to use the resulting structure to bootstrap model-free Reinforcement Learning. Our method is effective at learning various challenging dexterous motor control skills of higher difficulty than previously shown. In particular, we achieve dexterous in-hand manipulation of complex objects while simultaneously securing the object without the use of passive support surfaces. These policies also transfer effectively to real robots. A number of example videos can also be found on the project website: sbri.cs.columbia.edu

Keywords Exploration · Reinforcement learning · Dexterous manipulation · In-hand manipulation

1 Introduction

Reinforcement Learning (RL) of robot sensorimotor control policies has seen great advances in recent years, demonstrated for a wide range of motor tasks such as dexterous manipulation. This has translated into higher levels of dexterity than previously possible, typically demonstrated by the ability to reorient a grasped object in-hand using complex finger movements (Akkaya et al., 2019; Chen et al., 2021; Qi et al., 2022).

However, training a sensorimotor policy is still a difficult process, particularly for hard-exploration problems where the underlying state space exhibits complex structure, such

as "narrow passages" between parts of the space that are accessible or useful. Manipulation is indeed such a problem: even when starting with the object secured between the digits, a random action can easily lead to a drop, and thus to an irrecoverable state. Finger-gaiting (Ma et al., 2011) further implies transitions between different subsets of fingers used to hold the object, all while maintaining stability. This leads to difficulty in exploration during training, since random perturbations in the policy action space are unlikely to discover narrow passages in state space. Current studies address this difficulty through a variety of means: using simple, convex objects and non-convex objects of significantly reduced size to reduce the difficulty of the task, reliance on support surfaces to reduce the chances of a drop, object pose tracking through extrinsic sensing, etc.

The difficulty of exploring problems with labyrinthine state space structures is far from new in robotics. In fact, the large and highly effective family of Sampling-Based Planning (SBP) algorithms was developed to address this exact problem. By expanding a known structure towards targets randomly sampled in the state space of the problem, SBP methods can explore even very high-dimensional state spaces in ways that are probabilistically complete, or guaranteed to converge to optimal trajectories. However, SBP algo-

Gagan Khandate, Tristan L. Saidi and Siqi Shang have contributed equally to this work.

✉ Gagan Khandate
gagank@cs.columbia.edu

✉ Tristan L. Saidi
tls2160@columbia.edu

¹ Department of Computer Science, Columbia University, New York, NY 10027, USA

² Department of Mechanical Engineering, Columbia University, New York, NY 10027, USA

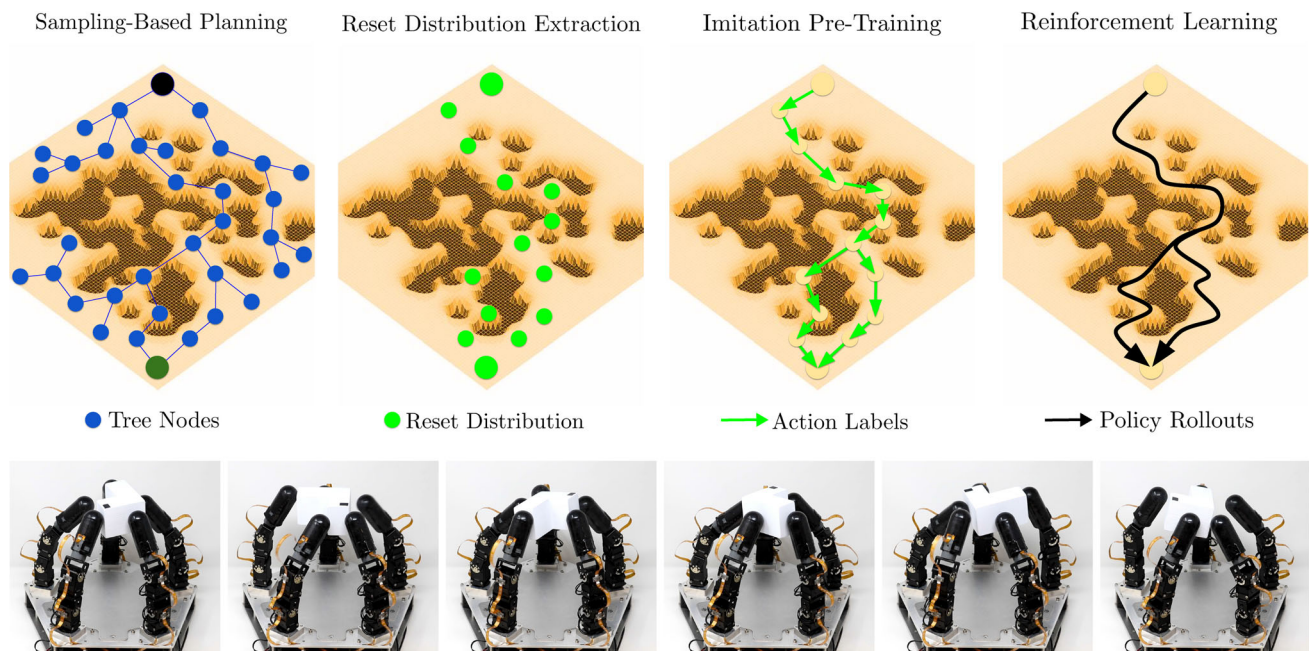


Fig. 1 Our method illustrated with an abstract state-space consisting of narrow stable regions (beige) between large unstable regions (black). The proposed two-stage approach uses sampling-based planning to

explore the challenging state-space and leverages the information within best paths by using state data as reset distribution and action data for imitation pre-training for efficient reinforcement learning

gorithms are traditionally designed to find trajectories rather than policies. For problems with computationally demanding dynamics, SBP cannot be used online for previously unseen start states or to quickly correct when unexpected perturbations are encountered along the way.

Broadly, this work draws on the strength of both RL and SBP methods in order to train motor control policies for in-hand manipulation with finger gaiting. We aim to manipulate more difficult objects, including concave shapes, while securing them at all times without relying on support surfaces. Furthermore, we aim to achieve large reorientation of the grasped object with purely intrinsic (tactile and proprioceptive) sensing. To achieve that we use a non-holonomic RRT algorithm with added constraints to find approximate trajectories that explore the useful parts of the problem state space. Then, we use these trajectories towards training complete RL policies based on the full dynamics of the problem. In particular, we use state data sampled from the tree to construct exploratory reset distribution and use action data from tree paths to obtain a warm-start policy through imitation learning and illustrate this in Fig. 1.

This extended version of our study contains a number of additions compared to the previous conference proceedings version (Khandate et al., 2023). We introduce Imitation Pre-training, a method for leveraging actions sampled from the RRT tree in order to further bootstrap learning and show that

it further improves training efficiency. We extend our applications to the more difficult in-hand manipulation task of reorientation to a desired pose, which can be either a canonical pose known at training time (for example, as needed by downstream steps in an assembly process) or arbitrary and specified as a goal at run-time. Finally, we provide a number of additional insights obtained from an extensive ablation of our approach.

Overall, the main contributions of this work include:

- To the best of our knowledge, we are the first to show that reset distributions generated via SBP with kinematic constraints can enable more efficient training of RL control policies for dexterous in-hand manipulation.
- We show that SBP can explore useful parts of the manipulation state space, allowing RL to later fill in the gaps between approximate trajectories by learning appropriate actions under more realistic dynamic constraints. A warm-start policy obtained by imitating the actions from sampling-based trajectories provides an additional boost in training speed.
- The exploration boost from SBP allows us to train policies for dexterous skills not previously shown, such as in-hand manipulation of concave shapes, with only intrinsic sensing and no support surfaces. We demonstrate these skills both in simulation and on real hardware.

2 Related work

2.1 Exploration in RL

Exploration methods for general RL operate under the strict assumption that the learning agent cannot teleport between states, mimicking the constraints of the real world. Under such constraints, proposed exploration methods include using intrinsic rewards (Pathak et al., 2017; Haarnoja et al., 2018) or improving action consistency via temporally correlated noise in policy actions (Amin et al., 2021) or parameter space noise (Plappert et al., 2017).

Fortunately, in cases where the policies are primarily trained in simulation, this requirement can be relaxed, and we can use our knowledge of the relevant state space to design effective exploration strategies. A number of these methods improve exploration by injecting useful states into the reset distribution during training. Nair et al. (2017) use states from human demonstrations in a block stacking task, while Ecoffet et al. (2019, 2021) use states previously visited by the learning agent itself for problems such as Atari games and robot motion planning. Tavakoli et al. (2018) evaluate various schemes for maintaining and resetting from the buffer of visited states. However, these schemes were evaluated only on benchmark continuous control tasks (Duan et al., 2016). From a theoretical perspective, Agarwal et al. (2020) show that a favorable reset state distribution provides a means to circumvent worst-case exploration issues using sample complexity analysis of policy gradients.

2.2 Combining SBP and RL

Finding feasible trajectories through a complex state space is a well-studied motion planning problem. Of particular interest to us are sampling-based methods such as Rapidly-exploring Random Trees (RRT) (LaValle, 1998; Karaman & Frazzoli, 2010; Webb & Berg, 2013) and Probabilistic Road Maps (PRM) (Kavraki et al., 1996, 1998). These families of methods have proven highly effective and are still being expanded. Stable Sparse-RRT (SST) and its optimal variant SST* (Li et al., 2021) are examples of recent sampling-based methods for high-dimensional motion planning with physics. However, the goal of these methods is finding (kinodynamic) trajectories between known start and goal states rather than closed-loop control policies which can handle deviations from the expected states.

Several approaches have tried to combine the exploratory ability of SBP with RL, leveraging planning for global exploration while learning a local control policy via RL (Chiang et al., 2019; Francis et al., 2020; Schramm & Boularias, 2022). These methods were primarily developed for and tested on navigation tasks, where nearby state space samples are generally easy to connect by an RL agent acting as a local planner.

The LeaPER algorithm (Pinto et al., 2018) also uses plans obtained by RRT as reset state distribution and learns policies for simple non-prehensile manipulation. However, the state space for the prehensile in-hand manipulation tasks we show here is highly constrained, with small useful regions and non-holonomic transitions. Other approaches use trajectories planned by SBP as expert demonstrations. Morere et al. (2020) recommend using a policy trained with SBP as expert demonstrations as an initial policy. Alternatively, Jurgenson and Tamar (2019) and Ha et al. (2020) use planned trajectories in the replay buffer of an off-policy RL agent for multi-arm motion planning. First, these methods requires that planned trajectories also include the actions used to achieve transitions, which SBP does not always provide. Next, it is unclear how off-policy RL can be combined with the extensive physics parallelism that has been vital in the recent success of on-policy methods for learning manipulation (Allshire et al., 2021; Makoviychuk et al., 2021; Chen et al., 2021).

2.3 Dexterous manipulation

Turning specifically to the problem of dexterous manipulation, a number of methods have been used to advance the state of the art, including planning, learning, and leveraging mechanical properties of the manipulator. Leveroni and Salisbury (1996) build a map of valid grasps and use search methods to generate gaits for planar reorientation, while (Han & Trinkle, 1998) consider finger-gaiting of a sphere and identify the non-holonomic nature of the problem. Some methods have also considered RRT for finger-gaiting in-hand manipulation (Yashima et al., 2003; Xu et al., 2007), but limited to simulation for a spherical object. More recently, Morgan et al. demonstrate robust finger-gaiting for object reorientation using actor-critic reinforcement learning (Morgan et al., 2021) and multi-modal motion planning (Morgan et al., 2022), both in conjunction with a compliant, highly underactuated hand designed explicitly for this task. Bhatt et al. (2022) also demonstrate robust finger-gaiting and finger-pivoting manipulation with a soft compliant hand, but these skills were hand-designed and executed in an open-loop fashion rather than autonomously learned.

Model-free RL has also led to significant progress in dexterous manipulation, starting with OpenAI's demonstration of finger-gaiting and finger-pivoting Akkaya et al. (2019) trained in simulation and translated to real hardware. However, this approach uses extensive extrinsic sensing infeasible outside a lab setting, and relies on support surfaces such as the palm underneath the object. Khandate et al. (2022) show dexterous finger-gaiting and finger-pivoting skills using only precision fingertip grasps to enable both palm-up and palm-down operation, but only on a range of simple convex shapes and in a simulated environment.

Recently (Makoviychuk et al., 2021) showed that GPU physics could be used to accelerate learning skills similar to OpenAI's. Allshire et al. (2021) used extensive domain randomization and sim-to-real transfer to re-orient a cube but used a table top as an external support surface. Chen et al. (2021), Chen et al. (2022), and Handa et al. (2022) demonstrated in-hand re-orientation for a wide range of objects under palm-up and palm-down orientations of the hand with extrinsic sensing providing dense object feedback.

Sievers et al. (2022), Pitz et al. (2023) and Röstel et al. (2023) demonstrated in-hand cube reorientation to desired pose with purely tactile feedback. Qi et al. (2022) and Yin et al. (2023) used rapid motor adaptation to achieve effective sim-to-real transfer of tactile only in-hand manipulation skills for continuous reorientation of small cylindrical and cube-like objects. Yuan et al. (2023) and Qi et al. (2023) use both visual and tactile sensing. We learn tactile in-hand manipulation skills critically, in our case, the exploration ability of SBP allows learning of policies for more difficult tasks, such as in-hand manipulation of non-convex and large shapes, with only intrinsic sensing. We also achieve successful and robust sim-to-real transfer without extensive domain randomization or domain adaptation by closing the sim-to-real gap via tactile feedback.

3 Problem description

We focus on the problem of achieving dexterous in-hand manipulation while simultaneously securing the manipulated object in a precision grasp. Keeping the object stable in the grasp during manipulation is needed in cases where a support surface is not available or the skill must be performed under different directions for gravity (i.e. palm up or palm down). However, it also creates a difficult class of manipulation problems, combining movement of both the fingers and the object with a constant requirement of maintaining stability.

Formally, our goal is to obtain a policy for issuing finger motor commands to achieve a desired object transformation. The state of our system at time t is denoted by $\mathbf{x}_t = (\mathbf{q}_t, \mathbf{p}_t)$, where $\mathbf{q} \in \mathcal{R}^d$ is a vector containing the positions of the hand's d degrees of freedom (joints), and $\mathbf{p} \in \mathcal{R}^6$ contains the position and orientation of the object with respect to the hand. An action (or command) is denoted by the vector $\mathbf{a} \in \mathcal{R}^d$ comprising new setpoints for the position controllers running at every joint.

For parts of our approach, we assume that a model of the forward dynamics of our environment (i.e. a physics simulator) is available for planning or training. We denote this model by $\mathbf{x}_{t+1} = F(\mathbf{x}_t, \mathbf{a}_t)$. We will show however that our results transfer to real robots using standard sim-to-real methods.

Algorithm 1 General-purpose non-holonomic RRT (G-RRT)

Input: Tree containing root node, G ; $N \leftarrow 1$

```

1: while  $N < N_{max}$  do
2:    $\mathbf{x}_{sample} \leftarrow$  random point in state space
3:    $\mathbf{x}_{node} \leftarrow$  node closest to  $\mathbf{x}_{sample}$  currently in  $G$ 
4:    $d_{min} \leftarrow \infty$ ;  $\mathbf{x}_{new} \leftarrow \text{NULL}$ 
5:   while  $k < K_{max}$  do
6:      $\mathbf{a} \leftarrow \mathcal{N}(0, \alpha \mathbf{I})$  random action
7:      $\mathbf{x}_a \leftarrow F(\mathbf{x}_{node}, \mathbf{a})$ 
8:     if Stable( $\mathbf{x}_a$ ) and  $dist(\mathbf{x}_{sample}, \mathbf{x}_a) < d_{min}$  then
9:        $d_{min} \leftarrow dist(\mathbf{x}_{sample}, \mathbf{x}_a)$ 
10:       $\mathbf{x}_{new} \leftarrow \mathbf{x}_a$ 
11:    end if
12:     $k \leftarrow k + 1$ 
13:  end while
14:  if  $\mathbf{x}_{new}$  is not NULL then
15:    Add  $\mathbf{x}_{new}$  to  $G$  with  $\mathbf{x}_{node}$  as parent
16:     $N \leftarrow N + 1$ 
17:  end if
18: end while
return  $G$ 

```

As discussed above, we also require that the hand maintains a stable precision grasp of the manipulated object at all times. Overall, this means that our problem is characterized by a high-dimensional state space, but only small parts of this state space are accessible to us: those where the hand is holding the object in a stable precision grasp. Furthermore, the transition function of our problem is non-holonomic: the subset of fingers that are tasked with holding the object at a specific moment, as well as the object itself, must move in a concerted fashion. Conceptually, the hand-object system must evolve on the complex union of high-dimensional manifolds that form our accessible states. Still, the problem state space must be effectively explored if we are to achieve dexterous manipulation with large object reorientation and finger gaing.

4 Sampling-based state space exploration

To effectively explore our high-dimensional state space characterized by non-holonomic transitions, we turn to the well-known Rapidly-exploring Random Trees (RRT) algorithm. We leverage our knowledge of the manipulation domain to induce tree growth along the desired manifolds in state space. In particular, we expect two conditions to be met for any state: (1) the hand is in contact with the object only via fingertips, (2) the distribution of these contacts must be such that a stable grasp is possible.

Assuming system dynamics $F()$ are available and fast to evaluate, we use here the general non-holonomic version of the RRT algorithm, which is able to determine an action that moves the agent towards a desired sample in state space via random sampling. We use the same version of this algorithm

as described, for example, by King et al. (2016), which we recapitulate here in Algorithm 1 and refer to as G-RRT.

The essence of this algorithm is the **while** loop in line 5: it is able to grow the tree in a desired direction by sampling a number K_{max} of random actions, then using the transition function $F()$ of our problem to evaluate which of these produces a new node that is as close as possible to a sampled target.

Our only addition to the general-purpose algorithm is the stability check in line 8: a new node gets added to the tree only if it passes a stability check. This check consists of advancing the simulation for an additional 2s with no change in the action; if, at the end of this interval, the object has not been dropped (i.e. the height of the object is above a threshold) then the new node is deemed stable and added to the tree. Assuming a typical simulation step of 2 ms, this implies 1000 additional calls to $F()$ for each sample.

Overall, the great advantage of this algorithm lies in its simplicity and generality. The only manipulation-specific component is the aforementioned stability check. However, its performance can be dependent on K_{max} (i.e. number of action samples at each iteration), and each of these samples requires a call to the transition function. This problem can be alleviated by the advent of highly efficient and massively parallel physics engines implementing the transition function, which is an important research direction complementary to our study.

We note that in previous work (Khandate et al., 2023) we used an additional condition on the state: the hand must maintain at least three fingers in contact with the object.¹ Empirically, we noticed that the three contact constraint had minimal impact on the results of G-RRT. Thus, we discontinued the contact constraint in this work.

In the same prior work, we also introduced a manipulation-specific version of the non-holonomic RRT algorithm dubbed M-RRT, which does not require the use of the transition function for stability checks. Instead, M-RRT uses manipulation kinematics alone to explore manifolds defined by the contact constraints. This version also required that the hand must maintain at least three fingers in contact with the object. While the M-RRT version has the potential to obtain significant speedups by foregoing dynamic simulation, G-RRT is appealing in its generality and simplicity. As both versions are able to effectively explore the state space, we focus this study on the G-RRT version of our algorithm and plan additional comparisons between these two versions for future work.

¹ Three contacts are the fewest that can achieve stable grasps without relying on torsional friction, which is highly sensitive to the material properties of the objects in contact.

Algorithm 2 R×R (+ IPT)

Input: Randomly initialized actor π_θ and critic V_ψ , learning rates η_1, v_1, η_2, v_2

```

# Execute Sampling-based Planning
1: Tree,  $G \leftarrow$  G-RRT() ▷ Alg 1
# Extract top trajectories
2: Let  $\mathbf{x}_{root}$  be the root state and  $\mathbf{x}_{goal}$  be the desired goal state towards
achieving the manipulation task of interest.
3:  $\mathcal{D} \leftarrow \{\}, P \leftarrow 1$ 
4: while  $P < P_{max}$  do
5: Sample goal  $\mathbf{x}_{goal}$  ▷ fixed or randomized
6:  $\mathbf{x}_{node} \leftarrow$  node closest to  $\mathbf{x}_{goal}$  now in  $G$ 
7:  $\tau = \{\mathbf{x}_{node}\}$ 
8: while  $\mathbf{x}_{node} \neq \mathbf{x}_{root}$  do
9: delete  $\mathbf{x}_{node}$  from  $G$ 
10: Let  $\mathbf{x}_{parent}$  be parent of  $\mathbf{x}_{node}$ 
11:  $\tau \leftarrow \mathbf{x}_{parent} \cup \tau$  ▷ nodes in reverse
12:  $\mathbf{x}_{node} \leftarrow \mathbf{x}_{parent}$ 
13: end while
14:  $\mathcal{D} \leftarrow \mathcal{D} \cup \tau, P \leftarrow P + 1$ 
15: end while
# Imitation Pre-training ▷ Optional
Get warm-start policy and critic
16:  $\pi_\theta, V_\psi \leftarrow$  IPT( $\mathcal{D}, \pi_\theta, V_\psi, \eta_1, v_1$ ) ▷ from Alg 3
# Reinforcement Learning
17:  $\pi_\theta, V_\psi \leftarrow$  RL( $\mathcal{D}, \pi_\theta, V_\psi, \eta_2, v_2$ )
18: for  $N$  iterations do
19: Initialize/clear rollout buffer  $\mathcal{R}$ 
20: for  $E$  episodes do
21: Collect rollout  $\tau$  with  $\pi_\theta$  from initial state  $\mathbf{x}_0$  where  $\mathbf{x}_0 \sim \mathcal{D}$ 
the buffer of best states from sampling-based planning.
22:  $\mathcal{R} \leftarrow \mathcal{R} \cup \tau$ 
23: end for
Update  $\pi_\theta$  w.r.t policy loss  $\mathcal{L}_\pi$  on  $\mathcal{R}$  with RL algorithm of choice
24:  $\theta \leftarrow \theta - \eta_2 \nabla_\theta \mathcal{L}_\pi(\theta)$ 
Update critic  $V_\psi$ 
25:  $\mathcal{L}_V = \mathbb{E}_{\mathbf{o}_t \sim \mathcal{R}} \left[ (V_\psi(\mathbf{o}_t) - V_t^{target})^2 \right]$ 
26:  $\psi \leftarrow \psi - v_2 \nabla_\psi \mathcal{L}_V$ 
27: end for

```

5 Reinforcement learning

While the exploration method we have discussed so far is capable of exploring the complex state space of in-hand manipulation and identifying approximate transitions that follow the complex manifold structure of this space, it does not provide directly usable policies. The state transitions themselves provide no mechanism to act in states that are not part of the tree, or to act under slightly different transition functions.

To generate closed-loop policies able to handle variability in the encountered states, we turn to RL algorithms. Critically, we rely on the trees generated by our sampling-based algorithms to ensure effective exploration of the state space during policy training.

5.1 Sampling-based reset distribution

One mechanism we use to transition information from the sampling-based tree to the policy training method is via the reset distribution: we select relevant paths from the planned tree and then use the nodes therein as reset states for policy training.

We note that the sampling-based trees as described here are task-agnostic. Their effectiveness lies in achieving good coverage of the state space (usually within pre-specified limits along each dimension). Once a specific task is prescribed (e.g. via a reward function), we must select states from paths through the tree that are relevant to the task. After selecting said states, we use a uniform distribution over these states as a reset distribution for RL. We describe this concretely in Algorithm 2. In particular, lines 2–15 show the heuristic approach we use to select paths through the tree to compose the set of reset states. Lines 18–27 show reinforcement learning from using these reset states for exploration. We note that other selection mechanisms are also possible; a promising and more general direction for future studies is to select tree branches that accumulate the highest reward.

This approach has a theoretical grounding in related work showing that derived reset distributions can be used to aid exploration. Let ρ denote the initial state distribution of the MDP to be solved with policy-gradient RL. Recent results (Agarwal et al., 2020) show that it is beneficial to compute policy gradients for policy π under a different initial distribution μ if it enables sufficient exploration. If d_ρ^π denotes the stationary state distribution induced by policy π under the initial state distribution ρ , then $d_\rho^{\pi^*}$ is the stationary distribution of the optimal policy π^* with initial state distribution ρ . Improved exploration can be achieved if μ sufficiently covers $d_\rho^{\pi^*}$. Here, we obtain such a favorable distribution μ from plans derived via sampling-based planning.

As this approach is compatible with online RL wherein policy rollouts are collected from a new set of states every episode, both off-policy and on-policy RL are equally feasible. However, we use on-policy learning due to its compatibility with GPU physics simulators and relative training stability.

5.2 Imitation pre-training

Relying on sampling-based trajectories exclusively for the reset state distribution disregards potentially valuable information concerning actions and the connectivity between states within exploration trees. Imitation learning, which maps states to actions via supervised learning, can be effective in using such action information embedded within the exploration trees. However, this is not straightforward to execute. The sampling-based trajectories cannot be treated as authentic demonstrations. Instead, they are quasi-static

Algorithm 3 Imitation Pre-training (IPT)

Input: Sampling-based trajectories (state-only) as demonstration \mathcal{D} , randomly initialized policy π and value network V with parameters θ, ψ , learning rates η and ν
Assemble buffer with actions \mathcal{D}'

- 1: **for** all consecutive state pairs $\mathbf{x}_k, \mathbf{x}_{k+1}$ **do**
- 2: Compute \mathbf{o}_k observation vector at \mathbf{x}_k
- 3: Get \mathbf{a}_k with Eq. (1)
- 4: $\mathcal{D}' \leftarrow \mathcal{D} \cup (\mathbf{o}_k, \mathbf{a}_k)$
- 5: **end for**
- # BC with MSE loss \mathcal{D}'
- 6: **for** E_1 epochs **do**
- 7: $L_\pi = \mathbb{E}_{\mathbf{o}, \mathbf{a} \sim \mathcal{D}'} [(a - \pi_\theta(\mathbf{o}))^2]$
- 8: $\theta \leftarrow \theta - \eta \nabla_\theta L_\pi(\theta)$
- 9: **end for**
- # Value pre-training
- 10: **for** E_2 epochs **do**
- 11: Collect rollouts with π_θ and store in \mathcal{R}
- 12: $L_V = \mathbb{E}_{\mathbf{o}_t \sim \mathcal{R}} [(V_\psi(\mathbf{o}_t) - V_t^{target})^2]$
- 13: $\psi \leftarrow \psi - \nu \nabla_\psi L_V$
- 14: **end for**

return π_θ, V_ψ

demonstrations composed of a sequence of stable states as a consequence of the stability constraint (Algorithm 1, line 8) we enforce during sampling-based planning. Due to this stability check, the actions are a distorted version of the desired actions under full dynamics. Thus, imitation learning, for instance, with Behavioral Cloning (BC), is unlikely to yield a successful policy due to the approximate nature of demonstrations and inherent challenges arising from distribution shift. Nonetheless, the resulting policies can be used as a warm-start policy for reinforcement learning.


Here we provide a description of this approach. First, we observe that the difference in joint angles between node transitions is a reliable approximation of the desired action, particularly its direction. Augmented by a scaling hyperparameter, these differences can serve as action labels. For illustration, let \mathbf{q}_k and \mathbf{q}_{k+1} be joint angles at successive states. The action labels \mathbf{a}_k can be obtained by simply scaling the difference with a scaling factor β , i.e. $\mathbf{a}_k = \beta(\mathbf{q}_{k+1} - \mathbf{q}_k)$. In cases where obtaining such action labels is more complex, an inverse dynamics model can be of assistance. Equation (1) restates this generally

$$\mathbf{a}_k = g(\mathbf{x}_k, \mathbf{x}_{k+1}) \quad (1)$$

where g is the inverse function to compute action labels between two successive states of a demonstration.

We then utilize these demonstrations via learning an imitation policy to bootstrap reinforcement learning. To retain the benefits from any generalization properties achieved via imitation learning, we pre-train the critic network on rollouts of the imitation policy to avoid washing it out with a randomly initialized critic (Hansen et al., 2022). This method aligns

Fig. 2 The object shapes for which we learn finger-gaiting. From left to right: the easy, medium and hard categories



with contemporary approaches for integrating reinforcement learning with imitation learning which go beyond simply augmenting the online replay buffer with demonstrations Hu et al. (2023). Our main approach, supplemented with imitation pre-training, is depicted in Algorithm 3. Lines 16 and 18 show imitation pre-training with sampling-based demonstrations and warm-starting reinforcement learning.

6 Experimental setup and tasks

6.1 Hardware

We use the robot hand shown in Fig. 1, consisting of five identical fingers. Each finger comprises a roll joint and two flexion joints for a total of 15 fully actuated position-controlled joints. For the real hardware setup, each joint is powered by a Dynamixel XM430-210T servo motor. The distal link of each finger consists of an optics-based tactile fingertip as introduced by Piacenza et al. (2020).

We test our methods on the object shapes illustrated in Fig. 2. We split this into categories: “easy” objects (sphere, cube, cylinder), “moderate” objects (cuboids with elongated aspect ratios), and “hard” objects (either concave L- or U-shapes) and train object-specific policies for each of these objects with the proposed methods and also the baseline methods to be discussed later. We note that in-hand manipulation of the objects in the “hard” category has not been previously demonstrated in the literature.

6.2 Exploration trees

The extensive sampling of possible actions, which is the main computational expense of G-RRT (line 5) lends itself well to parallelization. In practice, we use the IsaacGym Makoviy-chuk et al. (2021) simulator to parallelize this algorithm at both these levels (16 parallel evaluations of the main loop, and 1024 parallel evaluations of the action sampling loop). In general, given the advent of increasingly more powerful parallel architectures for physics simulation, we expect that more general methods that are easier to parallelize might win

out in the long term over more problem-specific solutions that are more sample efficient at the individual thread level.

6.3 Reinforcement learning

We train our policies using Asymmetric Actor Critic PPO Pinto et al. (2017); Schulman et al. (2017) where we separate actor and critic training for improved performance. All training is done in the IsaacGym simulator and the critic uses object pose p , object velocity \dot{p} , and net contact force on each fingertip $t_1 \dots t_m$ as feedback in addition to the feedback already provided to the policy network.

For Imitation Pre-training, we compute action labels from state transitions using $\beta = 2$ in Eq. (1). We train the critic with rollouts of the imitation policy for 2M steps to mitigate forgetting issues inherent in a randomly initialized critic network.

6.4 Tasks

We evaluate our method on a variety of challenging dexterous in-hand manipulation tasks.

6.4.1 Finger-gaiting

First, we focus on the task of achieving large in-hand object rotation about a desired axis. We, as others before Qi et al. (2022), believe this to be representative of this general class of problems, since it requires extensive finger gaiting and object reorientation.

We chose to focus on the case where the only sensing available is hand-centric, either tactile or proprioceptive. Achieving dexterity with only proprioceptive sensing, as biological organisms are clearly capable of, can lead to skills that are robust to occlusion and lighting and can operate in very constrained settings. With this directional goal in mind, the observation available to our policy consists of tactile and proprioceptive data collected by the hand, and no global object pose information. Formally, the observation vector is

$$o_t = [q_t, q_t^s, c_t] \quad (2)$$

where $\mathbf{q}_t, \mathbf{q}_t^s \in \mathcal{R}^d$ are the current positions and setpoints of the joints, and $\mathbf{c}_t \in [0, 1]^m$ is the vector representing binary (contact / no-contact) touch feedback for each of m fingers of the hand.

In this task, where our goal is finger-gaiting for z-axis object rotation, we plan trees where object rotation around the x- and y-axes was restricted to 0.2 radians. Then we select 2×10^4 nodes from the paths that exhibit the most rotation around the z-axis to construct our reset distribution and for imitation pre-training. On average, each such path comprises 20–30 nodes. We recall that all tree nodes are subjected to an explicit stability check under full system dynamics before being added to the tree; we can thus use each of them as is.

Similar to our previous work Khandate et al. (2022), we use a reward function that rewards object angular velocity about the z-axis. In addition, we include penalties for the object's translational velocity and its deviation from the initial position Qi et al. (2022).

6.4.2 Arbitrary reorientation

While finger-gaiting to achieve maximum rotation is a good proof-of-concept task with several applications, we now consider the canonically studied and versatile in-hand manipulation skill of reorientation to a desired pose. We continue with the requirement to achieve in-hand manipulation with only stable fingertip grasps as it enables reorientation in arbitrary orientations of the hand.

We take on the challenging task of orienting an object from a randomly assigned initial pose to a desired arbitrary orientation, a frequently studied task in dexterous in-hand manipulation Chen et al. (2021), Yuan et al. (2023). In addition to proprioceptive and tactile data collected by the hand, we include current and desired object orientations as inputs to the policy resulting in the following observation vector,

$$\mathbf{o}_t = [\mathbf{q}_t, \mathbf{q}_t^s, \mathbf{c}_t, \boldsymbol{\phi}_t, \boldsymbol{\phi}_g] \quad (3)$$

where $\boldsymbol{\phi}_t, \boldsymbol{\phi}_g \in \mathcal{R}^4$ are the current and desired object orientation respectively. Our reward function is a modification of the one proposed in Chen et al. (2021) as we found the original reward function fails in our setup. Equation 4 describes the reward function we use for this task,

$$r_t = c \cdot \max(\min(\Delta_t - \Delta_{t-1}, \epsilon), -\epsilon) + c_{\text{success}} \cdot \mathbb{1}[\text{success}] \quad (4)$$

where Δ_t is angular distance between current and desired object orientation, $\epsilon > 0$ is a threshold coefficient, and $c < 0$ is a scaling coefficient. A large positive reward, denoted c_{success} , is added if the agent successfully completes the task. Task completion is achieved through successful reorientation

plus satisfaction of heuristics adapted from previous work Chen et al. (2022). The criteria are described here:

1. Reorientation criterion: $\Delta_t < \theta_{\text{thresh}}$
2. Joint angular velocity criterion: $\|\dot{\mathbf{q}}_t\|_2 < \dot{q}_{\text{thresh}}$
3. Object linear velocity criterion: $\|\dot{\mathbf{x}}_t\|_2 < \dot{x}_{\text{thresh}}$
4. Object angular velocity criterion: $\|\boldsymbol{\omega}_t\|_2 < \omega_{\text{thresh}}$

Unlike the previous task where a subset of the tree needs to be extracted to aid exploration, here the full tree i.e all nodes of the tree can be used towards assembling the buffer of exploratory reset states. To ensure complete exploration of the state space we use G-RRT trees generated without any constraints on object orientation.

6.4.3 Go-to-root

While our method improves on sample efficiency for learning policies for arbitrary reorientation, obtaining such policies remains computationally expensive due to the requirement of a few tens of billion simulation training steps. Learning such policies can be overkill; often, in practice, it is sufficient to reorient the object to a fixed canonical pose starting from an arbitrary initial pose. In a packaging line, for example, items arrive at arbitrary orientations. The essential in-hand manipulation task is to reorient objects to a canonical pose before inserting them in the packaging container.

However, learning to reorient the object to a fixed desired orientation starting from an arbitrary initial pose is still difficult as it suffers the same exploration challenges seen in previous tasks. Fortunately, our method of using exploration trees is naturally well suited for this problem. The robot state with the desired canonical orientation can itself be used as the root node while generating the exploration tree. Thus, we refer to this as the ‘‘Go-to-root’’ task.

In this task, we aim to learn a policy to reorient the grasped object to reach a desired canonical / root orientation. The inputs to the policy network are similar to the arbitrary reorientation task, consisting of proprioceptive, tactile, and current object pose feedback but excluding the desired goal orientation as it is fixed, resulting in the observation vector in Eq. (5). The reward function and success criteria are identical to the arbitrary reorientation task except with fixed canonical orientation $\boldsymbol{\phi}_g$ as the goal.

$$\mathbf{o}_t = [\mathbf{q}_t, \mathbf{q}_t^s, \mathbf{c}_t, \boldsymbol{\phi}_t] \quad (5)$$

For imitation pre-training, we extract plans by backtracking from randomly selected nodes with large displacement from the root. In each task, this amounts about about 3K unique observation-action pairs for finger-gaiting and about 15k unique observation-action pairs for go-to-root task.

6.5 Algorithms and baselines

In our experiments, we compare the following approaches:

6.5.1 Ours, $R \times R$

In this variant, we use the method presented in this paper, relying on exploratory reset states obtained by growing the tree via G-RRT. In all cases, we use a tree comprising 10^5 nodes as informed the ablation study in Fig 14.

6.5.2 Ours, $R \times R + IPT$

In this variant of our method, we use Imitation Pre-training with labels obtained from the grown G-RRT tree, alongside using exploratory reset distribution extracted from the same tree.

6.5.3 Stable Grasp Sampler (SGS)

This baseline represents an alternative to the method presented in this paper: we use a reset distribution consisting of stable grasps generated by sampling random joint angles and object orientations. This approach has been effective in demonstrating precision in-hand manipulation with only intrinsic sensing Khandate et al. (2022); Qi et al. (2022) for simple shapes.

6.5.4 Explored Restarts (ER)

This method selects states explored by the policy itself during random exploration to use as reset states Tavakoli et al. (2018). It is highly general, with no manipulation-specific component, and requires no additional step on top of RL training. We implement the "uniform restart" scheme as it was shown to have superior performance on high dimensional continuous control tasks. However, we have found it to be insufficient for the complex state space of our problem: it fails to learn a viable policy even for simple objects.

6.5.5 Fixed Initialization (FI)

For completeness, we also tried resetting from a single fixed state.

6.5.6 Fixed Initialization (FI + IPT)

Towards understanding the effect of pre-training, we also tried FI baseline but with a warm start policy obtained from Imitation Pre-training using the G-RRT tree.

6.5.7 Gravity Curriculum (GC)

Additionally, we evaluated Fixed Initialization with gravity curriculum, linearly annealed from $0 \frac{m}{s^2}$ to $-9.81 \frac{m}{s^2}$ over the course of 50M timesteps.

7 Results

7.1 Evaluation in simulation

7.1.1 Finger-gaiting

Our training results aggregated over 3 seeds are summarized in Figs. 3 and 5. We find that our methods ($R \times R$, $R \times R + IPT$) converge within 50M timesteps across all objects for the Finger-gaiting task, consistently outperforming baselines.

On the easy object set, we find that all the methods except the ER baseline learn to gait. However, policies that use a random grasp reset distribution (SGS) and sampling-based reset distribution ($R \times R$) achieve higher performance with robust gaits. We also observe that warm starting with Imitation Pre-training ($R \times R + IPT$) results in faster convergence. Unlike our previous studies, we find that FI also learns finger-gaiting. We attribute this to the improved implementation of our asymmetric actor-critic PPO and increased training steps made possible with the help of GPU physics.

However, as we move to more difficult objects in the medium and hard categories, the performance gap between our methods and baselines broadens. For medium difficulty objects, we find that $R \times R$, $R \times R + IPT$, and SGS again all learn to gait, but the policies learned via our methods, $R \times R$ and $R \times R + IPT$, are more effective with higher returns. Again, we observe that warm starting with Imitation Pre-training ($R \times R + IPT$) results in faster convergence and also achieves improved final performance.

Finally, for the hardest object set, a random grasp-based reset distribution is no longer successful. Only $R \times R$ and $R \times R + IPT$ converge to stable and consistent gaits. Once again, $R \times R + IPT$ improves on $R \times R$ with respect to rate of convergence and final performance.

Interestingly, FI + IPT baseline that uses a warm start policy (obtained from Imitation Pre-training using the G-RRT tree) performs similarly to FI baseline, failing on all but easy objects. This suggests that exploratory reset distributions may be necessary to derive benefits from Imitation Pre-training.

To verify the scalability of the method to train policies for multiple objects, we successfully train object-agnostic policies with our $R \times R$ approach. To achieve this we simply use object-specific initial distributions in multi-object training. We found the training difficulty to be limited to the hardest




Fig. 3 Training performance of our methods and a number of baselines for the Finger-gaiting task on the object categories shown in Fig. 2

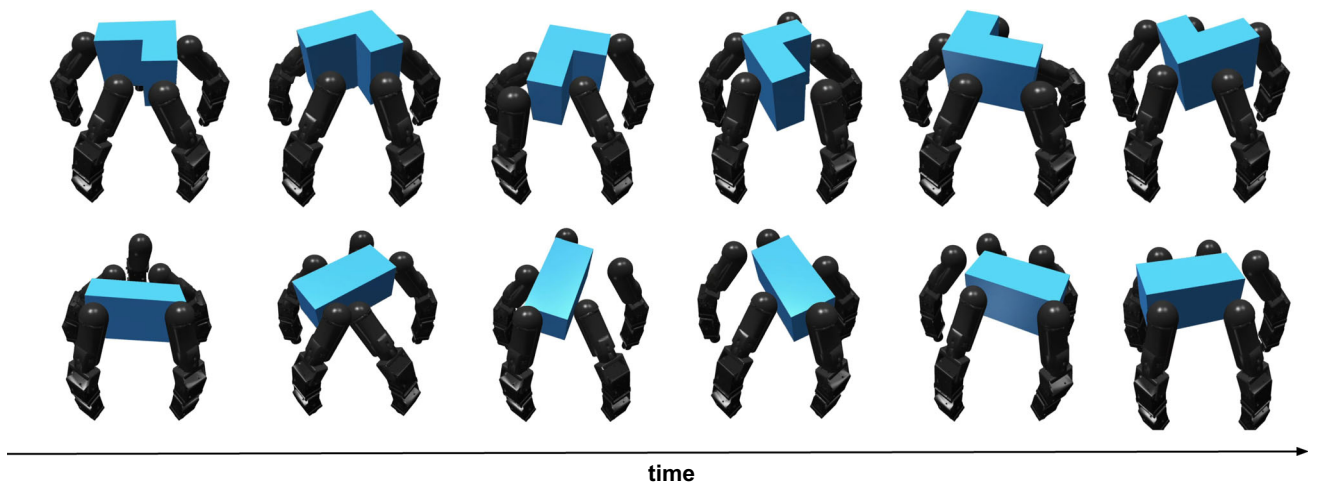


Fig. 4 Key frames of the policies for the Finger-gaiting achieved with our method $R \times R$ for representative objects in simulation

object in the set. Importantly, in Sec 7.2, we also verify our policies can be transferred to the real hand.

Next, we present our results in learning policies to reach a desired object orientation i.e. go-to-root and arbitrary reorientation tasks. Additionally, due to the poor performance of many baselines in the previous task, we continue with the best-performing ones. In particular, we continue with SGS and GC baselines. We drop the ER baseline as it uniformly fails across object classes and also drop the FI baseline as fixed initialization is incompatible with both tasks by definition.

7.1.2 Go-to-root

We evaluate $R \times R$ and $R \times R + IPT$ on the Go-to-root task and compare it with SGS and GC baselines. For this task we augment the Gravity Curriculum baseline to use 20 hand-crafted grasps as the reset distribution. These grasps are constructed to appropriately cover the orientation space of the object. We do this to ensure fair comparison between methods, as learn-

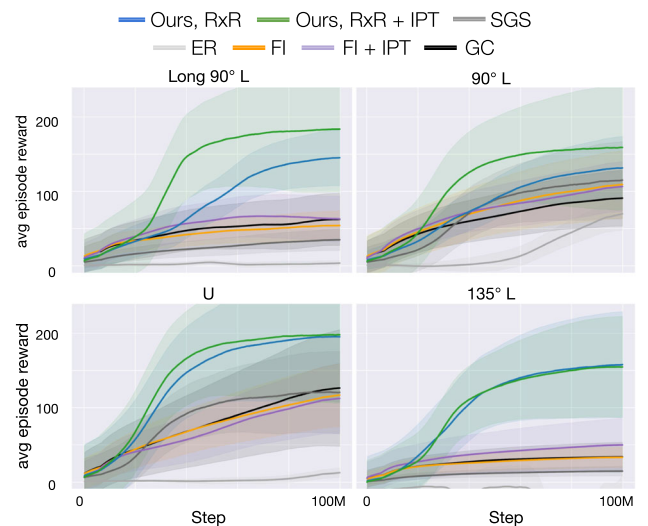


Fig. 5 Training performance of our methods and a number of baselines for the Finger-gaiting task on hard objects shown in Fig. 2




Fig. 6 Training performance for the Go-to-root task with our methods and other best-performing baselines




Fig. 7 Training performance for the Arbitrary Reorientation task comparing our methods (R×R, R×R + IPT) with Stable Grasp Sampler (SGS) baseline

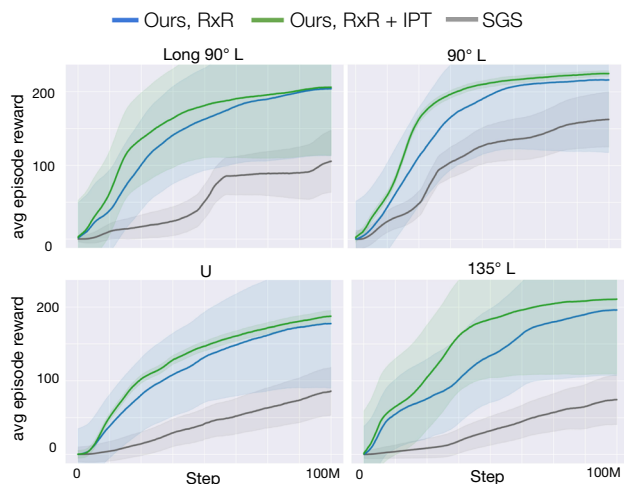


Fig. 8 Training performance for the Go-to-root task with our methods and other best-performing baselines for objects in hard category

ing reorientation from a fixed initial orientation to a fixed target is a significantly easier task.

Our results are summarized in Figs. 6 and 8. As shown by the curves, go-to-root is a harder task requiring between

100 and 200M steps. However, these results indicate that our methods consistently outperform baseline methods and produce effective control policies for the task. Again we find that R×R + IPT improves on R×R in both convergence and final performance across all object sets.

7.1.3 Arbitrary reorientation

We evaluate R×R on the Arbitrary Reorientation task. Our results in Fig. 7 plot the success rate. Note that we considered the hand in "palm-down" orientation as it tends to be the desired hand orientation in many applications. We demonstrate that our method outperforms SGS on all objects. Interestingly, our method significantly outperforms SGS baseline even for objects in the easy category, with increasing performance for objects in the medium and hard categories. The keyframes of our policy executing this task are shown in Fig. 9.

Overall, a common thread of all experiments presented here is that our methods (R×R, R×R + IPT) enable learning a range of challenging dexterous manipulation tasks, while




Fig. 9 Key frames of goal-reaching tasks, Go-to-root and Arbitrary Reorientation, in simulation

Table 1 Manipulation performance in real hardware

	Median Revolutions	Mean rotation Speed (rad/s)
Cylinder	5	0.42
Cube (s)	4.5	0.44
Cuboid	1.5	0.44
L-shape	1.5	0.24

We report the median number of object rotations achieved before dropping the object in ten consecutive trials, as well as the time needed to perform these rotations

none of the domain-agnostic methods (ER, FI, GC) are able to learn in-hand manipulation on objects beyond the easy set.

7.2 Evaluation on real hand

To test the applicability of our method on real hardware, we attempted to transfer the learned Finger-gaiting policies for a subset of representative objects: cylinder, cube, cuboid & L-shape. We chose these objects to span the range from simpler to more difficult manipulation skills.

We achieve transfer of the reinforcement learning policies to the real-hardware via domain randomization introduced via a curriculum. We adapt the policies achieved with nominal system parameters i.e. the system parameters used during prior tree generation and reinforcement learning phases with an additional fine-tuning phase where we introduce domain randomization of relevant system parameters. A thorough list of such parameters is detailed next. It is important to note that while randomizing these parameters, the exploratory tree generated with nominal parameters still remains valid as these parameter perturbations only mildly affected the stability of states. Hence, it is sufficient to use states obtained from this tree generated with with nominal system parameters.

First, we impose velocity and torque limits in the simulation, mirroring those used on the real motors (0.6 rad/s and 0.5 N·m, respectively). We found that our hardware has a significant latency of 0.05s, which we included in the simulation. In addition, we modified the angular velocity reward to maintain a desired velocity instead of maximizing the object's angular velocity. We also randomize joint origins (0.1 rad), friction coefficient (1–40), and train with perturbation forces (1 N). All these changes are introduced successively via a curriculum.

For observation, we used the current position and setpoint from the motor controllers with no additional changes. For tactile data, we found that information from our tactile fingers is most reliable for contact forces above 1 N. We thus did not use reported contact data below this threshold and imposed a similar cutoff in simulation. Overall, we believe that a key advantage of exclusively using proprioceptive data is a smaller sim-to-real gap compared to extrinsic sensors such as cameras. We note that an ablation study illustrating the importance of touch feedback is presented in Sect 8.

For the set of representative objects, we ran the respective policy ten consecutive times and counted the number of successful complete object revolutions achieved before a drop. In other words, five revolutions means the policy successfully rotated the object for 1,800° before dropping it. In addition, we also report the average object rotation speed observed during the trials.

The results of these trials are summarized in Table 1. Figures 10 and 11 show the keyframes of the real hand finger-gaiting we achieved with our policies. Finally, as our finger-gaiting policies do not rely upon vision feedback, our policies are robust to changes in lighting conditions. Examples of our method operating in dynamic lighting conditions can be found in the accompanying video.




Fig. 10 Key frames of the Finger-gaiting policies transferred to the real hand




Fig. 11 Key frames of a Finger-gaiting policy transferred to the hand in the “palm-down” orientation

We note that the starting policies used for sim-to-real transfer were trained with trees generated by a version of G-RRT that used the constraint that at least three fingers must be in contact with the object. However, as previously mentioned, we found that this constraint has minimal effect on the trees generated via G-RRT. Nevertheless, we expect similar sim-to-real performance of policies trained using trees generated after forgoing the three contact constraint. These policies are identical to the policies used for sim-to-real transfer, as per visual comparison in simulation. We also note that prelimi-

nary attempts to transfer policies for the cylinder and cube also show similar performance.

8 Ablations

8.1 Sampling-based exploration

First, we conduct ablation of G-RRT for our object set. We aim to discover how effectively the tree explores its available state space given the number of iterations through the main

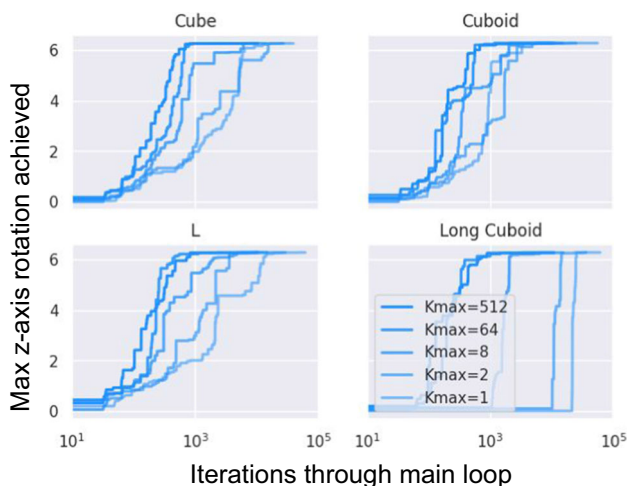


Fig. 12 Tree expansion performance for G-RRT. We plot the number of attempted tree expansions (i.e. iterations through the main loop, on a log scale) against the maximum object rotation achieved by any tree node so far. We plot performance for different values of K_{max} , the number of random actions tested at each iteration

loop (i.e. the number of attempted tree expansions towards a random sample). As a measure of tree growth, we look at the maximum object rotation achieved around our target axis. We note that any rotation beyond approximately $\pi/4$ radians can not be done in-grasp, and requires finger repositioning. Thus, we compare the maximum achieved object rotation vs. the number of expansions attempted (on a log scale). The results are shown in Fig. 12.

As expected, the performance of G-RRT improves with the number K_{max} of actions tested at each iteration. Interestingly, the algorithm performs well even with $K_{max} = 1$; this is equivalent to a tree node growing in a completely random direction, without any bias towards the intended sample. However, we note that, at each iteration, the node that grows is the closest to the state-space sample taken at the beginning of the loop. This encourages growth at the periphery of the tree and along the real constraint manifolds, and, as shown here, still leads to effective exploration.

We also found that G-RRT is sensitive to the action-scale parameter α . Figure 13 compares exploration as measured by angular distance of the farthest node from the root for varying α . Among the various settings evaluated, an $\alpha = 0.15$ was the fastest. Interestingly, higher values of action-scale adversely affects the performance of G-RRT. Furthermore, as we will discuss later, not only does α affect the rate of exploration it also impacts the quality of extracted paths.

8.2 Reinforcement learning

We now consider the ablation of training methods. We perform these ablations on the Finger-gaiting and Go-to-root tasks on a non-convex, L-shaped object.

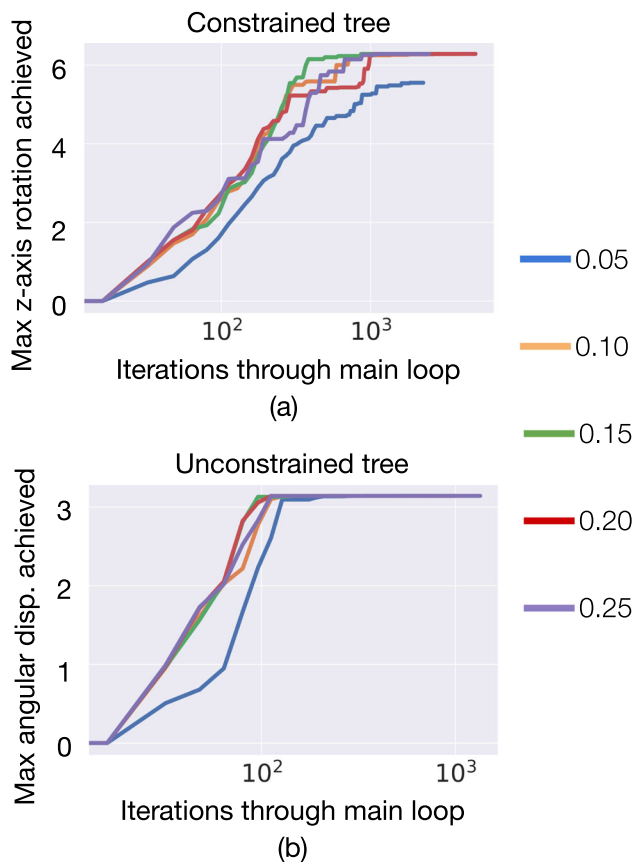


Fig. 13 G-RRT action-scale (α) ablation. $\alpha = 0.15$ is optimal

We begin by studying the impact of the size of the tree used in extracting reset states. Figure 14a summarizes our results for learning a Finger-gaiting policy using trees of different sizes grown via G-RRT. Qualitatively, we observe that, as the tree grows larger, the top 100–400 paths sampled from the tree contain increasingly more effective gaits, likely closer to the optimal policy. We see that we need a sufficiently large tree with at least 10^4 nodes to enable learning. However, training is most reliable with 10^5 nodes. This suggests a strong correlation between the quality of states used for reset distribution and sample efficiency of learning.

Figure 14b summarizes the tree-size ablation results for learning the Go-to-root task. For a very small tree, we see near 100% success rate at training time. This is rather unsurprising as, at these small tree sizes, the nodes are still close to the root orientation resulting in a trivial reorientation task. As we expect, validation performance on paths extracted from the largest tree increases as the size of the tree generating the initial state distribution increases.

We also reconsider the action-scale parameter. Besides affecting the rate of exploration it also impacts the quality of paths. We can infer this by the impact on training performance in Fig. 15. We see a similar trend, as the training performance




Fig. 14 The final performance of policies for Finger-gaiting (left) and Go-to-root (right) tasks using trees of increasing sizes. Validation performance is with reset states from paths extracted from the largest tree




Fig. 15 Action scale ablation for the Finger-gaiting and Go-to-root tasks with L

peaks again around $\alpha = 0.15$ for both Finger-gaiting and Go-to-root.

We also conducted an ablation study of policy feedback. Particularly, we aimed to compare intrinsically available tactile feedback vs. object pose feedback that would require external sensing. Figure 16 summarizes these results which demonstrate the importance of tactile feedback for both Finger-gaiting and Go-to-root tasks.

In the Finger-gaiting task, we found that touch feedback is essential in the absence of object pose feedback for all moderate and hard objects. For these objects, we also saw that replacing this tactile feedback with object pose feedback results in slower learning. Similarly, in the Go-to-root task, leaving out touch feedback from policy input also results in slower learning.

These results underscore the importance of touch feedback for in-hand manipulation skills. Richer tactile feedback such as contact position, normals, and force magnitude can be expected to provide even stronger improvements; we hope to explore this in future work.

9 Discussion and conclusions

The results we have presented show that sampling-based exploration methods make it possible to achieve difficult manipulation tasks via RL. In fact, these popular and widely used classes of algorithms are highly complementary in this case. RL is effective at learning closed-loop control policies

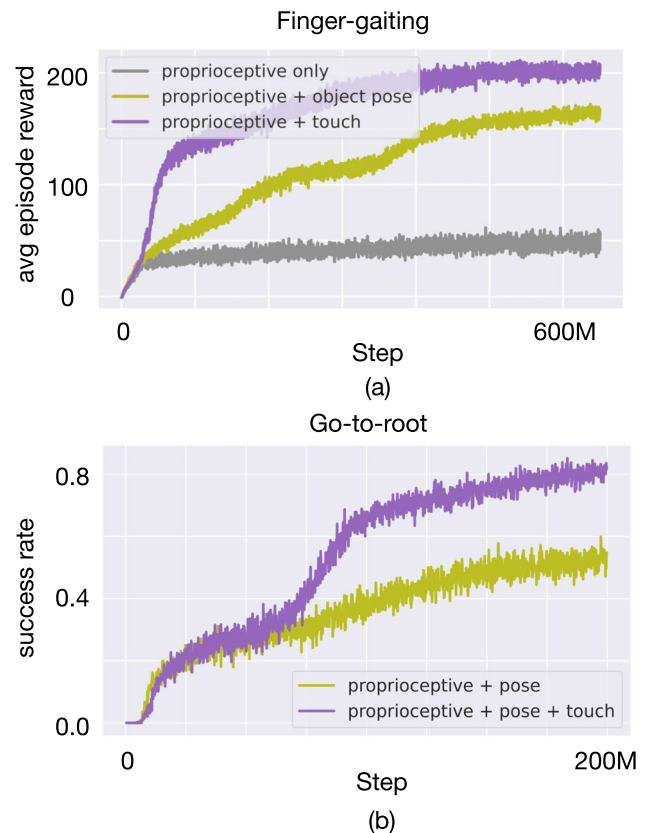


Fig. 16 Ablation of policy feedback components which highlights the importance of touch feedback

that maintain the local stability needed for manipulation, and, thanks to training on a large number of examples, are robust to variations in the encountered states. However, the standard RL exploration techniques (random perturbations in action space) are ineffective in the highly constrained state space with a complex manifold structure of manipulation tasks. Conversely, SBP methods, which rely on a fundamentally different approach to exploration, can effectively discover relevant regions of the state space.

We present multiple methods for conveying exploration information derived through SBP to RL training algorithms. In particular, explored states can be used as an effective reset distribution to enable learning. The transitions between states used during sample-based exploration are also useful surrogates for actions, and can thus be used in an imitation pre-training stage to boost learning. Imitation pre-training serves as an effective strategy to leverage the inherent structure accessible through sampling-based plans. This approach becomes particularly valuable in addressing the challenges of hard motor control tasks within reinforcement learning, not only by facilitating training but also by providing a favorable initialization of actor and critic networks. Recent work in this domain extends the use of the imitation policy beyond mere initialization by demonstrating further benefits when incor-

porated throughout the training process. This is a promising avenue for achieving further performance gains in reinforcement learning scenarios for hard motor control tasks.

We use our approach to demonstrate a number of dexterous manipulation skills: achieving large in-hand reorientation about a given axis via finger-gaiting and reorientation to a desired object pose, which is either pre-set or randomized. Importantly, we demonstrate finger gaiting precision manipulation of both convex and non-convex objects, using only tactile and proprioceptive sensing. Using only these types of intrinsic sensors makes manipulation skills insensitive to occlusion, illumination, or distractors, and reduces the sim-to-real gap. We take advantage of this by demonstrating our approach both in simulation and on real hardware. We note that, while some applications naturally preclude the use of vision (e.g. extracting an object from a bag), we expect that in many real-life situations, future robotic manipulators will achieve the best performance by combining touch, proprioception, and vision.

In future work, we believe that our approach can be scaled to tackle even more demanding motor control tasks with numerous degrees of freedom, including both prehensile and non-prehensile manipulation. Bi-manual manipulation seems like a natural candidate application, particularly if it involves coordination with multi-fingered dexterous hands. We also believe our approach can be applied beyond dexterous manipulation, such as in learning challenging agile locomotion skills with minimal changes. Towards this aim, we suggest expanding the sampling configuration space to the complete state of the robot with first order derivatives and also replacing the stability check with necessary boundary constraints on the state. With these changes in place, we believe it is a feasible recipe to accelerate learning an end-to-end policy for learning agile skills such as robot parkour Zhuang et al. (2023). Finally, sampling-based exploration techniques could be integrated directly into the policy training mechanisms, removing the need for two separate stages during training. We hope to explore all these ideas in the future.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10514-024-10170-8>.

Acknowledgements We would like to thank all the members of the Robotic Manipulation and Mobility Lab at Columbia University for helpful discussions and feedback on early versions of the manuscript.

Author Contributions Gagan Khandate led this research by conceptualizing the method, validating early versions, and developing the manuscript. Tristan L. Saidi contributed policies, new tasks, extensive ablations, and results on imitation pre-training and also assisted in writing the manuscript. Siqi Shang trained policies in simulation and policies for sim-to-real transfer for the conference version of the submission. Eric Chang trained contact detection models for our optics-based tactile sensor. Johnson Adams assisted in maintaining the robot

firmware and sim-to-real transfer. Yang Liu contributed to early versions of the hardware. Seth Dennis designed and fabricated the hardware used in our experiments. Matei Theodor Ciocarlie provided excellent guidance and mentorship throughout all the stages of this research and assisted with the manuscript.

Funding This work was supported in part by the Office of Naval Research Grant N00014-21-1-4010 and the National Science Foundation grant CMMI-2037101.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose. The authors have no Conflict of interest to declare that are relevant to the content of this article. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no financial or proprietary interests in any material discussed in this article.

References

- Agarwal, A., Kakade, S. M., Lee, J. D., & Mahajan, G. (2020). Optimality and approximation with policy gradient methods in markov decision processes. In Jacob, A. & Shivani A. (eds.), *Proceedings of thirty third conference on learning theory*, vol. 125 of *Proceedings of machine learning research*, pp. 64–66. PMLR.
- Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., & Zhang, L. (2019). Solving rubik’s cube with a robot hand
- Allshire, A., Mittal, M., Lodaya, V., Makoviychuk, V., Makoviichuk, D., Widmaier, F., Wüthrich, M., Bauer, S., Handa, A., & Garg, A. (2021). Transferring dexterous manipulation from GPU simulation to a remote real-world TriFinger.
- Amin, S., Gomrokchi, M., Satija, H., van Hoof, H., & Precup, D. (2021). A survey of exploration methods in reinforcement learning. [arXiv:2109.00157](https://arxiv.org/abs/2109.00157)
- Andrychowicz, O. M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., & Schneider, J. (2020). Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1), 3–20.
- Bhatt, A., Sieler, A., Puhlmann, S., & Brock, O. (2022). Surprisingly robust in-hand manipulation: An empirical study.
- Chen, T., Tippur, M., Wu, S., Kumar, V., Adelson, E. & Agrawal, P. (2022). Visual dexterity: In-hand dexterous manipulation from depth.
- Chen, T., Xu, J., & Agrawal, P. (2021). A system for general in-hand object re-orientation
- Chiang, H. T. L., Hsu, J., Fiser, M., Tapia, L., & Faust, A. (2019). RL-RRT: Kinodynamic motion planning via learning reachability estimators from RL policies. *IEEE Robotics and Automation Letters*, 4(4), 4298–4305.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., & Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control.
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2019). Go-explore: a new approach for hard-exploration problems.
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2021). First return, then explore. *Nature*, 590(7847), 580–586.

- Francis, A., Faust, A., Chiang, H. T. L., Hsu, J., Kew, J. C., Fiser, M., & Lee, T. W. E. (2020). Long-range indoor navigation with PRM-RL. *IEEE Transactions on Robotics*, 36(4), 1115–1134.
- Ha, H., Xu, J., & Song, S. (2020). Learning a decentralized multi-arm motion planner.
- Haarnoja, T., Zhou, A., & Abbeel, P. (2018). Soft actor-critic: Off-Policy maximum entropy deep reinforcement learning with a stochastic actor.
- Han, L., & Trinkle, J. C. (1998). Dextrous manipulation by rolling and finger gaiting. In *Proceedings 1998 IEEE international conference on robotics and automation (Cat. No. 98CH36146)*, vol. 1, pp. 730–735.
- Handa, A., Allshire, A., Makoviychuk, V., Petrenko, A., Singh, R., Liu, J., Makoviichuk, D., Van Wyk, K., Zhurkevich, A., Sundaralingam, B., & Narang, Y. (2022). DeXtreme: Transfer of agile in-hand manipulation from simulation to reality.
- Hansen, N., Lin, Y., Hao, S., Wang, X., Kumar, V. & Aravind, R. (2022). MoDem: Accelerating visual model-based reinforcement learning with demonstrations.
- Hu, H., Mirchandani, S., & Sadigh, D. (2023). Imitation bootstrapped reinforcement learning.
- Jurgenson, T., & Tamar, A. (2019). Harnessing reinforcement learning for neural motion planning.
- Karaman, S., & Frazzoli, E. (2010). Optimal kinodynamic motion planning using incremental sampling-based methods. In *49th IEEE conference on decision and control (CDC)*, pp. 7681–7687.
- Kavraki, L. E., Kolountzakis, M. N., & Latombe, J.-C. (1998). Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation*, 14(1), 166–171.
- Kavraki, L. E., Svestka, P., Latombe, J.-C., & Overmars, M. H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4), 566–580.
- Khandate, G., Haas-Heger, M., Ciocarlie, M. (2022). On the feasibility of learning finger-gaiting in-hand manipulation with intrinsic sensing. In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 2752–2758
- Khandate, S., Shang, S., Chang, E. T., Saidi, T. L., Liu, Y., Dennis, S. M., Adams, J., & Ciocarlie, M. (2023). Sampling-based exploration for reinforcement learning of dexterous manipulation.
- King, J. E., Cagnetti, M., & Srinivasa, S. S. (2016). Rearrangement planning using object-centric and robot-centric action spaces. In *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE
- LaValle, S. (1998). Rapidly-exploring random trees : A new tool for path planning. The annual research report.
- Leveroni, S., & Salisbury, K. (1996). Reorienting objects with a robot hand using grasp gaits. In *Robotics research*, pp 39–51. Springer London.
- Li, L., Miao, Y., Qureshi, A. H., & Yip, M. C. (2021). MPC-MPNet: Model-predictive motion planning networks for fast, near-optimal planning under kinodynamic constraints. *IEEE Robotics and Automation Letters*, 6(3), 4496–4503.
- Ma, R. R., & Dollar, A. M. (2011). On dexterity and dexterous manipulation. In *2011 15th International conference on advanced robotics (ICAR)*, pp. 1–7.
- Makoviychuk, V., Wawrzyniak, L., Guo, Y., Michelle, L., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., & Gavriel S. (2021). Isaac gym: High performance GPU-based physics simulation for robot learning.
- Morere, P., Francis, G., Blau, T., & Ramos, F. (2020). Reinforcement learning with probabilistically complete exploration.
- Morgan, A.S., Nandha, D., Chalvatzaki, G., D'Eramo, C., Dollar, A. M., & Peters, J. (2021). Model predictive actor-critic: Accelerating robot skill acquisition with deep reinforcement learning.
- Morgan, A. S., Hang, K., Wen, B., Bekris, K., & Dollar, A. M. (2022). Complex in-hand manipulation via compliance-enabled finger gaiting and multi-modal planning. *IEEE Robotics and Automation Letters*, 7(2), 4821–4828.
- Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2017). Overcoming exploration in reinforcement learning with demonstrations.
- Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction.
- Piacenza, P., Behrman, K., Schifferer, B., Kymissis, I., & Ciocarlie, M. (2020). A sensorized multicurved robot finger with Data-Driven touch sensing via overlapping light signals. *IEEE/ASME Transactions on Mechatronics*, 25(5), 2416–2427.
- Pinto, L., Andrychowicz, M., Welinder, P., Zaremba, W., & Abbeel, P. (2017). Asymmetric actor critic for image based robot learning.
- Pinto, L., Mandalika, A., Hou, B., & Srinivasa, S. (2018). Sample-efficient learning of nonprehensile manipulation policies via physics-based informed state distributions.
- Pitz, J., Röstel, L., Sievers, L., & Bäuml, B. (2023). Dextrous tactile in-hand manipulation using a modular reinforcement learning architecture.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., & Andrychowicz, M. (2017). Parameter space noise for exploration.
- Qi, H., Kumar, A., Calandra, R., Ma, Y., & Malik, J. (2022). In-hand object rotation via rapid motor adaptation.
- Qi, H., Yi, B., Suresh, S., Lambeta, M., Ma, Y., Calandra, R., & Malik, J. (2023). General in-hand object rotation with vision and touch. CoRL. [arXiv:2309.09979](https://arxiv.org/abs/2309.09979)
- Röstel, L., Pitz, J., Sievers, L., & Bäuml, B. (2023). Estimator-coupled reinforcement learning for robust purely tactile in-hand manipulation.
- Schramm, L., & Boularias, A. (2022). Learning-guided exploration for efficient sampling-based motion planning in high dimensions. In *2022 International conference on robotics and automation (ICRA)*. IEEE.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms.
- Sievers, L., Pitz, J., & Bäuml, B. (2022). Learning purely tactile In-Hand manipulation with a Torque-Controlled hand. In *2022 International conference on robotics and automation (ICRA)*, pp. 2745–2751.
- Tavakoli, A., Levdiq, V., Islam, R., Smith, C. M. & Kormushev, P. (2018). Exploring restart distributions.
- Webb, DJ. & van den Berg, J. (2013). Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In *2013 IEEE international conference on robotics and automation*, pp 5054–5061.
- Xu, J., Koo, TJ. & Li, Z. (2007). Finger gaits planning for multifingered manipulation. In *2007 IEEE/RSJ international conference on intelligent robots and systems*, pp. 2932–2937.
- Yashima, M, Shiina, Y, & Yamaguchi, H. (2003). Randomized manipulation planning for a multi-fingered hand by switching contact modes. In *2003 IEEE international conference on robotics and automation (Cat. No. 03CH37422)*, vol. 2, pp. 2689–2694.
- Yin, ZH., Huang, B., Qin, Y., Chen, Q. & Wang, X. (2023). Rotating without seeing: Towards in-hand dexterity through touch.
- Yuan, Y., Che, H., Qin, Y., Huang, B., Yin, ZH., Lee, KW., Yi, W., Lim, SC. & Wang, X. (2023). Robot synesthesia: In-hand manipulation with visuotactile sensing.
- Zhuang, Z., Fu, Z., Wang, J., Atkeson, C., Schwertfeger, S., Finn, C. & Zhao, H. (2023). Robot parkour learning.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.




Gagan Khandate received his B.Tech from Indian Institute of Technology Madras, India in 2015 and his M.S. from Columbia University in 2018. Currently, he is Ph.D. candidate in Dept. of Computer Science at Columbia University, New York, NY, USA. His research focuses on robot learning for robotic manipulation and complex motors skills with a particular focus on multi-fingered dexterity.



Tristan L. Saidi received his M.S. in Computer Science in 2024 from Columbia University, and is currently a doctoral student in Computer Science at Columbia University. His research focuses on machine learning and applied math, with a particular emphasis on geometric data analysis and manifold learning.



Siqi Shang received his M.S. in Computer Science from Columbia University and his B.S. in Computer Science from Chinese University of Hong Kong. Currently, he is a PhD candidate in Electrical and Computer Engineering at the University of Texas at Austin. His research interests center around advancing the dexterity and versatility of manipulation.




Eric T. Chang is a Ph.D. candidate in Mechanical Engineering at Columbia University. He received his B.S. in Mechanical Engineering and Computer Science at Duke University and his M.S. in Mechanical Engineering at Columbia University. His main research interests lie in tactile sensing and enabling robot dexterity through rich touch feedback.

Yang Liu received his B.S. from Zhejiang University, China in 2017 and his M.S. from Columbia University in 2020. His interests include design, fabrication and mechatronics for robotics.



Seth Dennis received his B.S. in Computer Science and M.S. in Mechanical Engineering from Columbia University. His interests include design, fabrication and system integration of robotic and mechatronic systems with experience in building robotic hands and micro-satellites.



Johnson Adams earned his M.S. in Mechanical Engineering from Columbia University, following his B.Tech from the Indian Institute of Technology, Madras. His interests include system integration for robotics with a focus on robust sensor fusion.



Matei Ciocarlie received the B.S. degree from the Polytechnic University of Bucharest, Bucharest, Romania, in 2003. He received the M.S. and Ph.D. degrees from Columbia University, New York, NY, USA, in 2005 and 2010 respectively. He is an Associate Professor with the Columbia University. His main research interests include reliable robotic performance in unstructured, human environments, looking to discover how artificial mechanisms can interact with the world as skill-

fully as biological organisms.