# ActivePusher: Active Learning and Planning with Residual Physics for Nonprehensile Manipulation

Zhuoyun Zhong Seyedali Golestaneh Constantinos Chamzas Department of Robotics Engineering, Worcester Polytechnic Institute {zzhong3, sgolestaneh, cchamzas}@wpi.edu



Figure 1: Two key challenges addressed by ACTIVEPUSHER. When learning, the robot must choose the most informative (active learning) skill parameters to efficiently improve its skills. When planning, the robot should select skill parameters with low model uncertainty (active planning) to ensure reliable task completion.

**Abstract:** Planning with learned dynamics models offers a promising approach toward real-world, long-horizon manipulation, particularly in nonprehensile settings such as pushing or rolling, where accurate analytical models are difficult to obtain. However, collecting training data for learning-based methods can be costly and inefficient, as it often relies on randomly sampled interactions that are not necessarily the most informative. Furthermore, learned models tend to exhibit high uncertainty in underexplored regions of the skill space, undermining the planning reliability. To address these challenges, we propose ACTIVEPUSHER, a novel framework that combines residual-physics modeling with kernel-based uncertainty-driven active learning to focus data acquisition on the most informative skill parameters. Additionally, ACTIVEPUSHER seamlessly integrates with model-based kinodynamic planners, leveraging uncertainty estimates to bias control sampling toward more reliable actions. We evaluate our approach in both simulation and real-world environments and demonstrate that it improves data efficiency and planning success rates compared to baseline methods.

Keywords: Active Learning, Nonprehensile Manipulation, Kinodynamic Planning

# **1** Introduction

Model-based planning methods offer a powerful framework for generalizing robotic behavior across tasks and enabling long-horizon decision making [1, 2]. However, their effectiveness critically depends on the accuracy of the underlying forward dynamics model. Inaccuracies in this model can cause cascading errors during execution, particularly in contact-rich settings such as nonprehensile manipulation (e.g., pushing, sliding, or rolling), where even minor deviations in predicted trajectories may lead to significant task failure.

Accurately modeling the dynamics for these tasks is challenging. Analytical physics-based models often rely on simplified assumptions about friction, contact geometry, and mass distribution, making

them brittle in practice [3]. As an alternative, data-driven approaches can learn dynamics directly from interaction data, either from scratch or by refining simplified analytical models through residual learning [4, 5]. However, these methods face two key limitations in real-world robotic settings:

- **Sample inefficiency:** Learning accurate models often requires large amounts of interaction data, which is costly and time-consuming to collect on physical systems.
- **Inaccuracy in underexplored regions:** Models may perform poorly in sparsely explored regions of the skill space, leading to unreliable predictions and failures during execution.

In this paper, we propose ACTIVEPUSHER, a framework that tightly integrates residual physics, active learning, and active kinodynamic planning, to address both challenges visually illustrated in Fig. 1. The core insight of our approach is to explicitly quantify epistemic uncertainty in a learned neural network model with residual physics. This uncertainty estimate allows the system to actively target informative skills for learning and reliable skills for planning.

To estimate the epistemic uncertainty of a learned neural network model, we adopt the Neural Tangent Kernel (NTK) [6]. During learning, rather than sampling pushing actions at random, ACTIVEPUSHER actively queries the NTK for uncertainty estimation and identify actions with the highest expected information gain using Batch Active learning via Information maTrices (BAIT) strategy [7]. This targeted exploration enables the model to improve rapidly with far fewer interactions. During planning, the uncertainty estimates are incorporated into a kinodynamic planner, biasing control selection toward high-confidence actions to maximize task success. By focusing on where the model is uncertain to learn, and where the model is certain to act, our approach tightly integrates learning and planning, enabling robust nonprehensile manipulation with few real-world interactions per task. Crucially, ACTIVEPUSHER operates without large offline datasets, high-fidelity simulation, or human demonstrations. Our main contributions are as follows:

- Active learning of skill models. We introduce a principled framework for data-efficient nonprehensile skill learning by selecting skill parameters that maximize expected information gain, enabling data collection in the most informative manner.
- Active uncertainty-aware planning. We propose a novel planning strategy that integrates model uncertainty into an asymptotically optimal kinodynamic planner, guiding action sampling toward reliable actions and improving overall task success rate.
- Empirical validation in simulation and the real world. We demonstrate the effectiveness of our approach with multiple objects and a manipulation task, showing improved data efficiency and planning success over baselines.

# 2 Related Work

ACTIVEPUSHER draws ideas from several areas, such as residual learning, active learning and kinodynamic planning. In this section, we briefly review each of these areas in the context of nonprehensile manipulation, with a focus on pushing.

**Residual Model Learning** aims to combine the strengths of analytical and data-driven approaches by training a neural network to predict corrections on top of an approximate physics-based model. This strategy increases predictive accuracy while leveraging the structure and inductive biases provided by analytical models. In robotic manipulation, purely analytical models and physics simulations can offer useful priors but are often coarse approximations of real-world dynamics, sensitive to assumptions about physical parameters [8]. Conversely, fully data-driven methods [9] can model complex behaviors without priors but typically require large amounts of real-world data. By learning only the residual error between a physics model and reality, residual learning approaches significantly reduce the data burden while improving real-world performance. Our work builds on residual learning approaches [4, 5]. In addition, ACTIVEPUSHER advances beyond prior work by actively selecting the most informative data points for model refinement, further improving sample efficiency.

Active Learning is a well-established topic in machine learning that aims to improve sample efficiency by actively selecting which data points to label [7, 10, 11, 12]. This paradigm naturally aligns with self-supervised robotic learning settings, where the robot can autonomously choose which experiences to collect. Given the high cost of real-world data acquisition, several robotic learning approaches have leveraged active learning to reduce the number of required interactions [13]. In the context of skill learning, recent methods [14, 15] have applied active learning strategies to accelerate the acquisition of binary success/failure skills, such as pouring [16] or scooping. However, existing approaches primarily focus on optimizing task success rates rather than predicting the detailed outcomes of actions. In contrast, our work applies active learning to improve predictive skill models, enabling more accurate outcome predictions and their integration into kinodynamic planners.

**Kinodynamic Non-Prehensile Planning** is a class of motion planning methods for non-prehensile tasks. Prior works use fixed physics-based models and primarily emphasize computational efficiency [17, 18]. The authors of [19] adapt to uncertainty online, while visual-dynamics methods [20, 21, 22] learned visual forward dynamics, but require extensive data collection. More closely related to our approach are methods such as [8, 23, 24], which learn dynamics models using GP and integrate them within a model predictive control (MPC) framework. Although these methods combine learned models with model-based planners, they rely on random data collection and do not reason about learned model uncertainty.

# **3** Problem Statement

**Kinodynamic Planning:** Let  $x \in \mathcal{X}$  denote the state and state-space, and  $u \in \mathcal{U}$  denote the control and control-space of a robotic system [1]. The true (unknown) system dynamics can be expressed as:

$$x(\mathcal{T}) = x(0) + \int_0^{\mathcal{T}} f(x(t), u(t)) dt$$

where  $\mathcal{T}$  denotes the trajectory duration, and f the unknown forward dynamics model of the system. Let  $\mathcal{X}_{obs} \subset \mathcal{X}$  denote the obstacle (invalid) state space, and  $\mathcal{X}_{free} = \mathcal{X} \setminus \mathcal{X}_{obs}$  denote the free (valid) space. The start state is  $x_{start} \in \mathcal{X}_{free}$ , and the goal region is  $X_{goal} \subseteq \mathcal{X}_{free}$ .

The kinodynamic planning problem is to determine a time  $\mathcal{T}$  and a control function  $u : [0, \mathcal{T}] \to \mathcal{U}$ such that the resulting trajectory satisfies  $x(0) = x_{\text{start}}, x(\mathcal{T}) \in X_{\text{goal}}$ , and  $x(t) \in \mathcal{X}_{\text{free}}$  for all  $t \in [0, \mathcal{T}]$ .

Active Learning of Forward Dynamics Models: We define an interaction as the application of a sequence of controls  $(u_1, u_2, u_3, ...)$ , with the corresponding observed states  $(x_1, x_2, x_3, ...)$ . The sequence of applied controls and observed states forms a dataset for model learning.

The active learning of forward dynamics models problem is to approximate the unknown f with a learned model  $\hat{f}$ , such that  $\hat{f}$  predicts the outcomes of applied controls as accurately as possible, while minimizing the number of interactions required for training with the help of uncertainty quantification.

Active Kinodynamic Planning: Borrowing the concept of active learning with uncertainty quantification, we define *active kinodynamic planning* as incorporating uncertainty estimation from the learned dynamics model in kinodynamic planning. The problem aims at finding a control trajectory  $u(\cdot)$  such that the resulting trajectory  $x(\cdot)$  satisfies the planning requirements with high probability.

#### 4 Methodology

We represent each object as a 2D oriented bounding box (OBB) and parameterize a 2D push by three variables (s, o, d), shown in Fig. 2, where  $s \in \{1, ..., 4\}$  selects one of the box's sides, o is the lateral offset along that side from the center, and d is the total push distance. The end-effector follows a straight-line fixed-duration (T = 3 s) sinusoidal velocity profile:

$$v(t) = \frac{d}{T} \left[ \sin\left(2\pi \, \frac{t}{T} - \frac{\pi}{2}\right) + 1 \right]$$

The effect of the push is defined as the SE(2) transform between the object's initial and final poses. In this work, we adopt the isotropic assumption, treating the push effect as invariant to the object's initial state by modeling it in an object-centric frame.

#### 4.1 Residual Physics

To effectively predict in a low-data setting, we adopt the approach of learning residual physics, which integrates a physics-based model with a neural network [4, 5]. Rather than replacing the physics-based model, the neural network is tasked with learning the residual error, i.e. deviations from the idealized model output to the real observations. This preserves the physical plausibility while allowing the learned component to correct and improve overall accuracy.



Figure 2: Neural network (NN) with residual physics architecture. The network takes both the control parameters and the output of the physics model to predict residuals, which are then added to the physics-based output to produce the final prediction.

For the analytical model, we follow the motion model proposed in [25] to predict object motion given pusher motion. In this framework, the object is treated as a rigid rectangle pushed under quasi-static assumption, with frictional forces obeying Coulomb's law. The model further requires knowledge of the object's shape and the ratio of frictional moment to frictional force. However, we do not assume having access to these exact parameters. Thus, the analytic prediction serves only as a coarse estimation on how the object will move. To keep the model compatible with our data-driven components, we apply additional simplification, details of which are provided in Appendix Sec. B.1.

As illustrated in Fig. 2, our neural network takes both the skill parameters and the output of the physics equation as input. This design enables the network to reason about both the nominal dynamics and the data-driven corrections required to account for object-specific and contact-specific variations. The network finally outputs the residual, and the model combines it with the physics equation output to provide the final estimate. We train the combined model by minimizing the SE(2) Mean Squared Error (MSE):

$$\mathcal{L}_{SE(2)} = \left\| \log \left( \hat{T}^{-1} T \right) \right\|^2$$

where  $\hat{T}, T \in SE(2)$  are the predicted and true poses. The logarithmic map  $\log(\cdot)$  maps the relative transform  $\hat{T}^{-1}T$  to its Lie algebra  $\mathfrak{se}(2)$  and gives a 3D vector of errors in the Lie tangent space.

#### 4.2 Uncertainty Quantification

Traditionally, neural network-based dynamics models produce only point estimates of action outcomes, lacking measure of their prediction uncertainty. By explicitly quantifying the epistemic uncertainty in the learned model, ACTIVEPUSHER enables both uncertainty-aware data acquisition and robust planning, as illustrated in Fig. 3. During learning, this uncertainty guides active data collection by prioritizing the most informative samples, thereby improving data efficiency (Sec. 4.3). At execution time, the planner leverages this uncertainty to select reliable actions from well-explored regions of the action space, resulting in more robust planning (Sec. 4.4).

ACTIVEPUSHER estimates model uncertainty by recasting a trained neural network as a Gaussian Process (GP) via its Neural Tangent Kernel (NTK) [6]. In the infinite-width limit setting, a fully-connected network trained by gradient descent evolves as a Gaussian Process with fixed covariance given by the NTK, defined as

$$x_{\text{NTK}}(x, x') = \langle \nabla_{\theta} f_{\theta}(x), \nabla_{\theta} f_{\theta}(x') \rangle$$

where  $f_{\theta}(x)$  denotes the network output given input x and  $\nabla_{\theta} f_{\theta}(x)$  is its gradient with respect to model's parameters. Empirically, for finite-width networks, the kernel changes during the training

process. But NTK with parameters after convergence (also known as empirical NTK) still yields accurate posterior uncertainty estimates in practice [6, 12, 26].

The gradient of the neural network  $\nabla_{\theta} f_{\theta}(x)$  reflects how sensitive the network's output  $f_{\theta}(x)$  is to small perturbations in parameter  $\theta$ . Intuitively, if the gradients at two input points x and x' are similar, these inputs contain comparable information from the perspective of the neural network.

Leveraging NTK as a prior covariance function allows for explicit posterior inference of unobserved data points within the data pool  $S_{pool}$ , conditioned on observed training data  $S_{train}$ . Formally, given a pre-defined data noise  $\sigma_d$ , the posterior can be represented as in Eq. 1.

$$Cov(\mathcal{S}_{pool}) = K_{NTK}(\mathcal{S}_{pool}, \mathcal{S}_{pool}) - K_{NTK}(\mathcal{S}_{pool}, \mathcal{S}_{train}) K_t^{-1} K_{NTK}(\mathcal{S}_{train}, \mathcal{S}_{pool}) K_t = K_{NTK}(\mathcal{S}_{train}, \mathcal{S}_{train}) + \sigma_d^2 I$$
(1)

where  $K_{\text{NTK}}(A, B)$  is the Gram matrix of pariwise NTK values between inputs in A and B.

In practice, we employ additional kernel transformations for computational efficiency and numerical stability. Specifically, a sketching kernel is applied to reduce the dimensionality of kernel matrices, and a scaling kernel is utilized for label normalization. Detailed discussions and implementations of these kernel transformations can be found in [12].

#### 4.3 Active Learning

Given the posterior estimation, one can apply different acquisition strategies to select next batch of data. The general active learning process is defined in Alg. 1. In each of the N training round, we perform uncertainty estimation over all unlabeled data in  $S_{pool}$  and select the B most informative samples  $S_{sel}$ . After quering their labels and moving them into the training set  $S_{train}$ , we retrain the model with the expanded  $S_{train}$  and proceed to the next round.

In this work, we adopt the BAIT algorithm [7] to actively select the most informative pushing parameter batch to execute during training. Intuitively, BAIT aims to select k points whose combined per-sample Fisher embeddings best approximate (in Frobenius norm) the global Fisher information, yielding a representative batch that jointly captures both model uncertainty and diversity. In regression problems with squared-error, BAIT is simplified to reduce the total pos-





b) Active Learning

c) Active Planning

Figure 3: a) ACTIVEPUSHER quantifies the model uncertainty of the learned model (Sec. 4.2). b) During the learning phase, ACTIVEPUSHER chooses the most *informative* push to apply to increase the learning efficiency (Sec. 4.3). c) During planning the most *reliable* pushes are chosen to maximize the task success rate (Sec. 4.4).

terior variance over the training and pool set after sample selection. Equivalently, in the kernel setting, BAIT acquisition function is given by:

$$S_{sel} = \operatorname{Acquire}_{BAIT}(k, S_{train}, S_{pool}, B) = \arg\min_{S_{sel} \subseteq S_{pool}} \sum_{x \in S_{train} \cup S_{pool}} k_{NTK}[S_{train} \cup S_{sel}](x, x)$$

$$(2)$$

where  $k_{\text{NTK}}[S]$  denotes the posterior after being conditioned on S. However, as discussed in [7], optimizing the Fisher objective Eq. 2 is intractable given the many potential different combinations for  $S_{sel}$ . To avoid this, BAIT employs a forward–backward greedy approximation to select a batch that closely approximates the solution. Details of the algorithm are provided in Appendix Sec. A.2.

## Algorithm 1: Active Learning

**Input:** Kernel k, training round N, batch size B, initial training set  $S_{train}$ , pool set  $S_{pool}$ 

1 for  $i \leftarrow 1$  to N do

- 2 |  $S_{sel} \leftarrow \text{Acquire}(k, S_{train}, S_{pool}, B);$
- 3  $S_{train} \leftarrow S_{train} \cup S_{sel};$
- 4  $S_{pool} \leftarrow S_{pool} \setminus S_{sel};$
- 5 Acquire labels of  $S_{sel}$ ;
- 6 Train model with  $S_{train}$  and corresponding labels;

#### Algorithm 2: Active Action Sampling

**Input:** Kernel k, training set  $S_{\text{train}}$ 

- 1  $\mathcal{S}_U \leftarrow \texttt{random\_sampling}();$
- 2 Var  $\leftarrow$  Diag(query\_uncertainty( $k, S_{\text{train}}, S_U$ )) (Eq. 1);

```
3 return \operatorname{argmin}_{u \in \mathcal{S}_U} \operatorname{Var}[u];
```

#### 4.4 Active Planning

Different from traditional trajectory tracking or closed-loop MPC approaches for pushing [8, 23, 24], we formulate nonprehensile pushing as a kinodynamic planning problem in the object's SE(2) state space. In this formulation, each parameterized push action becomes a discrete control that drives the object's pose. We use an asymptotically optimal kinodynamic planner, specifically SST [27], to explore the object's state space directly.

In the absence of model error, action sequences found by SST succeed by design; in practice, however, accumulated prediction errors can lead to execution failures. To improve robustness, we integrate epistemic uncertainty estimates into the action sampling step. Specifically, we use the method in Sec. 4.2 to provide uncertainty estimate on the potential pushing actions. This can bias action sampling away from uncertain regions of the skill space. At each planning step, we sample a batch of candidate pushing actions, query the model to evaluate their epistemic uncertainty, and select the action with the lowest predicted uncertainty. Alg. 2 summarizes our uncertainty-aware sampling procedure. By biasing exploration toward well-explored regions of the skill space, SST remains asymptotically optimal while avoiding poorly modeled dynamics.

## **5** Experiments

To evaluate the effectiveness of ACTIVEPUSHER, we design a nonprehensile manipulation scenario, as shown in Fig. 1. The same setup is replicated in the genesis [28] simulation, which enables parallel data collection for accelerated training. In both simulation and real-world experiments, we use a 6-DOF UR10 robotic arm equipped with a rigidly grasped cylinder tool to execute nonprehensile pushing actions. Our experiments focus on the task of pushing different objects, drawn from the YCB object dataset [29]. The selected objects vary in shape, mass, and frictional properties to test the robustness of our approach across a range of physical characteristics. To execute the push parameters in the robot, we use a global redundancy resolution method [30].

We evaluate the proposed method in two settings: (i) skill learning performance, which measures prediction accuracy (Sec. 5.1), and (ii) long-horizon kinodynamic planning, where we assess the task success effectiveness of uncertainty-guided planning (Sec. 5.2).

#### 5.1 Skill Learning

As described in Sec. 4.3, training our active learning method requires a pool of candidate samples for labeling. We randomly selected 1,000 push parameters to serve as the candidate set, denoted by  $S_{pool}$ . At each data acquisition stage, we collect a batch of 20 samples  $S_{sel}$  and repeat this process for a total of 10 batches. We repeat the whole experiment five times in simulation and three times in the real world. Four objects in simulation and one object in the real world are used, and we evaluate the following five methods:



Figure 4: Push-Skill Learning Results. The validation Loss for 1 real object, and 4 simulated objects in from the YCB Dataset: *Cracker Box, Mug, Mustard Bottle Banana*. The proposed active learning methods outperform random data collection for all objects. Models with residual physics performs better in low-data regime.

- Pure Physics: The analytical dynamics model described in Sec. 4.1.
- MLP Random: A fully connected neural network trained on random push parameters.
- *Residual Random:* The hybrid model as described in Sec. 4.1. The physics estimation is the same as *Pure Physics*. Meanwhile, the neural network architecture and data collection are the same as *MLP Random*.
- *MLP Active Learning:* The same MLP architecture as *MLP Random*, but trained via our NTK-driven active learning pipeline to select informative samples.
- *Residual Active Learning:* The hybrid residual model, trained using our active learning framework to assess the combined impact of residual physics and informed sample selection.

The summarized results for Root Mean Square Error (RMSE) loss on the validation set (1,000 data), shown in Fig. 4, demonstrate that modesl with residual physics features clear improvement when having a limited number of data. Also, active learning approach outperforms the random method for all objects. Therefore, our experiment shows that leveraging active learning and informing the model with the physics of the task would be beneficial in scenarios where data acquisition is challenging.

#### 5.2 Kinodynamic Planning

In this experiment we demonstrate the performance of a learned model in tandem with a kinodynamic planner for a downstream task. The planning was conducted in the object's state space, defined as X = SE(2), with a control space  $U = \mathbb{R}^3$  corresponding to the pushing parameters described in Sec. 4. The valid state space  $X_{free}$  is constrained to the table surface, and the robot base is treated as an obstacle that the object should not be pushed into.

We designed a Push-to-Grasp Task in which the robot must push a wide, non-graspable box toward the edge of the table to enable a feasible pick-up, as shown on the right side of Fig. 1. The goal region  $X_{goal}$  is defined via a geometry proxy: any object pose for which the center of the mass of the box is inside the table and at least one corner lies beyond a certain distance from the table boundary is treated as graspable and serves as a valid goal state.

We use SST[27], an asymptotically optimal kinodynamic planner, in the Open Motion Planning Library OMPL [31] to solve a set of 100 planning problems with varied start and goal configurations. A key advantage of asymptotic optimality is that it enables us to use path length as optimization objective, implicitly minimizing the number of control steps and reducing cumulative model error during execution. The produced sequence of controls are executed in open-loop.

Active Learning for Push-to-Grasp Task: We first evaluate the performance of the learned skill models, trained with different active learning strategies described in the previous section Sec. 5.1, in simulation. The original kinodynamic planner is used in this experiment. The whole experiment



Figure 5: Planning results for Push-to-Grasp Task in simulation, showing success rate over 100 different plans and the corresponding tracking error between planned and executed trajectories.



Figure 6: Results comparison of Active Planning and Random Planning in both simulation and real-world for Push-to-Grasp Task.

is repeated five times and we measure task success rate and trajectory tracking error in the SE(2) space, which are reported in Fig. 5.

Overall, the results are consistent with those of Sec. 5.1. In the low-data regime, models with residual physics consistently outperform those without, and incorporating active learning accelerates improvement in task success rate and reduction in tracking error.

Active Planning for Push-to-Grasp Task: We test the proposed active planning method, as described in Sec. 4.4, with the best trained model *Residual Active Learning*, both in simulation and in the real world. In simulation, we run the 100 planning problems with all ten models with different training data sizes. The whole experiment is repeated three times. For real-world models, we test five models at batches 2, 4, 6, 8, and 10, and the whole experiment is only run once. The results are presented in Fig. 6.

In both simulation and the real world, tracking error results show that the active action sampling algorithm would choose actions that models are more confident with, thus resulting in a lower execution error in push paths in general. Consequently, our active planning method is able to outperform the random method in terms of the success rate consistently.

# 6 Conclusion

In this paper, we present ACTIVEPUSHER, a framework that combines residual physics, active learning, and uncertainty-aware active kinodynamic planning to improve learning and planning of nonprehensile manipulation. By explicitly modeling epistemic uncertainty using NTK, our method selectively gathers informative training data with BAIT acquisition function, and biases kinodynamic planning toward reliable actions. Experimental results in both simulation and real-world pushing tasks demonstrate that ACTIVEPUSHER achieves higher prediction accuracy and planning success with fewer interactions compared to baseline approaches. This integration of learning and planning offers a promising path toward data-efficient and reliable robotic manipulation in nonprehensile setting.

# 7 Limitation

One limitation of our framework ACTIVEPUSHER is that we model only epistemic uncertainty through NTK-based posterior, and do not account for inherent aleatoric uncertainty. In many pushing scenarios, different actions exhibit inherently different noise characteristics. For instance, pushing a toy car forward may produce more variability than nudging it laterally, yet our planner treats all residual noise as if it were learnable. Without an aleatoric term, the system may avoid regions with high model uncertainty that are in fact low–noise in practice, or conversely over-commit to pushes with low model uncertainty that are intrinsically stochastic.

A second limitation is on task and object diversity. Although we evaluate across a variety of YCB objects with different shapes and mass, our study remains confined to planar pushes in SE(2). Future work could explore a broader array of geometries (e.g. non-convex or articulated objects), dynamics (e.g. uneven mass distribution or surface), or extend the framework to full 3D (SE(3)) pushing and other non-prehensile skills such as rolling and throwing.

A third limitation is that ACTIVEPUSHER operates in an open-loop fashion. During execution, we sample a fixed sequence of pushes without real-time feedback and replanning. In practice, closed-loop control and online replanning can substantially increase robustness to unmodeled disturbances or sensor noise.

## References

- A. Orthey, C. Chamzas, and L. E. Kavraki. Sampling-based motion planning: A comparative review. Annual Review of Control, Robotics, and Autonomous Systems, 7(1):285–310, July 2024. URL https://doi.org/10.1146/annurev-control-061623-094742.
- [2] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez. Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(Volume 4, 2021):265-293:, 2021. ISSN 2573-5144. URL https://www.annualreviews.org/content/journals/10.1146/ annurev-control-091420-084139.
- M. T. Mason. Toward robotic manipulation. Annual Review of Control, Robotics, and Autonomous Systems, 1(Volume 1, 2018):1-28, 2018. ISSN 2573-5144. URL https://www.annualreviews.org/content/journals/10.1146/ annurev-control-060117-104848.
- [4] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser. Tossingbot: Learning to throw arbitrary objects with residual physics. *IEEE Transactions on Robotics*, 36(4):1307–1319, 2020. URL https://ieeexplore.ieee.org/document/9104757.
- [5] A. Ajay, J. Wu, N. Fazeli, M. Bauza, L. P. Kaelbling, J. B. Tenenbaum, and A. Rodriguez. Augmenting physical simulators with stochastic neural networks: Case study of planar pushing and bouncing. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3066–3073, 2018. URL https://ieeexplore.ieee.org/document/8593995.
- [6] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper\_files/paper/ 2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf.
- [7] J. Ash, S. Goel, A. Krishnamurthy, and S. Kakade. Gone fishing: Neural active learning with fisher embeddings. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 8927–8939. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper\_files/ paper/2021/file/4afe044911ed2c247005912512ace23b-Paper.pdf.
- [8] F. R. Hogan and A. Rodriguez. Reactive planar non-prehensile manipulation with hybrid model predictive control. *The International Journal of Robotics Research*, 39(7):755–773, 2020. doi:10.1177/0278364920913938. URL https://doi.org/10.1177/0278364920913938.
- [9] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine. Learning to poke by poking: Experiential learning of intuitive physics. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper\_files/paper/ 2016/file/c203d8a151612acf12457e4d67635a95-Paper.pdf.
- [10] Y. Gal, R. Islam, and Z. Ghahramani. Deep Bayesian active learning with image data. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1183–1192. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/gal17a.html.
- [11] X. Li and Y. Guo. Adaptive active learning for image classification. In 2013 IEEE Conference on Computer Vision and Pattern Recognition, pages 859–866, 2013. doi:10.1109/CVPR.2013.116.
- [12] D. Holzmüller, V. Zaverkin, J. Kästner, and I. Steinwart. A framework and benchmark for deep batch active learning for regression. J. Mach. Learn. Res., 24(1), Jan. 2023. ISSN 1532-4435. URL https://dl.acm.org/doi/abs/10.5555/3648699.3648863.

- [13] A. T. Taylor, T. A. Berrueta, and T. D. Murphey. Active learning in robotics: A review of control principles. *Mechatronics*, 77:102576, 2021. ISSN 0957-4158. URL https://www. sciencedirect.com/science/article/pii/S0957415821000659.
- [14] Z. Wang, C. R. Garrett, L. P. Kaelbling, and T. Lozano-Pérez. Learning compositional models of robot skills for task and motion planning. *The International Journal of Robotics Research*, 40(6-7):866–894, 2021. URL https://doi.org/10.1177/02783649211004615.
- [15] A. LaGrassa, M. Lee, and O. Kroemer. Task-oriented active learning of model preconditions for inaccurate dynamics models. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pages 16445–16445. IEEE, 2024.
- [16] N. Kumar, T. Silver, W. McClinton, L. Zhao, S. Proulx, T. Lozano-Pérez, L. P. Kaelbling, and J. Barry. Practice makes perfect: Planning to learn skill parameter policies. *arXiv preprint* arXiv:2402.15025, 2024.
- [17] J. A. Haustein, I. Arnekvist, J. Stork, K. Hang, and D. Kragic. Learning manipulation states and actions for efficient non-prehensile rearrangement planning. *arXiv preprint arXiv:1901.03557*, 2019.
- [18] K. Ren, P. Chanrungmaneekul, L. E. Kavraki, and K. Hang. Kinodynamic rapidly-exploring random forest for rearrangement-based nonprehensile manipulation. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pages 8127–8133. IEEE, 2023.
- [19] M. Faroni and D. Berenson. Online adaptation of sampling-based motion planning with inaccurate models. In 2024 IEEE International Conference on Robotics and Automation (ICRA), pages 2382–2388. IEEE, 2024.
- [20] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In 2017 IEEE international conference on robotics and automation (ICRA), pages 2786–2793. IEEE, 2017.
- [21] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *Advances in neural information processing* systems, 28, 2015.
- [22] W. Zhou, B. Jiang, F. Yang, C. Paxton, and D. Held. Hacman: Learning hybrid actor-critic maps for 6d non-prehensile manipulation. In *Conference on Robot Learning*, pages 241–265. PMLR, 2023.
- [23] M. Bauza, F. R. Hogan, and A. Rodriguez. A data-efficient approach to precise and controlled pushing. In *Conference on Robot Learning*, pages 336–345. PMLR, 2018.
- [24] G. Wang, K. Ren, and K. Hang. Uno push: Unified nonprehensile object pushing via nonparametric estimation and model predictive control. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 9893–9900. IEEE, 2024.
- [25] K. Lynch, H. Maekawa, and K. Tanie. Manipulation and active sensing by pushing using tactile feedback. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 416–421, 1992. doi:10.1109/IROS.1992.587370.
- [26] M. A. Mohamadi, W. Bae, and D. J. Sutherland. Making look-ahead active learning strategies feasible with neural tangent kernels. *Advances in Neural Information Processing Systems*, 35: 12542–12553, 2022.
- [27] Y. Li, Z. Littlefield, and K. E. Bekris. Asymptotically optimal sampling-based kinodynamic planning. *The International Journal of Robotics Research*, 35(5):528–564, 2016. doi:10.1177/ 0278364915614386. URL https://doi.org/10.1177/0278364915614386.
- [28] G. Authors. Genesis: A universal and generative physics engine for robotics and beyond. December 2024. URL https://github.com/Genesis-Embodied-AI/Genesis.

- [29] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set. *IEEE Robotics Automation Magazine*, 22(3):36–52, 2015. doi:10.1109/MRA.2015.2448951.
- [30] Z. Zhong, Z. Li, and C. Chamzas. Expansion-grr: Efficient generation of smooth global redundancy resolution roadmaps. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 8854–8860, 2024. doi:10.1109/IROS58592.2024.10801917.
- [31] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. doi:10.1109/MRA.2012.2205651. https://ompl.kavrakilab.org.
- [32] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint* arXiv:2303.05499, 2023. URL https://arxiv.org/pdf/2303.05499.
- [33] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollar, and R. Girshick. Segment anything. In *Proceedings of* the IEEE/CVF International Conference on Computer Vision (ICCV), pages 4015–4026, October 2023. URL https://openaccess.thecvf.com/content/ICCV2023/papers/ Kirillov\_Segment\_Anything\_ICCV\_2023\_paper.pdf.

# Appendix

## A Active Learning Algorithms

In this section, we provide a more detailed and formal explanation of the active learning strategy used in ACTIVEPUSHER. The core idea is to leverage the Neural Tangent Kernel (NTK) [6] for posterior estimation and the Batch Active learning via Information maTrices (BAIT) [7] acquisition strategy.

#### A.1 Posterior Estimation

Given a fully-connected neural network  $f_{\theta}(x)$  with infinite width, parameterized by weights  $\theta$ , trained with dataset samples  $S_{train}$  and corresponding labels  $\mathcal{L}_{train}$ , its NTK is defined as

$$k_{\text{NTK}}(x, x') = \langle \nabla_{\theta} f_{\theta}(x), \nabla_{\theta} f_{\theta}(x') \rangle$$

where  $\nabla_{\theta} f_{\theta}(x)$  denotes the gradient of the network output with respect to the parameters  $\theta$ , x and x' are two different input samples, and  $\langle \cdot, \cdot \rangle$  is the inner product of two parameter gradient vectors.

After training to convergence under the infinite-width regime, the neural network  $f_{\theta}(x)$  becomes equivalent to a Gaussian Process(GP) governed by the NTK:

$$f_{\theta}(x) \sim \mathcal{GP}(0, k_{\text{NTK}}(x, x')).$$

Conditioned on the labeled training dataset  $S_{train}$ , we obtain the posterior kernel  $k_{\text{NTK}}[S_{train}]$ . The posterior allows us to estimate predictive uncertainty over a given unlabeled pool set  $S_{pool}$ :

$$\operatorname{Cov}(\mathcal{S}_{\text{pool}}) = k_{\text{NTK}}(\mathcal{S}_{\text{pool}}, \mathcal{S}_{\text{pool}}) - k_{\text{NTK}}(\mathcal{S}_{\text{pool}}, \mathcal{S}_{\text{train}}) k_{\text{NTK}}(\mathcal{S}_{\text{train}}, \mathcal{S}_{\text{train}} + \sigma^2 I)^{-1} k_{\text{NTK}}(\mathcal{S}_{\text{train}}, \mathcal{S}_{\text{pool}})$$

where  $k_{NTK}$  abbreviates  $k_{NTK}[S_{train}]$  for simplicity, and  $\sigma$  is the assumed inherent data noise (we set  $\sigma = 0.01$  in our experiments). By isolating the diagonal terms of the posterior covariance matrix, we obtain per-sample epistemic uncertainty estimates for  $S_{pool}$ :

$$\operatorname{Var}(\mathcal{S}_{\operatorname{pool}}) = \operatorname{Diag}(\operatorname{Cov}(\mathcal{S}_{\operatorname{pool}}))$$

Although real networks are of finite width, this empirical NTK approximation still provides a practical estimate of the model's uncertainty. Since the aleatoric uncertainty is captured by pre-defined  $\sigma_d$ , the posterior covariance primarily reflects the epistemic uncertainty of the model. These estimates on the unseen data  $S_{pool}$  will be used for both active skill learning and active skill planning.

#### A.2 BAIT

Given posterior estimation, we adopt BAIT acquisition strategy to choose informative samples [7]. BAIT builds upon a theoretically principled criterion from statistics: minimizing the expected error (or Bayes risk) of a maximum likelihood estimator by maximizing the Fisher information accumulated through labeling selected data points. Specifically, BAIT seeks to minimize the trace of the inverse Fisher information matrix of the selected batch (i.e., the model uncertainty after selecting a batch), pre-multiplied by the Fisher information of the entire unlabeled pool:

$$\underset{\mathcal{S}_{\text{sel}} \subseteq \mathcal{S}_{\text{pool}}}{\arg\min} \operatorname{tr} \left( \left( \sum_{x \in \mathcal{S}_{\text{sel}}} I(x; \theta) \right)^{-1} \left( \sum_{x \in \mathcal{S}_{\text{pool}}} I(x; \theta) \right) \right)$$
(3)

To reduce the necessary computation, BAIT chose to operate only on the last layer of the network. For regression, this can be further simplified with Kronecker product to:

$$\underset{\mathcal{S}_{\text{sel}} \subseteq \mathcal{S}_{\text{pool}}}{\arg\min c} \cdot \operatorname{tr} \left( \left( \sum_{x \in \mathcal{S}_{\text{sel}}} x^L (x^L)^\top \right)^{-1} \left( \sum_{x \in \mathcal{S}_{\text{pool}}} x^L (x^L)^\top \right) \right)$$

where Let  $(x^{L})$  is the penultimate layer representation induced by the neural network.

In the kernel setting, as shown in [12], one can prove that:

$$\sum_{x \in \mathcal{S}_{pool}} k \big[ \mathcal{S}_{sel} \big](x, x) = c \operatorname{tr} \left( G_{\mathcal{S}_{sel}}^{-1} \ G_{\mathcal{S}_{pool}} \right), \quad G_{\mathcal{S}} := \nabla_{\theta} f_{\theta}(\mathcal{S})^{\top} \nabla_{\theta} f_{\theta}(\mathcal{S})$$
(4)

In regression with Gaussian likelihood, the following is equivalent:

$$\sum_{x \in \mathcal{S}} I(x; \theta) = \nabla_{\theta} f_{\theta}(\mathcal{S})^{\top} \nabla_{\theta} f_{\theta}(\mathcal{S})$$

which shows that by optimizing Eq. 4, we are essentially optimizing the same fisher objective Eq. 3.

Different from the original BAIT [7], we made the following changes. First, considering that our neural network is relatively small, we use the full gradient NTK instead of the last-layer gradient. Second, instead of only computing the Fisher information of merely the selected set and pool set, we expand it by also considering the current training set. In combination, our new objective is:

$$\sum_{x \in \mathcal{S}_{train} \cup \mathcal{S}_{pool}} k_{\text{NTK}} \big[ \mathcal{S}_{train} \cup \mathcal{S}_{sel} \big](x, x)$$

As discussed in [7], optimizing such a Fisher objective is intractable given the many potential different combinations for  $S_{sel}$ . To address this, the same greedy forward–backward selection algorithms, proposed in [7], are adopted. Given *B* number of samples to collect in a batch, the algorithm first greedily selects 2*B* samples by:

$$\underset{x \in \mathcal{S}_{pool}}{\operatorname{arg\,min}} \sum_{x' \in \mathcal{S}_{train} \cup \mathcal{S}_{pool}} k_{\text{NTK}}[\mathcal{S}_{train} \cup \mathcal{S}_{sel} \cup \{x\}](x', x')$$

After collecting 2B samples, it greedily removes B samples from the selected set by:

$$\underset{x \in \mathcal{S}_{sel}}{\operatorname{arg\,min}} \sum_{x' \in \mathcal{S}_{train} \cup \mathcal{S}_{pool}} k_{\mathrm{NTK}} [\mathcal{S}_{train} \cup (\mathcal{S}_{sel} \setminus \{x\})](x', x')$$

## **B** Implementation Details

In this section, we provide a detailed description of our models, training procedure and the experiments conducted in both simulation and real-world environments. All the experiments are run with a 6-DOF UR10 robot, and on a workstation with an NVIDIA RTX 4070 Ti Super GPU and 32GB of RAM.

#### **B.1** Dynamics Model

To predict the object's final pose after a push, we use a neural network model with residual physics, where the analytical physics formulation is adapted from [25]. To integrate this analytical model into the neural network and enable efficient batch operations, we made the following simplifying assumptions:

- **Rectangular approximation.** As required by this mode, we approximate every object by a 2D Oriented Bounding Box (OBB). While this captures the overall geometry, it may introduce errors at highly irregular features (e.g., the neck of a mustard bottle).
- Fixed frictional ratio. The frictional moment to force ratio c is object-specific and generally unknown. We fix c = 0.05 for all objects, a value chosen empirically to best match our real-world pushing experiments.
- **Simplified contact and force application.** We assume the pusher follows a straight-line trajectory with a sinusoidal velocity profile. In reality, as the object rotates, the contact may slide along its edge and the force direction would tilt away from perpendicular. To simplify it, we hold the contact point fixed and assume the force is always perpendicular to that point, maintaining perfect sticking (no-slip) contact throughout the push. This oversimplification allows us to compute the final pose in a single vectorized integration step rather than simulating incremental, step-by-step dynamics (see Eq. 5).

We push the object with a rigid rod mounted on the robot's end effector. The push is defined with parameters (s, o, d), and follows a sinusoidal velocity profile v(t). We choose the sinusoidal velocity profile because it produces continuous and bounded acceleration and avoids sudden jerks. Additionally, it provides an analytically tractable form where the total displacement d determines the entire velocity profile.

Taking all of this into account, we derive the object's relative final pose by first adapting the following equation proposed in [25].

$$v_x = \frac{(c^2 + x_c^2) v_{px} + x_c y_c v_{py}}{c^2 + x_c^2 + y_c^2}, \quad v_y = \frac{(c^2 + y_c^2) v_{py} + x_c y_c v_{px}}{c^2 + x_c^2 + y_c^2}, \quad \omega = \frac{x_c v_y - y_c v_x}{c^2}$$

where  $(x_c, y_c)$  is the contact point and  $(v_{px}, v_{py})$  is the pusher speed. Under our assumption, they are:

$$x_c = L_{obj} \quad y_c = s, \quad v_{px} = -v(t), \quad v_{py} = 0$$

Given the local translational and rotational speed, we can compute the pose at time t by integration:

$$\theta(t) = \int_0^t \omega(t) dt,$$

$$x(t) = \int_0^t \left[\cos(\theta(t)) v_x(t) - \sin(\theta(t)) v_y(t)\right] dt,$$

$$y(t) = \int_0^t \left[\sin(\theta(t)) v_x(t) + \cos(\theta(t) v_y(t)\right] dt$$
(5)

For our neural network architecture, we employ a fully connected multilayer perceptron (MLP) consisting of five hidden layers with sizes [32, 64, 128, 64, 32]. The network is designed to predict the relative SE(2) pose of the object after a push. For standard MLPs, the input consists of the push parameters, resulting in an input dimensionality of 3. In contrast, for the residual model, the input additionally includes the output of the analytical physics model (as defined in Equation Eq. 5), increasing the input dimensionality to 6.

All models are trained using a batch size of 16 for 1000 epochs, or until convergence. We adopt the Adam optimizer in conjunction with a learning rate scheduler that reduces the learning rate on plateau, helping to mitigate overfitting. During the skill learning process, the model is retrained from scratch after each new data collection cycle—specifically, after acquiring an additional batch of 20 samples in simulation or 40 samples in the real world.

#### **B.2** Motion Planner

To solve the *Push-To-Grasp* motion planning problem, we use the Open Motion Planning Library (OMPL) [31] with the following specifications:

- State Space: SE(2), for an object's unconstrained motion on a plane (table),
- **Control Space:**  $\mathbb{R}^3$ , for the 3-dimensional push parameter (s, o, d),
- State Validity Checker: True or False, based on if it is within the boundary of  $C_{space}$  feasible for the robot and if it is collision-free.
- Control Propagator:  $(x_f, y_f, \theta_f)$ , propagating the SE(2) state based on the relative transformation output of the dynamics model,
- Control Sampler: Random (OMPL default) / Active (ours),
- **Goal Region:** A narrow zone alongside the tabletop edge, defined such that the object's center of mass remains on the table surface and at least one corner of the object extends by at least 2.5*cm* past the edge.
- **Planner:** SST, an asymptotically optimal kinodynamic planner, with a 10*s* planning time limit.

#### **B.3** Experiments

We conduct our simulation experiments using the Genesis simulator [28], which supports parallel execution of multiple environments. Our setup includes a 6-DOF UR10 robot and multiple test objects with different geometric shapes and physical characteristics from the YCB dataset [29]. The simulated scene in the Genesis environment (shown in Fig. 7) replicates our real-world setup. To accelerate both data collection and planning execution, we run 1000 environments in parallel during the data-collection phase and 100 environments in parallel during the planning phase, as in Fig. 7a.



(a) Genesis Simulation



(b) Real-World

Figure 7: Experiment Setup

We execute the same workflow on a physical robot as shown in Fig. 7b. For perception, we use an Intel RealSense Depth Camera D455 for an overhead view and an Intel RealSense Depth Camera D435 mounted on the end-effector for a more precise in-hand observation. Object pose is estimated with the LANGSAM model [32, 33] to detect the position and template-matching for orientation. Robot communication and control are handled via the robot's default Real-Time Data Exchange (RTDE) interface.

Additionally, in our real-world experiments we use a simple reset algorithm that automatically pushes the object back toward the center of the robot's high-manipulability workspace whenever it drifts outside that region. This mechanism enables autonomous, continual learning without any human intervention.