

Amortized Inference for Efficient Grasp Model Adaptation

Michael Noseworthy^{1*}, Seiji Shaw^{1*}, Chad Kessens², Nicholas Roy¹

Abstract—In robotic applications such as bin-picking or block-stacking, learned predictive models have been developed for manipulation of objects with varying but known dynamic properties (e.g., mass distributions and friction coefficients). When a robot encounters a new object, these properties are often difficult to observe and must be inferred through interaction, which can be expensive in both inference time and number of interactions. Our goal is to develop a predictive model that can efficiently adapt to new objects by estimating their unobserved properties through interaction. We propose an encoder/decoder action-feasibility model to efficiently predict grasp stability based on a variable, and limited, amount of interaction data. The encoder predicts a distribution over the unobserved parameters while the decoder predicts action feasibility which can be used in an uncertainty-aware planner. An explicit representation of uncertainty in the encoder enables information-gathering heuristics to minimize adaptation interactions. We also show that the amortized distribution is computationally efficient and performs comparably to particle-based distributions in a grasping domain. Finally, we deploy our method on a Panda robot to grasp heavy objects.

I. INTRODUCTION

In automated manipulation, tasks such as bin-picking or block-stacking are often repeated for a large set of objects that have a wide distribution of *geometric* and *dynamic* parameters (e.g., their masses, centers of mass, and friction coefficients). Our goal is for robots to efficiently interact with new objects to discover their unobserved dynamic parameters required for reliable prediction and planning.

Previous work has shown that robots are able to learn manipulation dynamics (e.g., stacking or throwing) entirely from online data when the dynamic properties of the objects are known [1] or have minimal impact on action outcome. It is unreasonable, however, to expect that a robot should know or be able to observe the dynamic properties when encountering a novel object. Robots without tactile or force feedback, for example, will not be able to directly observe inertial and frictional properties of objects. A robot may not even be able to know what or which dynamic properties will play a role in the manipulation dynamics.

A standard approach to handling unknown parameters is to use a Bayes filter [2] to maintain a probability distribution or belief over the unobserved properties. Every time a new action is executed, the resulting observation (e.g., grasp success or failure) is used to update the belief using an observation model that relates the unknown parameters to observations. When the observation is a simple function of the belief, such as a linear function with Gaussian noise, the



Fig. 1: Our method allows a Panda (right) to infer *unobserved properties* of objects based on interaction history. This allows the robot to efficiently grasp objects with non-uniform mass distributions such as weighted blocks (left) or ShapeNet objects (center).

belief update can be efficient. However, when the observation model is a complicated function, more expressive belief representations, such as sets of particles, are required.

Particle-based approaches suffer from the curse of dimensionality: the number of particles needed to represent the belief scales exponentially with the number of unobserved variables. For tasks such as grasping there are five unknowns: mass, static friction, and three for center of mass. Sampling and updating the belief can therefore be expensive.

When manipulating new objects with unknown properties, a robot will need to perform many belief updates as it interacts and collects observations. Further, while planning, it is common to simulate belief updates in order to choose informative actions. As such, the cost of the belief update is critical. Using particle-based representations can quickly become intractable and limit the amount of interaction possible within a fixed time budget.

To find a predictive dynamics model with an efficient belief representation and update procedure, we propose to *jointly learn* (1) inference networks that are trained to predict posterior distributions over the unobserved parameters given a history of interactions and their outcomes and (2) an action feasibility classifier that can be used for planning. The model uses an encoder/decoder architecture based on *Neural Processes* [3] for unsupervised learning of an object-level latent space shared between the inference network (encoder) and feasibility classifier (decoder). For our grasping task, the input to the inference network is a set of labeled grasps for a specific object. The posterior over the latent parameters is computed as a *single forward pass* through the inference network/encoder, which only scales linearly with the number of observed grasps irrespective of the number of unobservable parameters. The feasibility classifier/decoder can in turn be used for planning by consuming samples from the posterior and predicting the success probability for candidate grasps.

*Equal Contribution

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, ²Army Research Laboratory

We show that our model, called a Grasping Neural Process, can classify stable grasps comparably to particle filters at a fraction of the computational cost. Our inference network predicts a distribution over the learned-latent representation of the unobserved parameters, which we use to compute information-gathering heuristics for a robot to adapt with fewer grasps. In simulation, we verify the GNP can be used to find stable grasps for novel objects after observing a small set of interactions and can identify grasps that minimize the grasping force when deployed in an uncertainty-based planner. We also demonstrate the GNP, trained in simulation, can be deployed to successfully adapt and grasp a weighted block on a physical robot (see Figure 1).

II. PROBLEM FORMULATION

Given a previously unseen object with unknown dynamic properties, our goal is for a robot to reliably manipulate it (e.g., grasp or place) using a small number of adaptation interactions. In order to do so, we take a model-based approach where we aim to predict action feasibility in a way that allows the model to rapidly adapt to novel objects.

Formally, we consider a robot operating in a domain with K objects in a discrete-time setting, where each object's state can be divided into observed and unobserved properties. The observed state $\mathcal{X}_t = \{x_t^{(k)} | k = 1, \dots, K\}$ can be time-varying, whereas the unobserved state $\mathcal{Z} = \{z^{(k)} | k = 1, \dots, K\}$ is assumed to be static. Many time-varying properties (e.g., pose or velocity), can be extracted from a perception system while static properties often cannot be directly observed (e.g., friction coefficients, mass, or coefficients of restitution).

Our objective is to learn a model that predicts the feasibility of an action for a specific object¹ $y_t \in \{0, 1\}$ (e.g., grasp stability). This predictive model can then be integrated with a *task and motion planning* system for long-horizon manipulation tasks. Concretely, the goal is to estimate $p(y_t | x_t, z, a_t)$, where $a_t \in \mathcal{A}$ is the action space of the robot.

We assume the robot's operation is divided into three distinct phases: *training*, *adaptation*, and *testing*. In the *training* phase, the set of objects is fixed, and the robot has some finite amount of time to act in the environment without any task descriptions. In this work, we assume we already have a dataset of completed interactions provided to us for the training phase. In the *adaptation* phase, the robot is given a new object and a short amount of time to experiment and adapt by selecting information-gathering interactions. Finally, there is a *testing* phase, in which a robot is given a task to be performed, and interacts with the objects in order to complete the task. No further adaptation or learning is performed in the *testing* phase.

III. METHODS

We primarily focus on the *adaptation phase*. How can a robot *efficiently*, in both computational and sample efficiency, adapt to objects with unknown dynamic properties?

¹We will leave out the superscript k when referring to a single object.

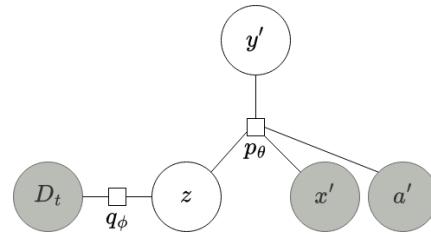


Fig. 2: After executing t actions and observing their outcomes, D_t , the robot can predict a posterior over the unobserved properties, $q_\phi(z|D_t)$. We aim to predict feasibility y' of unlabeled actions, a' , given an object's observed properties, x' , and current belief, z .

We take a probabilistic approach where we perform posterior estimation of the unknown properties, z , given a variable amount of observed data, $D_t = \{(x_1, a_1, y_1), \dots, (x_t, a_t, y_t)\}$. Rather than estimate the posterior distribution through an expensive online computation, we propose to amortize the inference procedure by first training an inference network, q_ϕ , to directly predict the parameters of the posterior distribution (Section III-A). This model is learned jointly with the feasibility model, $p_\theta(y|z, x, a)$. The resulting posterior distribution, $q_\phi(z|D_t)$ (which we will refer to as amortized distributions), can be used directly for informative action planning (Section III-B) or in conjunction with the feasibility model for uncertainty aware planning (Section III-C). See Figure 3 for an overview.

A. Amortized Posterior Updates

When interacting with a novel object, the robot must update its posterior after each interaction, which can be a time-intensive operation. To address this issue, we propose to use amortized inference for constant-cost posterior distributions. Instead of performing an expensive online posterior update, we will train a neural network to predict the posterior distribution given the online samples, D_t , observed so far. The inference network, $q_\phi(z|D_t)$, will need to approximate posterior distributions $p(z|D_t)$ for all possible timesteps, t , and observation realizations. We take $q_\phi(z|D_t)$ to be a multi-dimensional Gaussian distribution. The inference network will predict the mean and diagonal covariance.

At any timestep, t , during the adaptation phase, the graphical model in Figure 2 represents the data generation process. Note that this graphical model represents a single timestep of the adaptation phase where we have already observed the labels for t actions, D_t , but have not yet observed the label, y' , for a new action, a' , with observed state, x' . In order to learn both a feasibility model, $p_\theta(y'|z, x', a')$ that can be used for planning and an inference network, $q_\phi(z|D_t)$, we take an amortized variational inference approach.

For a single object, we seek to optimize:

$$\mathbb{E}_z [\log p_\theta(y'|x', a', z)] - D_{KL}(q_\phi(z|D_T) || q_\phi(z|D_t)), \quad (1)$$

where $D_T = D_t \cup \{(y', a', x')\}$. Note that we do not require labels for the unobserved properties as the latent space is learned (and need not be human-interpretable).

The proposed method is an instantiation of *Neural Processes* [3] for robotic manipulation. The objective is the

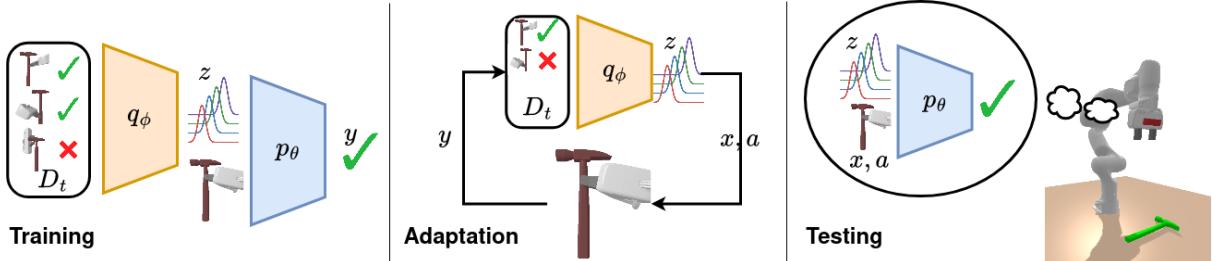


Fig. 3: Grasp Neural Processes use an offline *training phase* (left) to jointly learn an action feasibility model, p_θ , and an inference network, q_ϕ , that predicts a posterior distribution over unobserved properties. During the *adaptation phase* (center), the learned inference network can be used for efficient online posterior updates and action selection. Finally, in the *testing phase* (right), the robot can use the current belief, z , along with the feasibility model to perform manipulation tasks.

evidence lower bound (ELBO) for amortized variational inference over functions (i.e., each sample of the latent space represents a different action-feasibility classifier). The first term of this objective incentivizes the model to make accurate and object-consistent predictions, using the latent space when necessary. The second term is the KL-divergence between the *full posterior* and the *partial posterior*. During the adaptation phase, it is the *partial posterior*, $q_\phi(z|D_t)$, that will be used and the KL-divergence is used to make this posterior as accurate as possible given the reduced input set D_t .

During training, Eq. 1 is optimized using stochastic gradient descent. Mini-batches are sampled where each element of a batch corresponds to $D_T^{(k)}$ for many objects, k . Each $D_T^{(k)}$ is further divided into partial datasets $D_t^{(k)}$ by uniformly sampling t between 0 and T (a standard normal distribution $\mathcal{N}(0, 1)$ is used when $t = 0$). This ensures the encoder can represent posterior distributions for variable input sizes.

B. Informative Action Selection

We have shown so far how to perform efficient posterior updates at a single timestep: evaluate $q_\phi(z|D_t)$. However, we are also interested in minimizing the total number of time-consuming interactions used to collect observations. While selecting *random* actions may be computationally cheap, many interactions are likely required before reaching a sufficient performance level.

Instead, it is common to take an information-theoretic approach. Actions are selected which maximize the information gain at a single timestep (a greedy, but often good, approximation to optimizing over sequences of actions [4]):

$$\max_{a_{t+1}} H(q_\phi(z|D_t)) - \mathbb{E}_{y_{t+1}} [H(q_\phi(z|D_t, x_{t+1}, a_{t+1}, y_{t+1}))] \quad (2)$$

One way to solve this maximization is via a simple pool-based approach: first generate a collection of M unlabeled samples, then score each according to this objective. However, computing expected information gain requires two posterior updates, which can be expensive with particle-based methods: $O(MN)$ if N is the number of particles.

Using the inference network within the information gain computation has a much more favorable complexity, scaling linearly with the number of unlabeled samples: $O(M)$. Given a similar computational budget to baselines, more unlabeled

samples can be considered, leading to more informative actions and therefore fewer expensive interactions.

C. Uncertainty Aware Planning

After seeing a limited number of adaptation actions, captured by D_t , there will always be some level of remaining uncertainty about the unknown properties. This uncertainty, or model confidence, can be incorporated into task planning to ensure robust behaviour.

We consider two task formulations. The first is to simply find an action that is feasible. We use the learned networks, q_ϕ and p_θ , to find the most feasible action:

$$a^* = \operatorname{argmax}_{a \in \mathcal{A}} p(y|x, a) \quad (3)$$

We also consider bandit-style tasks specified by a reward function, $R(a)$, where we must trade off maximizing the reward with current uncertainty:

$$a^* = \operatorname{argmax}_{a \in \mathcal{A}} \mathbb{E}_{y \sim p(y|x, a)} [R(a)] \quad (4)$$

The predictive posterior used in both task formulations, $p(y|x, a)$, is computed using the learned distributions:

$$p(y|x, a) = \int p_\theta(y|x, z, a) q_\phi(z|D_t) dz \quad (5)$$

This is approximated using Monte-Carlo sampling to account for uncertainty arising from limited data.

IV. GRASPING TASK

We focus on grasping objects with unknown physical parameters. Grasping is important for many manipulation tasks and success depends on understanding mass and friction properties. For example, to grasp a hammer with a heavy head, the robot must grasp it as close to its center of mass as possible to avoid slip. Concretely, we aim to predict grasp stability (i.e., force closure) of candidate grasps given an object's observable properties, x_t , which consist of object geometry and grasp pose. The amortized posterior will need to infer center of mass (CoM), friction, and mass from a sequence of interactions to form an accurate prediction. We refer to our model as a *Grasp Neural Process* (GNP).

A. Dataset Generation

All experiments are run on two synthetic datasets: box primitives (*Boxes*), and a post-processed set of objects derived from 191 ShapeNet classes (see Figure 1, center) [5].² For the training phase, we use 1000 unique object geometries with 5 dynamic property samples each. This leads to datasets with 5000 objects and 50 labeled grasps per object.

Object Generation: The CoM is uniformly sampled within the convex hull of the object and lateral friction/mass is chosen from a uniform distribution over [0.1, 0.3].

Grasp Generation: We assume access to object meshes and a floating Panda gripper to check grasp stability. The mesh is first used to sample grasp points that are within a specified antipodal tolerance (30 degrees). We then sample a random gripper orientation around the line connecting the grasp points and ensure there are no collisions with the gripper. Finally a grasp force is uniformly sampled in the range $(5, 20)N$. The stability label is generated by closing the gripper in simulation and applying perturbation forces in random directions (a force closure approximation).

B. Network Architectures

Grasp Neural Processes (GNPs) use an encoder/decoder architecture (representing the inference net and learned feasibility model respectively) with domain-specific structure for grasping data. The encoder accepts an arbitrary number of grasps representing the history of grasps tried so far.

Input Features: Each object is represented by a *global point cloud* of 256 points sampled uniformly from the surface of the mesh. Each grasp is further represented by a local point cloud. The *local point cloud* is in the reference frame of the gripper and only includes points within 3cm of either grasp point. The *global point cloud* allows the network to reason about object-level properties like *volume* and *moments of inertia* while the *local point cloud* allows reasoning about local features like *curvature* and *surface normals*. Each grasp is further represented by the grasp points as well as the grasping force. The *encoder* will have access to the grasp’s label while the *decoder* will not.

Auxiliary Networks: We use PointNets [7] to encode both local and global point clouds. A separate PointNet instance is used for each point cloud type and maps the set of points into a single fixed-length embedding.

Encoder: The GNP encoder takes as input a collection of grasps. It outputs the parameters of a d -dimensional diagonal Gaussian distribution, $\mu_z, \sigma_z \in \mathbb{R}^d$. We repurpose the PointNet architecture to operate over a set of grasps.

Decoder: Finally, the NP decoder takes as input a set of unlabeled grasps, and a latent sample, $z \sim \mathcal{N}(\mu_z, \sigma_z^2)$, to predict grasp stability using an MLP.

V. EXPERIMENTS

In our experiments, we evaluate whether GNPs yield well-performing grasp-detectors with faithful uncertainty representations of the unknown dynamic parameters in the

²The ShapeNet objects were post-processed to be watertight for simulation and rescaled so they could fit in the robot gripper as in [6].

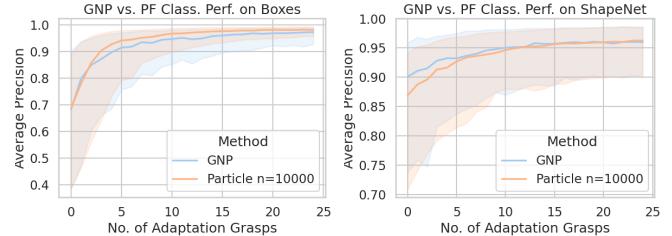


Fig. 4: **GNP vs. PF** The proposed method achieves comparable adaptation performance to an expensive particle filter baseline with many particles. Lines show the median *Average Precision* of the grasp classifier across 500 novel objects. Shaded regions are the top and bottom quartiles.

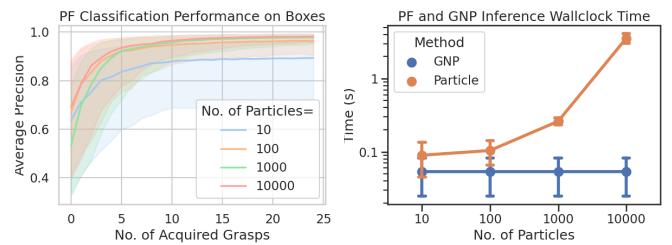


Fig. 5: **Particle Filter Analysis (Boxes)** The particle filter performance is dependent on the number of particles (left) yet the inference time scales poorly with the number of particles (right), compared to the proposed amortized inference method. Timing results represent the mean/std across 2500 inference runs.

ShapeNet and *Boxes* datasets. We first compare the proposed method to an accurate (but expensive) particle-based baseline, achieving comparable performance at a fraction of the cost (Section V-A). Then, we show how the learned posterior can be leveraged for faster information-gathering (Section V-B). Finally, we evaluate the model as part of an uncertainty-aware planner for a grasping task (Sections V-C and V-D) and deploy the model on a real Franka Panda robot.

A. Particle Filter Baseline

To understand the trade-off between efficiency and prediction performance of GNPs, we compare our approach to a particle filter. We reuse our decoder architecture and train it to predict grasp stability using the known dynamic parameters with a standard cross-entropy loss. This model is then used during the *fitting* phase as the observation model of a standard particle filter found in Chopin et al. [8]. The filter is initialized with a varying number of particles ($N = 10, 100, 1000, 10000$) sampled from a uniform distribution over the range of valid dynamic parameters.

Figure 4 shows the performance of our method compared to the particle filter using a large number of particles ($N = 10,000$). We report *average precision*³ of the grasp stability classifier evaluated over time as the model collects more adaptation grasps. We find that GNPs perform comparably to the particle filter at a fraction of the online inference cost

³Average precision characterizes how well we classify positive grasps (balancing precision and recall), which is indicative of how the classifier would perform in downstream tasks discussed later in this section.

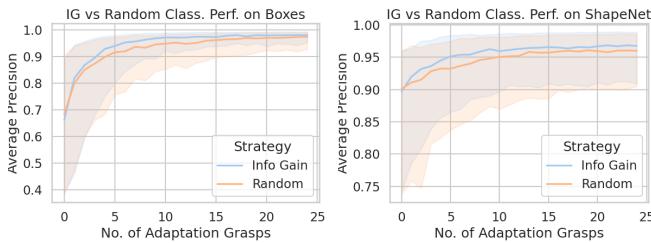


Fig. 6: **Information Gain** GNP encoders can be exploited for faster information-gathering. *Average Precision* per grasp for IG and random adaptation strategies. Solid lines are medians, and error ranges are the upper and lower quartiles across 500 test objects.

used by the baseline as shown in Figure 5. Our amortized prediction is a constant-pass time through the encoder which is significantly faster than the best-performing filters with large numbers of particles.

B. Using Heuristics for Efficient Information Gathering

We evaluate if the amortized posteriors can be used to gather data more efficiently using an information-gain (IG) acquisition function (Section III-B). A useful uncertainty representation should permit the model to reason about what inputs are more informative than others.

Given a trained GNP, we simulate the adaptation phase under the random and IG strategies. At each timestep, 25 unlabeled grasps from the novel object are sampled. The random data-collection strategy chooses a single grasp from this pool. The information-gain strategy computes the expected information-gain metric for each grasp, and greedily chooses the grasp with maximal IG. To evaluate performance, we again report the average precision of the grasp classifier. Our results are shown in Figure 6.

On the Boxes dataset, our results show that a greedy IG heuristic produces better performing classifiers with less acquisition data. This suggest that amortized posteriors are still useful for probabilistic reasoning. On the ShapeNet dataset, the IG heuristic has a smaller performance improvement over random. We attribute this to the significant influence of object geometry on grasp stability. Some objects have irregular geometry that is difficult to learn. For other objects, grasp stability is governed primarily by their geometry (e.g. an hourglass that is only graspable at the center), rather than their underlying dynamic parameters. Since all boxes have similar geometry, the stability of a grasp is far more sensitive to the dynamic parameters than those of ShapeNet.

C. Finding Stable Grasps for Novel Objects

To evaluate the utility of GNPs for practical grasping tasks, we integrate the model into the grasp planner discussed in Section III-C. In simulation, we first generate a novel object and perform the adaptation phase using both IG and random strategies. After each observation, we use a sampling based approximation of Eq. 3 to select the most likely grasp.

In Figure 7, we report how many adaptation grasps are required to reach a desired *grasp success rate* (how often the chosen grasp is actually successful) on ShapeNet and Boxes

datasets (with 500 novel objects per dataset). We consider two task difficulty levels: where the chosen grasp must be either $< 20N$ (easy) or $< 10N$ (hard). Using less force requires better understanding of dynamic properties.

Across both datasets and adaptation strategies, GNPs are able to successfully adapt to novel objects: performance increases with the number of adaptation grasps (requiring as little as 5 grasps to achieve high success rates). The IG strategy shows small improvements for easy grasps ($< 20N$ force; smaller bars are better). For hard grasps ($< 10N$ force), IG shows a larger gain, sometimes needing as few as half the required adaptation grasps for similar performance to random (e.g., to achieve 85% grasp success rate).

D. Uncertainty-Aware Planning

Another task that requires understanding dynamic properties is that of grasping an object with minimal force. Succeeding at this task requires grasping as close to the center of mass as possible to avoid slip. We formalize this task using the reward function:

$$R(G) = \begin{cases} F_{max} - F_G, & \text{grasp is stable,} \\ R_{failure}, & \text{otherwise.} \end{cases}$$

The robot can choose forces in the range $(5, 20)N$ and smaller forces lead to larger reward. F_G is the force of the chosen grasp and we choose a large negative reward ($R_{failure} = -F_{max}$) to prefer robust behaviour where the robot uses larger forces in the face of uncertainty (as opposed to dropping the object).

We perform this task using an uncertainty aware planner (Eq. 4), with sampling-based optimization (using 200 samples). We evaluate task performance with 500 objects after each adaptation grasp. In Figure 8, we report how often the chosen grasp was successful and the normalized regret achieved for successful grasps.

We compare our approach that uses the uncertainty from the amortized posterior distributions (Monte Carlo) to an ablation that only uses the mean of the posterior distribution when planning (Most Likely). This ablation evaluates the utility of learned posteriors for robust grasp planning under uncertainty. Monte Carlo achieves higher success rates, showing robust behaviour even with few adaptation grasps. This comes at a cost of higher initial regrets but task performance improves as the model collects more information, which is the desired robust behaviour.

E. Panda Grasping Demo

We deploy the trained model on a Franka Emika Panda robot. The *training* phase is performed solely in simulation but we perform the *adaptation* and *testing* phases on the real robot (Figure 9) using a heavy block with an offset center of mass (i.e., only grasps on one side of the object will succeed). In the real world, we only sample grasps of $5N$ that are kinematically feasible and grasps are labeled based on whether they slip in the gripper after lifting the arm.

We compare to an antipodal grasp sampler that chooses grasps based solely on object geometry. For our method, we

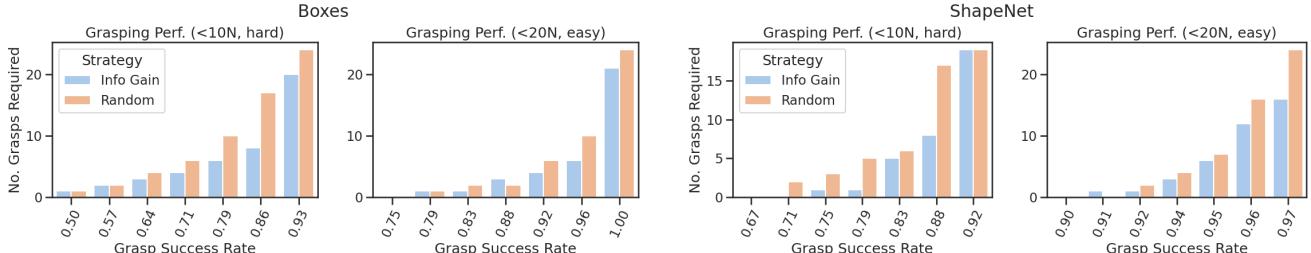


Fig. 7: **GNP Performance as a Grasp Detector** Number of adaptation grasps needed to reach a desired *Grasp Success Rate*. The GNP chooses the most likely grasp considering the current uncertainty arising from limited data. We consider *easy* and *hard* versions of the task that use different force limits. The IG adaptation strategy tends to produce higher success rates with fewer adaptation grasps.

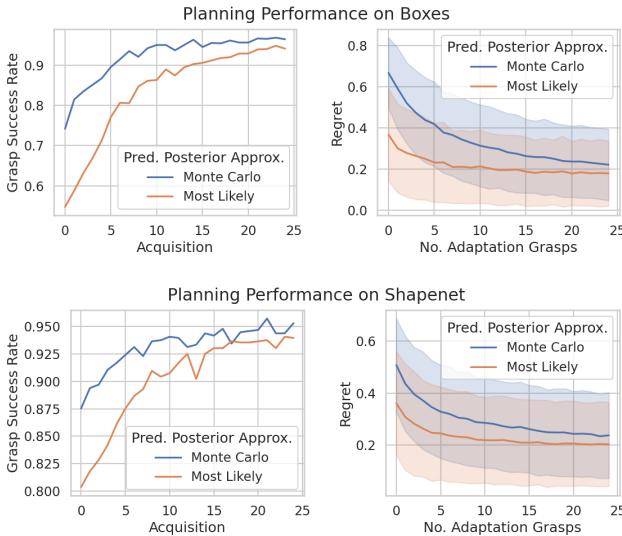


Fig. 8: **Planning** Success and Normalized Regret for the *minimum force grasping task*. We ablate our GNP model to show the impact of the explicit uncertainty representation. Solid lines are means and shaded regions are a standard deviation.

allow the robot 10 adaptation grasps chosen using the IG strategy. The baseline achieves 8/20 successful grasps while our adapted model has 18/20 successes, highlighting the ability to successfully adapt to non-visual properties. Please see the accompanying video for more robot results.

VI. RELATED WORK

Several previous works have introduced Bayesian models for rapid adaptation. For example, [9] and [10] both use latent variables to parameterize a task and Gaussian Processes as the global dynamics. [11] has a similarly structured model but use a BNN as the dynamics model. [12] and [13] use both deep ensembles and low-dimensional latent variables within their models. For grasping, [14] developed adaptive grasp classifiers using a similar approach to ours. However, we extend this approach by considering probabilistic latent spaces and showing that this learned latent space can be used for robust planning and efficient adaptation.

Other work has focused on learning object dynamics with unobservable object properties. In some works, the authors assume the global dynamics are known *a priori* which can inform the values of object-specific properties [15], [16], [17],

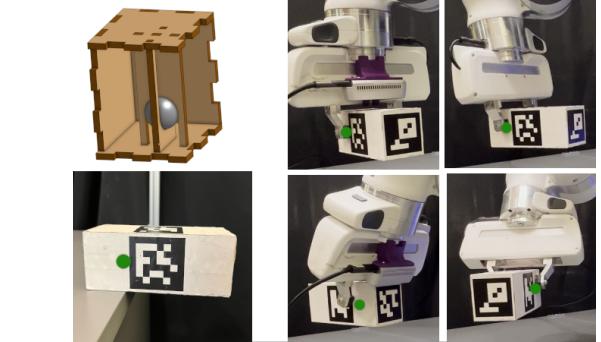


Fig. 9: **Robot Demo** We test our method on a Panda robot by grasping a block with an offset CoM (left). After 10 adaptation grasps, the robot consistently grasps the object near the CoM (green dot). The right 4 grasps show 4 successful evaluation grasps.

[18]. In our work, we desire to jointly infer the dynamics and object properties. Related work which does not assume the dynamics are known often rely on either a fixed dataset [19], [20], [21], [22] at adaption time or a task definition in order to infer the latent properties [23], [9], [24], [10]. Instead, we propose to use information-gain heuristics with respect to model uncertainty which is applicable even when the task is not yet known. Our approach can yield zero-shot performance on a new task, because all the experimentation was performed in a task-agnostic adaptation phase.

Dissentangling object properties through interaction, termed interactive and active perception, [25] is a common strategy. Applications include understanding object kinematics [26], [27], or object geometry [28], [29]. Bandit-style approaches to grasp selection require a large number of grasps to be effective [30], [31]. Unlike our approach, very few methods except for [27], [29] address model uncertainty over dynamic or geometric information, and only [29] leverages uncertainty for efficient information-gathering.

VII. CONCLUSIONS

We present GNPs: an action feasibility model with an explicit representation of uncertainty over a novel object's dynamic parameters. We experimentally verify that the amortized predicted distributions are more computationally efficient than particle filters and lead to robust grasp planners. In the future, we are interested in training dynamic models for *multiple* actions that share the same latent parameter space to enable more robust long-horizon manipulation planning.

REFERENCES

[1] M. Noseworthy, C. Moses, I. Brand, S. Castro, L. Kaelbling, T. Lozano-Pérez, and N. Roy, "Active learning of abstract plan feasibility," in *Proceedings of Robotics Science and Systems (RSS)*, 2021.

[2] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.

[3] M. Garnelo, J. Schwarz, D. Rosenbaum, F. Viola, D. J. Rezende, S. M. A. Eslami, and Y. W. Teh, "Neural Processes," in *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.

[4] N. Houlsby, F. Huszar, Z. Ghahramani, and M. Lengyel, "Bayesian Active Learning for Classification and Preference Learning," in *NeurIPS Workshop on Bayesian optimization, experimental design and bandits: Theory and applications*, 2011.

[5] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An Information-Rich 3D Model Repository," Stanford University — Princeton University — Toyota Technological Institute at Chicago, Tech. Rep. arXiv:1512.03012 [cs.GR], 2015.

[6] C. Eppner, A. Mousavian, and D. Fox, "Acronym: A large-scale grasp dataset based on simulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.

[7] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017, pp. 652–660.

[8] N. Chopin, "A sequential particle filter method for static models," *Biometrika*, vol. 89, no. 3, pp. 539–552, 2002.

[9] F. Doshi-Velez and G. Konidaris, "Hidden parameter markov decision processes: A semiparametric regression approach for discovering latent task parametrizations," in *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, 2016.

[10] S. Sæmundsson, K. Hofmann, and M. P. Deisenroth, "Meta reinforcement learning with latent variable gaussian processes," in *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, 2018.

[11] T. W. Killian, S. Daulton, G. Konidaris, and F. Doshi-Velez, "Robust and Efficient Transfer Learning with Hidden Parameter Markov Decision Processes," in *Proceedings of the 31st Conference on Neural Information Processing Systems (NeurIPS)*, 2017.

[12] C. F. Perez, F. P. Such, and T. Karaletsos, "Generalized Hidden Parameter MDPs Transferable Model-based RL in a Handful of Trials," in *Proceedings of The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, 2020.

[13] S. Belkhale, R. Li, G. Kahn, R. McAllister, R. Calandra, and S. Levine, "Model-Based Meta-Reinforcement Learning for Flight with Suspended Payloads," *IEEE Robotics and Automation Letters (RAL)*, vol. 6, no. 2, 2021.

[14] R. Chen, N. Gao, N. A. Vien, H. Ziesche, and G. Neumann, "Meta-Learning Regrasping Strategies for Physical-Agnostic Objects," in *ICRA Workshop on Scaling Robot Learning*, 2022.

[15] P. W. Battaglia, J. B. Hamrick, and J. B. Tenenbaum, "Simulation as an engine of physical scene understanding," *Proceedings of the National Academy of Sciences (PNAS)*, vol. 110, no. 45, 2013.

[16] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum, "Galileo: Perceiving Physical Object Properties by Integrating a Physics Engine with Deep Learning," 2015, p. 9.

[17] J. Wu, J. Lim, H. Zhang, J. Tenenbaum, and W. Freeman, "Physics 101: Learning Physical Object Properties from Unlabeled Videos," in *Proceedings of the British Machine Vision Conference 2016*, 2016.

[18] L. Zhang and J. C. Trinkle, "The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing," in *2012 IEEE International Conference on Robotics and Automation (ICRA)*, 2012.

[19] P. Y. Lu, S. Kim, and M. Soljačić, "Extracting interpretable physical parameters from spatiotemporal systems using unsupervised learning," *Physical Review X*, vol. 10, no. 3, p. 031056, 2020.

[20] R. Veerapaneni, J. D. Co-Reyes, M. Chang, M. Janner, C. Finn, J. Wu, J. Tenenbaum, and S. Levine, "Entity Abstraction in Visual Model-Based Reinforcement Learning," in *Conference on Robot Learning (CoRL)*, 2019.

[21] Z. Xu, J. Wu, A. Zeng, J. B. Tenenbaum, and S. Song, "DensePhysNet: Learning Dense Physical Object Representations via Multi-step Dynamic Interactions," in *Proceedings of Robotics: Science and Systems (RSS)*, 2019.

[22] D. Zheng, V. Luo, J. Wu, and J. B. Tenenbaum, "Unsupervised Learning of Latent Physical Properties Using Perception-Prediction Networks," in *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, 2018.

[23] J. Denil, P. Agrawal, T. D. Kulkarni, T. Erez, P. Battaglia, and N. de Freitas, "Learning to Perform Physics Experiments via Deep Reinforcement Learning," in *Proceedings of The International Conference on Learning Representations (ICLR)*, 2017.

[24] K. Rakelly, A. Zhou, D. Quillen, C. Finn, and S. Levine, "Efficient Off-Policy Meta-Reinforcement Learning via Probabilistic Context Variables," in *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.

[25] J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, and G. S. Sukhatme, "Interactive Perception: Leveraging Action in Perception and Perception in Action," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1273–1291, Dec. 2017.

[26] S. Y. Gadre, K. Ehsani, and S. Song, "Act the part: Learning interaction strategies for articulated object part discovery," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

[27] R. Martín-Martín and O. Brock, "Coupled recursive estimation for online interactive perception of articulated objects," *The International Journal of Robotics Research (IJRR)*, vol. 41, no. 8, pp. 741–777, July 2022.

[28] A. Herzog, P. Pastor, M. Kalakrishnan, L. Righetti, J. Bohg, T. Asfour, and S. Schaal, "Learning of grasp selection based on shape-templates," *Autonomous Robots*, vol. 36, no. 1, pp. 51–65, Jan. 2014.

[29] S. Dragiev, M. Toussaint, and M. Gienger, "Uncertainty aware grasping and tactile exploration," in *2013 IEEE International Conference on Robotics and Automation (ICRA)*, May 2013.

[30] M. Danielczuk, A. Balakrishna, D. Brown, and K. Goldberg, "Exploratory Grasping: Asymptotically Optimal Algorithms for Grasping Challenging Polyhedral Objects," in *Proceedings of the 2020 Conference on Robot Learning (CoRL)*, 2021.

[31] L. Fu, M. Danielczuk, A. Balakrishna, D. S. Brown, J. Ichnowski, E. Solowjow, and K. Goldberg, "LEGS: Learning Efficient Grasp Sets for Exploratory Grasping," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022.