SIMULATION-BASED BAYESIAN INFERENCE FOR MULTI-FINGERED ROBOTIC GRASPING

Norman Marlier
University of Liège
norman.marlier@uliege.be

Olivier Brüls
University of Liège
o.bruls@uliege.be

Gilles Louppe
University of Liège
g.louppe@uliege.be

ABSTRACT

Multi-fingered robotic grasping is an undeniable stepping stone to universal picking and dexterous manipulation. Yet, multi-fingered grippers remain challenging to control because of their rich nonsmooth contact dynamics or because of sensor noise. In this work, we aim to plan hand configurations by performing Bayesian posterior inference through the full stochastic forward simulation of the robot in its environment, hence robustly accounting for many of the uncertainties in the system. While previous methods either relied on simplified surrogates of the likelihood function or attempted to learn to directly predict maximum likelihood estimates, we bring a novel simulation-based approach for full Bayesian inference based on a deep neural network surrogate of the likelihood-to-evidence ratio. Hand configurations are found by directly optimizing through the resulting amortized and differentiable expression for the posterior. The geometry of the configuration space is accounted for by proposing a Riemannian manifold optimization procedure through the neural posterior. Simulation and physical benchmarks demonstrate the high success rate of the procedure.

*K*eywords Multi-fingered grasping, Bayesian inference, Robot learning

Two of the grand challenges for the deployment of robots in warehouses, assembly lines or homes are universal picking, *i.e* the ability to robustly grasp a large variety of objects in diverse environments, and dexterous manipulation, *i.e* the ability to manipulate objects to perform useful actions. Multi-fingered robotic grasping represents a promising avenue towards these objectives and has the potential to greatly improve and facilitate human-machine interactions. However, due to the wide variety of object shapes, sensor noise and nonsmooth contact dynamics, multi-fingered grippers remain challenging to control. In addition, multi-fingered grasps entail high dimensional configuration spaces compared to two-fingered grasps, making them difficult to optimize and plan.

Early analytical approaches [1, 2] for planning multi-fingered grasps rely on force analysis where a metric is optimized based on the laws of mechanics and 3D models. These analytical methods require the accurate knowledge of many model and world parameters, which is practically difficult to achieve in real settings. By contrast, learning-based methods are getting more and more established because they can extract useful information from raw sensor data [3]. These methods typically make use of machine learning or deep learning models for learning a mapping between grasp configurations and their success. This strategy has been applied on top-down grasping with a parallel jaw gripper and has demonstrated excellent results [4, 5, 6, 7, 8], including its extension to 6 DOF [9, 10].

In this work, we consider the more challenging setting of robust multi-fingered grasping plans, including the 6 DOF grasp poses and the finger configuration. By framing the problem as an inference task, we demonstrate the generic usefulness and applicability of likelihood-free Bayesian inference algorithms to difficult robotic tasks such as multi-fingered grasping considered here. We summarize our contributions as follow:

- We bring simulation-based Bayesian inference methods [11] to multi-fingered grasping. By learning a model for the likelihood-to-evidence ratio and using an analytical prior, we derive an amortized and differentiable posterior for the hand configurations.

- We make use of Riemannian manifold optimization to deal with the nonlinearity of the configuration space, in particular of the 3D rotation group tied to the hand pose.

Figure 1: Our benchmark scene for multi-fingered grasping. The pose of the hand $(\mathbf{x}, \mathbf{R})$ is defined in the local object frame. The depth camera produces an image i of the scene.

- We validate our approach on a realistic and challenging 3-finger gripper setup, through both simulation and real-world physical benchmarks. Results demonstrate its high success rate.

1 Problem statement

Description We consider the problem of grasping a rigid and stable body with a multi-fingered gripper, as illustrated in Fig. 1. The object $\mathcal{O}$ is modelled as a 3D surface mesh and its centroid stands on a table at a location $(x_\mathcal{O}, y_\mathcal{O}, z_\mathcal{O})$ with a rotation $\varphi_{z,\mathcal{O}}$ around the z-axis in the world reference frame $\mathcal{F}_W$. We refer to its 2D pose $(x_\mathcal{O}, y_\mathcal{O}, \varphi_{z,\mathcal{O}})$ as $\mathbf{p}_\mathcal{O} \in \mathbb{R}^2 \times \mathrm{SO}(2)$. The hand configuration $\mathbf{h} \in \mathcal{H} = \mathbb{R}^3 \times \mathrm{SO}(3) \times \mathcal{G}$ is defined as the combination of the pose $(\mathbf{x}, \mathbf{R}) \in \mathbb{R}^3 \times \mathrm{SO}(3)$ of the hand and the type $g \in \mathcal{G} = \{\text{basic}, \text{wide}, \text{pinch}\}$ of the grasp. The hand pose $(\mathbf{x}, \mathbf{R})$ is defined with respect to the world frame coordinate. The robot evolves in a 3D workspace observed with a fixed depth camera producing images $i \in \mathcal{I}$. The goal is to find a robust hand configuration $\mathbf{h}^*$ with respect to a given binary metric $S = \{0, 1\}$, where $S = 1$ indicates a successful grasp.

Probabilistic modeling We model the scene and the grasping task according to the Bayesian network shown in Fig. 2a. The variables $S, i, \mathbf{h}, \mathcal{O}$ and $\mathbf{p}_\mathcal{O}$ are modelled as random variables in order to capture the noise in the robot or in the depth camera, as well as our prior beliefs about the hand configuration, the object or its pose. The structure of the Bayesian network is motivated by the fact that $\mathbf{h}, \mathcal{O}$ and $\mathbf{p}_\mathcal{O}$ are independent, while S is dependent on $\mathbf{h}, \mathcal{O}$ and $\mathbf{p}_\mathcal{O}$ and i is dependent on $\mathcal{O}$ and $\mathbf{p}_\mathcal{O}$. This structure also enables a direct and straightforward way to generate data: $\mathbf{h}, \mathcal{O}$ and $\mathbf{p}_\mathcal{O}$ are sampled from their respective prior distributions while S and i can be generated using forward physical simulators for the grasping and the camera.

The prior distribution $p(\mathbf{x})$ of the spatial position is uniformly distributed between the extreme values $\mathbf{x}_{\text{lim}} = (x_{\text{low}}, y_{\text{low}}, z_{\text{low}}, x_{\text{high}}, y_{\text{high}}, z_{\text{high}})$, chosen to be within the range of physical dimensions of the gripper and the biggest object. It emphasizes our ignorance about interesting regions of space for grasping. The rotation $\mathbf{R}$ is parameterized with a quaternion. A quaternion $\mathbf{q}$ is an element of the quaternion group $\mathbb{H}$, in the form $\mathbf{q} = q_0 \mathbf{1} + q_1 \mathbf{i} + q_2 \mathbf{j} + q_3 \mathbf{k} = (q_0, q_1, q_2, q_3)^T$ with $(q_0, q_1, q_2, q_3)^T \in \mathbb{R}^4$ and $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$. The conjugate $\bar{\mathbf{q}}$ of quaterion $\mathbf{q}$ is given by $\bar{\mathbf{q}} := q_0 \mathbf{1} - q_1 \mathbf{i} - q_2 \mathbf{j} - q_3 \mathbf{k}$. A unit quaternion, called *versor*, $\mathbf{q}_1 \in \mathbb{H}_1$ has a unit norm defined as $\|\mathbf{q}\| = \sqrt{\mathbf{q}\bar{\mathbf{q}}} = 1$. They give a more compact representation than rotation matrices and avoid gimbal lock and singularities. Unit quaternions can be identified with the elements of a hyperspherical manifold $\mathbb{S}^3$ embedded into $\mathbb{R}^4$. Moreover, $\mathbb{S}^3$ is a double covering of $\mathrm{SO}(3)$, meaning that antipodal points $\pm \mathbf{q}$ represent the same rotation, which implies that $p(\mathbf{q}; \cdot) = p(-\mathbf{q}; \cdot)$. We define the prior $p(\mathbf{q})$ as a mixture of *power-spherical* distributions [12] with 4 modes μ_i. Each mode is a mixture that satisfies $p(\mathbf{q}; \cdot) = p(-\mathbf{q}; \cdot)$. In total, we have

$$p(\mathbf{q}) = \frac{1}{N} \sum_{i=1}^{N=4} \frac{\text{PowerSpherical}(\mathbf{q}; \mu_i, \kappa)}{2} + \frac{\text{PowerSpherical}(\mathbf{q}; -\mu_i, \kappa)}{2}. \tag{1}$$

These modes μ_i encode information about the orientation of the gripper and share the same concentration factor $\kappa = 30$. To grasp an object, the gripper point toward the table and thus toward the object – an informed prior which indeed

Simulation-based Bayesian inference for multi-fingered robotic grasping

| | (a) | | | (b) |
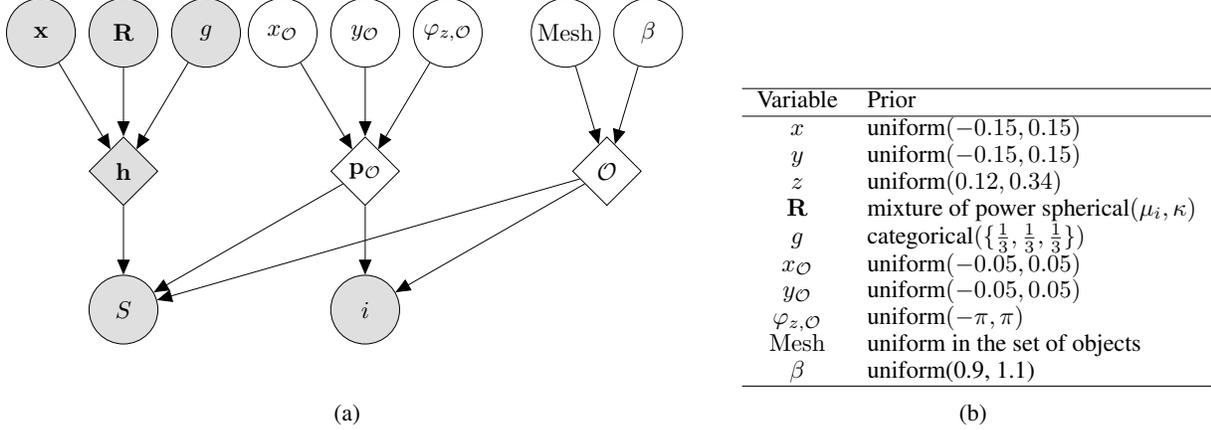
Figure 2: (a) Probabilistic graphical model of the environment. Gray nodes correspond to observed variables and white nodes to unobserved variables. (b) Prior distributions.

results in sufficiently many successful grasps. We then define four rotations, separated by a rotation of $\frac{\pi}{2}$ around the z-axis (see Fig. 5 in Appendix B). In this way, our prior covers a large part of the rotation space and is sufficiently informative by contrast to a uniform prior over the unit sphere $\mathbb{S}^3$. The grasp type g is uniformly distributed between the three types basic, wide and pinch. These three modes modulates the spacing between the fingers in the opposite side of the thumb. Finally, $p(\mathbf{h}) = p(\mathbf{x})p(\mathbf{R})p(g)$.

The prior $p(\mathcal{O}) = p(\text{Mesh})p(\beta)$ is a discrete uniform distribution over a fixed set of object meshes and a uniform distribution for the scaling factor β. Finally, the prior $p(\mathbf{p}_{\mathcal{O}})$ captures our belief that the object can be everywhere on the table with any rotation around the vertical axis. For this reason, uniform distributions are used for all three parameters $x_{\mathcal{O}}, y_{\mathcal{O}}, \varphi_{z,\mathcal{O}}$. Table 2b summarizes the prior distributions.

Given our probabilistic graphical model, we finally formulate the problem of grasping as the Bayesian inference of the hand configuration $\mathbf{h}^*$ that is a posteriori the most likely given a successful grasp and an observation i. That is, we are seeking for the maximum a posteriori (MAP) estimate

$$\mathbf{h}^* = \arg\max_{\mathbf{h}} \ p(\mathbf{h}|S = 1, i). \tag{2}$$

2 Likelihood-free Bayesian inference for multi-fingered grasping

2.1 Density ratio estimation

From the Bayes's rule, the posterior of the hand configuration is

$$p(\mathbf{h}|S, i) = \frac{p(S, i|\mathbf{h})}{p(S, i)} p(\mathbf{h}). \tag{3}$$

The likelihood function $p(S, i|\mathbf{h})$ and the evidence $p(S, i)$ are both intractable, which makes standard Bayesian inference procedures such as Markov chain Monte Carlo unusable. However, drawing samples from forward models remains feasible with physical simulators, hence enabling likelihood-free Bayesian inference algorithms.

First, we express the likelihood-to-evidence ratio as a product of two individual ratios,

$$r(S, i|\mathbf{h}) = \frac{p(S, i|\mathbf{h})}{p(S, i)} = \frac{p(S|\mathbf{h})}{p(S)} \frac{p(i|S, \mathbf{h})}{p(i|S)} = r(S|\mathbf{h})r(i|S, \mathbf{h}). \tag{4}$$

By adapting the approach described in [13, 14] for likelihood ratio estimation, we train two neural network classifiers d_ϕ and d_θ that we will use to approximate $r(S|\mathbf{h})$ and $r(i|S, \mathbf{h})$. The first network d_ϕ is trained to distinguish positive tuples $(S, \mathbf{h})$ (labeled $y = 1$) sampled from the joint distribution $p(S, \mathbf{h})$ against negative tuples (labeled $y = 0$) sampled from the product of marginals $p(S)p(\mathbf{h})$. The Bayes optimal classifier $d^*(S, \mathbf{h})$ that minimizes the cross-entropy loss is given by

$$d^*(S, \mathbf{h}) = \frac{p(S, \mathbf{h})}{p(S)p(\mathbf{h}) + p(S, \mathbf{h})}, \tag{5}$$

3

Simulation-based Bayesian inference for multi-fingered robotic grasping
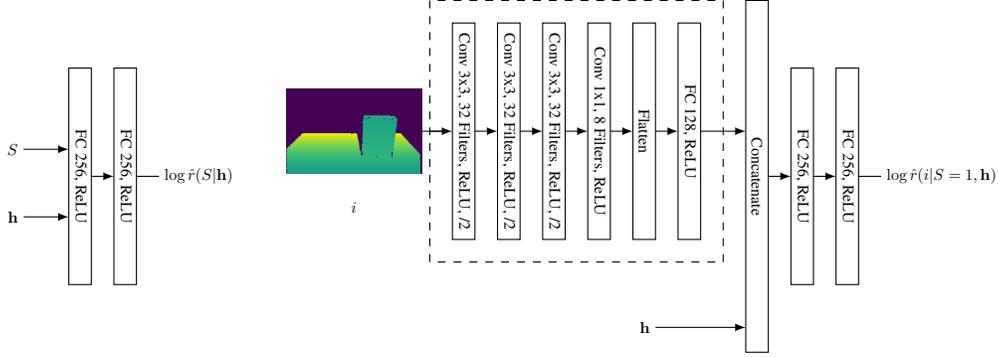
Figure 3: Neural network architectures of the classifiers d_ϕ and d_θ used to respectively approximate the likelihood ratios $r(S|\mathbf{h})$ and $r(i|S=1,\mathbf{h})$.

which recovers the likelihood ratio $r(S|\mathbf{h})$ as

$$\frac{d^*(S,\mathbf{h})}{1-d^*(S,\mathbf{h})} = \frac{p(S,\mathbf{h})}{p(S)p(\mathbf{h})} = \frac{p(S|\mathbf{h})}{p(S)}. \tag{6}$$

Therefore, by modelling the classifier with a neural network d_ϕ trained on the binary classification problem, we obtain an approximate but amortized and differentiable likelihood ratio

$$\hat{r}(S|\mathbf{h}) = \frac{d_\phi(S,\mathbf{h})}{1-d_\phi(S,\mathbf{h})}. \tag{7}$$

The second network d_θ is trained similarly, over positive tuples $(i,\mathbf{h})$ (labeled $(y=1)$ sampled from the conditional joint distribution $p(i,\mathbf{h}|S=1)$ against negative tuples $(i,\mathbf{h})$ (labeled $y=0$) sampled from the product of marginals $p(i|S=1)p(\mathbf{h}|S=1)$. Using the same likelihood ratio trick, we obtain

$$\hat{r}(i|S=1,\mathbf{h}) = \frac{d_\theta(i,\mathbf{h})}{1-d_\theta(i,\mathbf{h})}. \tag{8}$$

Finally, the likelihood ratios are combined with the prior to approximate the posterior as

$$\hat{p}(\mathbf{h}|S=1,i) = \hat{r}(i|S=1,\mathbf{h})\hat{r}(S=1|\mathbf{h})p(\mathbf{h}), \tag{9}$$

which enables immediate posterior inference despite the initial intractability of the likelihood function $p(S,i|\mathbf{h})$ and of the evidence $p(S,i)$.

The neural network classifiers d_ϕ and d_θ are architectured as in Fig. 3. In d_θ, the camera image i of size $640 \times 480 \times 1$ is pre-processed by scaling the depths in the interval $\{0\} \cup [0.45,1]$ and resized to $256 \times 160 \times 1$. Then, i is fed to a convolutional network made of four convolutional layers followed by a fully connected layer and which goal is to produce a vector embedding of the image. The image embedding and $\mathbf{h}$ are then both fed to a subsequent network made of 2 fully connected layers. The hand configuration $\mathbf{h}$ enters the neural network as a 1×13 vector where the rotation matrix $\mathbf{R}$ is flattened [15, 16] and the grasp type g is passed through an embedding. In d_ϕ, both S and $\mathbf{h}$ are directly fed to a network made of 2 fully connected layers. The parameters ϕ and θ are optimized by following Algorithm 1 (Appendix A), using Adam as optimizer.

Finally, we note that the factorization of the likelihood-to-evidence ratio forces the two ratio estimators to focus their respective capacity on the information brought by S and i. Because of the high discriminative power of S, training instead a single ratio taking both S and i as inputs would indeed lead to an estimator that usually discards the smaller information brought in i.

2.2 Maximum a posteriori estimation

Due to the intractability of the likelihood function and of the evidence, Eq. (2) cannot be solved analytically nor numerically. We rely instead on the approximation given by the likelihood-to-evidence ratio $\hat{r}$ to find an approximation of the maximum a posteriori (MAP) estimate as

$$\hat{\mathbf{h}}^* = \arg\max_{\mathbf{h}} \hat{r}(S=1,i|\mathbf{h})p(\mathbf{h}) \tag{10}$$

$$= \arg\min_{\mathbf{h}} -\log \hat{r}(S=1,i|\mathbf{h})p(\mathbf{h}), \tag{11}$$

4

which we solve using gradient descent. For a given g, the gradient of Eq. (11) decomposes as

$$-\nabla_{(\mathbf{x},\mathbf{R})} \log \hat{r}(S,i|\mathbf{h})p(\mathbf{h}) = -\nabla_{(\mathbf{x},\mathbf{R})} \log \hat{r}(S,i|\mathbf{h}) - \nabla_{(\mathbf{x},\mathbf{R})} \log p(\mathbf{h}). \tag{12}$$

Our closed-form prior $p(\mathbf{h})$ has analytical gradients. In fact, uniform distributions are set to have null gradient everywhere in the domain. Therefore, $\nabla_{\mathbf{x}}p(\mathbf{h}) = \mathbf{0}$. By contrast, $p(\mathbf{R})$ is a weakly informative prior and has a non null gradient from the power spherical distribution. Its derivative with respect to $\mathbf{q}$ is

$$\begin{aligned}
\nabla_{\mathbf{q}}p(\mathbf{q};\mu,\kappa) &= C(\kappa)\kappa(1+\mu^T\mathbf{q})^{\kappa-1}\nabla_{\mathbf{q}}(1+\mu^T\mathbf{q}) \\
&= C(\kappa)\kappa\mu(1+\mu^T\mathbf{q})^{\kappa-1},
\end{aligned} \tag{13}$$

where $C(\kappa)$ is the normalization term. Since the likelihood-to-evidence ratio estimator $\hat{r}$ is modelled by a neural network, it is fully differentiable with respect to its inputs and its gradients can be computed by automatic differentiation. However, not all variables of the problem are Euclidean variables and naively performing gradient descent would violate our geometric assumptions. Let us consider a variable $\mathcal{Z}$ on the smooth Riemannian manifold $\mathcal{M} = \mathbb{R}^3 \times \mathrm{SO}(3)$ with tangent space $\mathcal{T}_{\mathcal{Z}}\mathcal{M}$ and a function $f : \mathcal{M} \to \mathbb{R}$. Since $\mathrm{SO}(3)$ is embedded in the set of 3×3 matrices $\mathbb{R}^{3\times3}$, f can be evaluated on $\mathbb{R}^3 \times \mathbb{R}^{3\times3}$, leading to the definition of the Euclidean gradients $\nabla f(\mathcal{Z}) \in \mathbb{R}^3 \times \mathbb{R}^{3\times3}$. In turn, these Euclidean gradients can be transformed into their Riemannian counterparts $\mathrm{grad} f(\mathcal{Z})$ via orthogonal projection $\mathbf{P}_{\mathcal{Z}}$ into the tangent space $\mathcal{T}_{\mathcal{Z}}\mathcal{M}$ [17, 18]. Therefore,

$$\mathrm{grad} f(\mathcal{Z}) = \mathbf{P}_{\mathcal{Z}}(\nabla f(\mathcal{Z})) \tag{14}$$

where the orthogonal projection onto $\mathbb{R}^3$ is the identity $\mathbb{I}_3$ and the orthogonal projection onto $\mathrm{SO}(3)$ at $\xi \in \mathrm{SO}(3)$ is $\xi \mathrm{skew}(\xi^T \nabla f(\xi))$ where $\mathrm{skew}(A) := \frac{1}{2}(A - A^T)$. Thus, we can solve Eq. (11) by projecting Euclidean gradients of Eq. (12) to the tangent space $\mathcal{T}_{\mathcal{Z}}\mathcal{M}$ and plug them to a manifold optimization procedure. In our experiments, we use the geometrical conjugate gradient method [17] implemented in Pymanopt [19] to perform 20 optimization steps and we scan for the best value of g. For completeness, the full optimization algorithm is provided in Algorithm 2 (Appendix A).

3 Experiments

For training, we use 19 objects from the YCB data [20] (see Fig. 6 in Appendix C) together with 5 objects from the ShapeNet dataset [21], for a total of 24 types of objects. We selected a diverse range of objects compatible with the geometry of our gripper. In simulation, the success rate was evaluated on the 19 objects used for training, as well as on 5 new unseen objects from the YCB data (see Fig. 7 in Appendix C). Only the 19 objects from YCB are used in the real setup.

3.1 Data generation

Grasp generative model A physical simulator is used to sample from $p(S|\mathbf{h},\mathcal{O},\mathbf{p}_{\mathcal{O}})$. Hand configurations, objects and object poses are sampled from their priors $p(\mathbf{h})$, $p(\mathcal{O})$ and $p(\mathbf{p}_{\mathcal{O}})$, and are then submitted to a lift test. First, a planner generates a trajectory in the joint space to bring the gripper to the hand configuration $\mathbf{h}$. If the pose is not reachable, the test fails. Otherwise, the gripper closes its fingers until contact with the object or the gripper itself. Then, the robot lifts possibly the object to a given height. If the object is held in the gripper, the grasp is considered as successful. Simulations are performed using Pybullet [22]. We use volumetric hierarchical approximate decomposition to get convex meshes of objects from *obj* files for collision detection [23].

Sensor generative model The sensor generative model $p(i|\mathbf{p}_{\mathcal{O}},\mathcal{O})$ is implemented in Pybullet, with an approach similar to the Blensor sensor framework [24] used to render depth images from a Kinect model sensor. Simulating a structured-light sensor allows a better transfer to the real setup. Objects and poses are sampled from their priors, $\mathcal{O} \sim p(\mathcal{O}), \mathbf{p}_{\mathcal{O}} \sim p(\mathbf{p}_{\mathcal{O}})$. Then, the object is placed and an image $i \sim p(i|\mathbf{p}_{\mathcal{O}},\mathcal{O})$ is generated.

Domain randomization for sim-to-real transfer Generative models in simulation differ from their real world counterparts due to incorrect physical modelling and inaccurate physical parameters. This *reality gap* may lead to failures from our model because the synthetic data and the real data distributions are different. To overcome this issue, we use *domain randomization* [25] with nuisance parameters on the position and the orientation of the camera, the minimum and maximum distance of the depth, the field of view, and the coefficient of lateral and spinning frictions μ and γ. Domain randomization avoids the precise calibration of both the grasp and the image simulators, which can be very difficult and costly. We use uniform distributions for the nuisance parameters which are difficult to measure with $\pm 2\%$ error and Gaussian distributions for easily measurable parameters. For the orientation of the camera, a multivariate normal variable $\eta \sim \mathcal{N}(\mathbf{0},\Sigma), \Sigma = \mathrm{diag}(\sigma_\alpha = 0.002, \sigma_\beta = 0.01, \sigma_\gamma = 0.002)$ is drawn and then mapped to $\mathrm{SO}(3)$ using the exponential map.

Grasping inference strategy		Success rate	
		Sim	*Real*
Prior based	$\mathbf{h} \sim p(\mathbf{h})$	0.8%	-
	$\mathbf{h} = \arg\max p(\mathbf{h})$	0.6%	-
Metric based	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(S = 1|\mathbf{h})$	43%	-
	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(\mathbf{h}|S = 1)$	44%	46%
Partial observation based	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(S = 1, i|\mathbf{h})$	64%/71%	-
	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(\mathbf{h}|S = 1, i)$	71%/**75**%	70%
Full observation based (ideal)	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(S = 1, \mathcal{O}, \mu, \beta|\mathbf{h})$	80%	-
	$\mathbf{h} = \arg\max_{\mathbf{h}} \hat{p}(\mathbf{h}|S = 1, \mathcal{O}, \mu, \beta)$	85%	-

Table 1: Grasping success rate for various inference strategies of the hand configuration. The success rate obtained by performing Bayesian posterior inference through the full forward simulation reaches 71% for objects seen during training and 75% for 5 new objects. In real experiments, the success rate reaches 70%.

3.2 Simulation benchmarks

We evaluate the performance of our approach incrementally, adding algorithmic components one by one to assess their respective marginal benefits. For each inference strategy, we estimate the success rate over 1000 grasping attempts for randomly sampled objects and camera images. Nuisances parameters are resampled in simulation when evaluating the success. Our general results are summarized in Table 1, while supplementary results for each individual category of objects can be found in Appendix D. We first report results for strategies maximizing the (conditional) densities $p(\mathbf{h}|\cdot)$ of hand configurations. Optimizing for the maximum a priori estimate $\mathbf{h} = \arg\max p(\mathbf{h})$, without conditioning on success or an observation of the scene, leads to a very low success rate of 0.6%. As expected, these results are too poor to be usable but they should underline the informativeness of the prior. Sampling hand configurations from a uniform prior would instead result in a much smaller success rate, by about one order of magnitude (less than 0.1%). When conditioning on the expected success $S = 1$, performance increases to 44%. Taking both the expected success $S = 1$ and the image i into account and following the inference methodology proposed in Section 2, leads to an even larger success rate of 71% for the maximum a posteriori estimates. For the 5 new objects, we reach a comparable success rate of 75%, which demonstrates the good generalization of the approach. In comparison, had the properties $\mathcal{O}$, μ, and β of the object been perfectly known, the success rate would reach 85%, which shows that the convolutional network manages to extract most of the relevant information from the observation i. Table 1 also reports results for maximum likelihood estimates, achieving success rates of 43%, 64% and 80% when maximizing the likelihoods $p(S = 1|\mathbf{h})$, $p(S = 1, i|\mathbf{h})$, and $p(S = 1, \mathcal{O}, \mu, \beta|\mathbf{h})$ respectively. Note that maximizing $p(S = 1|i, \mathbf{h})$ and $p(S = 1|\mathcal{O}, \mu, \beta, \mathbf{h})$ would give the same result since i and $\mathcal{O}, \mu, \beta$ are independent from $\mathbf{h}$. Our informative prior can explain the difference in success rates between the MAP and the MLE estimates and motivates the use of a Bayesian approach.

Not only our framework can be used for grasp planning, it also provides immediate access to the full posterior $p(\mathbf{h}|S = 1, i)$ of hand configurations. As an illustrating example, we extract the marginal posterior densities $p(\mathbf{x}|S = 1, u)$, $p(\mathbf{R}|S = 1, i)$ and $p(g|S = 1, i)$ for the simulated scene of Fig. 1a, with the box centered at $(0, 0)$ without any rotation. The resulting posterior is shown in Fig. 4. First, $p(\mathbf{x}|S = 1, i)$ shows the distribution in space of the hand configuration $\mathbf{h}$. The concentration along the x-axis and z-axis are high, meaning that high density regions are located slightly behind and in front of the box, at a given height related to the geometrical dimensions of the box. Concerning the y-axis, the posterior fails to capture the symmetry and places all the density at the right of the box. Overall, the positions $\mathbf{x}$ which are the most likely to give a successful grasp are on the right corner of the box. This is underlined by the posterior of $p(\mathbf{R}|S = 1, i)$. The red dots correspond to the density of the x-axis, the green dots to the y-axis and the blue dots to the z-axis. The x-axis has one mode, directed toward the table, inherited from the prior and slightly deviates to the right. The y-axis, however, has only two antipodal modes by contrast to the prior. These modes correspond to the situation in which the fingers are placed on the front surface or the back surface. The z-axis can be constructed by taking the cross product between x and y. Uncertainties from x and y are propagated, leading to two antipodal modes with lower concentration than y. Finally, the posterior $p(g|S = 1, i)$ for the grasp type indicates a preference towards pinch and wide modes over the basic mode. Pinch mode is preferred when the position $\mathbf{x}$ is far from the right corner while the wide mode is mainly used when $\mathbf{x}$ is located near the right corner.

3.3 Physical benchmarks

We carried out experiments with a Robotiq 3-finger gripper attached to a UR5 robotic arm, as shown in Figure 1b. A Kinect v1 provides depth images and is rigidly linked to the robot base. The robotic arm is controlled in position in

Simulation-based Bayesian inference for multi-fingered robotic grasping
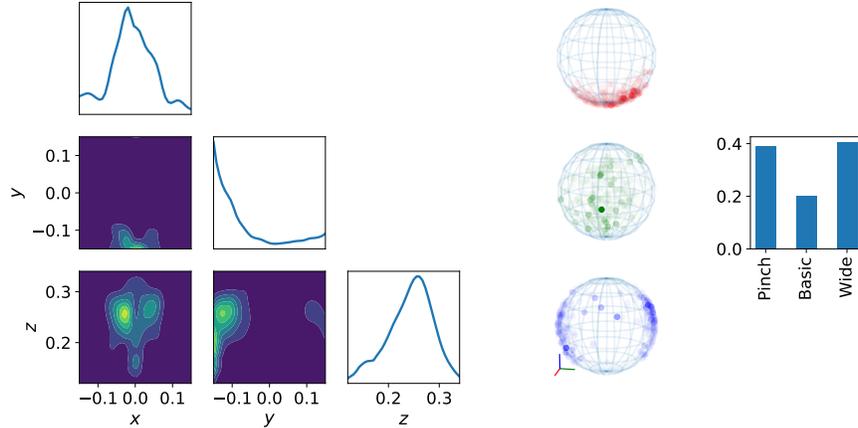
Figure 4: Posterior $p(\mathbf{h}|S=1,i)$ for the setup in Fig. 1a with the box centered at $(0,0)$. (Left) Posterior $p(\mathbf{x}|S=1,i)$, (Middle) Posterior $p(\mathbf{R}|S=1,i)$, (Right) Posterior $p(g|S=1,i)$.

the joint space. Communications are performed within the ROS framework. We calibrate the centre of the table by computing the trajectory in the simulator and then sending it to the real robot. We perform 10 trials per object, for a total of 190 grasps. Objects are placed at best at $(x_{\mathcal{O}}=0, y_{\mathcal{O}}=0, \theta_{\mathcal{O}}=0)$. As shown in Table 1, our success rate of 70% is similar than in simulation, which indicates that the simulation-to-reality transfer works well, at least on average. These results also demonstrate competitive performance with respect to related works (see Section 4), although this remains difficult to establish because of the distinct hardware setups. We observe that failures are mainly behavioral and geometric. Behavioral failures arise when the simulation does not model the physics correctly. For example, in the real setup, the bowl slides on the table when the gripper closes its fingers, while in simulation, the bowl is just lifted. We could reduce these errors by using a more accurate simulator. Geometric failures arise when there is a shift in the location or in the orientation of the object. Most of the time, the robot either collides with the object or misses it for smaller objects. These failures could be avoided using a more precise calibration, additional sensors, or more extensive domain randomization. Finally, the time of computation is reasonable (from 5 to 10s), but could be decreased by tuning the architecture of the neural network, lowering the number of optimization steps, or using more powerful hardware. We leave this for future work.

4   Related work

Over the last decade, progress in multi-fingered robotic grasping has been steady [26, 27, 28, 29, 30, 31, 32] thanks to differentiable models such as neural networks. Unfortunately, the variety of robotic hands, their actuation modes and sensor inputs, and the lack of standard benchmarks, make it difficult to compare these advances fairly against one another.

From early works, [26] identify poses and fingertip positions of stable grasps with a deep neural network from RGB-D images and use a planner (GraspIt!) based on simulated annealing to determine the best hand configurations. They reach a success rate of 75% over a set of 8 objects but suffer from slow execution times (16.6s on average). By contrast, our method is faster and reaches comparable performance over a larger set of objects. [29] use generative adversarial networks to sample both grasp poses and finger joints efficiently based on RGB-D images. While being a fast approach, they reach only 60% of success rate in real experiments. The work of [27] is the most similar to ours. They perform grasp planning as probabilistic inference via a classifier trained to predict the success of a grasp. They retrieve maximum a posteriori estimates using gradient ascent and the fact that the classifier is fully differentiable. The prior distribution is fitted with a Gaussian mixture model from the dataset. In contrast, our method uses an analytical prior based on power-spherical distributions, does not require an external grasp sampler, and relies on a neural classifier to approximate the likelihood-to-evidence ratio. Similarly, [28] compute maximum likelihood estimates of the hand configuration by making use of gradients provided by a neural network. Finally, both of these works treat the rotations with Euler angles and optimize them as real numbers with boundary constraints. This representation is not suitable for a neural network according to [16]. Instead, our optimization relies on Riemannian conjugate gradients, which preserve the geometrical structure of the rotation group. Other interesting approaches to multi-fingered grasping include the use of deep reinforcement learning based on vision and tactile sensors [31], or the use of tactile information only for learning a closed-loop controller [32].

7

From a statistical perspective, several Bayesian likelihood-free inference algorithms [33, 34, 35, 36, 37, 38, 13] have been developed to carry out inference when the likelihood function is implicit and intractable. These methods operate by approximating the posterior through rejection sampling or by learning parts of the Bayes' rule, such as the likelihood function, the likelihood-to-evidence ratio, or the posterior itself. These algorithms have been used across a wide range of scientific disciplines such as particle physics, neuroscience, biology, or cosmology [11]. To the best of our knowledge, our work is one of the first to apply one of those for the direct planning successful grasps. More specifically, we rely here on the amortized inference approach of [13] to carry out inference within seconds for any new observation i. In contrast, an approach such as ABC [33, 34] could take up to hours to determine a single hand configuration $\mathbf{h}$ since data would need to be simulated on-the-fly for each observation i due to the lack of amortization of ABC. Neural posterior estimation [38] is also amortizable but would have required new methodological developments to be applicable on distributions defined on manifolds, such as those needed here for the rotational part of the pose.

5   Summary and future work

We demonstrate the usefulness and applicability of simulation-based Bayesian inference to multi-fingered grasping. The approach is generic yet powerful because it can work with any simulator, thereby incorporating from the simplest to the more sophisticated piece of domain knowledge, while leveraging all recent developments from deep learning to solve the Bayesian inference problem. Maximum a posteriori hand configurations are found by directly optimizing through the resulting amortized and differentiable expression for the posterior. The geometry of the configuration space is accounted for by proposing a Riemannian manifold optimization procedure through the neural posterior. We demonstrate a working proof-of-concept achieving robust multi-fingered grasping, both in simulation and in real experiments thanks to domain randomization. Our success rate is comparable to previous works.

Acknowledgments

References

[1] Domenico Prattichizzo and Jeffrey C. Trinkle. *Grasping*, pages 671–700. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[2] Carlo Ferrari and John F Canny. Planning optimal grasps. In *ICRA*, volume 3, pages 2290–2295, 1992.

[3] Kilian Kleeberger, Richard Bormann, Werner Kraus, and Marco F Huber. A survey on learning-based robotic grasping. *Current Robotics Reports*, pages 1–11, 2020.

[4] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio, and Ken Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. In *Proceedings of Robotics: Science and Systems*, Cambridge, Massachusetts, July 2017.

[5] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, pages 651–673, 2018.

[6] Ulrich Viereck, Andreas Pas, Kate Saenko, and Robert Platt. Learning a visuomotor controller for real world robotic grasping using simulated depth images. In *Conference on Robot Learning*, pages 291–300. PMLR, 2017.

[7] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 3406–3413. IEEE, 2016.

[8] Douglas Morrison, Peter Corke, and Jurgen Leitner. Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach. *Robotics: Science and Systems XIV*, pages 1–10, 2018.

[9] Andreas ten Pas, Marcus Gualtieri, Kate Saenko, and Robert Platt. Grasp pose detection in point clouds. *The International Journal of Robotics Research*, 36(13-14):1455–1473, 2017.

[10] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2901–2910, 2019.

[11] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020.

[12] Nicola De Cao and Wilker Aziz. The power spherical distribution. *arXiv preprint arXiv:2006.04437*, 2020.

[13] Joeri Hermans, Volodimir Begy, and Gilles Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4239–4248. PMLR, 13–18 Jul 2020.

[14] Johann Brehmer, Siddharth Mishra-Sharma, Joeri Hermans, Gilles Louppe, and Kyle Cranmer. Mining for Dark Matter Substructure: Inferring subhalo population properties from strong lenses with machine learning. *Astrophys. J.*, 886(1):49, 2019.

[15] Kieran Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia. Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold. *arXiv preprint arXiv:2106.05965*, 2021.

[16] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019.

[17] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.

[18] Jiang Hu, Xin Liu, Zai-Wen Wen, and Ya-Xiang Yuan. A brief introduction to manifold optimization. *Journal of the Operations Research Society of China*, pages 1–50, 2019.

[19] James Townsend, Niklas Koep, and Sebastian Weichwald. Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation. *The Journal of Machine Learning Research*, 17(1):4755–4759, 2016.

[20] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 International Conference on Advanced Robotics (ICAR)*, pages 510–517, 2015.

[21] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.

[22] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. `http://pybullet.org`, 2016–2020.

[23] Khaled Mamou and Faouzi Ghorbel. A simple and efficient approach for 3d mesh approximate convex decomposition. In *2009 16th IEEE international conference on image processing (ICIP)*, pages 3501–3504. IEEE, 2009.

[24] Michael Gschwandtner, Roland Kwitt, Andreas Uhl, and Wolfgang Pree. Blensor: Blender sensor simulation toolbox. In *International Symposium on Visual Computing*, pages 199–208. Springer, 2011.

[25] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.

[26] Jacob Varley, Jonathan Weisz, Jared Weiss, and Peter Allen. Generating multi-fingered robotic grasps via deep learning. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4415–4420. IEEE, 2015.

[27] Qingkai Lu and Tucker Hermans. Modeling grasp type improves learning-based grasp planning. *IEEE Robotics and Automation Letters*, 4(2):784–791, 2019.

[28] Qingkai Lu, Kautilya Chenna, Balakumar Sundaralingam, and Tucker Hermans. Planning multi-fingered grasps as probabilistic inference in a learned deep network. In *Robotics Research*, pages 455–472. Springer, 2020.

[29] Jens Lundell, Enric Corona, Tran Nguyen Le, Francesco Verdoja, Philippe Weinzaepfel, Gregory Rogez, Francesc Moreno-Noguer, and Ville Kyrki. Multi-fingan: Generative coarse-to-fine sampling of multi-finger grasps. *arXiv preprint arXiv:2012.09696*, 2020.

[30] Jens Lundell, Francesco Verdoja, and Ville Kyrki. Ddgc: Generative deep dexterous grasping in clutter. *arXiv preprint arXiv:2103.04783*, 2021.

[31] Bohan Wu, Iretiayo Akinola, Abhi Gupta, Feng Xu, Jacob Varley, David Watkins-Valls, and Peter K Allen. Generative attention learning: a "general" framework for high-performance multi-fingered grasping in clutter. *Autonomous Robots*, pages 1–20, 2020.

[32] Hamza Merzić, Miroslav Bogdanović, Daniel Kappler, Ludovic Righetti, and Jeannette Bohg. Leveraging contact forces for learning to grasp. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3615–3621. IEEE, 2019.

[33] Jean-Michel Marin, Pierre Pudlo, Christian P Robert, and Robin J Ryder. Approximate bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.

[34] Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.

[35] George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848. PMLR, 2019.

[36] George Papamakarios and Iain Murray. Fast ϵ-free inference of simulation models with bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*. 2016.

[37] Jan-Matthis Lueckmann, Pedro J. Goncalves, Giacomo Bassetto, Kaan Öcal, Marcel Nonnenmacher, and Jakob H. Macke. Flexible statistical inference for mechanistic models of neural dynamics. In *Advances in Neural Information Processing Systems*, 2017.

[38] David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pages 2404–2414. PMLR, 2019.

A   Algorithms

Algorithm 1: Training procedure for $d_\phi(S, \mathbf{h})$ and $d_\theta(i, \mathbf{h})$.

    Input: Priors $p(\mathbf{h})$, $p(\mathcal{O})$, $p(\mathbf{p}_\mathcal{O})$
             Sensor generative model $p(i|\mathcal{O}, \mathbf{p}_\mathcal{O})$
             Grasp generative model $p(S|\mathbf{h}, \mathcal{O}, \mathbf{p}_\mathcal{O})$
             Criterion ℓ (e.g, the binary cross-entropy)
    Output: Trained classifiers $d_\phi(S, \mathbf{h})$, $d_\theta(i, \mathbf{h})$

1 **while** *not converged* **do**
2   **Sample** $\mathbf{h} \leftarrow \{\mathbf{h}_m \sim p(\mathbf{h})\}_{m=1}^M$
3   **Sample** $\mathbf{h}' \leftarrow \{\mathbf{h}'_m \sim p(\mathbf{h})\}_{m=1}^M$
4   **Sample** $\mathcal{O}, \mathbf{p}_\mathcal{O} \leftarrow \{\mathcal{O}_m, \mathbf{p}_{\mathcal{O},m} \sim p(\mathcal{O})p(\mathbf{p}_\mathcal{O})\}_{m=1}^M$
5   **Simulate** $S \leftarrow \{S_m \sim p(S|\mathbf{h}_m, \mathcal{O}_m, \mathbf{p}_{\mathcal{O},m})\}_{m=1}^M$
6      $\mathcal{L} \leftarrow \ell(d_\phi(S, \mathbf{h}), 1) + \ell(d_\phi(S, \mathbf{h}'), 0)$
7      $\phi \leftarrow \texttt{OPTIMIZER}(\phi, \nabla_\phi \mathcal{L})$
8 **end**

9 **while** *not converged* **do**
10   **Sample** $\mathbf{h} \leftarrow \{\mathbf{h}_m \sim p(\mathbf{h}|S = 1)\}_{m=1}^M$
11   **Sample** $\mathbf{h}' \leftarrow \{\mathbf{h}'_m \sim p(\mathbf{h}|S = 1)\}_{m=1}^M$
12   **Sample** $\mathcal{O}, \mathbf{p}_\mathcal{O} \leftarrow \{\mathcal{O}_m, \mathbf{p}_{\mathcal{O},m} \sim p(\mathcal{O}, \mathbf{p}_\mathcal{O}|S = 1, \mathbf{h})\}_{m=1}^M$
13   **Simulate** $i \leftarrow \{i_m \sim p(i|\mathcal{O}_m, \mathbf{p}_{\mathcal{O}m})\}_{m=1}^M$
14      $\mathcal{L} \leftarrow \ell(d_\theta(i, \mathbf{h}), 1) + \ell(d_\theta(i, \mathbf{h}'), 0)$
15      $\theta \leftarrow \texttt{OPTIMIZER}(\theta, \nabla_\theta \mathcal{L})$
16 **end**
17 **return** d_ϕ, d_θ

Algorithm 2: Manifold optimization procedure to obtain the MAP estimate $\hat{\mathbf{h}}^*$

    Input: Differentiable priors $p(\mathbf{h})$
             Differentiable likelihood-to-evidence ratio $\hat{r}$
             Depth image i
    Output: Approximate MAP estimate $\hat{\mathbf{h}}^*$

1 MAP cost function $f(i, \mathbf{h}) = -\log \hat{r}(S = 1, i|\mathbf{h}) - \log p(\mathbf{h})$

2 Sample an initial subset $\mathcal{S} = \{\mathbf{x}, \mathbf{q}\} \sim \{p(\mathbf{x})p(\mathbf{q})\}_1^{1000}$

3 **foreach** $g_k \in \mathcal{G}$ **do**
4      $\mathbf{h} = (\mathbf{x}, \mathbf{R}(\mathbf{q}), g_k)$
5   Initial iterate $\mathbf{h}_0 = \arg\min_{\mathbf{h} \in \mathcal{S}} f(i, \mathbf{h})$
6      $\hat{\mathbf{x}}_k^*, \hat{\mathbf{R}}_k^* = \text{Geometric CG method}(f, i, \mathbf{h}_0)$
7      $\hat{\mathbf{h}}_k^* = (\hat{\mathbf{x}}_k^*, \hat{\mathbf{R}}_k^*, g_k)$
8 **end**
9 **return** $\arg\min_{\mathbf{h}_k} f(i, \hat{\mathbf{h}}_k^*)$

B Prior Distribution

(a)

(b)

(c)

(d)
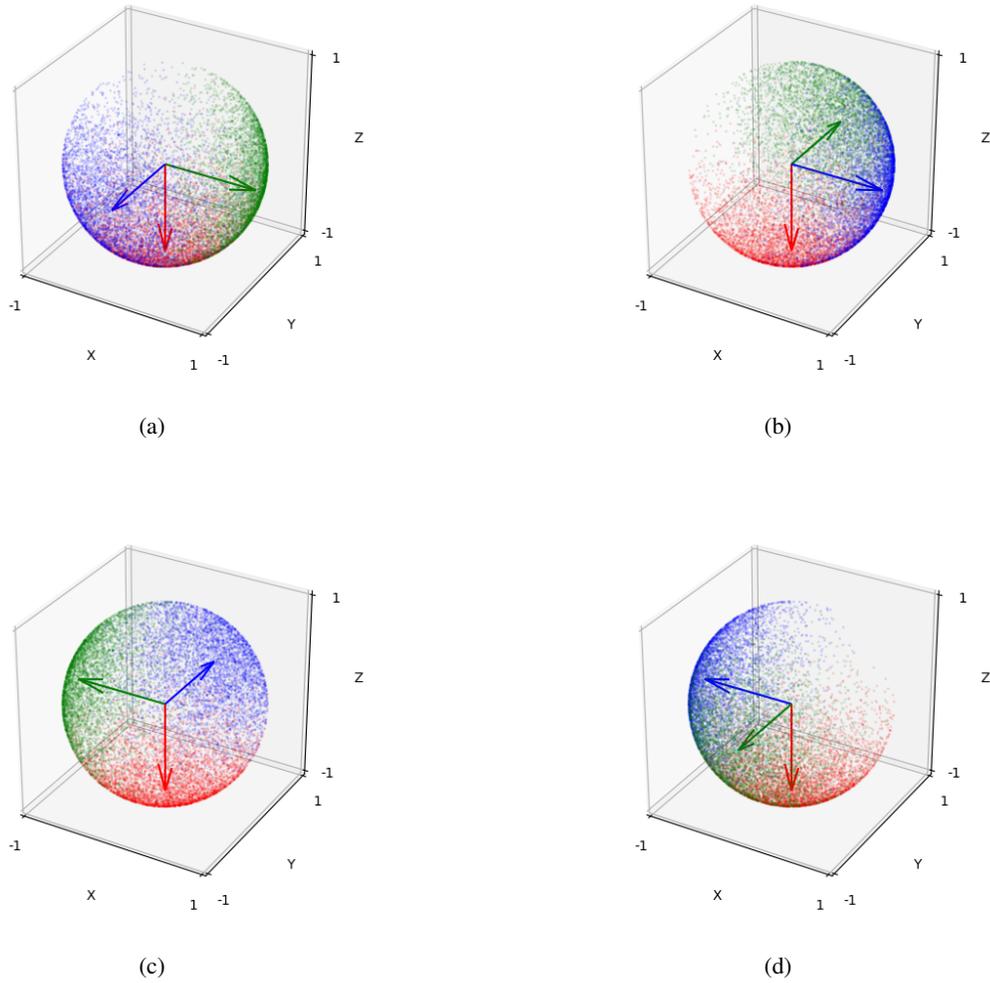
Figure 5: Prior distribution $p(\mathbf{q})$. (a)-(d) correspond to the four modes of the mixture.

C   YCB objects

Figure 6: Objects from YCB used for training and testing.
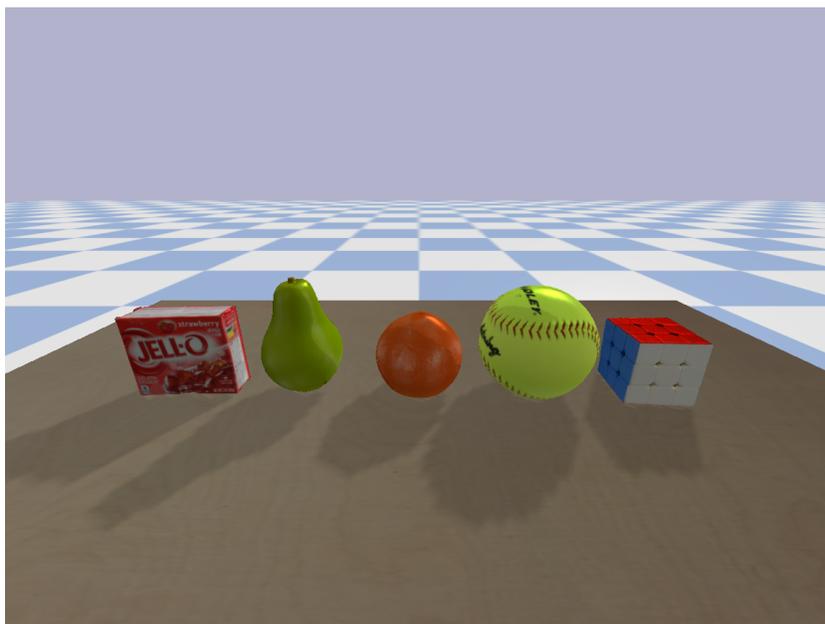
Figure 7: Additional unseen objects from YCB used for testing in simulation.

D   Success rate by object category
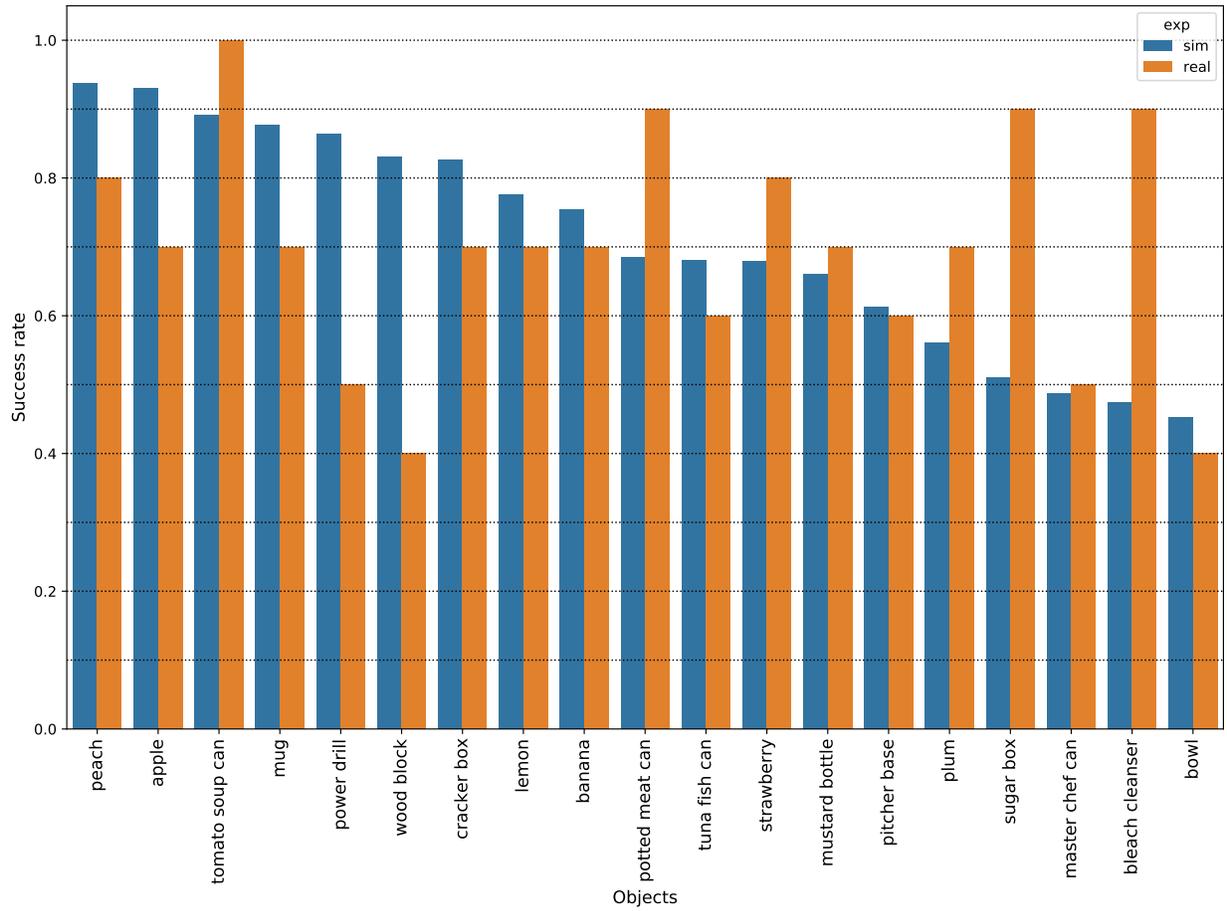
Figure 8: Success rate obtained from simulation and real experiments for each object.