

Implicit representation priors meet Riemannian geometry for Bayesian robotic grasping

Norman Marlier^{1*} Julien Gustin² Olivier Brüs³ Gilles Louppe⁴

Abstract—Robotic grasping in highly noisy environments presents complex challenges, especially with limited prior knowledge about the scene. In particular, identifying good grasping poses with Bayesian inference becomes difficult due to two reasons: i) generating data from uninformative priors proves to be inefficient, and ii) the posterior often entails a complex distribution defined on a Riemannian manifold. In this study, we explore the use of implicit representations to construct scene-dependent priors, thereby enabling the application of efficient simulation-based Bayesian inference algorithms for determining successful grasp poses in unstructured environments. Results from both simulation and physical benchmarks showcase the high success rate and promising potential of this approach.

I. INTRODUCTION

Grasping is a fundamental skill for any robotic system. While current methods are effective for highly constrained tasks in structured environments, new and complex applications require increased flexibility and more advanced algorithms to account for the uncertainties that emerge in unstructured and noisy environments. Bayesian inference offers a well-principled approach to address these uncertainties; however, robotic tasks present unique challenges that make Bayesian inference difficult to apply, particularly for sampling-based algorithms. Firstly, many Bayesian approaches assume that the likelihood is tractable and can be evaluated, which is seldom the case in robotics. Secondly, parameters may span a vast space, leading to inefficient sampling strategies. Lastly, parameters of interest often belong to smooth Riemannian manifolds, further complicating the inference procedure. In this paper, we address these challenges by designing informative scene-dependent priors and using simulation-based inference algorithms combined with geometric sampling methods. Our contributions are summarized as follows:

- We integrate simulation-based Bayesian inference methods [1] with 3D implicit representations for robotic grasping.
- We adapt geodesic Monte Carlo [2] with a neural ratio estimator to sample on Riemannian manifolds with an intractable likelihood.
- We validate our method through simulated and real experiments, demonstrating promising grasping performance.



Fig. 1: Our benchmark scene. (left) The simulated environment. (right) The real setup.

II. PROBLEM STATEMENT

We consider the problem of planning 4-DoF hand configurations for a robotic gripper handling various unknown objects on a table, observed with a depth camera. A benchmark scene is shown in Fig. 1.

A. Notations

Frames We use several reference frames in our work. The world frame \mathcal{F}_W and the workspace frame \mathcal{F}_S can be chosen freely and are not tied to a physical location. The world frame is used for the robot and the sensor, while the workspace frame is used for our inference system. \mathcal{F}_C and \mathcal{F}_E correspond respectively to the camera and the tool centre point.

Hand configuration The hand configuration $\mathbf{h} \in \mathcal{H} = \mathbb{R}^3 \times \mathbb{S}^1$ is defined as the pose $(\mathbf{x}, \mathbf{q}) \in \mathbb{R}^3 \times \mathbb{S}^1$ of the hand, where \mathbf{x} is the vector SE expressed in \mathcal{F}_S and \mathbf{q} is the planar rotation represented with complex numbers defined in \mathcal{F}_S .

Binary metric A binary variable $S \in \{0, 1\}$ indicates if the grasp fails ($S = 0$) or succeeds ($S = 1$).

Observation Given the depth image I with its corresponding transformation camera to world \mathbf{T}_{WC} and camera intrinsic matrix K , we construct a point cloud $\mathbf{P} \in \mathbb{R}^{2048 \times 3}$ expressed in \mathcal{F}_S .

Occupancy A binary variable $o \in \{0, 1\}$ indicates if a point $\mathbf{p} \in \mathbb{R}^3$ is occupied by any object of the scene.

Latent variables Unobserved variables \mathbf{z} capture uncertainties about the nonsmooth dynamics of contact, the sensor noise, as well as the number of objects and their geometry.

*The authors come from the University of Liège, Belgium

¹norman.marlier@uliege.be

B. Grasping as inference

We 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, an occupancy o and a point cloud \mathbf{P} . That is, we are seeking the maximum a posteriori (MAP) estimate

$$\mathbf{h}^* = \arg \max_{\mathbf{h}} p(\mathbf{h} | S = 1, o = 1, \mathbf{P}), \quad (1)$$

from which we then compute the joint trajectory

$$\tau_{1:m} = \Lambda(\tau_0, \text{IK}(\mathbf{h}^*), \mathbf{P}), \quad (2)$$

where IK is an inverse kinematic solver, $\tau_{1:m}$ are waypoints in the joint space, $\tau_m = \text{IK}(\mathbf{h}^*)$ with \mathbf{h}^* expressed in $\underline{\mathcal{F}}_W$ and Λ is a path planner.

III. IMPLICIT REPRESENTATION OF PRIORS AND POSTERIOR FOR ROBOTIC GRASPING

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

$$p(\mathbf{h} | S, o, \mathbf{P}) = \frac{p(S | \mathbf{h}, o, \mathbf{P})}{p(S | o, \mathbf{P})} p(\mathbf{h} | o, \mathbf{P}), \quad (3)$$

which can be rewritten as the product of the likelihood-to-evidence ratio r and a scene-dependent prior

$$p(\mathbf{h} | S, o, \mathbf{P}) = r(S | \mathbf{h}, o, \mathbf{P}) p(\mathbf{h} | o, \mathbf{P}). \quad (4)$$

A. Priors

Position The scene-dependent prior over the position \mathbf{x} is the distribution $p(\mathbf{x} | o, \mathbf{P}) = \frac{p(o | \mathbf{x}, \mathbf{P})}{p(o | \mathbf{P})} p(\mathbf{x})$, where $p(o | \mathbf{x}, \mathbf{P})$ is the likelihood of the occupancy o , $p(\mathbf{x})$ is uniform over the workspace, and $p(\mathbf{x} | \mathbf{P})$ is simplified to $p(\mathbf{x})$ by independence.

We model the occupancy likelihood $p(o | \mathbf{x}, \mathbf{P})$ using a Convolutional Occupancy Network [3]. This network computes the occupancy by first producing three canonical features planes $\mathbf{c}_{xy}(\mathbf{P})$, $\mathbf{c}_{xz}(\mathbf{P})$ and $\mathbf{c}_{yz}(\mathbf{P})$. Then, the bilinear interpolations of the three planes are used to compute $\psi(\mathbf{P}, \mathbf{x}) = \mathbf{c}_{xy}(\mathbf{P})(\mathbf{x}) + \mathbf{c}_{xz}(\mathbf{P})(\mathbf{x}) + \mathbf{c}_{yz}(\mathbf{P})(\mathbf{x})$. These point-wise features at point \mathbf{x} are finally processed by a fully connected network, outputting the occupancy probability.

This implicit representation allows us to sample interesting grasping positions from $p(\mathbf{x} | o, \mathbf{P}) \propto p(o | \mathbf{x}, \mathbf{P}) p(\mathbf{x})$. We use Hamiltonian Monte Carlo (HMC) to take advantage of the differentiability of the occupancy network.

Orientation The prior of the orientation \mathbf{q} is a uniform distribution over the unit circle \mathbb{S}^1 . This prior is *invariant* to any rotation $\mathbf{R} \in \text{SO}(2)$ applied to \mathbf{q} , satisfying $p(\mathbf{q}) = p(\mathbf{R}\mathbf{q})$. This property enables free selection of the reference frame on the table. Additionally, the prior can be extended to $\text{SO}(3)$ by using quaternions on \mathbb{S}^3 .

Hand configuration Finally, the prior of the hand configuration is $p(\mathbf{h} | o, \mathbf{P}) = p(\mathbf{x} | o, \mathbf{P}) p(\mathbf{q})$.

B. Ratio

The likelihood function $p(S | \mathbf{h}, o, \mathbf{P})$ and the evidence $p(S | o, \mathbf{P})$ are both intractable. However, drawing samples from forward models remains feasible with physical simulators, hence enabling likelihood-free Bayesian inference algorithms. In particular, the likelihood-to-evidence ratio $r(S | \mathbf{h}, o, \mathbf{P})$ (Eq. (4)) can be approximated by a neural network $r_\phi(S | \mathbf{h}, o, \mathbf{P})$ using amortized neural ratio estimation [5]. Here, instead of using only point-wise features $\psi(\mathbf{P}, \mathbf{x})$ as input for the ratio, we add a crop of the features plane $\mathbf{c}_{xy}(\mathbf{P})$, centred at \mathbf{x} , sized by the gripper and rotated by \mathbf{q} . These features $\Psi(\mathbf{P}, \mathbf{h})$ local in the neighbourhood of the grasping point are nearly equivariant to a 2D transformation applied to the object, *i.e.* $\mathbf{T}\Psi(\mathbf{P}, \mathbf{h}) \approx \Psi(\mathbf{TP}, \mathbf{h})$, $\mathbf{T} \in \text{SE}(2)$.

C. Posteriors

Given our scene-dependent prior and our likelihood-to-evidence ratio, we approximate the posterior over the hand configurations as

$$\hat{p}(\mathbf{h} | S, o, \mathbf{P}) = r_\phi(S | \mathbf{h}, o, \mathbf{P}) p(\mathbf{h} | o, \mathbf{P}). \quad (5)$$

This approximation defines an implicit function [6] on the product of manifolds $\mathbb{R}^3 \times \mathbb{S}^1$ that is both fully tractable and differentiable, allowing the use of gradient-based methods for computing the MAP and sampling. Therefore, we can use Markov Chain Monte Carlo methods to sample from our posterior approximation $\hat{p}(\mathbf{h} | S, o, \mathbf{P})$. In particular, based on [5], we use a likelihood-free version of HMC by replacing the intractable likelihood with the ratio. The potential energy function is defined as $U(\mathbf{h}) \triangleq -\log p(S | \mathbf{h}, o, \mathbf{P})$ and its difference is $U(\mathbf{h}_t) - U(\mathbf{h}') = \log r(S | \mathbf{h}_t, \mathbf{h}')$. The gradient used in the integration step is given by $\nabla_{\mathbf{h}} U(\mathbf{h}) = -\nabla_{\mathbf{h}} \log r(S | \mathbf{h}, o, \mathbf{P})$. To account for the geometry of the parameter space, we then further extend our likelihood-free HMC with a geodesic integrator. The geodesic Monte Carlo scheme uses geodesic flow to perform the integration while staying on the manifold. To this end, orthogonal projections and geodesics are needed in a closed form. Finally, geodesic Monte Carlo can be applied to a product of manifolds $\mathcal{M}_1 \times \mathcal{M}_2 : \{(x_1, x_2) : x_1 \in \mathcal{M}_1, x_2 \in \mathcal{M}_2\}$, such as $\mathbb{R}^3 \times \mathbb{S}^1$ in our specific case. Geodesic flow can be executed in parallel; only the evaluation of the ratio requires both variables \mathbf{x} and \mathbf{q} . In this manner, we can sample from the posterior density defined on a smooth manifold with closed-form geodesic. The full sampling procedure is summarized in Algorithm 1 of Appendix A.

IV. EXPERIMENTS

We assess our approach on a robotic grasping task in both simulation and real-world settings. We generate data in a *packed* scenario, as defined in [7]. Additional experiments can be found in Appendix B and C.

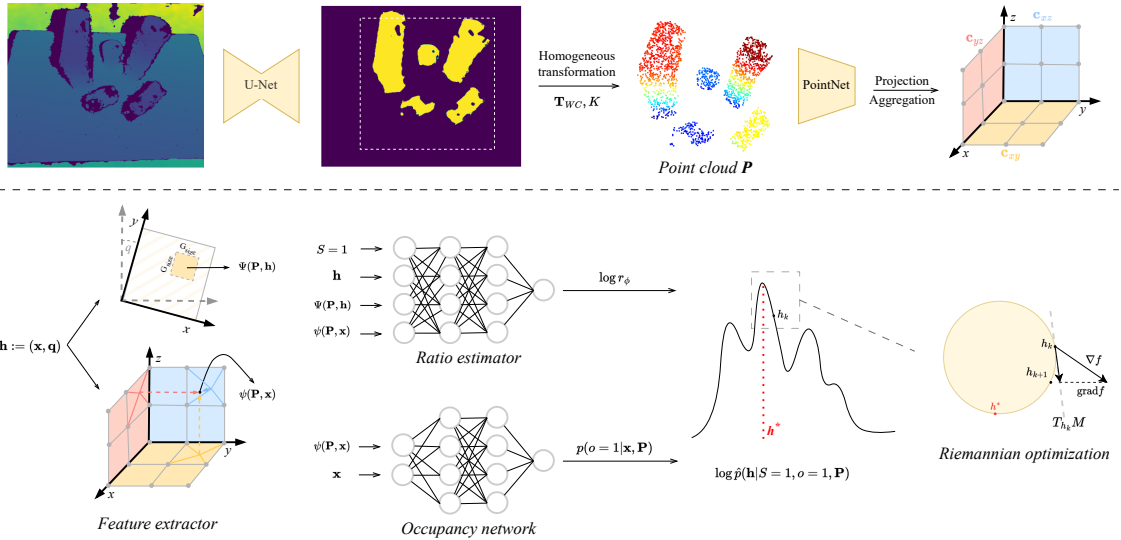


Fig. 2: Our grasp inference pipeline, as run on the scene of Fig. 1. It begins with a noisy depth image of the scene, from which we first separate the objects from the background using a U-Net [4]. We then generate three canonical feature planes following the approach in [3]. To evaluate a given \mathbf{h} , we extract point-wise $\psi(\mathbf{P}, \mathbf{x})$ and local $\Psi(\mathbf{P}, \mathbf{h})$ features and feed them to the ratio and occupancy networks. Using the resulting differentiable posterior and Riemannian optimization, we finally identify the most plausible hand configuration \mathbf{h}^* .

A. Grasp inference pipeline

Starting from the depth image I , we remove the background and extract only pixels of the objects with a segmentation model based on a U-Net architecture [4]. Then, we convert I to a point cloud \mathbf{P} with 2048 points, which passes through an encoder and produces three canonical feature planes. We then extract point-wise $\psi(\mathbf{P}, \mathbf{x})$ and local $\Psi(\mathbf{P}, \mathbf{h})$ features to evaluate the occupancy and the ratio networks. To smooth the posterior approximation, we use an ensemble of 6 ratio models. Finally, we compute the MAP by maximizing the log posterior density [8]. To this end, we use a Riemannian gradient ascent which preserves the nonlinearity of \mathbb{S}^1 . A visual summary of our method is given in Fig. 2.

Our likelihood-free geodesic Monte Carlo is used to sample plausible hand configurations $\mathbf{h} \sim \hat{p}(\mathbf{h} \mid S = 1, o = 1, \mathbf{P})$ for successful grasps, as shown in Fig. 3. Although our conditional prior distributes density across everywhere on the objects, the posterior assigns minimal density to the bottom of objects when multiple objects are present on the table. This occurs due to potential collisions between the gripper and the table, or the gripper and other objects. Regarding rotation, the posterior resembles the prior because multiple objects on the table, some with axial symmetry, allow for a wide range of grasp orientations. When only a single object is used, distinct modes can be observed, indicating that our posterior captures meaningful orientations, as shown in Fig. 4.

B. Simulation results

To compare our approach, we evaluated it against Grasp Pose Detection (GPD) [9] and Volumetric Grasp-

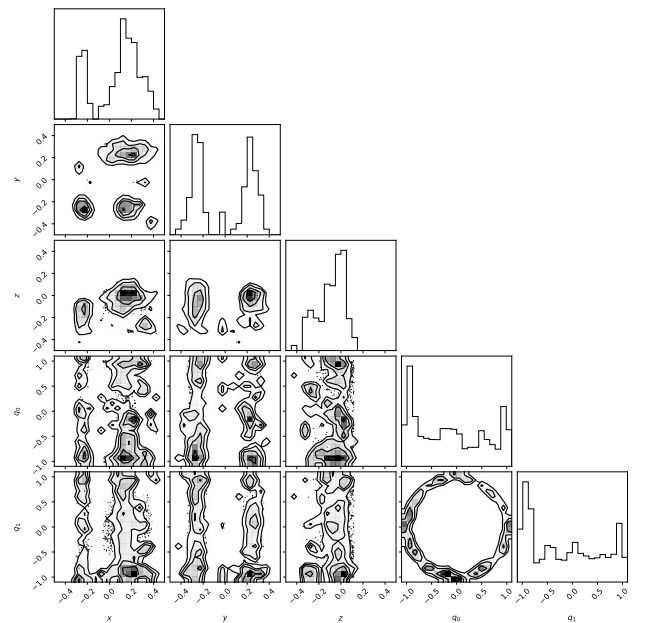


Fig. 3: Estimated posterior distribution $\hat{p}(\mathbf{h} \mid S = 1, o = 1, \mathbf{P})$ of plausible successful hand configurations, for the scene shown in Fig. 1.

ing Network (VGN) [7] in terms of success rate and percent cleared [7] using the same dataset and a similar scenario (Table I). Our model achieved a high success rate of 91.1%, which is very close to VGN’s best success rate of 91.5%. However, our model operated in a more constrained setup with 4 DoF instead of 6, limiting

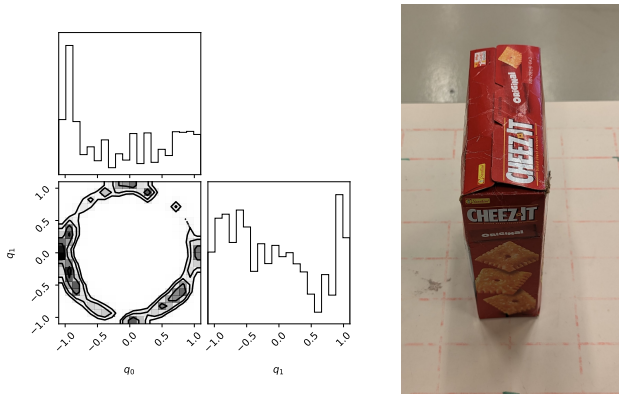


Fig. 4: (left) Estimated posterior distribution of the orientation $\hat{p}(\mathbf{q} \mid S = 1, o = 1, \mathbf{P})$. (right) Single object.

TABLE I: Success rates (%) and % cleared for picking experiments for the packed scenario with 5 objects over 200 rounds.

Method	Success rate	% cleared
<i>Simulation results</i>		
GPD [9]	73.7	72.8
VGN ($\epsilon = 0.95$) [7]	91.5	79
VGN ($\epsilon = 0.9$) [7]	87.6	80.4
VGN ($\epsilon = 0.85$) [7]	80.4	79.9
Ours	91.1	77
<i>Real-world results</i>		
Ours	95.6	88

the gripper’s movement and increasing the risk of collisions. This limitation resulted in a slightly lower percent cleared compared to VGN. Furthermore, the gripper used was primarily designed for larger objects, making it challenging to handle smaller ones.

C. Real-world results

We adopt the setup from [8]. We perform 20 rounds with 5 objects among 13 (see Fig. 5), using a protocol similar to the simulation experiments. The objects are chosen based on their availability in the lab and whether they were seen or unseen during training. For novel objects, we achieve a success rate of **95.6%** and a percent cleared of **88%**, showing the strong adaptability and performance of our approach. The discrepancy between the simulation and real-world setup is overcome without any decrease in performance. In both simulation and real-world settings, the majority of failure cases are due to insufficient friction forces, causing the objects to slip.

V. RELATED WORK

Grasp sampling strategies that generate data for machine learning methods can be categorized based on their coverage of the hand configuration space \mathcal{H} [10]. Simple strategies like uniform sampling provide direct density estimation but are highly inefficient. Heuristic methods, on the other hand, mainly rely on object geometry. Notably, the most efficient strategies [11], [12] are not suitable



Fig. 5: (left) Object assets used in the real setup. (right) Successful grasp of a mug.

for complex settings such as multi-fingered grippers. Our occupancy networks-based approach offers direct density estimation, efficient sampling, and does not depend on specific object or gripper assumptions.

Representing a 3D scene as a parameterized function using neural networks has recently gained substantial interest. Occupancy networks [13] determine whether a point is occupied or not. However, they lack *equivariance* for translations and rotations. To overcome this limitation, translation-equivariant feature planes are derived from convolutional networks, resulting in Convolutional Occupancy Networks [3]. Achieving rotation-equivariance is more challenging, but innovative solutions have recently emerged [14], [15].

Various methods exist for sampling from a distribution defined on a Riemannian manifold. Adapting Markov Chain Monte Carlo methods [16] to Riemannian manifolds allows sampling from a differentiable and tractable likelihood. Our approach eliminates the need for an explicit likelihood. Normalizing flows [17] which target density defined on manifolds can rapidly sample from the posterior distribution. However, it remains unclear how to use their gradients for Riemannian optimization.

Finally, probabilistic approaches for grasping problems typically depend on explicit likelihood functions that model the probability of success or a grasp quality metric related to an observation and a grasp pose [7], [18]–[20]. Closer to our work, a similar study [8] uses simulation-based inference to compute the maximum a posteriori through Riemannian gradient ascent. However, this approach uses a heuristic prior and cannot accommodate multiple objects.

VI. CONCLUSION

We have shown that Bayesian inference can be effectively applied to robotic grasping in complex and noisy environments. Through innovative enhancements, we have improved the sample efficiency of the inference pipeline. Our approach can manage tasks with escalating complexity and proves valuable for real-world robotic applications. Future research will focus on controlling the complete 6-DoF of the robotic hand.

ACKNOWLEDGEMENT

Norman Marlier acknowledges the Belgian Fund for Research training in Industry and Agriculture for its financial support (FRIA grant). Computational resources have been provided by the Consortium des Équipements de Calcul Intensif (CÉCI), funded by the Fonds de la Recherche Scientifique de Belgique (F.R.S.-FNRS) under Grant No. 2.5020.11 and by the Walloon Region.

REFERENCES

- [1] K. Cranmer, J. Brehmer, and G. Louppe, “The frontier of simulation-based inference,” *Proceedings of the National Academy of Sciences*, 2020.
- [2] S. Byrne and M. Girolami, “Geodesic monte carlo on embedded manifolds,” *Scandinavian Journal of Statistics*, vol. 40, no. 4, pp. 825–845, 2013.
- [3] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, “Convolutional occupancy networks,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 523–540.
- [4] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” 2015.
- [5] J. Hermans, V. Begy, and G. Louppe, “Likelihood-free MCMC with amortized approximate ratio estimators,” in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 4239–4248. [Online]. Available: <http://proceedings.mlr.press/v119/hermans20a.html>
- [6] K. Murphy, C. Esteves, V. Jampani, S. Ramalingam, and A. Makadia, “Implicit-pdf: Non-parametric representation of probability distributions on the rotation manifold,” *arXiv preprint arXiv:2106.05965*, 2021.
- [7] M. Breyer, J. J. Chung, L. Ott, S. Roland, and N. Juan, “Volumetric grasping network: Real-time 6 dof grasp detection in clutter,” in *Conference on Robot Learning*, 2020.
- [8] N. Marlier, O. Bröls, and G. Louppe, “Simulation-based bayesian inference for robotic grasping,” *arXiv preprint arXiv:2303.05873*, 2023.
- [9] A. Ten Pas, M. Gualtieri, K. Saenko, and R. Platt, “Grasp pose detection in point clouds,” *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1455–1473, 2017.
- [10] C. Eppner, A. Mousavian, and D. Fox, “A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set,” in *Robotics Research: The 19th International Symposium ISRR*. Springer, 2022, pp. 890–905.
- [11] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, “Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics,” 2017.
- [12] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee, “Learning 6-dof grasping interaction via deep geometry-aware 3d representations,” 2018.
- [13] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, “Occupancy networks: Learning 3d reconstruction in function space,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.
- [14] Y. Chen, B. Fernando, H. Bilen, M. Nießner, and E. Gavves, “3d equivariant graph implicit functions,” in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*. Springer, 2022, pp. 485–502.
- [15] C. Deng, O. Litany, Y. Duan, A. Poulencard, A. Tagliasacchi, and L. J. Guibas, “Vector neurons: A general framework for so (3)-equivariant networks,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 200–12 209.
- [16] C. Liua and J. Zhub, “Geometry in sampling methods: A review on manifold mcmc and particle-based variational inference methods,” *Advancements in Bayesian Methods and Implementations*, vol. 47, p. 239, 2022.
- [17] D. J. Rezende, G. Papamakarios, S. Racaniere, M. Albergo, G. Kanwar, P. Shanahan, and K. Cranmer, “Normalizing flows on tori and spheres,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 8083–8092.
- [18] Q. Lu, K. Chenna, B. Sundaralingam, and T. Hermans, “Planning multi-fingered grasps as probabilistic inference in a learned deep network,” in *Robotics Research*. Springer, 2020, pp. 455–472.
- [19] J. Cai, J. Cen, H. Wang, and M. Y. Wang, “Real-time collision-free grasp pose detection with geometry-aware refinement using high-resolution volume,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1888–1895, 2022.
- [20] M. Van der Merwe, Q. Lu, B. Sundaralingam, M. Matak, and T. Hermans, “Learning continuous 3d reconstructions for geometrically aware grasping,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 11 516–11 522.

A. Likelihood-free geodesic Monte Carlo

Algorithm 1: Likelihood-free geodesic Hamiltonian Monte Carlo

Input: A Manifold \mathcal{M} Initial parameters \mathbf{h}_0 Prior $p(\mathbf{h} \mid \mathbf{P})$ Momentum distribution $p(\mathbf{v})$ Trained classifier $d_\phi(S, \mathbf{h}, \mathbf{P})$ Observations S, \mathbf{P} **Output:** Markov chain $\mathbf{h}_{1:T}$

```

1  $t \leftarrow 0$ 
2  $\mathbf{h}_t \leftarrow \mathbf{h}_0$ 
3 for  $t < T$  do
4    $\mathbf{v}_t \sim p(\mathbf{v})$ 
5    $\mathbf{v}_t \leftarrow \pi_{\mathbf{h}_t}(\mathbf{v}_t)$ 
6    $k \leftarrow 0$ 
7    $\mathbf{v}_k \leftarrow \mathbf{v}_t$ 
8    $\mathbf{h}_k \leftarrow \mathbf{h}_t$ 
9   for  $k < L$  do
10     $\mathbf{v}_k \leftarrow \mathbf{v}_k + \frac{\epsilon}{2} \nabla_{\mathbf{h}_k} \log r(S \mid \mathbf{h}_k, \mathbf{P})$ 
11     $\mathbf{v}_k \leftarrow \pi_{\mathbf{h}_k}(\mathbf{v}_k)$ 
12     $\mathbf{h}_k \leftarrow \gamma(\epsilon), \gamma(0) = \mathbf{h}_k$ 
13     $\mathbf{v}_k \leftarrow \dot{\gamma}(\epsilon), \dot{\gamma}(0) = \mathbf{v}_k$ 
14     $\mathbf{v}_k \leftarrow \mathbf{v}_k + \frac{\epsilon}{2} \nabla_{\mathbf{h}_k} \log r(S \mid \mathbf{h}_k, \mathbf{P})$ 
15     $\mathbf{v}_k \leftarrow \pi_{\mathbf{h}_k}(\mathbf{v}_k)$ 
16     $k \leftarrow k + 1$ 
17  end
18   $\lambda_k \leftarrow \log r(S \mid \mathbf{h}_k, \mathbf{P}) + \log p(\mathbf{h}_k \mid \mathbf{P}) - \frac{1}{2} \mathbf{v}_k^T \mathbf{v}_k$ 
19   $\lambda_t \leftarrow \log r(S \mid \mathbf{h}_t, \mathbf{P}) + \log p(\mathbf{h}_t \mid \mathbf{P}) - \frac{1}{2} \mathbf{v}_t^T \mathbf{v}_t$ 
20   $\rho \leftarrow \min(\exp(\lambda_k - \lambda_t), 1)$ 
21   $\mathbf{h}_{t+1} \leftarrow \begin{cases} \mathbf{h}_k & \text{with a probability } \rho \\ \mathbf{h}_t & \text{with a probability } 1 - \rho \end{cases}$ 
22   $t \leftarrow t + 1$ 
23 end
24 return  $\mathbf{h}_{1:T}$ 

```

B. Sampling the orientation: toy problem

Given a model parameter sample $q_\theta \in \mathbb{S}^d$, the forward generative process is defined as:

$$\nu = q_\theta \tag{6}$$

$$\kappa = 20 \tag{7}$$

$$q_x \sim \exp(\kappa \nu^T q_x) \tag{8}$$

with the prior $p(q_\theta) \triangleq \text{SphericalUniform}(d)$. It follows the true posterior $p(q_\theta | q_x) \propto \exp(\kappa q_x^T q_\theta)$. We use a MLP of 3 layers with 64 neurons to approximate the likelihood-to-evidence ratio. All the activation functions are ReLU except the last one which is linear. We train the ratio with 1000000 samples with a batch size of 8000 for 50 epochs. For the geodesic HMC, we use 100 chains and 2000 transitions with a burn in of 1000. The integration parameters are $\epsilon = 0.01, L = 20$. The approximate posterior shares the same structure that the true posterior, demonstrating its accuracy (Fig. 6). Moreover, we conduct a quantitative analysis by computing the Maximal Mean Discrepancy (MMD) for \mathbb{S}^1 and \mathbb{S}^3 . We obtain respectively $0.0028 \pm 5.84e^{-6}$ and $0.01 \pm 3e^{-5}$ for an identity kernel between the two geodesic means for 10 different observations q_x .

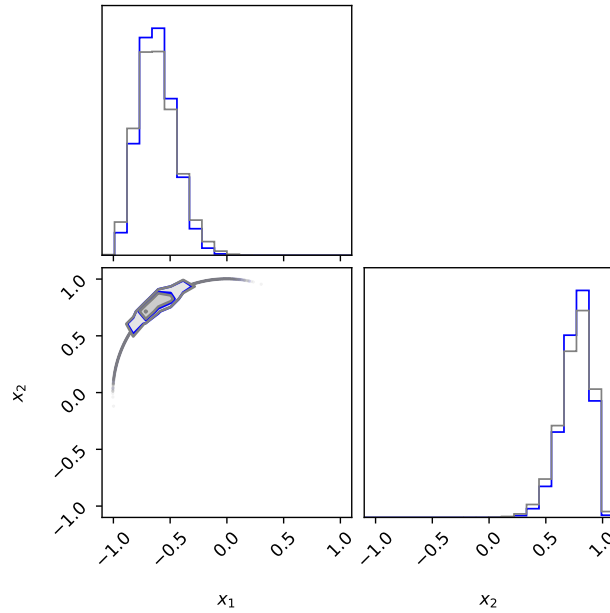


Fig. 6: Posterior for \mathbb{S}^1 . In blue, the distribution obtained through GMC with the ratio and in grey, the ground truth posterior.

C. Sampling the position: multiple objects scene

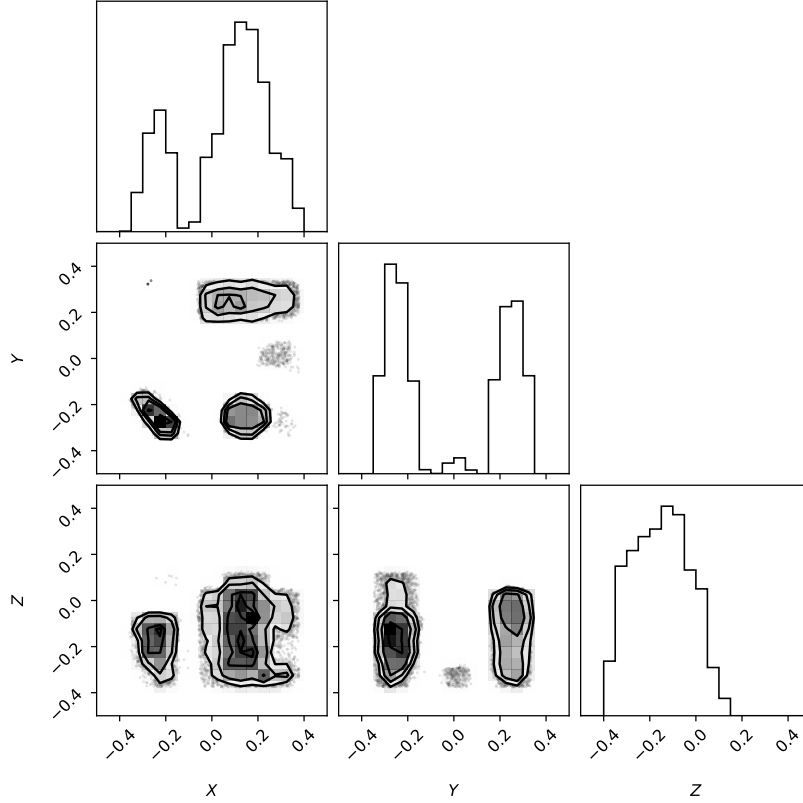


Fig. 7: Approximate posterior $p(\mathbf{x}|o = 1, \mathbf{P})$ of a scene with 5 different objects.

In this experiment, we have a scene containing five objects with varying levels of difficulty, as shown in Fig. 2. We use a convolutional occupancy network [3] that is trained for 120,000 iterations with a batch size of 32 samples. The network has a resolution of 128 and a feature dimension of 32. To sample from the posterior $p(\mathbf{x}|o = 1, \mathbf{P}) \propto p(o = 1|\mathbf{x}, \mathbf{P})p(\mathbf{x})$, as we previously explained, we employed HMC with specific hyper-parameters. These include 100 chains of 5000 transitions and a burn-in of 1000. The integration parameters are $\epsilon = 0.01$, $L = 20$. The chains' initial points are sampled uniformly in the bounding box of the objects. The corner plot in Fig. 7 illustrates the resulting posterior, which aligns with the location and shape of four out of the five objects and the potential grasping point. However, the occupancy of the fifth object cannot be recovered due to either being regarded as noise or having a much smaller density compared to the other objects.