

Supplementary material: Bayesian Object Models for Robotic Interaction with Differentiable Probabilistic Programming

Anonymous Author(s)

Affiliation

Address

email

1 In this brief document, we provide additional details about our approach.

2 1 Differentiable tree sampler

3 First, we prove that our differentiable tree sampler is guaranteed to produce a tree structured random
4 variable. Recall that the DP-DAG strategy for sampling a DAG represents the graph structured
5 random variable \mathcal{G} by its factorized adjacency matrix $A = PUP^T$. Here, P is a permutation matrix
6 and U is an upper-triangular matrix. For A to characterize a tree, we must explicitly forbid any edge
7 among siblings, or from a descendant to an ancestor in U .

8 For a proof sketch, recall that in our opportunistic rounding scheme, we disallow all edges (k, j) (for
9 $k > i$) for all edges (i, j) in the canonical graph (i.e., with adjacency matrix U). By contradiction,
10 let us assume that an edge (k, j) is added to the graph while another edge (i, j) exists. This would
11 imply that a new path $i \mapsto k \mapsto j$ is created in the graph. This violates a basic property of a tree;
12 only a unique path may exist across each pair of nodes (whereas we now have two such paths). This
13 proves that our sampling strategy is guaranteed to return a tree.

14 *Caveat:* The sampling strategy we present is not an unbiased strategy (the opportunistic rounding
15 step strongly favours edges between earlier entries in the canonical graph).

16 2 Implementation details

17 For model specification and gradient-based probabilistic inference, we leverage the pyro [1] (and
18 numpyro [2]) probabilistic programming language (PPL). Pyro allows us to specify generative
19 (probabilistic) models in python code using popular automatic differentiation API (e.g., pytorch [3],
20 JAX [4]). Additionally, the use of a modern PPL enables us to use the wide variety of (Bayesian) in-
21 ference engines (including ones that do not leverage autodifferentiation). In our initial experiments,
22 we found gradient-free inference engines to work only for low-dimensional systems (3-5 param-
23 eters) and for likelihood functions involving only a small number of timesteps (less than 10). We,
24 therefore, focus only on the gradient-based inference engines provided by pyro.

25 2.1 Inference engines

26 **SVI:** The stochastic variational inference (SVI) solver proceeds by maximizing the evidence lower
27 bound (ELBO). The *stochasticity* in variational inference arises from the number of samples from
28 the variational distribution used to compute the likelihood (and by extension the ELBO). In our
29 experiments, we use 4 samples to compute the expected ELBO.

30 **HMC and NUTS:** For our gradient-based Monte Carlo solvers (Hamiltonian Monte Carlo and No-
31 U-Turn Sampler), we leverage numpyro for efficient inference. Pyro with a pytorch backend has
32 been extensively benchmarked against PyMC3 [5] and numpyro [2], and has been found to be sig-
33 nificantly slower.

34 2.2 Prior distributions

35 **Choice of priors for continuous variables:** For the continuous-valued random variables in our
36 probabilistic model, recall that we use a standard normal distribution as our prior. We also experi-

Submitted to the 6th Conference on Robot Learning (CoRL 2022). Do not distribute.

37 mented with the standard uniform distribution $\text{Unif}(0, 1)$, but found marginal gains in inference time
38 when using Gaussian priors.

39 **Reparameterization trick:** Several physical variables of interest in our model have varying con-
40 straints on parameter ranges. For instance, elasticity and stiffness parameters are on the scale of
41 $[1000, 10000]$, while coefficients of static friction are in the $[0, 1]$ range. As such both variational
42 and Monte Carlo inference engines are numerically unstable for such large variations in parameter
43 ranges. We therefore reparameterize all parameters to a canonical range (-1 through 1).

44 **Computing statistics over reparameterization ranges:** To carry out the above reparameterization,
45 we use the held out instances (e.g., the *train* split comprising PartNet [6] 3D assets).

46 **Hierarchical priors:** Our framework, optionally, allows for the flexibility of auto-tuning the ob-
47 servation noise variable if needed. This may be done by introducing a half-normal distribution that
48 samples the standard-deviation of the observation noise parameter.

49 **References**

- 50 [1] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip,
51 P. Horsfall, and N. D. Goodman. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine*
52 *Learning Research*, 2018.
- 53 [2] D. Phan, N. Pradhan, and M. Jankowiak. Composable effects for flexible and accelerated probabilistic
54 programming in numpyro. *International Conference on Probabilistic Programming*, 2020.
- 55 [3] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and
56 A. Lerer. Automatic differentiation in pytorch. 2017.
- 57 [4] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. Van-
58 derPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy pro-
59 grams, 2018. URL <http://github.com/google/jax>.
- 60 [5] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck. Probabilistic programming in python using pymc3. *PeerJ*
61 *Computer Science*, 2:e55, 2016.
- 62 [6] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su. PartNet: A large-scale benchmark
63 for fine-grained and hierarchical part-level 3D object understanding. In *CVPR*, 2019.