

1 Appendices

2 A Terminal Losses

3 This section provides detailed derivations of the different terminal loss functions, $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G})$,
4 compared in our generalized Bayesian inference planner. Recall, the loss function $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G})$
5 measures the discrepancy between the terminal state distribution and the distribution of goal states,
6 represented implicitly by samples \tilde{X}_T and \tilde{G} respectively.

7 **Kernel Maximum Mean Discrepancy (MMD).** The squared MMD [1] measures discrepancy
8 between two distributions $p(x)$ and $q(y)$, and is defined as:

$$\text{MMD}^2(p, q) = \mathbb{E}[k(x, x')] - 2\mathbb{E}[k(x, y)] + \mathbb{E}[k(y, y')] \quad (1)$$

9 where $k(\cdot, \cdot)$ is a kernel function. In the case where distributions p and q are implicit, an unbiased
10 two-sample approximation of the squared MMD is given by:

$$\text{MMD}_u^2(p, q) = \frac{1}{m-1} \sum_{i=1}^m \sum_{j \neq i}^m k(x_i, x_j) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j) + \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(y_i, y_j) \quad (2)$$

11 where $x_i \sim p$, $1 \leq i \leq m$ and $y_i \sim q$, $1 \leq i \leq n$. When applied to our setting we simply define
12 $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G}) = \text{MMD}^2(\tilde{X}_T, \tilde{G})$ where $\tilde{X}_T = \{x_T^{(i)}\}_{i=1}^M$ correspond to samples from the terminal
13 state associated with the corresponding trajectory particle $\tau^{(i)}$.

14 **KL Divergence.** A differentiable two-sample KL divergence approximation can be obtained
15 through density ratio estimation via classification [2, 3]. The KL divergence between distributions p
16 and q is defined as:

$$D^{\text{KL}}(p, q) = \mathbb{E}_{x \sim p(x)} \left[\log \frac{p(x)}{q(x)} \right]. \quad (3)$$

17 This quantity can be approximated over samples $x^{(i)} \sim p(x)$ if the ratio $r(x) = p(x)/q(x)$ can be
18 computed analytically:

$$\mathbb{E}_{x \sim p(x)} [\log r(x)] \approx \frac{1}{M} \sum_{i=1}^M \log r(x^{(i)}) \quad (4)$$

19 In the case where the goal distribution and posterior distribution are implicit, the ratio $r(x)$ cannot
20 be evaluated. Instead, the ratio can be approximated through classification, where binary label $y = 1$
21 indicates that a sample x was drawn from p and $y = 0$ indicates a sample was drawn from q . The
22 ratio can then be approximated as:

$$r^*(x) = \frac{P(x|y=1)}{P(x|y=0)}, \quad (5)$$

23 which can be further simplified to:

$$r^*(x) = \exp(\sigma^{-1}(P(y=1|x))). \quad (6)$$

24 This approximation reduces the density ratio estimation to a classification problem. Thus, the ratio
25 is computed by evaluating a binary classifier for each sample which can be performed efficiently.

26 Intuitively, this approach assumes that the KL divergence should be highest when the classifier
27 cannot distinguish between the samples from each distribution. The classifier must be trained in-
28 dividually for each estimated trajectory distribution in order to use this divergence. To improve
29 computational efficiency in practice, the classifier can be initialized with the result from the pre-
30 vious iteration. Finally, we can define $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G}) = D^{\text{KL}}(\tilde{X}_T, \tilde{G})$ where the KL divergence is
31 approximated following Eq. (4).

32 **Smooth K-Nearest Neighbor.** It is possible to define a differentiable two-sample test based on the
33 well known k-NN algorithm as demonstrated in [4]. The Smooth K-Nearest Neighbor test possesses

34 important statistical properties such as consistency and convergence of its statistics to f -divergence.
 35 This is despite the complexity of having to solve a combinatorial optimization problem (nearest
 36 neighbor match) required by the k-NN method. Let n_1 be the number of samples of p , and n_2 be the
 37 number of samples of q . The k-NN divergence can be defined as

$$D_\alpha^{\text{NN}} = \int \frac{\alpha^2 p^2(x) + (1 - \alpha)^2 q^2(x)}{\alpha p(x) + (1 - \alpha)q(x)} dx, \quad (7)$$

38 for $n_1/(n_1 + n_2) \rightarrow \alpha \in (0, 1)$.

39 As proved in [5], the statistic

$$1 - \frac{T(X_1, X_2)}{(n_1 + n_2)^k}, \quad (8)$$

40 where $T(X_1, X_2)$ refers to the number of edges connecting samples in a set $X_1 = \{x_i\}_{i=1}^{n_1}$ to a
 41 set $X_2 = \{x_i\}_{i=1}^{n_2}$ from a k -neighborhood graph created with points in X_1 and X_2 , converges in
 42 probability to the D_α^{NN} divergence, and can be used as an efficient approximation. To make the
 43 computation differentiable, the authors of [6] define

$$T(X_1, X_2) = \sum_{i=1}^n \sum_{j=1}^n s_i(\{x_m\}_{m=1}^{n_1+n_2})_j, \quad (9)$$

44 where $s_i(\{x_m\}_{m=1}^{n_1+n_2})$ denotes the softmax function computed on the Euclidean distances between
 45 all points, except point i . As the softmax function is differentiable, the statistic in Eq. (8) becomes
 46 differentiable and can be used directly as our loss function $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G})$. To avoid specifying a
 47 particular value for α , our implementation computes the statistic for several values and averages
 48 them as the final result.

49 **Energy Statistic.** This two-sample test is based on Newton’s gravitational potential energy which
 50 relates two entities by the Euclidean distance between them [7]. Given two distributions $p(x)$ and
 51 $q(y)$, the energy distance is defined as:

$$D^E(p, q) = 2\mathbb{E}[\|x - y\|^2] - \mathbb{E}[\|x - x'\|^2] - \mathbb{E}[\|y - y'\|^2], \quad (10)$$

52 where x and y are independent random variables. The corresponding two-sample statistic given two
 53 sets of samples $X = \{x_i\}_{i=1}^{n_1}$ and $Y = \{y_j\}_{j=1}^{n_2}$ can then be written as:

$$D^E(X, Y) = \frac{2}{n_1 n_2} \sum_{i=1}^{n_1} \sum_{m=1}^{n_2} \|x_i - y_m\|^2 - \frac{1}{n_1^2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_1} \|x_i - x_j\|^2 - \frac{1}{n_2^2} \sum_{l=1}^{n_2} \sum_{m=1}^{n_2} \|y_l - y_m\|^2. \quad (11)$$

54 This provides a computationally efficient statistic which can be directly used as our loss
 55 $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G}) = D^E(\tilde{X}_T, \tilde{G})$.

56 A.1 Practical Considerations

57 The statistics considered in this work are good local approximations of distribution divergences.
 58 In the case of trajectory optimization, when the terminal states in early planning iterations are far
 59 from the goal set, the goal loss gradients can be uninformative. We therefore include a prior in our
 60 set planning method consisting of a smooth uniform distribution constructed by placing a bounding
 61 box around the goal samples. This can be included in our framework by multiplying a prior over
 62 the terminal state $p(x_T)$ with the goal likelihood in Eq. (6). This mitigates the poor divergence
 63 approximation in early iterations. Note that the uniform prior is insufficiently informative on its
 64 own, particularly in cases where the goal set is multi-modal. Furthermore, the prior needs to be
 65 differentiable, which is not the case for a standard uniform distribution. Therefore we define a
 66 smooth uniform prior in the region $R = x_T: a \leq x_T \leq b$ as

$$p(x_T) \propto \exp\left(-d(x_T, R)^2 / \sqrt{(2\sigma^2)}\right) \quad (12)$$

67 where $d(x, R) = \min |x - x'|$, $x' \in R$ is a distance function, and σ controls the *sharpness* of the
 68 approximation.

69 Once inference over the trajectory distributions converges, we must select a single trajectory esti-
70 mate to execute. A common approach to accomplish this is by taking the mean, or weighted mean,
71 of the particle set. This method is ineffective when the trajectory distribution is multi-modal. An
72 alternative approach is to pick the maximum weighted particle. Our proposed set-based terminal
73 losses yield a single score over the whole distribution, which does not enable weighing individual
74 particles based on terminal loss. To select our final sample, we instead select the lowest cost trajec-
75 tory. In practice, we also include the prior in the weight computation to avoid local minima with
76 very low-cost trajectories.

77 B Experiment Details

78 B.1 Planar Navigation

79 The agent state x_t is composed of a 2D position and velocity, and the control signal u_t is a 2D
80 acceleration. We use known, linear dynamics in a fully observed environment for the controller
81 rollouts.

82 **Losses.** For each method, the running cost is summed over each timestep, where the cost for one
83 timestep is:

$$c_t(x_t, u_t, z) = x_t^\top Q x_t + u_t^\top R u_t + \alpha c_{\text{SDF}}(x_t, z) \quad (13)$$

84 where Q and R are quadratic cost parameters for the state and action, and $c_{\text{SDF}}(x_t, z)$ is the obstacle
85 avoidance term, computed using the Signed Distance Function (SDF) over the environment z . The
86 goal loss $\mathcal{L}_{\text{goal}}(\tilde{X}_T, \tilde{G})$ is a set goal loss which is differentiable with respect to the trajectory τ .

87 **Implementation Details.** For each of our goal set planner ablations, $N = 50$ samples are ran-
88 domly selected from goal distribution, except for *KL (Ratio Estimation)*, which uses $N = 100$
89 samples. This method involves training a learned classifier so is aided by a higher sample size. For
90 the closest point methods, the goal sample with the smallest Euclidean distance from the start state
91 is selected. For all the methods, $M = 50$ particles are used to represent the trajectory distribution.
92 The particles represent the discrete control signals, u_t , which are 2D accelerations at each 0.1 sec-
93 ond timestep over a horizon of 3 seconds. All planners are initialized with the distribution from
94 the previous timestep, shifted to the current timestep, and run for 50 iterations. We use the Adam
95 optimizer [8] to select the step size in the SVGD update rule.

96 The KL divergence uses a 3 layer fully-connected network as the classifier, retrained at each
97 timestep. To mitigate computational complexity, we warm start the training with the weights from
98 the previous timestep. The Kernel MMD uses an RBF kernel, with a bandwidth selected by applying
99 the median heuristic over the goal samples [8]. The Smooth k-NN loss uses a value of $k = 1$.

100 We use the RBF kernel for SVGD, and set the bandwidth using the median heuristic, a popular tech-
101 nique for choosing the kernel bandwidth which yields a good estimate in many cases [9]. Without
102 access to data consisting of trajectory samples, we assume that they are normally distributed with
103 covariance $\sigma = 1$. Under these assumptions, it follows that the expected distance between samples
104 drawn from the distribution is $2D$, where D is the trajectory dimension.

105 B.2 Grasping

106 Figure 1 shows a case where the goal set planner finds a reachable point from the set, but the closest
107 point method may select goal points which are in collision in some of the example scenes. This
108 phenomenon helps explain the discrepancy in grasp success at the cost of increasing distance error.

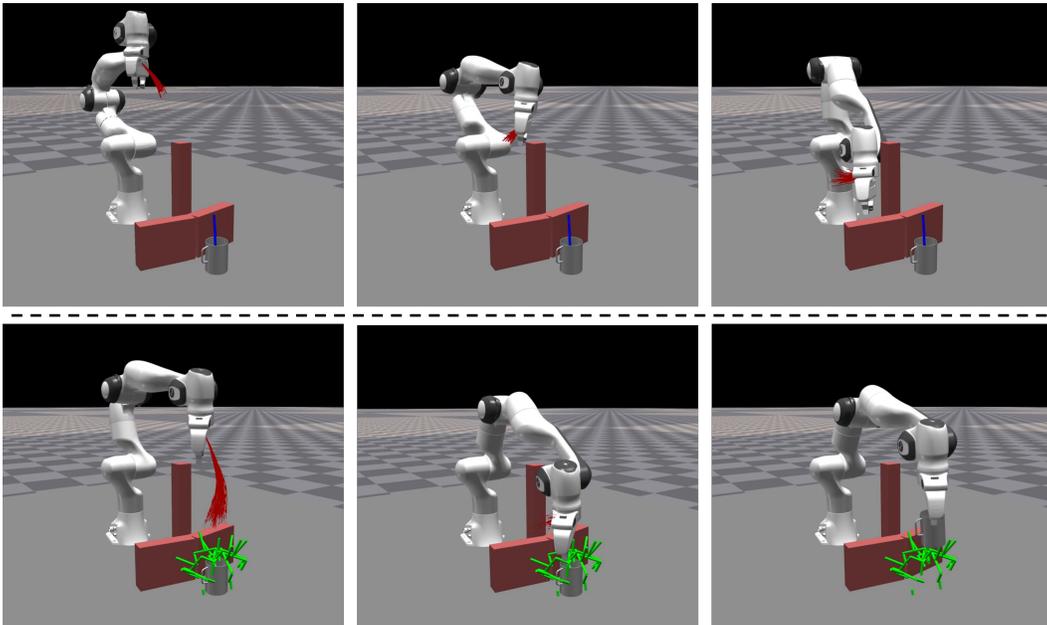


Figure 1: Example failure execution for the closest point planner (top) compared the the goal set planner (bottom). The closest point selected is shown in blue. The point is not reachable in the environment, causing the robot to fail to reach it (top right). The set planner considers all the grasp samples, shown in green. It finds a reachable grasp point and grasps successfully (bottom right).

109 **References**

- 110 [1] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample
111 test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- 112 [2] M. Sugiyama, T. Suzuki, and T. Kanamori. *Density Ratio Estimation in Machine Learning*.
113 Cambridge University Press, 2012.
- 114 [3] A. Menon and C. S. Ong. Linking losses for density ratio and class-probability estimation. In
115 *International Conference on Machine Learning*, pages 304–313. PMLR, 2016.
- 116 [4] M. F. Shilling. Multivariate two-sample tests based on nearest neighbours. *Journal of the*
117 *American Statistical Association*, 81(395):799–806, 1986.
- 118 [5] N. Henze. A multivariate two-sample test based on the number of nearest neighbor type coinci-
119 dences. *Annals of Statistics*, pages 772–783, 1988.
- 120 [6] J. Djolonga and A. Krause. Learning implicit generative models using differentiable graph tests.
121 *arXiv preprint arXiv:1709.01006*, 2017.
- 122 [7] G. J. Székely and M. L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal*
123 *of statistical planning and inference*, 143(8):1249–1272, 2013.
- 124 [8] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the*
125 *3rd International Conference on Learning Representations (ICLR)*, 2014.
- 126 [9] D. Garreau, W. Jitkrittum, and M. Kanagawa. Large sample analysis of the median heuristic.
127 *arXiv preprint arXiv:1707.07269*, 2017.