Geometrically-Aware Goal Inference: Leveraging Motion Planning as Inference

Kay Pompetzki
Intelligent Autonomous Systems Group
Technical University of Darmstadt
kay@robot-learning.de

An T. Le
Intelligent Autonomous Systems Group
Technical University of Darmstadt
an@robot-learning.de

Theo Gruner
Intelligent Autonomous Systems Group
Technical University of Darmstadt
Hessian Center for Artificial Intelligence (hessian.AI)
theo@robot-learning.de

Joe Watson
Intelligent Autonomous Systems Group
Technical University of Darmstadt
German Research Center for AI (DFKI)
joe.watson@robot-learning.de

Georgia Chalvatzaki
Interactive Robot Perception & Learning Lab
Technical University of Darmstadt
Hessian Center for Artificial Intelligence (hessian.AI)
georgia.chalvatzaki@tu-darmstadt.de

Jan Peters
Intelligent Autonomous Systems Group
Technical University of Darmstadt
Centre for Cognitive Science
German Research Center for AI (DFKI)
Hessian Center for Artificial Intelligence (hessian.AI)
jan@robot-learning.de

Keywords: Goal Inference, Probabilistic Inference, Motion Planning, Robotics

Acknowledgements

This work receives funding from the European Union's Horizon Europe program under Grant Agreement No. 101135959 (project ARISE), the German Research Foundation project CHIRON (PE 2315/8-1), and the German Research Foundation project METRIC4IMITATION (PE 2315/11-1).

Abstract

Goal inference is crucial in robotics, enabling effective collaboration in Human-Robot Interaction (HRI) and assisted teleoperation. Current approaches often rely on Markov Decision Processs (MDPs) and maximum entropy principles to infer intentions by integrating over trajectory space. However, these methods commonly employ local approximations around optimal trajectories, which oversimplify the integration and result in unimodal trajectory predictions. They predominantly consider straight-line paths or user input-related costs, neglecting geometric and contextual constraints such as obstacles. This paper proposes a Geometrically-aware goal inference framework that integrates motion planning with Bayesian inference. By leveraging batch-wise motion planning to generate multimodal trajectory distributions and employing belief updates through Sequential Monte Carlo methods, our approach demonstrates the efficiency of capturing goal-directed behavior in complex environments. This preliminary study highlights the promise of combining motion planning with goal inference and motivates future research toward more comprehensive evaluations.

1 Introduction

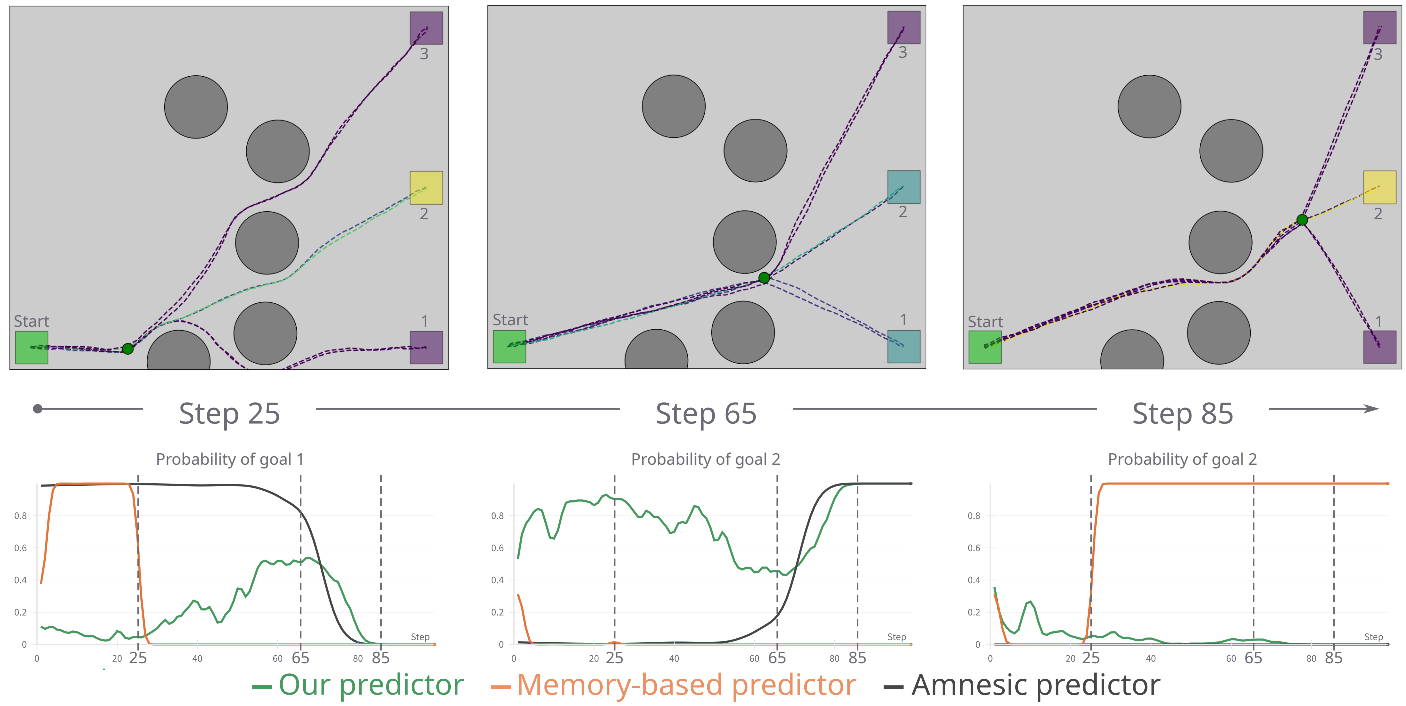

Figure 1: Goal inference example in a 2D navigation task. The agent (green dot) starts at the "Start" location and moves toward one of three goals, following an unknown suboptimal policy that avoids obstacles. **- Top row:** Planned paths generated by Gaussian Process Motion Planning (GPMP). **- Bottom row:** Evolution of goal probabilities over time for three methods. The amnesic predictor (black) relies only on Euclidean distance to the goals. The memory-based predictor (orange) employs a direct, Euclidean-optimal path between points without accounting for obstacles, leading to incorrect probability updates. Our predictor (green) integrates motion planning to infer goals based on path smoothness and distance-to-go, dynamically updating beliefs as the agent progresses. This conceptual demonstration highlights how leveraging motion planning can improve goal inference accuracy.

Goal inference is one of the foundational problems in robotics, particularly in human-robot interaction and assisted teleoperation. In such scenarios, robots must collaborate with humans safely and efficiently, requiring an understanding of the human's objectives or goals. Inverse planning has emerged as a widely-used framework to tackle this challenge, assuming that observed agent movements arise from an underlying intention [1, 2]. By leveraging this assumption, inverse planning aims to predict the intention behind observed actions.

A key finding from the early work on inverse planning [1] is that humans consider geometric information, such as obstacles, when forming beliefs about goals. Their experiments demonstrate that humans do not merely rely on straight-line distances but integrate contextual information from the environment to infer an agent's intention. This insight underscores the importance of incorporating geometric and environmental constraints into computational models of goal inference.

A substantial body of work has explored goal inference within the MDP, modeling it as a rational decision-making problem [3–7]. These methods assume the agent optimizes its behavior to minimize cost or maximize utility. Researchers have adopted the Boltzmann agent model to better reflect real-world scenarios, where observed behavior is rarely optimal, first introduced by Baker and Ziebart [1, 2]. This model relaxes the strict assumption of perfect rationality by treating actions probabilistically, assigning higher probabilities to lower-cost actions while allowing for suboptimal behavior. This probabilistic framing transforms intention prediction into an inference problem, enabling a more realistic representation of human behavior.

Despite the progress, many existing approaches rely on simplified assumptions that limit their applicability. For example, methods using Laplace approximations [3, 6] or Euclidean metrics often approximate the distribution over trajectories with straight-line paths, ignoring contextual information such as obstacles or environmental constraints. Similarly, Bayesian inference methods [8–11] and deep learning approaches [12–14] refine goal inference by modeling temporal dependencies or encoding task-specific dynamics. However, these methods frequently assume overly simplistic state transitions or require extensive data, and they struggle to capture the inherent multi-modal nature of real-world trajectories, where diverse paths may lead to different goals.

Recent advancements in motion planning and machine learning provide a unique opportunity to address these limitations. In particular, batch motion planning has become increasingly potent due to advancements in CPU/GPU vectorization, as demonstrated in [15–18]. These methods significantly accelerate computation, making it practical to handle multiple trajectory hypotheses in parallel. Furthermore, the increasing computational power of GPUs and frameworks like JAX [19] allows for scalable and efficient batch optimization. By integrating motion planning into goal inference, it becomes possible to better model the complex interplay between goals, environmental constraints, and the multi-modal nature of trajectories.

In this paper, we propose a geometrically-aware goal inference framework that leverages motion planning to address the limitations of existing approaches. Specifically, we:

- Use batch-wise GPMP [20] to generate multi-modal trajectory distributions by sampling diverse initial conditions [21].
- Incorporate path smoothness and obstacle-aware costs into a Bayesian inference framework for dynamic goal prediction.
- Evaluate the method using two baseline predictors: an amnesic predictor, which relies solely on Euclidean distance, and a memory-based predictor inspired by [3].

Our method demonstrates the potential of combining motion planning with Bayesian inference to capture the complexity of goal-directed behavior. By accounting for geometric constraints and leveraging optimal planning manifolds, this approach improves predictive accuracy in scenarios where ambiguity between goals is prevalent. Furthermore, this framework is computationally efficient, taking advantage of GPU-accelerated batch optimization to handle complex, high-dimensional settings.

This work serves as a conceptual step toward bridging the gap between motion planning and goal inference, with implications for shared control in human-robot collaboration and extensions to multi-agent systems where goal negotiation is critical.

2 Geometrically-Aware Goal Inference

In this work, we aim to infer the agent's intended goal $\boldsymbol{g} \in \mathcal{G}$, where $\mathcal{G} \subseteq \mathcal{S}$ is a known set of possible goals, based on observations $\boldsymbol{o}_t = \{\boldsymbol{s}_i, \boldsymbol{a}_i, \boldsymbol{c}_i\}_{i=1}^t$ up to the current time t. The observation includes the agent's states $\boldsymbol{s} \in \mathcal{S}$, agent's actions $\boldsymbol{a} \in \mathcal{A}$, and additional environmental information $\boldsymbol{c}$. The agent is assumed to follow an *unknown, suboptimal policy* that avoids environmental constraints, such as obstacles. We model a pseudo-likelihood of the agent's goal $\boldsymbol{g}$ using a Boltzmann distribution over trajectories

$$p(\mathcal{O}_{\boldsymbol{o}_t} = 1 | \boldsymbol{g}, \boldsymbol{\tau}) \propto \exp(-\mathcal{L}(\boldsymbol{g}, \boldsymbol{\tau}, \boldsymbol{o}_t)).$$

The cost function $\mathcal{L} : \mathcal{G} \times \mathcal{T} \times \mathcal{O} \to \mathbb{R}$ accounts for path feasibility, smoothness, and other constraints. Using Bayes' rule, the posterior over goals becomes

$$p(\boldsymbol{g} | \mathcal{O}_{\boldsymbol{o}_t} = 1) \propto \int_{\mathcal{T}} p(\mathcal{O}_{\boldsymbol{o}_t} = 1 | \boldsymbol{g}, \boldsymbol{\tau}) p(\boldsymbol{\tau} | \boldsymbol{g}) p(\boldsymbol{g}) \, \mathrm{d}\boldsymbol{\tau} \propto p(\mathcal{O}_{\boldsymbol{g}} = 1 \mid \boldsymbol{g}) \int p(\mathcal{O}_{\boldsymbol{\tau}} = 1 \mid \boldsymbol{\tau}) p(\boldsymbol{\tau} \mid \boldsymbol{g}) \, \mathrm{d}\boldsymbol{\tau} \; p(\boldsymbol{g}), \tag{1}$$

assuming an independence between goal and trajectory related costs $\mathcal{L}_{\boldsymbol{o}}(\cdot) = \mathcal{L}_{\boldsymbol{g}}(\cdot) + \mathcal{L}_{\boldsymbol{\tau}}(\cdot)$. Thus, $p(\boldsymbol{g})$ is the prior over goals, $p(\boldsymbol{\tau} \mid \boldsymbol{g})$ a probability over possible trajectories and the pseudo-likelihoods $p(\mathcal{O}_{\boldsymbol{g}} = 1 \mid \boldsymbol{g})$ and $p(\mathcal{O}_{\boldsymbol{\tau}} = 1 \mid \boldsymbol{\tau})$.

Marginalizing over the entire trajectory space $\mathcal{T}$ in equation 1 is computationally intractable for continuous state and action spaces [3,6]. We do not follow a Laplace approximation around an optimal trajectory to address this issue. Rather, we approximate the posterior by leveraging motion planning to compute a batch of *maximum a-posteriori* (MAP) trajectories. These trajectories are conditioned on the observed portion $\boldsymbol{o}_t$ and represent the most likely completions of $\boldsymbol{\tau}$ for each goal $\boldsymbol{g}$. In this paper, we use batch-wise GPMP [20] to optimize these MAP trajectories, incorporating geometric constraints such as obstacle avoidance. To capture the multimodality of $\mathcal{T}$, we initialize the GPMP instances with diverse initial states, sampling them in a manner similar [21]. For each goal $\boldsymbol{g}$, the MAP cost is given

$$\mathcal{L}_{\mathrm{MAP}}(\boldsymbol{g}) = \mathcal{L}(\boldsymbol{g}) + \mathbb{E}_{\boldsymbol{\tau} \sim p(\boldsymbol{\tau} | \boldsymbol{g}, \boldsymbol{o}_t)}[\mathcal{L}(\boldsymbol{\tau}, \boldsymbol{o}_t)] \approx \mathcal{L}(\boldsymbol{g}) + \frac{1}{B} \sum_{\mathcal{T}_{\mathrm{MAP}}} \mathcal{L}(\boldsymbol{\tau}_{\mathrm{MAP}}, \boldsymbol{o}_t) \tag{2}$$

with $\mathcal{L}(\boldsymbol{g})$ and $\mathcal{L}(\boldsymbol{\tau}, \boldsymbol{o}_t)$ being goal-specific costs and the trajectory-related costs, respectively. This approximation reduces the need to integrate over the entire trajectory space while still capturing the most relevant information for inference.

To refine the belief over goals, we consider a pseudo-posterior problem that minimizes the divergence between the current belief and the MAP-based cost distribution

$$\pi^*(\boldsymbol{g}) = \operatorname*{arg\,min}_{\pi(\cdot) \in \mathcal{P}(\mathcal{G})} \; \mathbb{D}_{\mathrm{KL}}(\pi(\boldsymbol{g}), \exp(-\alpha \mathcal{L}_{\mathrm{MAP}}(\boldsymbol{g})) p(\boldsymbol{g}))$$
$$\text{s.t. } \mathbb{D}_{\mathrm{KL}}(\pi(\boldsymbol{g}), \pi_i(\boldsymbol{g})) \leq \epsilon, \; \textstyle\sum_{j=1}^n \pi(\boldsymbol{g}_j) = 1$$

The update rule includes a relative entropy term to ensure smooth belief updates between iterations [22, 23]. Following [22], we know that a closed-form update for $\pi(g)$ is

$$\pi(\boldsymbol{g})^* = Q^{-1} \exp(-\alpha\mathcal{L}(\boldsymbol{g}))^\eta p(\boldsymbol{g})^\eta \pi_i(\boldsymbol{g})^{1-\eta},$$

where Q is the normalization constant, and $\eta = 1/(1 + \lambda_1)$ controls the learning rate.

Since the direct computation of the posterior remains intractable, we implement an *iterative approximation scheme* using Sequential Monte Carlo (SMC) [24, 25]. The update process is as follows: (i) Start with the observed trajectory $\boldsymbol{o}$, a prior $p(\boldsymbol{g})$ and the current belief distribution $\pi_i(\boldsymbol{g})$; (ii) For each goal $\boldsymbol{g}$, use GPMP to compute the MAP cost $\mathcal{L}_{\text{MAP}}(\boldsymbol{g})$, conditioned on the observed trajectory $\hat{\boldsymbol{\tau}}$; (iii) Then update the belief distribution over goals using a

$$\pi_{i+1}(\boldsymbol{g}) = \textstyle\sum_{j=1}^n a_j^{i+1} \delta(\boldsymbol{g}_j)$$
$$a_j^{i+1} := \frac{a_j^i \exp(-\alpha\mathcal{L}_{\text{MAP}}(\boldsymbol{g}_j))^\eta p(\boldsymbol{g}_j)^\eta \pi_i(\boldsymbol{g}_j)^{-\eta}}{\sum_{k=1}^n a_k^i \exp(-\alpha\mathcal{L}_{\text{MAP}}(\boldsymbol{g}_k))^\eta p(\boldsymbol{g}_k)^\eta \pi_i(\boldsymbol{g}_k)^{-\eta}}.$$

(iv) Iterate steps 2 - 4 until the belief converges or additional observations are made. Thus, our proposed approach systematically incorporates optimal trajectories into goal inference, effectively incorporating geometric constraints to infer the most likely goal.

3 Experiment

We conducted a preliminary study to explore the potential of geometrically-aware goal inference in a 2D navigation task (see figure 1). In this task, the agent moves from an initial state toward one of three possible goal locations, following an unknown, suboptimal policy that avoids constraints like obstacles. Our method aims to infer the agent's intended target based on its observed trajectory. All computations and updates were implemented in JAX [19], leveraging its efficient automatic differentiation and GPU acceleration capabilities to scale the method effectively.

Our predictor leverages batch-wise motion planning to approximate the trajectory distribution from the start state through the agent's current state to each potential goal. Specifically, we sample initial trajectory particles from an initial high-variance GP for GPMP optimization, resulting in diverse initial conditions discovering multiple solutions [21]. The motion planner generates paths assuming the initial state, the current agent state, and each goal state as inputs. These planned paths are used to compute two key costs: (i) the path length from the current state to the goal and (ii) the smoothness of the whole trajectory. These costs are incorporated into our framework to update the belief distribution over potential goals.

We compared our method to two heuristic-based baselines: an amnesic predictor, which only considers the squared Euclidean distance to each goal, and a memory-based predictor inspired by [3], which uses the sum of squared velocities as the cost. These baselines were selected to align with our current cost-based framework, which does not yet incorporate learned priors. The results illustrate how our predictor dynamically adjusts goal probabilities based on the observed trajectory. Initially, it assigns a high probability to a likely goal, adapts as the agent's path evolves, and eventually converges to the true goal as the agent approaches it. These results suggest that integrating motion planning with Bayesian inference provides a promising framework for goal inference.

This experiment serves as a conceptual demonstration of the method, illustrating its potential while acknowledging current limitations. Future studies will explore incorporating learned priors into the optimization process and extending the approach to more complex environments and higher-dimensional tasks.

4 Conclusion

In this work, we introduced a goal inference framework that seamlessly integrates motion planning as a subproblem. By approximating the trajectory distribution, we consider key geometric and topological factors in the framework. Although recent advances suggest computationally efficient motion planning approaches [15,16,18], the scalability of our approach remains to be further investigated in high-dimensional spaces. Nonetheless, our approach lays a basis for practical applications in areas such as assisted teleoperation and HRI, and enables a geometrically-aware goal inference pipeline for real-world scenarios.

References

[1] C. L. Baker, J. B. Tenenbaum, and R. R. Saxe, "Goal inference as inverse planning," in *Proceedings of the annual meeting of the cognitive science society*, vol. 29, no. 29, 2007.

[2] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey *et al.*, "Maximum entropy inverse reinforcement learning." in *Aaai*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.

[3] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.

[4] K. Muelling, A. Venkatraman, J.-S. Valois, J. E. Downey, J. Weiss, S. Javdani, M. Hebert, A. B. Schwartz, J. L. Collinger, and J. A. Bagnell, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Autonomous Robots*, vol. 41, pp. 1401–1422, 2017.

[5] S. Javdani, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization," *Robotics science and systems: online proceedings*, vol. 2015, 2015.

[6] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization for teleoperation and teaming," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, 2018.

[7] A. Gottardi, S. Tortora, E. Tosello, and E. Menegatti, "Shared control in robot teleoperation with improved potential fields," *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 3, pp. 410–422, 2022.

[8] Z. Wang, K. Mülling, M. P. Deisenroth, H. Ben Amor, D. Vogt, B. Schölkopf, and J. Peters, "Probabilistic movement modeling for intention inference in human–robot interaction," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 841–858, 2013.

[9] K. Hauser, "Recognition, prediction, and planning for assisted teleoperation of freeform tasks," *Autonomous Robots*, vol. 35, pp. 241–254, 2013.

[10] O. Dermy, A. Paraschos, M. Ewerton, J. Peters, F. Charpillet, and S. Ivaldi, "Prediction of intention during interaction with icub with probabilistic movement primitives," *Frontiers in Robotics and AI*, vol. 4, p. 45, 2017.

[11] S. Jain and B. Argall, "Probabilistic human intent recognition for shared autonomy in assistive robotics," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 9, no. 1, pp. 1–23, 2019.

[12] P. Kratzer, N. B. Midlagajni, M. Toussaint, and J. Mainprice, "Anticipating human intention for full-body motion prediction in object grasping and placing tasks," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2020, pp. 1157–1163.

[13] X. Zhao, H. Li, T. Miao, X. Zhu, Z. Wei, L. Tan, and A. Song, "Learning multimodal confidence for intention recognition in human-robot interaction," *IEEE Robotics and Automation Letters*, 2024.

[14] X. Zheng, C. Tang, Z. Wan, C. Hu, and W. Zhang, "Multi-level confidence learning for trustworthy multimodal classification," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 9, 2023, pp. 11 381–11 389.

[15] W. Thomason, Z. Kingston, and L. E. Kavraki, "Motions in microseconds via vectorized sampling-based planning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 8749–8756.

[16] A. T. Le, K. Hansel, J. Carvalho, J. Watson, J. Urain, A. Biess, G. Chalvatzaki, and J. Peters, "Global tensor motion planning," *arXiv preprint arXiv:2411.19393*, 2024.

[17] C. W. Ramsey, Z. Kingston, W. Thomason, and L. E. Kavraki, "Collision-affording point trees: Simd-amenable nearest neighbors for fast collision checking," *arXiv preprint arXiv:2406.02807*, 2024.

[18] A. T. Le, G. Chalvatzaki, A. Biess, and J. R. Peters, "Accelerating motion planning via optimal transport," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[19] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018. [Online]. Available: http://github.com/jax-ml/jax

[20] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuous-time gaussian process motion planning via probabilistic inference," *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.

[21] J. Urain, A. T. Le, A. Lambert, G. Chalvatzaki, B. Boots, and J. Peters, "Learning implicit priors for motion optimization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 7672–7679.

[22] J. Peters, K. Mulling, and Y. Altun, "Relative entropy policy search," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 24, no. 1, 2010, pp. 1607–1612.

[23] J. Schulman, "Trust region policy optimization," *arXiv preprint arXiv:1502.05477*, 2015.

[24] P. Del Moral, A. Doucet, and A. Jasra, "Sequential monte carlo samplers," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 68, no. 3, pp. 411–436, 2006.

[25] B. Dai, N. He, H. Dai, and L. Song, "Provable bayesian inference via particle mirror descent," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 985–994.