Geometrically-Aware Goal Inference: Leveraging Motion Planning as Inference

Kay Pompetzki[1], An T. Le[1], Theo Gruner[1,4], Joe Watson[5], Georgia Chalvatzaki[1,4], Jan Peters[1−4]

***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 motion planning as inference 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.**

I. Introduction

Goal inference is a fundamental problem in robotics, particularly in human-robot interaction and assisted teleoperation, where robots must interpret human intentions for safe and efficient collaboration. Inverse planning provides a well-established framework assuming observed movements stem from an underlying objective [1], [2]. By leveraging this assumption, inverse planning aims to predict the intention behind observed actions.

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 reflect better 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 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. However, incorporate such contextual factors rather than relying solely on direct distances when inferring an agent's intention [1]. 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.

In this paper, we propose a geometrically-aware goal inference framework to address the limitations of existing approaches. Specifically, we: (i) Approximate trajectory distributions using motion planning as inference [15]–[19], enabling multi-modal predictions; (ii) Incorporate path smoothness and obstacle-aware costs into the goal inference framework for dynamic goal prediction; (iii) Conducting a preliminary study comparing our method against two baselines. Our method demonstrates the potential of combining motion planning with Bayesian inference to capture the complexity of goal-directed behavior. It improves predictive accuracy in scenarios where ambiguity between goals is prevalent by accounting for geometric constraints and leveraging optimal planning manifolds. This work bridges the gap between motion planning and goal inference, with implications for shared control in human-robot collaboration and multi-agent systems where goal negotiation is critical.

II. 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

$$\begin{aligned} p(\boldsymbol{g}|\mathcal{O}_{\boldsymbol{o}_t} = 1) &\propto \textstyle\int_{\mathcal{T}} p(\mathcal{O}_{\boldsymbol{o}_t} = 1|\boldsymbol{g}, \boldsymbol{\tau}) p(\boldsymbol{\tau}|\boldsymbol{g}) p(\boldsymbol{g}) \ \mathrm{d}\boldsymbol{\tau} \\ &\propto p(\mathcal{O}_{\boldsymbol{g}} = 1 \mid \boldsymbol{g}) \textstyle\int_{\mathcal{T}} p(\mathcal{O}_{\boldsymbol{\tau}} = 1 \mid \boldsymbol{\tau}) p(\boldsymbol{\tau} \mid \boldsymbol{g}) \ \mathrm{d}\boldsymbol{\tau} \ p(\boldsymbol{g}), \end{aligned} \tag{1}$$

This work receives funding from the European Union's Horizon Europe program under Grant Agreement No. 101135959 (project ARISE).

[1]Computer Science Department, TU Darmstadt; [2]Centre for Cognitive Science, TU Darmstadt; [3]German Research Center for AI (DFKI), Research Department: SAIROL; [4]Hessian.AI; [5]Department of Engineering Science, University of Oxford;

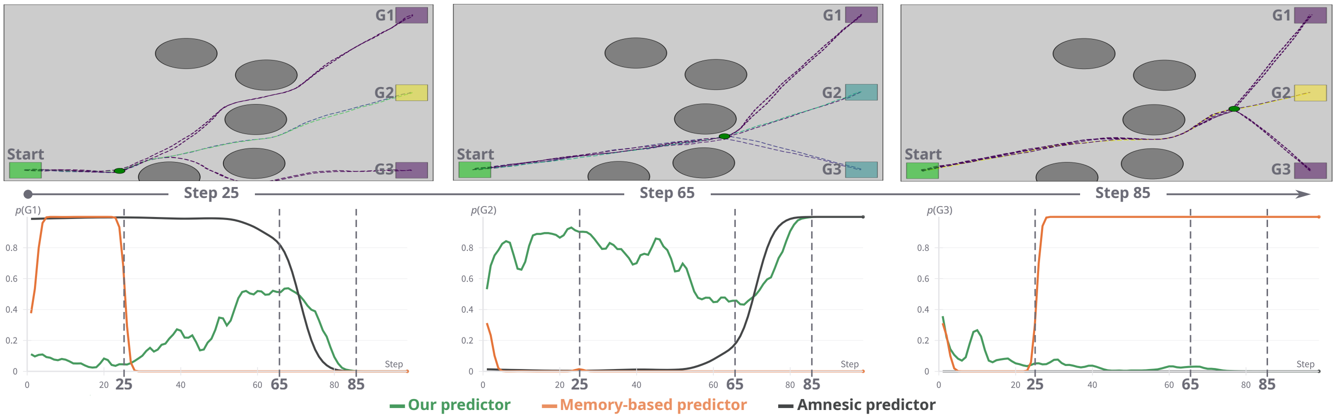

Fig. 1: Goal inference example in a 2D navigation task. The agent (green dot) starts at the "Start" location and moves toward one of three goals, following an unknown suboptimal policy that avoids obstacles. **- Top row:** Planned paths generated by Gaussian Process Motion Planning (GPMP). **- Bottom 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.

assuming an independence between goal and trajectory related costs $\mathcal{L}(\cdot) = \mathcal{L}_{\boldsymbol{g}}(\cdot) + \mathcal{L}_{\boldsymbol{\tau}}(\cdot)$. 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]. To address this issue, we approximate the posterior using motion planning as inference [15]–[20]. Thus, we either approximate the posterior using a batch of local *maximum a-posteriori* (MAP) estimations [21]–[23] or employ variational inference approaches [24]–[26].

To refine the belief over goals, we consider a pseudo-posterior problem [27], [28]that minimizes the divergence between the current belief and Gibbs posterior

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

The update rule includes a relative entropy term to ensure smooth belief updates between iterations [29], [30]. Following [29], we know the closed-form update solution $\pi(\boldsymbol{g})^* = Q^{-1}\exp(-\alpha\mathcal{L}(\boldsymbol{g}))^{\eta} p(\boldsymbol{g})^{\eta} \pi_i(\boldsymbol{g})^{1-\eta}$, 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) [31], [32].

III. 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 intention based on its observed trajectory. All implementations were done in JAX [33], leveraging its efficient automatic differentiation and GPU acceleration capabilities to scale the method effectively.

Our predictor leverages batch-wise GPMP [22], [23] to approximate the trajectory distribution. Specifically, we sample initial trajectory particles from an initial high-variance GP prior, resulting in diverse initial conditions discovering multiple MAP solutions [34]. GPMP generates paths conditioned on the initial state, the current agent state, and each goal state. 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: (i) 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 the highest probability to the second goal, adapts as the agent's path evolves, and eventually converges to the true third goal. These results serve as a conceptual demonstration, illustrating the potential of integrating motion planning within the goal inference problem. Bayesian inference provides a promising framework for goal inference. Future studies will explore incorporating learned priors into the optimization process and extending the approach to more complex environments and higher-dimensional tasks.

IV. 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 [35]–[37], 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] K. Rawlik, M. Toussaint, and S. Vijayakumar, "On stochastic optimal control and reinforcement learning by approximate inference," 2013.

[16] M. Toussaint, "Robot trajectory optimization using approximate inference," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 1049–1056.

[17] M. Toussaint and C. Goerick, "A bayesian view on motor control and planning," in *From Motor Learning to Interaction Learning in Robots*. Springer, 2010, pp. 227–252.

[18] M. Botvinick and M. Toussaint, "Planning as inference," *Trends in cognitive sciences*, vol. 16, no. 10, pp. 485–488, 2012.

[19] K. Hansel, J. Urain, J. Peters, and G. Chalvatzaki, "Hierarchical policy blending as inference for reactive robot control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10 181–10 188.

[20] J. Watson, H. Abdulsamad, and J. Peters, "Stochastic optimal control as approximate input inference," in *Conference on Robot Learning*. PMLR, 2020, pp. 697–716.

[21] J. Dong, M. Mukadam, F. Dellaert, and B. Boots, "Motion planning as probabilistic inference using gaussian processes and factor graphs." in *Robotics: Science and Systems*, vol. 12, no. 4, 2016.

[22] M. Mukadam, X. Yan, and B. Boots, "Gaussian process motion planning," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 9–15.

[23] 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.

[24] A. Lambert and B. Boots, "Entropy regularized motion planning via stein variational inference," *arXiv preprint arXiv:2107.05146*, 2021.

[25] H. Yu and Y. Chen, "A gaussian variational inference approach to motion planning," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2518–2525, 2023.

[26] L. C. Cosier, R. Iordan, S. N. Zwane, G. Franzese, J. T. Wilson, M. Deisenroth, A. Terenin, and Y. Bekiroglu, "A unifying variational framework for gaussian process motion planning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2024, pp. 1315–1323.

[27] J. Watson and J. Peters, "Inferring smooth control: Monte carlo posterior policy iteration with gaussian processes," in *Conference on Robot Learning*. PMLR, 2023, pp. 67–79.

[28] T. Gruner, B. Belousov, F. Muratore, D. Palenicek, and J. R. Peters, "Pseudo-likelihood inference," *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[29] 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.

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

[31] 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.

[32] 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.

[33] 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

[34] 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.

[35] 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.

[36] 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.

[37] 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.