

Extended Abstract: Motion Planners Learned from Geometric Hallucination

Xuesu Xiao[†], Bo Liu[†], and Peter Stone[†]

Abstract—Learning motion planners to move robot from one point to another within an obstacle-occupied space in a collision-free manner requires either an extensive amount of data or high-quality demonstrations. This requirement is caused by the fact that among the variety of maneuvers the robot can perform, it is difficult to find the single optimal plan without many trial-and-error or an expert who is already capable of doing so. However, given a plan performed in obstacle-free space, it is relatively easy to find an obstacle geometry, where this plan is optimal. We consider this “dual” problem of classical motion planning and name this process of finding appropriate obstacle geometry as *hallucination*. In this work, we present two different approaches to hallucinate (1) the *most constrained* and (2) a *minimal* obstacle space where a given plan executed during an exploration phase in a completely safe obstacle-free environment remains optimal. We then train an end-to-end motion planner that can produce motions to move through realistic obstacles during deployment. Both methods are tested on a physical mobile robot in real-world cluttered environments.

I. INTRODUCTION

While classical motion planners, such as Dynamic Window Approach (DWA) [1], can reliably navigate the robot in cluttered spaces with properly tuned parameters, recent machine learning techniques have also been applied to the motion planning problem [2]. Those approaches either learn from a classical motion planner [2] or a human expert [3], or through an extensive amount of trial-and-error, such as Reinforcement Learning (RL) [4]. However, most learned motion planners still under-perform their classical counterparts. Despite the advantage of learning motion planners without hand-crafted rules [1] or in-situ adjustment [5] in a data-driven fashion, the performance is still bottlenecked by the requirement of good-quality training data [6].

Consider the “dual” problem of motion planning: instead of finding the optimal motion plan for a specific obstacle configuration, either using hand-crafted rules or training data, we seek to find the obstacle configuration(s) where a specific motion plan is guaranteed to be optimal. We name this process *hallucination*. Solving this problem gives us the freedom to allow random exploration in a completely safe obstacle-free space and collect an extensive amount of motion plans, whose optimality will be assured by a class of *hallucination* techniques. In this work, we introduce two of those techniques: to hallucinate (1) the (unique) *most constrained* and (2) a (not unique) *minimal* obstacle configuration.

[†]Xuesu Xiao, Bo Liu, and Peter Stone are with Department of Computer Science, University of Texas at Austin, Austin, TX 78712 {xiao, bliu, pstone}@cs.utexas.edu

II. GEOMETRIC HALLUCINATION

Given a robot’s configuration space (C-space) partitioned by unreachable (obstacle) and reachable (free) configurations, $C = C_{obst} \cup C_{free}$, we define the classical motion planning problem as to find a function $f(\cdot)$ that can be used to produce optimal plans $p = f(C_{obst} | c_c, c_g)$ that results in the robot moving from the robot’s current configuration c_c to a specified goal configuration c_g without intersecting the interior of C_{obst} . Here, a plan $p \in \mathcal{P}$ is a sequence of low-level actions $\{u_i\}_{i=1}^t$, where $u_i \in \mathcal{U}$. This work introduces two methods to approach the “dual” problem of finding optimal $f(\cdot)$.

A. Hallucinating Most Constrained Obstacle Space [7]

Since different C_{obst} can lead to the same plan, the left inverse of f , f^{-1} , is not well defined (see Fig. 1 left). However, we can instead define a similar function $g(\cdot)$ such that $C_{obst}^* = g(p | c_c, c_g)$, where C_{obst}^* denotes the C-space’s *most constrained* unreachable set corresponding to p .¹ Formally, given a plan p and the set of all unreachable sets \mathcal{E}_{obst} , we say

$$\begin{aligned} C_{obst}^* = g(p | c_c, c_g) & \text{ iff } \forall C_{obst} \in \mathcal{E}_{obst}, \\ f(C_{obst} | c_c, c_g) = p & \implies C_{obst} \subseteq C_{obst}^*, \end{aligned} \quad (1)$$

We denote the corresponding reachable set of C as $C_{free}^* = C \setminus C_{obst}^*$. We call $g(\cdot)$ the *most constrained* hallucination function and the output of $g(\cdot)$ a *most constrained* hallucination. This hallucination can be projected onto the robot’s sensors. For example, for a LiDAR sensor, we perform ray casting from the sensor to the boundary between C_{obst}^* and C_{free}^* in order to project the hallucination onto the range readings (Fig. 1 right). Given the hallucination C_{obst}^* for p , the only viable (and therefore optimal) plan is $p = g^{-1}(C_{obst}^* | c_c, c_g)$. Note that $g(\cdot)$ is bijective and its inverse $g^{-1}(\cdot)$ is well defined. Leveraging machine learning, $g^{-1}(\cdot)$ is represented using a function approximator $g_\theta^{-1}(\cdot)$. Note that we aim to approximate $g_\theta^{-1}(\cdot)$ instead of the original $f(\cdot)$ due to the vastly different domain size: the *most constrained* (\mathcal{E}_{obst}^*) vs. *all* (\mathcal{E}_{obst}) unreachable sets.

During deployment, we use a smoothed coarse global path from a global planner to generate runtime hallucination so $g_\theta^{-1}(\cdot)$ does not need to generalize to unseen scenarios. Other components, including a Turn in Place, Recovery Behavior, and Speed Modulation modules, are used in conjunction with $g_\theta^{-1}(\cdot)$ to address inevitable out-of-distribution scenarios and adapt to the real C-space.

¹Technically, c_g can be uniquely determined by p and c_c , but we include it as an input to $g(\cdot)$ for notational symmetry with $f(\cdot)$.




Fig. 1: $p = f(C_{obst} | c_c, c_g)$ and $C_{obst}^* = g(p | c_c, c_g)$




Fig. 2: Left: $\overline{C_{obst}^{min}}$ is defined by three consecutive configurations with c_c, c_m, c_g , and symmetry point c'_m . Right: LiDAR reading is randomly sampled between min and max range.

B. Hallucinating Minimal Obstacle Space [8]

Learning from hallucinated C_{obst}^* can efficiently reduce input space, and therefore learning complexity, but requires runtime hallucination and other components during deployment. Hallucination of a *minimal* obstacle space generates C_{obst}^{min} , which is a minimal set of obstacle configurations required to cause the plan p to be optimal. We then randomly sample augmentations to the minimal unreachable set. Formally, we define the set of C_{obst}^{min} as:

$$\mathcal{E}_{obst}^{min} = \{C_{obst}^{min} | \forall c \in C_{obst}^{min}, f(C_{obst}^{min} \setminus \{c\} | c_c, c_g) \neq f(C_{obst}^{min} | c_c, c_g)\} \quad (2)$$

We use a special $\overline{C_{obst}^{min}}$ to approximate any $C_{obst}^{min} \in \mathcal{E}_{obst}^{min}$ (Fig. 2 left). This approximation is sufficient when the robot trajectory is composed of a dense sequence of configurations and C_{obst}^{min} is instantiated on discrete LiDAR beams, which will be shown empirically. As shown in Fig. 2 right, the max range of a LiDAR beam is determined by $\overline{C_{obst}^{min}}$ (if the beam intersects $\overline{C_{obst}^{min}}$) or the sensor's physical limit (if not), while the min range for each beam is determined by the boundary of the robot path. A random range is sampled between the min and max values, considering possible continuity among neighboring beams and being offset for uncertainty/safety induced by the optimal plan p . Therefore, many C_{obst} can be augmented based on $\overline{C_{obst}^{min}}$. We then train a parameterized policy $f_\theta(\cdot)$ to approximate classical motion planner $f(\cdot)$.

The advantage of augmenting $\overline{C_{obst}^{min}}$ and generating many C_{obst} is, during deployment, no runtime hallucination with the help of a global path and other extra components are required. The learned $f_\theta(\cdot)$ can plan in response to the real perception and adapt to the actual scenarios on its own.

C. Physical Experiments²

Two datasets are collected by two random exploration policies in an obstacle-free space in simulation: one with

²Videos: <https://www.youtube.com/watch?v=T72Z6rz9ges&t=1s> and <https://www.youtube.com/watch?v=xtLaSF0kiB0&t=49s>.

mostly constant 0.4m/s linear velocity ($v \approx 0.4\text{m/s}$) and varying angular velocity ($\omega \in [-1.57, 1.57]\text{rad/s}$), the other with varying $v \in [0, 1.0]\text{m/s}$ and $\omega \in [-1.57, 1.57]\text{rad/s}$. If trained on the first dataset, the speed of the planner output is modulated by a Model Predictive Control based collision probability checker, achieving a max $v = 0.6\text{m/s}$. Four neural network based planners are trained using the two datasets and two hallucination techniques. Simulated [9] and physical experiments are performed. While the *minimal* hallucination works well on both datasets and outperforms all other variants, and even a classical motion planner [1], the *most constrained* hallucination only performs well on the 0.4m/s dataset. This is because learning from varying speed while hallucinating only the most constrained space causes ambiguity for the learner.

III. CONCLUSIONS

We present a class of two geometric hallucination techniques that approach the classical motion planning problem from the opposite direction. Instead of seeking an optimal motion plan for an obstacle configuration, we find the obstacle configuration(s), where a motion plan is optimal. The first approach hallucinates the *most constrained* C-space, where the plan is the only feasible, and therefore optimal, plan. It largely reduces the learning complexity, since the learned motion planner $g_\theta^{-1}(\cdot)$ only plans in the most constrained C-spaces, instead of any C-spaces. However, the downside of this approach is during deployment, runtime hallucination along with other extra components are necessary. The second approach finds a *minimal* obstacle set to make a plan optimal, augments this minimal set to generate a large body of training data, and therefore does not require any extra components during deployment.

REFERENCES

- [1] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [2] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, and C. Cadena, "From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1527–1533.
- [3] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [4] A. Faust, K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson, "Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5113–5120.
- [5] X. Xiao, J. Dufek, T. Woodbury, and R. Murphy, "Uav assisted usv visual navigation for marine mass casualty incident response," in *2017 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6105–6110.
- [6] B. Liu, X. Xiao, and P. Stone, "Lifelong navigation," *arXiv preprint arXiv:2007.14486*, 2020.
- [7] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *arXiv preprint arXiv:2007.14479*, 2020.
- [8] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," *arXiv preprint arXiv:2010.08098*, 2020.
- [9] P. Daniel, T. Abigail, X. Xiao, and P. Stone, "Benchmarking metric ground navigation," *arXiv preprint arXiv:2008.13315*, 2020.