

Intent-based Deep Reinforcement Learning for Multi-agent Informative Path Planning

Tianze Yang¹, Yuhong Cao¹, Guillaume Sartoretti^{1†}

Abstract—In multi-agent informative path planning (MAIPP), agents must collectively construct a global *belief* map of an underlying distribution of interest (e.g., gas concentration, light intensity, or pollution levels) over a given domain, based on measurements taken along their trajectory. They must frequently replan their path to balance the exploration of new areas with the exploitation of known high-interest areas, to maximize information gain within a predefined budget. Traditional approaches rely on reactive path planning conditioned on other agents’ predicted future actions. However, as the belief is continuously updated, the predicted actions may not match the executed actions, introducing noise and reducing performance. We propose a decentralized, deep reinforcement learning (DRL) approach using an attention-based neural network, where agents optimize long-term individual and cooperative objectives by sharing their *intent*, represented as a distribution of medium-/long-term future positions obtained from their own policy. Intent sharing enables agents to learn to claim or avoid broader areas, while the use of attention mechanisms allows them to identify useful portions of imperfect predictions, maximizing cooperation even based on imperfect information. Our experiments compare the performance of our approach, its variants, and high-quality baselines across various MAIPP scenarios. We finally demonstrate the effectiveness of our approach under limited communication ranges, towards deployments under realistic communication constraints.

I. INTRODUCTION

Many robotic deployments rely on information gathering, such as autonomous exploration [1], environment monitoring [2], search-and-rescue [3], and surface inspection [4]. Starting from a (usually empty) initial belief of the environment, the agent needs to reason about known, informative areas as well as unknown ones to plan an informative path and collect information to update its belief under a given budget constraint (e.g., maximum path length or mission duration). Such deployments are known as adaptive *informative path planning* (IPP) [5], [6].

As a natural extension of such single-agent IPP, the *Multi-Agent Informative Path Planning* (MAIPP) aims to further improve efficiency through distributed cooperation. There, multiple agents collect information individually but share their measurements in real-time to construct a *global agents’ belief*, which estimates the true, underlying interest map from a finite set of discrete measurements. To achieve good cooperation in MAIPP, agents are expected to make *non-myopic* decisions that balance the trade-off between distribut-

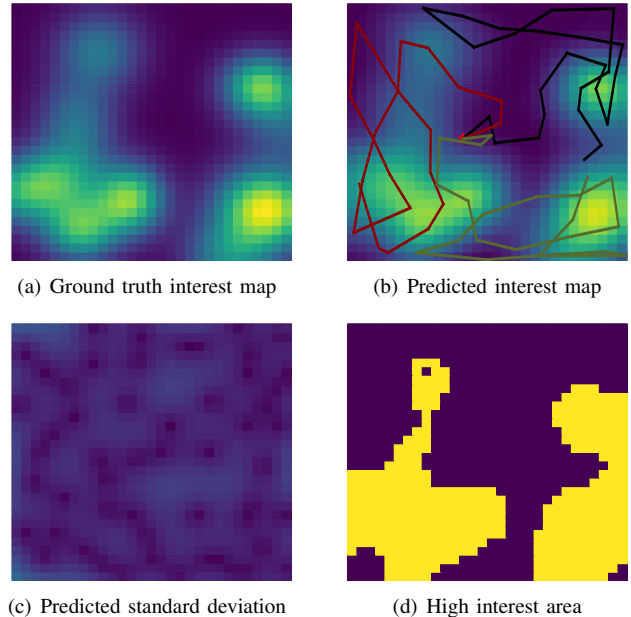


Fig. 1. Example MAIPP instance with 3 agents. The black, green, and red lines represent the trajectories of these 3 agents, seen spreading out initially to distributedly cover the domain, and then regrouping (in particular, the black and green agents) to cover the lower right high-interest area into more details. Note that agents start without any prior knowledge over the ground truth.

ing to explore the environment and regrouping to exploit a known, high-interest area meticulously. For example, Fig. 1 illustrates an episode of an MAIPP instance, to visualize this trade-off. Starting from the same position, the three agents start by distributing to cover the environment; later, the black and green agents regroup to collectively cover the lower-right high-interest area in detail.

Cooperation in MAIPP can be achieved through sequential greedy assignment (SGA) [7]. In SGA, agents plan their paths sequentially, considering the paths planned by all other agents before them. This allows agents to be aware of others’ future actions. As the interest map is updated online, agents often plan paths in a receding-horizon manner to remain responsive to map updates. In scenarios where the interest map remains stable, agents can replan long-horizon paths based on existing plans from other agents, leading to efficient long-term cooperation. However, when the interest map changes frequently and drastically, shared paths can become volatile and confuse agents.

To address this issue, we propose a framework that utilizes an attention-based model trained with distributed deep reinforcement learning (DRL). Each agent plans its path based on the global belief and the *intent* of other agents, represented

[†] Corresponding author, to whom correspondence should be addressed

¹Authors are with the Department of Mechanical Engineering, College of Design and Engineering, National University of Singapore. yangtianze@u.nus.edu, caoyuhong@u.nus.edu, mpegas@nus.edu.sg

as a spatial distribution of their predicted medium-/long-term positions. This framework facilitates efficient cooperation in MAIPP: at each decision step, agents use their learned policy and knowledge of other agents' intents to plan multiple fixed-horizon trajectories and construct a spatial *intent distribution*. This distribution is then shared with all agents at low communication cost. Each agent combines the intent distributions of other agents into a single map, which serves as an additional input to our model, along with the agent's individual belief. The intent distribution allows agents to learn to claim/avoid areas based on other agents' ongoing strategies. Furthermore, our attention-based network enables agents to recognize the useful portions of imperfect intent, unlike conventional methods that assume perfect intent and struggle with frequent plan changes.

Through numerical experiments, we demonstrate the effectiveness of our approach compared to two SGA-based baselines. We highlight the importance and performance benefits of each technique in our framework. Additionally, we evaluate our models (trained under global communications) under limited communication ranges, simulating practical deployments (where global communication may not be a practical/feasible assumption). Remarkably, our learning-based approach maintains performance without requiring retraining or fine-tuning, showcasing its potential for deployments under realistic communication constraints.

II. RELATED WORKS

Single-agent IPP has drawn the community's attention for more than ten years [8], [9]. Recent works have mainly focused on adaptive IPP, where the agent starts with an empty initial belief and where frequent replanning is necessary as this belief is updated online with new measurements. There, the real-time requirements drastically lead state-of-the-art IPP planners to focus on cutting down planning times through meta-heuristic and deep reinforcement learning [10], [11], [12], [13]. Notably, our recent work [12] recently proposed a DRL-based single-agent IPP approach that learns to utilize the full agent's belief to make reactive, yet non-myopic decisions. Multi-agent IPP has been relatively less studied so far. Most relevant, Viseras et al. [14] proposed to approach multi-agent IPP by combining a greedy strategy and collision avoidance. However, lessons from single-agent IPP planners have shown that greedy strategies will often lead to myopic decisions and thus poorer long-term efficiency [15], [10], [16], [11]. This issue is further aggravated in a cooperative multi-agent system. A promising way to extend the recent successes in single-agent IPP to multi-agent IPP (in particular, in a distributed manner) is through sequential greedy assignment (SGA) [7], which has been designed and validated in other information-gathering scenarios. There, each agent sequentially plans its own individual path, conditioned on all paths planned by other agents before itself, thus leading to scalable, priority-based cooperation. SGA can essentially be implemented on top of any conventional single-agent IPP planner, and the final

performance will ultimately depend on the choice of this underlying planner [7].

III. BACKGROUND

Following [8], [14], [17], [10], [12], in this work, we rely on Gaussian Processes (GPs) to construct and update the agents' beliefs and formulate adaptive MAIPP based on GPs.

A. Gaussian Process

In IPP, the underlying distribution of *interest* over the 2D environment $\mathcal{E} \subset \mathbb{R}^2$ is typically modeled as a continuous function $\zeta : \mathcal{E} \rightarrow \mathbb{R}$ (and is not available to the agent(s)). Gaussian Processes offer us a natural means to interpolate between discrete measurements, in order to build an approximation $\mathcal{GP}(\mu, P) \approx \zeta$ of this ground truth. Given a set of n' locations $\mathcal{X}^* \subset \mathcal{E}$ where the interest is to be inferred, a set of n measured locations $\mathcal{X} \subset \mathcal{E}$ and the sampled measurement set \mathcal{Y} , the mean and covariance of the GP are obtained as:

$$\begin{aligned} \mu &= \mu(\mathcal{X}^*) + K(\mathcal{X}^*, \mathcal{X})[K(\mathcal{X}, \mathcal{X}) + \sigma_n^2 I]^{-1}(\mathcal{Y} - \mu(\mathcal{X})), \\ P &= K(\mathcal{X}^*, \mathcal{X}^*) - K(\mathcal{X}^*, \mathcal{X})[K(\mathcal{X}, \mathcal{X}) + \sigma_n^2 I]^{-1} \times K(\mathcal{X}^*, \mathcal{X})^T \end{aligned}$$

where $K(\cdot)$ represents a pre-defined kernel function, σ_n^2 is a parameter representing the measurement noise, and I is a $n \times n$ identity matrix. Specifically, our work uses the Matérn 3/2 kernel following [11], [18].

B. Multi-Agent Adaptive Informative Path Planning

A solution to an MAIPP instance is a set of agent trajectories $\psi = \{\psi^1, \dots, \psi^m\}$. MAIPP seeks an optimal trajectory set ψ^* from the set of all available trajectories Ψ , which maximizes the information gain from all measurements obtained from all agent trajectories while satisfying some budget constraint for each agent:

$$\psi^* = \operatorname{argmax}_{\psi \in \Psi} \sum_{i=1}^m I(\psi^i), \text{ s.t. } C(\psi^i) \leq B, 1 \leq i \leq m, \quad (1)$$

where I denotes the information gained by agent i , $C(\psi^i)$ the trajectory length of agent i , B the individual budget (i.e., max trajectory length) for each agent (here, assumed identical for all agents, but this assumption can be lifted).

In line with recent works [11], [12], in this work, we define the information gain as the uncertainty reduction over the high-interest areas: $I(\psi^i) = \operatorname{Tr}(P_I^-) - \operatorname{Tr}(P_I^+)$, where $\operatorname{Tr}(\cdot)$ denotes the trace of a matrix, and P_I^- and P_I^+ are the prior and posterior covariance of the high-interest areas. These high-interest areas are defined by an *upper confidence bound*: $\mathcal{X}_I = \{x_i \in \mathcal{X}^* \mid \mu_i^- + \beta P_{i,i}^- \geq \mu_{th}\}$, where μ_i^- and $P_{i,i}^-$ are the mean and variance of the GP at the location x_i , and $\mu_{th}, \beta \in \mathbb{R}^+$ control the threshold and confidence interval, respectively. That is, by replacing \mathcal{X}^* with \mathcal{X}_I when calculating the covariance, the information gain in the objective function Eq.(1) is restricted to the high-interest areas predicted by the GP. Consequently, the planner needs to be reactive and adaptive to the change of predicted high-interest areas, which are updated along with new measurements.

IV. METHOD

A. MAIPP as a Sequential Decision-making problem

We formulate MAIPP as a sequential decision-making problem on graphs to decrease the complexity of a continuous environment. Each agent i maintains its individual *waypoint graph* G^i , which is constructed by a probabilistic roadmap (PRM) [19] at the beginning of the mission and covers the whole search domain. Specifically, n nodes are randomly sampled and connected to their k nearest neighbors, such that $G^i = (V^i, E^i), i \in \{1, \dots, m\}$, where $V^i = \{v_1^i, \dots, v_n^i\}$ is a set of nodes and E^i a set of edges. We do not assume the presence of any obstacles over the domain, but note that adding obstacles would be easy, as it would only require pruning any edge that collides with an obstacle in each agent's graph (similar to [20]).

All agents start from the same node at the beginning of the mission. At each of its decision step (i.e., upon reaching its next waypoint), an agent selects one of its neighboring nodes as its next waypoint according to its individual policy p and immediately sets course for it along a straight line and at a constant speed. Note that this parallel decision-making is done asynchronously, since agents typically reach their next waypoints at non-synchronous times during the mission. Agents continue constructing their path iteratively until their budget is exhausted, constructing the final trajectory $\psi^i = \{\psi_1^i, \dots, \psi_c^i\}, \forall \psi_n^i \in V^i, (\psi_j^i, \psi_{j+1}^i) \in E^i$, with trajectory length $C(\psi_c^i) = \sum_{j=1}^{n_i-1} L_2(\psi_j^i, \psi_{j+1}^i)$, where $L_2(\cdot, \cdot)$ denotes the Euclidean distance between two nodes. As a result, in our work, the MAIPP problem can be seen as a set of cooperative, sequential decision-making problems on a set of parallel graphs $G = \{G^1, G^2, \dots, G^m\}$.

B. Agent's Intent

We term the probabilistic distribution of a given agent's future positions as the *intent* of that agent. Each agent will often update and share its most recent intent and receive that of all other agents, which it will fuse into an accumulated intent map to use in its own decision-making. In our approach, each agent generates its own intent by sampling trajectories from its individual policy, which is conditioned on other agents' most recent intent. At each decision step, agent i plans multiple trajectories using its policy p . During this predictive sampling process, agent i virtually updates its position and the covariance matrix output by the Gaussian process (GP), assuming other agents are stationary. From these sampled predicted trajectories, agent i fits a Gaussian distribution $GD(\mu^i(t), \Sigma^i(t))$ over some (or all, see Sect. V for more details) of the nodes from these sampled predicted trajectories, with mean and covariance $\mu^i(t)$ and $\Sigma^i(t)$. We also set a minimum bound for the resulting covariance matrix to avoid singularities.

Agent i updates its own interest map after reaching each node along its path. It receives updated intent maps from the other $m - 1$ agents and fuses them into an accumulated intent distribution by summing the corresponding $m - 1$

Gaussian distributions and normalizing the result. Agent i uses its own belief and accumulated intent distribution to sample future paths, update its intent distribution, and select its next action. It then broadcasts its new $\mu^i(t)$ and $\Sigma^i(t)$ (describing its updated intent distribution) to all other agents. It is important to note that agent i only needs to broadcast a vector $[\mu^i(t), \Sigma^i(t)]$ to other agents, and the Gaussian distribution $GD(\mu^i(t), \Sigma^i(t))$ is constructed by other agents when they need to make a decision. This low communication cost allows for scalability and real-life implementation.

C. RL Cast

Observation: The observation for agent i at decision step t is $s_t^i = \{G^i, v_c^i, B_c^i, \psi^i\}$, which consists of the combined graph and the current state. In MAIPP, the combined graph $G^i = (V^i, E^i)$ integrates information from the agent's own waypoint graph, the current agent's belief predicted by the GP, and the intent of other $m - 1$ agents at each discrete node location: specifically, each node $v_j^i = (x_j^i, y_j^i, \mu(v_j^i), P(v_j^i), f(v_j^i)) \in V^i$. The current state of the agent is $\{v_c^i, B_c^i, \psi^i\}$, where $v_c^i = (x_c^i, y_c^i) \in V^i$ is the current position of the agent i , $\psi^i = (\psi_1^i, \dots, \psi_c^i)$ is the agent's executed trajectory, and $B_c^i = B - C(\psi^i)$ its remaining budget.

Action: When agent i arrives at its most recently selected node, its belief is updated according to all measurements taken by all m agents since agent i made its last decision (i.e., at its previous node). The intent is also updated based on all other agents' future predictions made since then. Then, based on this updated information, the agent's attention-based neural network outputs a policy over its neighboring nodes. Finally, agent i selects an action, i.e., the next node to visit from its policy, and starts moving in a straight line to that next waypoint.

Reward: At decision step t , agent i is given a positive reward according to its most recent action, proportional to the amount of uncertainty reduction: $r_t = (\text{Tr}(P^{t-1}) - \text{Tr}(P^t)) / \text{Tr}(P^{t-1})$, $\text{Tr}(P^{t-1})$ is used to scale the reward. Note that other agents may have also made decisions at step $t - 1$ that can have influenced this uncertainty reduction. Second, a final negative reward $r_f = -\beta \cdot \text{Tr}(P^f)$ is assigned to agent i 's when its budget is exhausted, where P^f is the remaining covariance in high-interest areas, and β is a scaling factor (0.02 in practice).

D. Neural Network

Inspired by the success of [12] in single agent IPP, we rely on attention layers [21] as the fundamental modules of our neural network for their powerful ability to capture dependencies between high-interest areas at multiple scales. Our network consists of two parts: an encoder that learns to reason about the relationships over the agent's belief and others' intents between nodes in the waypoint graph, and a decoder that outputs a policy over the next node to visit based on the encoder's output, the agent position, high-interest threshold, remaining budget, and the executed trajectory so far.

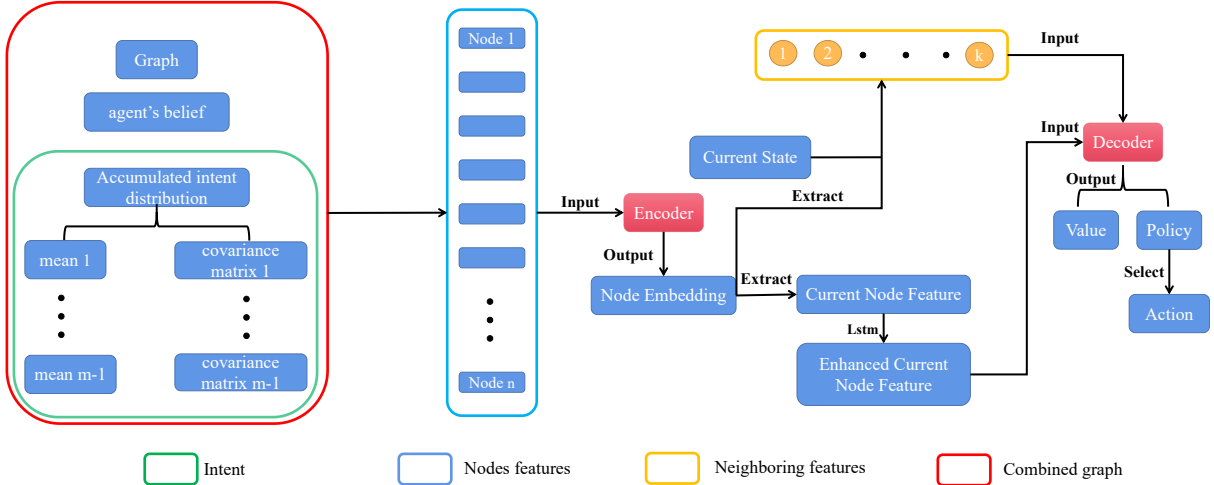


Fig. 2. **Attention-based neural network structure.** The (augmented) nodes of the combined graph (circled in red) are the first input to the network, and contain information about: 1) their location, 2) the agents’ belief (predicted by the GP) at their location, and 3) the intent level, obtained from the accumulated intent distribution (received from the other $m - 1$ agents) at their location. Relying on the self-attention mechanism in the encoder, each node reasons about the dependencies between all other nodes and itself (i.e., learning the relationship between each node position, the global belief, and the intent distribution). After that, the decoder takes the output of this encoder (the *nodes embedding*), and combines it with the current state (agent’s position, remaining budget, and trajectory executed so far), to finally select one of the neighboring nodes as the next position to visit.

Attention Layer: The input of an attention layer consists of a query source h^q and a key-and-value source $h^{k,v}$. The main principle for this general (cross-)attention mechanism is: $q_i = W^Q h_i^q$, $k_i = W^K h_i^{k,v}$, $v_i = W^V h_i^{k,v}$, $u_{ij} = \frac{q_i^T \cdot k_j}{\sqrt{d}}$, $a_{ij} = \frac{e^{u_{ij}}}{\sum_{j=1}^n e^{u_{ij}}}$, $h'_i = \sum_{j=1}^n a_{ij} v_j$, Where all weight matrices are learnable.

Encoder In the encoder, we first project each node in V^i into a higher-dimensional (128 in practice) feature vector, and add a graph positional embedding to each node feature as in [22] to represent edge connectivity. We then input these high-dimensional node features as both the query source and key-and-value source for the following attention layer, which is trained to let each node adaptively merge other nodes’ features into its own feature vector. We term the final encoder’s output the *Node Embeddings*.

Decoder In the decoder, we first extract the node embeddings of the robot’s current node (termed *current features*) as well as its neighboring node embeddings (termed *neighboring features*). We then enhance the current features by adding information about the high-interest threshold, remaining budget, and executed trajectory so far. Specifically, we concatenate the current feature with the high-interest threshold and remaining budget, then project it back to its previous dimension (128). This updated current feature is then passed through an LSTM [23] layer to yield the *enhanced current node features*, where the hidden and cell states are input from the previous enhanced current feature along the executed trajectory. Together with the neighboring features as the Key, we input the enhanced current node features as the Query to a pointer layer [24], a special attention layer that directly uses normalized attention weights as output. We regard these attention weights as our final policy π , a probability distribution over neighboring nodes whose dimensions automatically depend on the number of neighboring nodes.

E. Training

We train our model using PPO [25]. At the beginning of each training episode, we randomly generate 8 to 12 2D Gaussian distributions in the $[0, 1]^2$ space, to construct the ground truth that agents will explore. The resolution of the ground truth and the agent’s belief is 30×30 , and thus the covariance matrix trace (i.e., the uncertainty over the agent’s belief, initially empty), starts at 900. The start position of all agents is randomly drawn within $[0, 1]^2$. A PRM randomly samples 200 nodes to build a graph for each agent, in which we connect each node to its nearest $k = 20$ neighbors. We train with $m = 3$ agents under a budget of $B = 3$ for each agent. An agent takes a new measurement every time it has traveled 0.2 from its previous measurement, and an episode terminates when all agents exhaust their budget. We use the Adam optimizer with a learning rate 5×10^{-5} , which decays by a factor 0.96 every 32 steps, and train with batch size 512. PPO updates the neural network for 8 iterations at each training step. Our model is trained using one i9-10980XE CPU and one NVIDIA GeForce RTX 3090 GPU, and needs about 20 hours to fully converge. Our code is available at <https://github.com/marmotlab/MAIPP>.

V. EXPERIMENTS

In our experiments, all agents also share the same start position v_1 and the same budget B . We further investigate two methods to let each agent i aggregate its set of sampled predicted trajectories $\{g^i(t)\}$ into an intent distribution:

- *destination intent* (DI), where the intent distribution is built by fitting a Gaussian distribution (GD) over the set of all *end nodes* of the sampled trajectories only, and
- *trajectory intent* (TI), where the intent distribution is generated by fitting a GD over all nodes along all sampled trajectories. See Fig. 3 for a visual comparison.

A. Results

We compare our method with two sequential greedy assignment (SGA) based baselines: one using an underlying RRT [15] based single-agent IPP planner, while the other relies on the state-of-the-art single-agent DRL-based planner CATNIPP [12] (greedy variant, using the purely reactive model). In SGA+RRT, agents plan their individual paths sequentially, considering the current belief and paths selected by previous agents. Specifically, each agent samples numerous paths and finally selects the best one to be executed (i.e., the one that minimizes $\text{Tr}(P_f)$). They execute a portion of the best path (0.2, in practice), update their belief, and then replan the next sub-path. SGA+CATNIPP uses a reactive planner (CATNIPP) where the agent only selects the next node to visit at each decision step.

Table I shows the comparison results between our methods and these SGA baselines. We vary the agents’ budget B (same for each agent) and present results for $m = 3$ agents. $\text{RRT}(a, b)$ indicates that agents select the best path from all branches of the generated tree considering path lengths in the range $[a, b]$ (typically more than 8 paths). For our DI/TI variants, (a, j) means agents sample a paths with j nodes to generate their intent distribution. In $\text{TI}(8, 5)^*$, agents select the first step of the best path sampled during intent generation as their next node to visit. We also extend the evaluation to $m = 10$ agents in Table II by reducing the sensor range.

In many scenarios, such as disaster relief operations or ocean exploration, agents may need to accommodate limited communication ranges. We conducted additional experiments with fixed communication ranges of 0.3 and 0.6. Agents acquire beliefs and intent only from agents within the communication range during decision-making. The results are summarized in Table IV.

B. Discussion

From our results in Table I and Table II, we first observe that our variants outperform both SGA+RRT and SGA+CATNIPP across all experiments with varying budgets. Second, we also observe that SGA+CATNIPP consistently outperforms SGA+RRT. We believe that the reason behind these performance differences is that the learning-based approaches (SGA+CATNIPP and our variants) allow the agent to sequence reactive, non-myopic decisions based on its entire belief, whereas RRT may likely remain short-sighted since it can only rely on a finite sampling horizon. Second, we note that SGA+CATNIPP is equivalent to our method without the additional intent input, thus making this comparison an ablation study as well. There, the fact that any of our variants outperforms SGA+CATNIPP highlights the added performance obtained from allowing agents to learn to cooperate through these medium-/long-term predictions, even when those may be inaccurate (e.g., when longer paths are used to generate individual intent maps).

We then investigated how the performance of SGA+RRT is affected by the planning horizon, and results can be found in Table I. There, we show that, unsurprisingly, SGA+RRT sees a sharp decline in performance as the planning horizon




Fig. 3. **Our two intent generation methods, for the same sampled trajectories (colored lines, here composed of 5 waypoints each) from the agent’s current position (star):** Destination intent (DI) fits a Gaussian distribution over the final node of each trajectory only (colored squares), which can result in a wider Gaussian when these points are more spread out, whereas trajectory intent (TI) relies on all waypoints (colored circles and squares) along all trajectories, often resulting in a more peaked Gaussian since many of these nodes are more clumped together.

increases. That is, as agents share longer paths, which they often do not execute as they replan along the way (likely to very different new paths), this type of path sharing is actually hurting the team performance by introducing noise/confusion in the sampling process. Therefore, the longer the sampling path, the less stable the predicted future intent.

However, when we look at the average performance (final $\text{Tr}(P_f)$, reported in Table III) of these variants, the story is very different. There, for our DI variant, we observe the same correlation between prediction stability and final performance as for the SGA+RRT baseline. That is, DI seems to perform best when agents share shorter-term, more stable predicted intents. However, and most interestingly, our TI variant seems to break free from this trend: its best performance is found when agents share intent maps created using medium-length trajectories (composed of 6 nodes in our case), despite these maps exhibiting lower stability in time. We believe that this confirms our hypothesis, about the attention-based structure of our neural network allowing agents to dynamically learn which portion(s) of the shared intent maps to trust/rely on to maximize their performance, even/especially in cases where these intent may not be fully accurate. That is, agents in our TI variant are able to extract more information from intent maps constructed from medium-length sampled trajectories, which are inherently less stable/accurate, leading to improved overall performance. That being said, we still note that this performance comes back down if the length of the sampled trajectories is further increased to 9 nodes, where the accuracy/stability of the intent map decreases below some critical level, introducing too much noise/confusion in the team and leading to worse final performances.

Table IV shows that traditional methods like RRT-based approaches have high communication requirements. When the communication range is small (e.g., 0.3), their performance significantly drops. However, increasing the range to 0.6 improves their performance, bringing it close to the global communication scenario. This is because, without acquiring the beliefs of other agents, the SGA+RRT method allows agents to select paths solely based on their own

TABLE I

COMPARISON BETWEEN OUR TWO SAMPLING VARIANTS, IN TERMS OF $\text{Tr}(P_f)$, FOR 3 AGENTS (10 TRIALS ON 30 INSTANCES).

Method	Budget 2 for each agent		Budget 3 for each agent		Budget 4 for each agent		Budget 5 for each agent	
	mean	std	mean	std	mean	std	mean	std
SGA+RRT(0.9, 1.0)	178.37	6.66	151.59	3.44	145.15	6.47	135.41	6.74
SGA+RRT(0.6, 0.7)	102.03	3.80	55.95	3.51	45.44	1.80	41.47	3.21
SGA+RRT(0.3, 0.4)	39.42	1.91	16.90	1.10	9.16	0.65	5.46	0.41
SGA+CatNIPP	38.75	3.05	13.39	0.91	5.61	0.34	3.27	0.29
Destination Intent(8,3)	37.43	1.27	11.55	0.89	4.56	0.41	2.82	0.11
Trajectory Intent(8,5)	32.85	2.07	8.25	0.61	3.66	0.27	2.34	0.10
Trajectory Intent(8,5)*	30.30	1.26	7.76	0.34	3.32	0.24	1.99	0.16

TABLE II

PERFORMANCE COMPARISON IN TERMS OF FINAL AVERAGE $\text{Tr}(P_f)$, FOR 10 AGENTS (10 TRIALS ON 30 INSTANCES).

Method	Budget 2 for each agent		Budget 3 for each agent		Budget 4 for each agent	
	mean	std	mean	std	mean	std
different graph						
SGA+RRT	30.50	1.79	16.12	1.33	10.85	0.88
SGA+CatNIPP	27.83	3.45	8.90	1.17	4.51	0.30
DI (8,3)	23.86	3.09	7.14	0.39	3.81	0.44
TI (8,5)*	19.39	2.34	5.87	0.46	3.07	0.34

TABLE III

AVERAGE FINAL $\text{Tr}(P_f)$ FOR DIFFERENT NUMBERS OF SAMPLED WAYPOINTS DURING THE INDIVIDUAL INTENT GENERATION IN DI/TI.

	Budget 3 for each agent		Budget 4 for each agent	
	mean	std	mean	std
	DI (8,3)	11.55	0.89	4.56
DI (8,6)	12.53	0.45	4.81	0.39
DI (8,9)	12.61	0.91	5.24	0.45
TI (8,3)	9.65	0.72	3.94	0.25
TI (8,6)	8.43	0.48	3.67	0.21
TI (8,9)	8.92	0.78	3.91	0.30

belief, resulting in wasted budget from redundant work. Even with a large budget, SGA+RRT performs poorly as agents lack accurate enough information to distribute well and cooperatively cover the domain.

In contrast, learning-based methods such as SGA+CatNIPP and our intent-based method excel at making non-myopic decisions even under limited communication. Although we notice a slight decrease in performance ($\sim 20\%$) between performances under global and 0.6 communication ranges, the learning-based methods do not experience a significant decline like SGA+RRT. This gap between traditional and learning-based approaches is further magnified when the communication range is decreased to 0.3. We believe that these results highlight how learning-based approaches enable agents to identify relevant/accurate portions of the incomplete/inaccurate information available to them, particularly through the use of attention mechanisms, leading to improved decision-making. More importantly, we note that our model trained with global communications keeps promising performance even under limited communication ranges, showcasing its generalizability. We believe that training models specifically under known communication constraints in real-world applications could yield even better performance.

TABLE IV

PERFORMANCE UNDER LIMITED COMMUNICATION RANGES, IN TERMS OF $\text{Tr}(P_f)$, FOR 3 AGENTS (10 TRIALS ON 30 INSTANCES)

Communication Range = 0.3	Budget 2		Budget 3		Budget 4	
	mean	std	mean	std	mean	std
SGA+RRT(0.3,0.4)	99.53	6.33	52.29	3.95	26.82	3.61
SGA+CatNIPP	65.62	4.59	25.35	2.24	11.11	0.80
DI (8,3)	66.59	3.57	24.34	1.99	9.74	0.59
TI (8,5)*	50.95	2.34	16.04	1.82	6.21	0.27
Communication Range = 0.6	mean	std	mean	std	mean	std
SGA+RRT(0.3,0.4)	54.96	6.31	26.75	1.62	17.86	2.17
SGA+CatNIPP	54.02	2.91	20.16	2.54	8.87	0.75
DI(8,3)	51.70	3.72	18.00	1.73	7.61	1.07
TI(8,5)*	40.99	3.07	11.39	0.65	5.12	0.65
Global Communication	mean	std	mean	std	mean	std
TI(8,5)*	30.30	1.26	7.76	0.34	3.32	0.24

VI. CONCLUSION

This paper introduces a decentralized, policy-based deep reinforcement learning model for adaptive MAIPP, where each agent utilizes its own policy network to predict its medium-/long-term behavior, and shares this information with others to enhance cooperative decision-making. We propose two methods to aggregate shared information into individual decision-making and compare their performance. Our approach demonstrates advantages across various budgets and team sizes without any retraining, where it outperforms two SGA-based MAIPP baselines: a conventional single-agent IPP solver, and a state-of-the-art learning-based approach. Our experiments indicate that the improved performance of our approach is likely due to learned attention mechanisms, enabling agents to determine which portions of shared intent they can trust and rely on, even in cases of inaccuracy or rapidly changing information. Additionally, we present results under limited communication ranges, demonstrating that our trained model maintains significant performance improvements in realistic constraint scenarios not encountered during training.

We emphasize that our approach to generating, sharing, and utilizing intent for cooperative decision-making extends beyond MAIPP and can be applied or extended to other multi-agent tasks, such as multi-agent pathfinding, task allocation, or collective robotic construction. Future work will explore these extensions and applications, as well as implement our models on hardware, leveraging the reactive nature of our approach to meet real-time deployment requirements.

ACKNOWLEDGMENT

This work was supported by Temasek Laboratories (TL@NUS) under grants TL/SRP/21/19 and TL/FS/2022/01.

REFERENCES

- [1] F. Mascariich, S. Khattak, C. Papachristos, and K. Alexis, "A multi-modal mapping unit for autonomous exploration and mapping of underground tunnels," in *2018 IEEE aerospace conference*. IEEE, 2018, pp. 1–7.
- [2] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 24–39, 2012.
- [3] J. L. Baxter, E. Burke, J. M. Garibaldi, and M. Norman, "Multi-robot search and rescue: A potential field based approach," *Autonomous robots and agents*, pp. 9–16, 2007.
- [4] Q. Wu, J. Lu, W. Zou, and D. Xu, "Path planning for surface inspection on a robot-based scanning system," in *2015 IEEE international conference on mechatronics and automation (ICMA)*. IEEE, 2015, pp. 2284–2289.
- [5] Z. W. Lim, D. Hsu, and W. S. Lee, "Adaptive informative path planning in metric spaces," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 585–598, 2016.
- [6] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.
- [7] M. Corah and N. Michael, "Efficient online multi-robot exploration via distributed sequential greedy assignment," in *Robotics: Science and Systems*, vol. 13, 2017.
- [8] R. Marchant and F. Ramos, "Bayesian optimisation for informative continuous path planning," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6136–6143.
- [9] J. Binney and G. S. Sukhatme, "Branch and bound for informative path planning," in *2012 IEEE international conference on robotics and automation*. IEEE, 2012, pp. 2147–2154.
- [10] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017.
- [11] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for uav-based terrain monitoring," *Autonomous Robots*, vol. 44, no. 6, pp. 889–911, 2020.
- [12] Y. Cao, Y. Wang, A. Vashisth, H. Fan, and G. A. Sartoretti, "CAAtNIPP: Context-aware attention-based network for informative path planning," in *6th Annual Conference on Robot Learning*, 2022. [Online]. Available: <https://openreview.net/forum?id=cAIIbdNAeNa>
- [13] J. Rückin, L. Jin, and M. Popović, "Adaptive informative path planning using deep reinforcement learning for uav-based active sensing," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4473–4479.
- [14] A. Viseras, T. Wiedemann, C. Manss, L. Magel, J. Mueller, D. Shutin, and L. Merino, "Decentralized multi-agent exploration with online-learning of gaussian processes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4222–4229.
- [15] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [16] S. Arora and S. Scherer, "Randomized algorithm for informative path planning with budget constraints," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4997–5004.
- [17] R. Mishra, M. Chitre, and S. Swarup, "Online informative path planning using sparse gaussian processes," in *2018 OCEANS-MTS/IEEE Kobe Techno-Oceans (OTO)*. IEEE, 2018, pp. 1–5.
- [18] M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake, "Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring," *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 658–685, 2019.
- [19] L. E. K. J.-C. Latombe, "Probabilistic roadmaps for robot path planning," *Practical motion planning in robotics: current approaches and future challenges*, pp. 33–53, 1998.
- [20] Y. Cao, T. Hou, Y. Wang, X. Yi, and G. Sartoretti, "Ariadne: A reinforcement learning approach using attention-based deep networks for exploration," *arXiv preprint arXiv:2301.11575*, 2023.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [22] V. P. Dwivedi and X. Bresson, "A generalization of transformer networks to graphs," *arXiv preprint arXiv:2012.09699*, 2020.
- [23] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [24] O. Vinyals, M. Fortunato, and N. Jaitly, "Pointer networks," *arXiv preprint arXiv:1506.03134*, 2015.
- [25] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *CoRR*, vol. abs/1707.06347, 2017. [Online]. Available: <http://arxiv.org/abs/1707.06347>