AFRL-AFOSR-VA-TR-2024-0214

Information-geometric path planning

Tanaka, Takashi
UNIVERSITY OF TEXAS AT AUSTIN
110 INNER CAMPUS DR
AUSTIN, TX, 78712
USA

04/30/2024
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	**2. REPORT TYPE**	**3. DATES COVERED**	
20240430	Final	**START DATE** 20200901	**END DATE** 20241231

4. TITLE AND SUBTITLE
Information-geometric path planning

5a. CONTRACT NUMBER	**5b. GRANT NUMBER**	**5c. PROGRAM ELEMENT NUMBER**
	FA9550-20-1-0101	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Takashi Tanaka

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
UNIVERSITY OF TEXAS AT AUSTIN
110 INNER CAMPUS DR
AUSTIN, TX 78712
USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**	**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**
Air Force Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203	AFRL/AFOSR RTA2	AFRL-AFOSR-VA-TR-2024-0214

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The purpose of this proposal is to introduce the new concept of information-geometric path planning (IGPP), where we seek the shortest path in the configuration space, not in terms of the conventional Euclidean distance functions, but in terms of more general distance functions that can quantify the path complexity in an information-theoretic sense. The proposed development is motivated by realistic motion planning tasks in probabilistic environments (e.g., autonomous navigation of UAVs) in which simple and long paths are sometimes preferred to short and complex paths in view of safety and reliability. Inspired by Sim's notion of rational inattention, the IGPP framework incorporates a distance function that is proportional to the information gain required to follow the given path. We equip the IGPP with highly efficient algorithms for path search, path smoothing and path following, all leveraged by the recent advancements on information gain maximization algorithms in the networked control systems theory literature. This project contains three thrusts. Thrust 1 develops mathematical bases for IGPP, including its characterization using the Finsler manifold theory. Thrust 2 focuses on algorithmic developments, including the modifications of the standard path planning algorithms such as RRT and A-star. Thrust 3 assesses how IGPP can advance the frontiers of the current path-planning technologies. In particular, we consider (i) path planning under sensing resource constraints, (ii) decentralized path planning, and (iii) path planning mimicking human experts. The expected outcomes of this project include a new theoretical framework, algorithms, and insights that add a new dimension to the existing autonomous motion planning technologies. The added dimension will help overcome the existing challenges in unmanned/manned Air Force platforms and enhance their autonomous capabilities.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	UU	13
U	U	U		

19a. NAME OF RESPONSIBLE PERSON	**19b. PHONE NUMBER** *(Include area code)*
FREDERICK LEVE	696-9730

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

Final Report

Award Number:
AFOSR FA9550-20-1-0101

Report Type:
Final Report

Reporting Periods:
09/01/2020 – 12/31/2023

Distribution Statement:
Ex: Distribution A – Approved For Public Release
Distribution B – Not Approved For Public Release

Program Officer Name:
Dr. Frederick A. Leve
(Dynamical Systems and Control Theory)

Project Title:
Information-Geometric Path Planning: Perception-based Motion Planning Algorithms for Highly Autonomous Systems

Principal Investigator:
Takashi Tanaka
Department of Aerospace Engineering and Engineering Mechanics
University of Texas at Austin
Email: ttanaka@utexas.edu
Address: 2617 Wichita Street, Austin, Texas 78712-1221

1 Abstract

The purpose of this proposal is to introduce the new concept of information-geometric path planning (IGPP), where we seek the shortest path in the configuration space, not in terms of the conventional Euclidean distance functions, but in terms of more general distance functions that can quantify the path complexity in an information-theoretic sense. The proposed development is motivated by realistic motion planning tasks in probabilistic environments (e.g., autonomous navigation of UAVs) in which simple and long paths are sometimes preferred to short and complex paths in view of safety and reliability. Inspired by Sim's notion of rational inattention, the IGPP framework incorporates a distance function that is proportional to the information gain required to follow the given path. We equip the IGPP with highly efficient algorithms for path search, path smoothing and path following, all leveraged by the recent advancements on information gain maximization algorithms in the networked control systems theory literature.

This project contains three thrusts. Thrust 1 develops mathematical bases for IGPP, including its characterization using the Finsler manifold theory. Thrust 2 focuses on algorithmic developments, including the modifications of the standard path planning algorithms such as RRT and A-star. Thrust 3 assesses how IGPP can advance the frontiers of the current path-planning technologies. In particular, we consider (i) path planning under sensing resource constraints, (ii) decentralized path planning, and (iii) path planning mimicking human experts. The expected outcomes of this project include a new theoretical framework, algorithms, and insights that add a new dimension to the existing autonomous motion planning technologies. The added dimension will help overcome the existing challenges in unmanned/manned Air Force platforms and enhance their autonomous capabilities.

2 Research Objectives

The purpose of this research is to introduce the new concept of information-geometric path planning (IGPP), where we seek the shortest path in the configuration space, not in terms of the conventional Euclidean distance functionals, but in terms of more general distance functionals that can quantify the path complexity in an information-theoretic sense. The new framework has a built-in capability to penalize both path length and path complexities and thus serves as an ideal tool to balance these two design objectives. The proposed development is motivated by path planning tasks in stochastic environments (e.g., autonomous navigation of UAVs, Figure 1) in which simple and long paths (Path B) are sometimes preferred to short and complex paths (Path A) in view of safety.

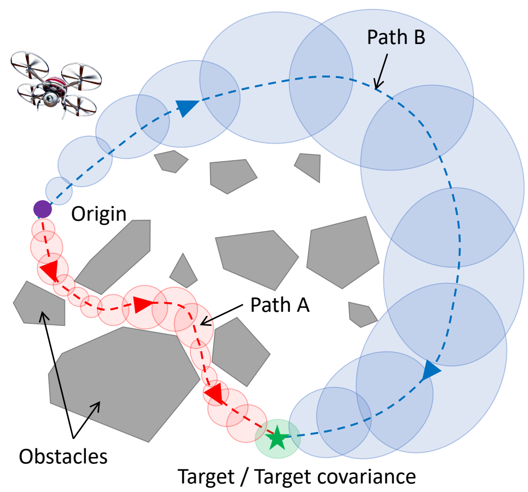

Figure 1: Short and complex path A vs. long and simple path B. Confidence ellipses are shown along each path.

To characterize the "simplicity" concept in mathematical terms, we equip the underlying configuration space with a novel information-geometric distance function. In particular, the distance from the origin to the destination is the weighted sum of the Euclidean distance between these two points and the information gain required to make this transition. Here, the information gain is the amount of sensor data (measured in bits) that the decision-maker has to collect to follow the path. In this standard, a path is considered "simple" if navigation through it does not require many sensing actions. A path that is traceable by an open-loop control policy is thus the simplest. This idea of simplicity is related to Sim's notion of rational inattention for bounded rationality [1]. In general, the introduced distance function is non-Euclidean and asymmetric.

The shortest path problem we formulate is motivated by the increasing need for simultaneous perception and action planning in modern information-rich autonomy. Owing to the wide availability of low-cost and high-performance sensing devices, obtaining a large amount of sensor data has become easier in many applications. Nevertheless, operating a sensor at its full capacity may not be the best strategy for resource-constrained robots, especially if it drains the robot's scarce power or computational resources with little benefit. As sensor modalities increase, how to achieve a given task with minimum perceptual resources (e.g., with reduced sensing frequencies or sensor gains) becomes an increasingly relevant question. For instance, planetary rovers need to estimate wheel slippage using visual odometry (VO) when traversing harsh and unknown terrains [2]. However, using VO reduces the navigation speed as the rover needs to stop frequently to capture images [3] and drive slowly due to its limited computational capability [4]. The Mars Science Laboratory rover reaches a maximum speed of 140 m/h in blind-drive mode (no VO update) and 45 m/h in hazard avoidance mode (VO update every 10 m) [5]. Another example is the vision-based navigation of micro aerial vehicles (MAV), where the visual data are sent to a ground station (e.g., see [6]) or are processed by computationally constrained onboard processors (e.g., see [7]). In the navigation of MAVs, the available resources (computation speed, memory, power, and communication bandwidth) are limited and require specific consideration. These examples showcase the importance of perception effort management in autonomous navigation.

3 Results

The research outcomes of our three-year-long investigations are summarized in research papers (see the "Publication" section below). From a methodological perspective, it is convenient to categorize our results into two classes. In the first category, our research focus is centered around the concept of Gaussian belief space planning. While this approach is applicable to continuous-time continuous-space path planning scenarios, it critically relies on the Gaussianity of uncertainty and the robot dynamics must be assumed linear. In the second category, we formulate the minimum sensing planning concept in finite state systems. While Gaussianity/linearity assumptions are not necessary for such a formulation, it is inherently restricted to finite state systems. While it was inevitable for us to take mathematically different methodologies in each category, they are both derived from the same governing principle (motivated by the discussion above), and they will complement each other's shortcomings. Below, we summarize the results of our effort in each category.

3.1 Gaussian belief space planning for minimum sensing navigation.

In the series of research publications [C6], [C7], [C11], [C12], and [J2], we adopted the Gaussian belief space formalism to develop theory and algorithms for minimum sensing navigation. Problem formulation, overview of the developed planning algorithms, and numerical validations are briefly summarized below.

3.1.1 Problem formulation

Let $x(t_k)$ be the random vector representing the robot's actual position at time t_k. In an open-loop control scenario, it is assumed to satisfy

$$x(t_{k+1}) = x(t_k) + (t_{k+1} - t_k)v_k + n_k \tag{1}$$

where v_k is the velocity input and $n_k \sim \mathcal{N}(0, \|x_{k+1} - x_k\|W)$ is a Gaussian disturbance whose covariance matrix is proportional to the commanded travel distance.

In Gaussian belief space planning, a reference trajectory is given as a sequence of belief waypoints $b_k = (x_k, P_k), k = 0, 1, ..., K$, where $x_k \in \mathbb{R}^d$ and $P_k \in \mathbb{S}^d_{++}$ are planned mean and covariance of the random vector $x(t_k)$. In the sequel, we call $\mathbb{B} := \mathbb{R}^d \times \mathbb{S}^d_{++}$ the *Gaussian belief space* or simply the *belief space*. We first introduce an appropriate "distance" function from a belief point $b_k = (x_k, P_k)$ to another $b_{k+1} = (x_{k+1}, P_{k+1})$. The distance function is the cost of steering the Gaussian probability density characterized by b_k to the one characterized by b_{k+1}. We assume that the distance function is a weighted sum of the travel cost $\mathcal{D}_{\text{travel}}(b_k, b_{k+1})$ and the information cost $\mathcal{D}_{\text{info}}(b_k, b_{k+1})$.

- Travel cost: We assume that the travel cost is simply the commanded travel distance:

$$\mathcal{D}_{\text{travel}}(b_k, b_{k+1}) := \|x_{k+1} - x_k\|.$$

- Information cost: Assuming that no sensor measurement is utilized while the deterministic control input v_k is applied to (1), the covariance at time step $k+1$ is computed as

$$\hat{P}_{k+1} := P_k + \|x_{k+1} - x_k\| W. \tag{2}$$

  We refer to $\hat{P}_{k+1}$ as the prior covariance at time step $k+1$. Suppose that the prior covariance is updated to the posterior $P_{k+1} (\preceq \hat{P}_{k+1})$ by a sensor measurement y_{k+1} at time step $k+1$. The minimum information gain required for this transition is given by the entropy reduction:

$$\begin{aligned}\mathcal{D}_{\text{info}}(b_k, b_{k+1}) &= h(x_{k+1}|y_0, \cdots, y_k) - h(x_{k+1}|y_0, \cdots, y_{k+1}) \\ &= \frac{1}{2}\log\det \hat{P}_{k+1} - \frac{1}{2}\log\det P_{k+1}.\end{aligned}$$

  Here, $h(\cdot|\cdot)$ denotes conditional differential entropy. Intuitively, $\mathcal{D}_{\text{info}}(b_k, b_{k+1})$ represents the minimum bits of information required to reduce the uncertainty from $\hat{P}_{k+1}$ to P_{k+1}.

Introducing a weight factor $\alpha \geq 0$, the total cost is defined as

$$\mathcal{D}(b_k, b_{k+1}) := \mathcal{D}_{\text{travel}}(b_k, b_{k+1}) + \alpha \mathcal{D}_{\text{info}}(b_k, b_{k+1}). \tag{3}$$

The total cost function (3) satisfies the following properties:

1. $\mathcal{D}(b_1, b_2) \geq 0 \;\; \forall b_1, b_2 \in \mathbb{B}$;
2. $\mathcal{D}(b, b) = 0 \;\; \forall b \in \mathbb{B}$; and
3. $\mathcal{D}(b_1, b_2) \leq \mathcal{D}(b_1, b_3) + \mathcal{D}(b_3, b_2) \;\; \forall b_1, b_2, b_3 \in \mathbb{B}$.

The function (3) is not symmetric in that $\mathcal{D}(b_1, b_2) \neq \mathcal{D}(b_2, b_1)$ in general. Consequently, the notion of the path length we introduce is direction-dependent. Also, $\mathcal{D}(b_1, b_2) = 0$ does not necessarily imply $b_1 = b_2$. Because of these properties, the function (3) is formally not a distance metric. In the literature, such a function is called a *quasi-pseudometric* or *Lawvere metric*.

Despite (3) not being a metric, it turns out that the function (3) still allows us to formulate a shortest path problem of our interest. First, for a given continuous path $\gamma(t) = (x(t), P(t))$ in the Gaussian belief space and a partition $\mathcal{P} = (0 = t_1 < t_2 < \cdots < t_K = T)$, the path length with respect to $\mathcal{P}$ is defined by $c(\gamma, \mathcal{P}) = \sum_{k=1}^{K} D(\gamma(t_k), \gamma(t_{k+1}))$. Taking a supremum, the length of a path γ is defined by $c(\gamma) = \sup_{\mathcal{P}} c(\gamma, \mathcal{P})$. We consider that a belief path $\gamma(t) = (x(t), P(t))$ is collision-free if

$$(x(t) - x_{\text{obs}})^\top P^{-1}(t)(x(t) - x_{\text{obs}}) \geq \chi^2 \;\; \forall t \in [0, T], \forall x_{\text{obs}} \in \mathcal{X}_{\text{obs}}$$

where χ^2 is a user-defined confidence parameter. With this setup, the shortest belief path problem

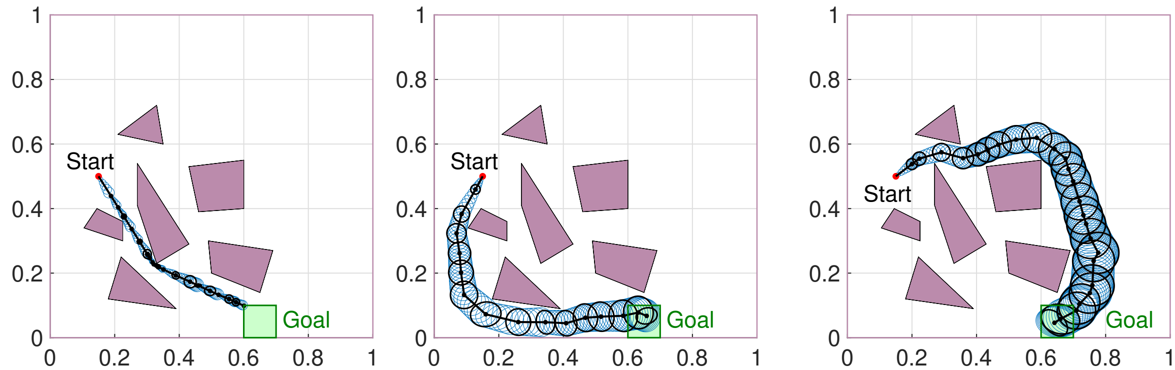

(a) $\alpha = 0$ (b) $\alpha = 0.3$ (c) $\alpha = 0.7$

Figure 2: Outputs of IG-RRT* with $\alpha = 0, 0.3, 0.7$ under the existence of multiple obstacles. Confidence ellipses representing $\Pr = 90\%$ certainty regions.

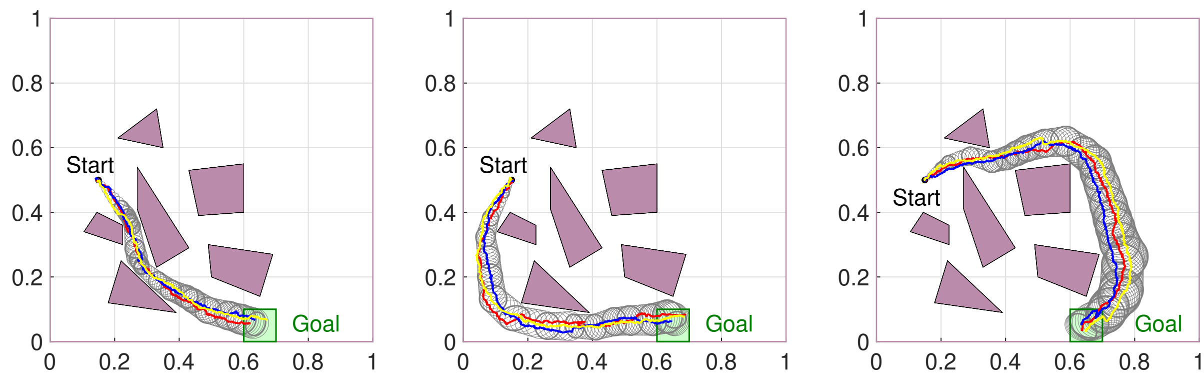

(a) $\alpha = 0.1$ (b) $\alpha = 0.3$ (c) $\alpha = 0.7$

Figure 3: The reference paths generated with $\alpha = 0.1, 0.3$ and 0.7 are followed by an event-based LQG controller using the high-precision sensor with $Z = 10^{-4} I_2$. Three sample trajectories (shown in red, blue, and yellow) are plotted in each case.

is formulated as

$$\begin{aligned} \min_{\gamma} \quad & c(\gamma) \\ \text{s.t.} \quad & \gamma(0) = b_0, \quad \gamma(T) \in \mathcal{B}_{\text{target}} \\ & (x(t) - x_{\text{obs}})^\top P^{-1}(t)(x(t) - x_{\text{obs}}) \geq \chi^2 \;\; \forall t \in [0, T], \forall x_{\text{obs}} \in \mathcal{X}_{\text{obs}}. \end{aligned} \tag{4}$$

Now, one can attempt to solve the shortest path problem (4) using various numerical solvers. In what follows, we summarize our approach based on the RRT* and PRM*-based algorithms.

3.1.2 IG-RRT* and IG-PRM* Algorithms

In [J2], we developed several versions of RRT*-based numerical algorithms dedicated to the Gaussian belief space planning problem (4). We refer to our algorithms as the IG-RRT* algorithm. Details of the algorithm, its variations, and tricks to improve its sampling efficiency can be found in the reference [J2]. In [P1], we developed a PRM*-based belief state planning algorithm, referred to as IG-PRM*.

Figure 2 shows the path plans obtained by IG-RRT* with $\alpha = 0, 0.3, 0.7$ under the existence of multiple obstacles. It demonstrates that the proposed cost function can induce both a short and complex path (like path A in Figure 1) and a simple and long path (like path B in Figure 1) by properly adjusting the tuning parameter $\alpha > 0$.

3.1.3 Effectiveness of IGPP

We also performed extensive numerical studies to demonstrate the practical benefits of the IGPP concept. In Figure 3, we considered the situation in which the traveling agent (noisy double integrator) tries to follow the planned path in Figure 2 using an event-based sensing strategy. Roughly, sensing actions are triggered only when the confidence ellipses computed in real-time by Kalman filtering "escape" the planned path tube computed offline – see [J2] for the details about the triggering rule for the sensing actions. Figure 4 shows the number of required measurements for the robot to travel from the start position to the goal region. It shows that the required frequency of measurements tends to decrease as we increase α. This result indicates that the proposed IGPP is an effective framework for minimum sensing navigation.

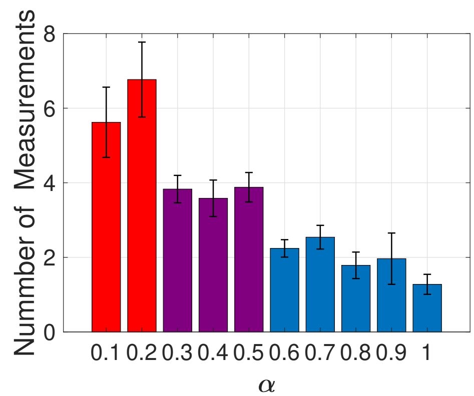

Figure 4: The number of required measurements for a double integrator robot using event-based LQG controller.

In [J2], we also conducted similar numerical experiments using a high-fidelity quadrotor simulator (Figure 5). In the off-line phase, two belief space paths were generated using the IG-RRT* algorithm using two different values of α. In the online (real-time) implementation phase, each path is followed by the event-based sensing strategy. The number of sensing actions (vision-based localization) required to follow the planned path from the starting point to the goal region is monitored for each case. We performed many simulation trials for each case. We found a similar trend to Figure 4 – the nominal path generated by a larger value of α tends to be easier to follow with fewer sensing actions.

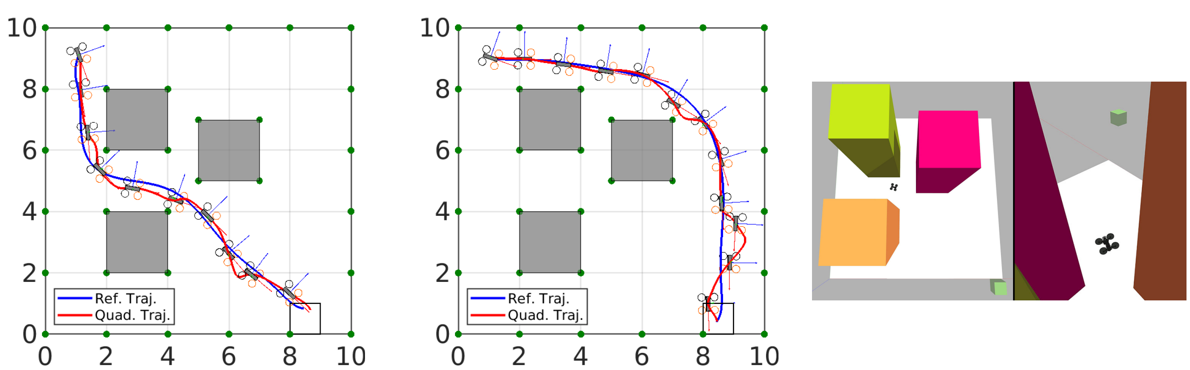

(a) Horizontal position of CM (X-Y plane) of the quadrotor following the path generated by α = 0.2.

(b) Horizontal position of CM (X-Y plane) of the quadrotor following the path generated by α = 2.0.

(c) Snapshots of the quadrotor simulation conducted in the Gazebo simulator. The small green box shows the goal region.

Figure 5: Smoothed reference trajectories (blue), sampled paths (red) of the quadrotor, and snapshots of the simulation in Gazebo.

3.1.4 Asymptotic optimality

A sampling-based path planning algorithm is said to be asymptotically optimal if the length of the best path discovered by the algorithm converges to the length of the optimal path with probability one as the sample size tends to infinity. It is well understood that some sampling-based path planning algorithms, including RRT* and PRM*, attain asymptotic optimality with respect to the Euclidean metric. However, it is not trivial that RRT* and PRM* remain asymptotically optimal even if it is evaluated under the Lawvere metric . Generally speaking, despite the extensive use of sampling-based algorithms and their rigorous analysis in the deterministic setting, there has been little formal analysis of their asymptotic optimality in Gaussian belief space.

We aimed to address this lack of research by examining the asymptotic behavior of the proposed IG-RRT* and IG-PRM* algorithms. Even though we discovered that the asymptotic behavior of the IG-RRT* algorithm was cumbersome to analyze, in [P1], we were able to provide a formal proof of the asymptotic optimality of the IG-PRM* algorithm. While we conjecture that the IG-RRT* algorithm also attains asymptotic optimality, a formal proof must be postponed as future work.

3.2 Belief space planning over a finite state space.

In [J1], we generalized the belief space planning concept for minimum sensing navigation discussed above to a non-Gaussian regime.

3.2.1 Problem formulation

We consider a variation of the standard Markov decision process (MDP) referred to as a *perception MDP*, which consists of a finite set of states S and actions A as well as a countably infinite set of observations Z. Furthermore, the perception MDP additionally consists of an environmental cost function $C : S \times A \to \mathbb{R}$ mapping state-action pairs to an associated cost and a transition function $T : S \times A \times S \to [0, 1]$ mapping state-action pairs to a probability distribution over successor states.

In contrast to the standard partially observable MDP (POMDP), in a perception MDP, the decision-maker is able to synthesize its own perception strategy, represented in terms of a belief-dependent observation function, in addition to its strategy for action selection. We refer to this joint synthesis procedure as the *simultaneous perception-action design* (SPADE) framework. A visualization of the SPADE framework is shown in Figure 6. To capture the cost associated with a given perception strategy, we impose an information-theoretic perception cost that penalizes information flow from the environment to the decision-maker. In this work, we specifically focus on using the *directed information* to model this perception cost, which we subsequently decompose in terms of the sum of the stage-wise mutual information gains. In conjunction with the environmental

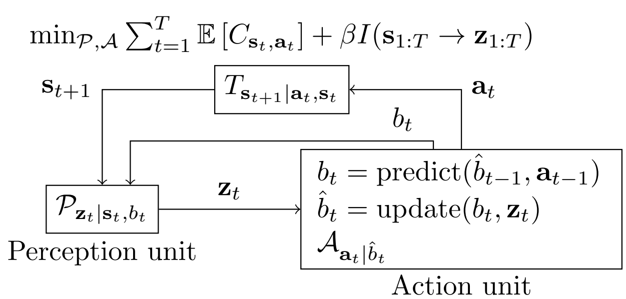

Figure 6: Visualization of the simultaneous perception-action design (SPADE) framework.

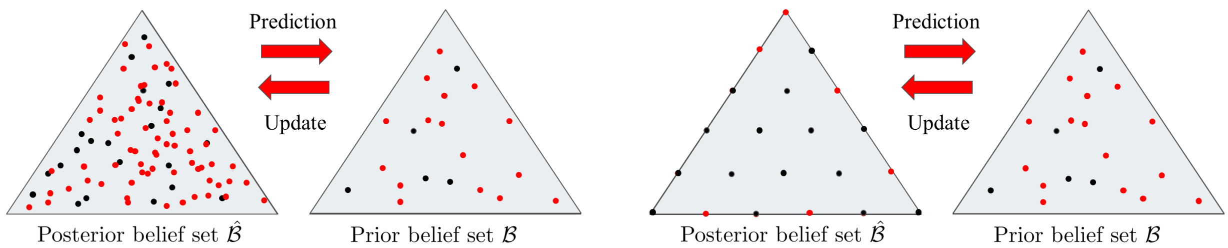

(a) Unconstrained value iteration over belief space. (b) Value iteration over IFBS.

Figure 7: Left: Unconstrained value iteration over belief space. Without imposing any constraints on admissible perception policies, the value iteration does not keep a finite set of belief states invariant. Consequently, one must evaluate the value function on the continuum of the belief states. Right: A finite set of belief states can be made invariant by imposing linear constraints on admissible perception policies. The value iteration can be performed only on a finite set of belief states.

cost function, we can then express the objective of the decision-maker as follows:

$$\min_{\mathcal{A},\mathcal{P}} \sum_{t=0}^{\infty} \gamma^{t-1}(\beta I(\mathbf{s}_t; \mathbf{z}_t | \mathbf{z}_{1:t-1}) + \mathbb{E}[C_{s_t,a_t}]), \tag{5}$$

where $\gamma \in (0, 1]$ is a discount factor, $\beta \in \mathbb{R}_{\geq 0}$ is a relative cost weighting factor, $I(\mathbf{s}_t; \mathbf{z}_t | \mathbf{z}_{1:t-1})$ is the mutual information, and C_{s_t,a_t} is the environmental cost of taking action $a \in A$ while in state $s \in S$ at time step t.

3.2.2 Solution via the invariant finite belief sets (IFBS)

Due to the perception costs, the decision-maker does not generally know its true state at a given time step. As in the continuous state-space version of the problem, the decision-maker instead maintains a *belief state*, an $|S|$-dimensional vector consisting of the probabilities that the decision-maker resides in each state $s \in S$ at time t. We show that, in principle, the SPADE problem (i.e., minimizing (5)) can be solved via a value iteration procedure over the continuous set of belief states. However, this value iteration is intractable. To overcome this difficulty, we propose a novel *method of invariant, finite belief sets* (IFBS), in which the decision-maker operates exclusively on a representative sample of belief states (Figure 7). We provide conditions under which this sample of belief states remains invariant, and present a computationally-efficient algorithm for solving the SPADE problem over this representative sample of belief states. Furthermore, we provide theoretical results proving that, as the resolution of the sampling increases, the resulting solution approaches that of the SPADE problem over the continuous belief space.

3.2.3 Application to dynamic sensor selection

One particularly noteworthy application of the proposed SPADE framework is in the problem of strategic sensor selection. In this problem, due to sensing constraints such as limited communication bandwidth or processing capability, only a subset of a set of sensors can be used for observation at a given time step. The objective is then to dynamically choose this subset of sensors such that the cumulative sensing error is minimized. Although this problem can naturally be expressed in terms of a POMDP, its solution quickly becomes intractable. In our work, we alternatively use the perception strategy obtained through solving the SPADE problem (which makes no assumption on the structure of the set of sensors) as a reference observation function. The decision-maker then chooses the subset of sensors whose combined observation function most closely matches that of

Table 1: Comparison of the PBVI-based approach and our algorithm's offline solution times, online solution times, cumulative MAP estimation error, and average number of chosen sensors.

Method	Offline [s]	Online [s]	Error	$\lvert\mathcal{S}^t\rvert$
PBVI, $k = 1$	24.95	8.1×10^{-4}	23.6191	1.00
PBVI, $k = 2$	8.6×10^{2}	6.0×10^{-4}	15.84	2.00
PBVI, $k = 3$	1.5×10^{4}	7.1×10^{-4}	11.4284	3.00
PBVI, $k = 4$	TO	–	–	–
$\beta = 0.3$	39.75	0.98	5.51	7.96
$\beta = 0.4$	42.73	0.95	5.96	7.19
$\beta = 0.5$	46.44	1.05	10.62	4.91
$\beta = 0.6$	50.43	1.15	16.33	3.18
$\beta = 0.7$	50.37	1.14	21.57	2.12
$\beta = 0.8$	54.13	1.13	25.23	1.35

this reference value at the given time step. Figure 8 shows the resulting sensor selection strategies for an 8-state environment with a total of 8 sensors, wherein one sensor is centered on each of the states. Intuitively, as the relative cost of information increases, the number of selected sensors at each time step decreases. Furthermore, as shown in Table 1, although requiring slightly more online computation time, our proposed sensor selection strategy is able to scale to sensor selection problems with greater allowed cardinalities at only a modest performance loss. Note that "PBVI" refers to the point-based value iteration algorithm, a common algorithm for solving for an optimal strategy in a POMDP.

4 Task achievements

The overall achievement levels of the proposed tasks are summarized below. Most of the tasks were executed as planned. Tasks 3-2 and 3-3 were left largely unexplored, which must be postponed as future work.

Thrust 1: Mathematical foundation: We provide IGPP with a rigorous mathematical foundation based on information geometry and Finsler manifold theory.

- Task 1-1: Finsler manifolds and geodesics. (100% complete)
- Task 1-2: Non-holonomic constraints. (100% complete)
- Task 1-3: Partially observable configurations. (80% complete)

Thrust 2: Algorithmic foundation: We develop a software package for IGPP, containing the shortest path search, path-smoothing and path-following (feedback control) algorithms.

- Task 2-1: Path smoothing by SDP/ADMM algorithm. (100% complete)
- Task 2-2: Path smoothing by reverse water-filling algorithm. (80% complete)
- Task 2-3: Rationally inattentive path following. (80% complete)
- Task 2-4: Path following by Feynman-Kac sampling. (100% complete)

Thrust 3: Path planning frontiers: We use the IGPP framework to tackle well-recognized open challenges in path planning research.

- Task 3-1: Simultaneous perception and motion planning. (100% complete)
- Task 3-2: Distributed path planning by multiple agents. (30% complete)
- Task 3-3: Human-aware path planning. (10% complete)

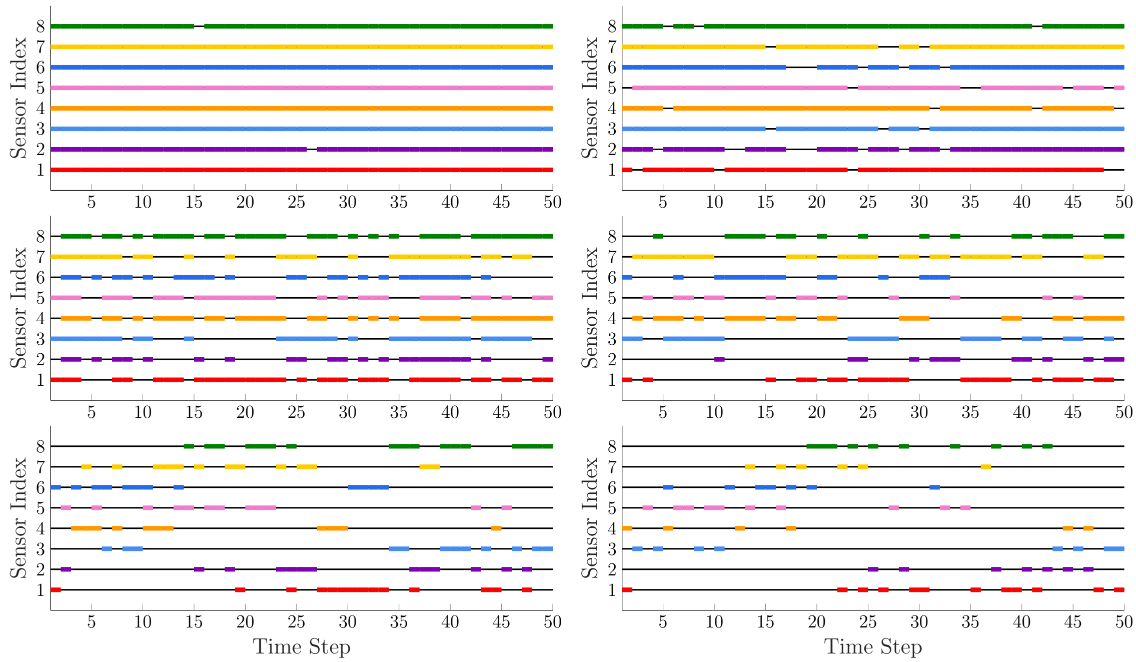

Figure 8: Comparison of the sensor selection at each time step for various values of β. Clockwise from top left are the plots for $\beta = 0.3$, $\beta = 0.4$, $\beta = 0.5$, $\beta = 0.6$, $\beta = 0.7$, and $\beta = 0.8$, respectively.

5 Future work

Besides the main results described above, our three-year-long research effort supported by this YIP project has also identified several new research directions. The research directions summarized below are both natural extensions of this YIP project and are worthwhile for further investigations in the future.

5.1 Path integral control

In motion planning under stochastic uncertainties or disturbances, correct evaluation of the probability of collision with obstacles is crucial. Often in practice, we are more concerned with end-to-end safety (the probability that the robot can travel from the initial configuration to the target configuration without collision) rather than instantaneous safety (the probability that the robot is in a safe state at a specific point in time). Unfortunately, the former is more difficult to compute, especially in continuous-time settings. In [C3, C5, C9, C10], we conducted an in-depth study on chance-constrained motion planning problems. As part of the proposed Task 2-4, we applied the concept of path integral control to chance-constrained motion planning problems [C3, C9, P6].

Path integral control (PIC) [8], a control algorithm inspired by the path integral formulation of quantum mechanics, has been gaining popularity in autonomy research [9]. It provides a Monte-Carlo-based algorithm to numerically compute the optimal control actions in real time. Model predictive path integral (MPPI) [10], a GPU-friendly receding horizon version of PIC, has been widely adopted in various nonlinear and stochastic control scenarios. Despite its growing popularity, the PIC theory is still immature in several aspects. First, the PIC theory for handling chance constraints (e.g., evaluation of rare event probabilities) is immature. Second, its sample complexity is not well understood, and it often requires prohibitive computational costs for real-time applications. There is significant room for improvement in its sample efficiency by adopting "smart" Monte Carlo algorithms, such as multi-level Monte Carlo. Third, currently, PIC theory is largely restricted to control systems with fully observable state spaces. These limitations must be addressed in the future for the PIC framework to be applicable to motion planning problems satisfactorily.

5.2 Real-time path planning using trained neural network models

The IG-RRT* and IG-PRM* algorithms developed in this YIP effort provide us with the desired path planning capability described in the "Research Objective" section above. However, the computational cost for executing these algorithms are often too high for real-time implementations in many application contexts. Therefore, we have explored an idea to use deep learning models that allow for real-time path planning.

First, we prepared a training data set comprising a large number of pairs $(I_{\text{in}}, I_{\text{out}})$ of an input image I_{in} and an output image I_{out}. The input image I_{in} shows a 2D configuration space (e.g., obstacle locations) describing the path planning problem of interest. The output image I_{out} describes a solution to the problem encoded by I_{in} obtained by IG-RRT* or IG-PRM* algorithms. Second, we trained a U-net model until it is able to predict an output image from a given input image. Finally, a trained U-net model is tested on unseen input images to see if it can predict the corresponding path plan. Figure 9 shows some of the output images predicted by the trained U-net and the ground truth solution obtained by IG-RRT* algorithm. They often exhibit a good match. Even though the solution predicted by a trained U-net is currently not fully reliable, we expect that the solution accuracy will improve as we increase the training data size.

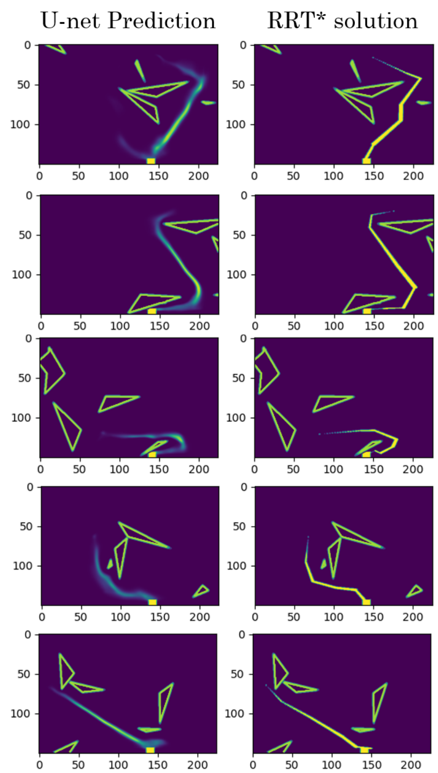

Figure 9: Comparison of the solution predicted by a trained U-net and the solution produced by IG-RRT*.

While the preparation of a training data set and the training of a U-net are time-consuming processes, a forward execution of the trained U-net is instantaneous. Hence, this solution approach is valuable in applications where real-time path planning is necessary. Future research will be concerned with how to improve the reliability of such a deep-learning-based approaches, how to ensure robustness against a large variety of input data, and how to generalize the approach to high-dimensional problems.

6 Acknowledgement

PI Tanaka would like to thank the Dynamical Systems and Control Theory program at AFOSR and Dr. Frederick A. Leve for supporting this project.

7 Products

Journal Publication

- [J1] M. Hibbard, T. Tanaka and U. Topcu, "Simultaneous Perception-Action Design via Invariant Finite Belief Sets," Automatica, vol. 155, pp. 111140, Sep. 2023.
- [J2] A. Pedram, R. Funada and T. Tanaka, "Gaussian Belief Space Path Planning for Minimum Sensing Navigation," IEEE Transactions on Robotics, vol. 39, no. 3, pp. 2040-2059, June 2023.

[J3] Y. Savas, M. Hibbard, B. Wu, T. Tanaka and U. Topcu, "Entropy Maximization for Partially Observable Markov Decision Processes," IEEE Transactions on Automatic Control, v.67, no. 12, (2022): 6948 – 6955

Preprints (Journal Papers Under Review)

[P1] V. Zinage, A. Pedram and T. Tanaka, "Optimal Sampling-based Motion Planning in Gaussian Belief Space for Minimum Sensing Navigation." The manuscript is under preparation for resubmission.

[P2] B. He and T. Tanaka, "Safety Control of Uncertain Lagrangian Systems Using Dynamic Output Feedback Barrier Pairs." IEEE Transactions on Automatic Control (Under review).

[P3] A. Patil, G. A. Hanasusanto, T. Tanaka, "Discrete-Time Stochastic LQR via Path Integral Control and Its Sample Complexity Analysis." The IEEE Control Systems Letters (Under review).

[P4] K. Hoshino, T. Tanaka, Y. Chen, "A Path Integral Algorithm for Partially Observed Control Problems." The IEEE Control Systems Letters (Under review).

[P5] A. Patil, R. Funada, T. Tanaka, L. Sentis, "Task Hierarchical Control via Null-Space Projection and Path Integral Approach." The IEEE Control Systems Letters (Under review).

[P6] A. Patil, A. Duarte, F. Bisetti, T. Tanaka, "Strong Duality and Dual Ascent Approach to Continuous-Time Chance-Constrained Stochastic Optimal Control" The manuscript is under preparation for resubmission.

Conference Papers

[C1] H. Park, D. Zhou, G. A. Hanasusanto, and T. Tanaka, Distributionally Robust Path Integral Control, The 2024 American Control Conference, 2024.

[C2] Z. Wang, R. Keller, X. Deng, K. Hoshino, T. Tanaka, and Y. Nakahira, Physics-Informed Representation and Learning: Control and Risk Quantification, The 38th Annual AAAI Conference on Artificial Intelligence, 2024.

[C3] A. Patil, Y. Zhou, D. Fridovich-Keil, T. Tanaka, Risk-Minimizing Two-Player Zero-Sum Stochastic Differential Game via Path Integral Control, The 62nd IEEE Conference on Decision and Control, 2023.

[C4] A. Patil, M. O. Karabag, T. Tanaka, U. Topcu, Simulator-Driven Deceptive Control via Path Integral Approach, The 62nd IEEE Conference on Decision and Control, 2023.

[C5] A. Patil and T. Tanaka, Upper and Lower Bounds for End-to-End Risks in Stochastic Robot Navigation, The 22nd IFAC World Congress, 2023.

[C6] R. Funada, K. Miyama, T. Toyooka, T. Tanaka and M. Sampei, Minimum Sensing Strategy for a Path-following Problem via Discrete-time Control Barrier Functions, The 22nd IFAC World Congress, 2023.

[C7] A. Pedram and T. Tanaka, A Smoothing Algorithm for Minimum Sensing Path Plans in Gaussian Belief Space, The 2023 American Control Conference, 2023.

[C8] B. He and T. Tanaka, Barrier Pairs for Safety Control of Uncertain Output Feedback Systems, The 2023 American Control Conference, 2023.

[C9] A. Patil, A. Duarte, A. Smith, T. Tanaka and F. Bisetti, Chance-Constrained Stochastic Optimal Control via Path Integral and Finite Difference Methods, The 61st IEEE Conference on Decision and Control, 2022.

[C10] A. Patil and T. Tanaka, Upper Bounds for Continuous-Time End-to-End Risks in Stochastic Robot Navigation, The 2022 European Control Conference, 2022.

[C11] A. Pedram, R. Funada and T. Tanaka, Dynamic Allocation of Visual Attention for Vision-based Autonomous Navigation under Data Rate Constraints, The 60th IEEE Conference on Decision and Control, 2021.

[C12] A. Pedram, J. Stefan, R. Funada and T. Tanaka, Rationally Inattentive Path-Planning via

RRT*, The 2021 American Control Conference, 2021.

[C13] H. Jung and T. Tanaka, Optimal Data Rate Allocation for Dynamic Sensor Fusion over Resource Constrained Communication Networks, The 2021 American Control Conference, 2021.

Thesis

[T1] A. Pedram, "Information-Theoretic Path Planning and Navigation" Ph.D. dissertation, The University of Texas at Austin, 2023.

References

[1] Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

[2] Cagri Kilic, Nicholas Ohi, Yu Gu, and Jason N Gross. Slip-based autonomous zupt through gaussian process to improve planetary rover localization. *IEEE Robotics and Automation Letters*, 6(3):4782–4789, 2021.

[3] Olivier Toupet, Jeffrey Biesiadecki, Arturo Rankin, Amanda Steffy, Gareth Meirion-Griffith, Dan Levine, Maximilian Schadegg, and Mark Maimone. Terrain-adaptive wheel speed control on the curiosity mars rover: Algorithm and flight results. *Journal of Field Robotics*, 37(5):699–728, 2020.

[4] Rongxing Li, Bo Wu, Kaichang Di, Anelia Angelova, Raymond E Arvidson, I-Chieh Lee, Mark Maimone, Larry H Matthies, Lutz Richer, Robert Sullivan, et al. Characterization of traverse slippage experienced by spirit rover on husband hill at gusev crater. *Journal of Geophysical Research: Planets*, 113(E12), 2008.

[5] John P Grotzinger, Joy Crisp, Ashwin R Vasavada, Robert C Anderson, Charles J Baker, Robert Barry, David F Blake, Pamela Conrad, Kenneth S Edgett, Bobak Ferdowski, et al. Mars science laboratory mission and science investigation. *Space science reviews*, 170:5–56, 2012.

[6] Michael Blösch, Stephan Weiss, Davide Scaramuzza, and Roland Siegwart. Vision based mav navigation in unknown and unstructured environments. In *2010 IEEE International Conference on Robotics and Automation*, pages 21–28. IEEE, 2010.

[7] Shaojie Shen, Nathan Michael, and Vijay Kumar. Autonomous multi-floor indoor navigation with a computationally constrained mav. In *2011 IEEE International Conference on Robotics and Automation*, pages 20–25. IEEE, 2011.

[8] Hilbert J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11):P11011, 2005.

[9] Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.

[10] Grady Williams, Andrew Aldrich, and Evangelos A Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40(2):344–357, 2017.