

A Gaussian Variational Inference Approach to Motion Planning

Hongzhe Yu[✉], *Graduate Student Member, IEEE*, and Yongxin Chen[✉], *Senior Member, IEEE*

Abstract—We propose a Gaussian variational inference framework for the motion planning problem. In this framework, motion planning is formulated as an optimization over the distribution of the trajectories to approximate the desired trajectory distribution by a tractable Gaussian distribution. Equivalently, the proposed framework can be viewed as a standard motion planning with an entropy regularization. Thus, the solution obtained is a transition from an optimal deterministic solution to a stochastic one, and the proposed framework can recover the deterministic solution by controlling the level of stochasticity. To solve this optimization, we adopt the natural gradient descent scheme. The sparsity structure of the proposed formulation induced by factorized objective functions is further leveraged to improve the scalability of the algorithm. We evaluate our method on several robot systems in simulated environments, and show that it achieves collision avoidance with smooth trajectories, and meanwhile brings robustness to the deterministic baseline results, especially in challenging environments and tasks.

Index Terms—Motion and path planning, planning under uncertainty, optimization and optimal control.

I. INTRODUCTION

MOTION planning [1] is a fundamental problem in robotics where the goal is to obtain a sequence of states in the space such that it connects a start and goal state while remaining feasible along the plan. When considering motion planning problems, ubiquitous uncertainties arise from imperfect system modeling and measurement noise. Robust motion planning under uncertainties has attracted attentions in the community. Guaranteed robustness was achieved by control and verification design [2], [3] where uncertainties are implicit in the formulation. Stochasticity can also be explicitly brought into the formulation [4]. Planning in belief space [5], [6] models states and measurements as distributions named ‘*belief*’, and planning and control are conducted in these spaces over distributions. Explicitly encoding stochasticity in motion planning has been shown [4] helpful in overcoming locally minimum deterministic solution for non-convex and multimodal [7] optimization problem.

Manuscript received 12 September 2022; accepted 28 February 2023. Date of publication 13 March 2023; date of current version 20 March 2023. This letter was recommended for publication by Associate Editor T. Bandyopadhyay and Editor A. Bera upon evaluation of the reviewers’ comments. This work was supported by the National Science Foundation (NSF) under Grants 1942523 and 2008513. (Corresponding author: Hongzhe Yu.)

The authors are with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: hyu419@gatech.edu; yongchen@gatech.edu).

Digital Object Identifier 10.1109/LRA.2023.3256134

In this work we propose a Gaussian variational inference (GVI) approach to solve motion planning as a probability inference. [8] solved this inference problem using maximum a priori (MAP) estimation. Variational inference (VI) [9] used in this paper, on the other hand, approaches inference problems by solving an optimization within a proposed distribution family. Operating on distributions, VI naturally accounts for stochasticity in an explicit way. A natural gradient descent scheme is used to solve the optimization. The linear Gaussian process (GP) representation of the trajectory used in this paper has gained its popularity in planning [8] and estimation [10] since it encodes smoothness and enjoys a sparsity pattern. Our framework takes into account uncertainties on top of Gaussian Process Motion Planning (GPMP2) [8]. We show that the proposed method is equivalently motion planning with entropy regularization. Entropy maximization in motion planning and reinforcement learning have been studied in [11], [12] and was shown to increase system’s robustness to disturbances [13]. Different from the existing works, our proposed method (1) uses a Newton-style optimization scheme which does not need a sampling scheme or learning process, and (2) is scalable by leveraging the sparsity. (3) The proposed method is shown to be an interpolation from a deterministic solution to a stochastic one. It recovers the deterministic solution by controlling the uncertainty level. (4) We show by experiment that the entropy term encodes the level of risk, which then serves as a metric measuring robustness in decision-making among multiple candidate plans. The optimization scheme for GVI in this paper was first proposed in [14], and has been applied in the robot estimation problems in [15], where the factorized property of the problem was leveraged. To the best knowledge of the authors this is the first work that GVI is used in robot motion planning.

This paper is organized as follows. Section II discusses the related works. In Section III we formulate the motion planning problem as variational inference. The method to solve this inference problem is presented in Section IV. Our framework is illustrated in Section V through numerical experiments.

II. RELATED WORK

The study of motion planning has a long history in robotics community. Sampling based methods such as Rapidly-exploring random tree (RRT) and Probabilistic road map (PRM) [16] [17] provide with optimal yet course paths as graphs or trees connecting start and goal configurations. However they do not

consider dynamical feasibility of the system in their formulations. Trajectory optimization [18], [19], [20] uses optimal control framework to generate trajectories by formulating the problem as a constrained optimization. Direct or collocation methods [4], [8], [18], [21] operate in control and trajectory space while indirect methods [22] optimize only on control inputs, both of which have gained successes in obtaining locally optimal solutions.

The most relevant works to our work is [8] and [15]. In [8] the authors formulated planning as a inference problem and solved it using MAP. [15] proposed a sparse Gaussian variational inference method to solve inference problem in robot estimation. Variational inference has been used in motion plannings also in [7], [23]. We leverage the connection between motion planning and estimation problems, and we also find interesting connections between the GPMP formulation [8], stochastic control problem, and variational inference problem, as discussed in [24], [25].

Robustness to uncertainties is also one of the main motivations of this formulation. Robust motion planning seeks robustness against environment uncertainties. In [2], [3] the authors compute verifiably safe reachable sets using Lyapunov analysis, where the robustness is measured by the volume of the reachable sets around a nominal trajectory. In this work robustness is encoded in the system entropy which is also proportional to the volume of the covariance matrix in Gaussian case.

III. PROBLEM FORMULATION

In this section we formulate motion planning as a variational inference problem. Our formulation generalizes the Gaussian process motion planning [21] that casts motion planning as a MAP task.

A. Gaussian Process Motion Planning

Trajectory optimization formulates the motion planning problem as an optimization of the form

$$\begin{aligned} \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot)} & \mathcal{F}(\mathbf{x}, \mathbf{u}) \\ \text{s.t. } & \mathcal{G}_i(\mathbf{x}, \mathbf{u}) \leq 0, \quad i = 1, \dots, m \\ & \mathcal{H}_i(\mathbf{x}, \mathbf{u}) = 0, \quad i = 1, \dots, r, \end{aligned} \quad (1)$$

where \mathcal{F} is the cost function and \mathcal{G}_i 's, \mathcal{H}_i 's are constraints often related to system dynamics, collision avoidance, or actuation limits. The optimization is over the trajectory $\mathbf{x}(\cdot)$ and the control input $\mathbf{u}(\cdot)$ jointly.

The GPMP framework, alternatively, formulates the motion planning as a MAP problem

$$\begin{aligned} \mathbf{x}^* &= \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{z}) \\ &= \arg \max_{\mathbf{x}} p(\mathbf{z}|\mathbf{x})p(\mathbf{x}), \end{aligned} \quad (2)$$

where the prior distribution $p(\mathbf{x})$ promotes smoothness of the solution, and the likelihood $p(\mathbf{z}|\mathbf{x})$ of some desired behavior encoded by event \mathbf{z} enforces collision avoidance. In particular,

the prior distribution is associated with a linear Gaussian process

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{F}(t)\mathbf{w}(t) + \mathbf{b}(t), \quad (3)$$

where \mathbf{w} denotes standard white noise with covariance \mathbf{Q}_c .

After discretization over time steps $\mathbf{t} = [t_0, \dots, t_N]$, the trajectory becomes a vector $\mathbf{x} = [x_0, \dots, x_N]^T$ and the prior becomes a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{K})$ where the inverse covariance matrix $\mathbf{K}^{-1} = \mathbf{B}^T \mathbf{Q}^{-1} \mathbf{B}$ has an interesting sparse structure [10] with

$$\mathbf{B} = \begin{bmatrix} \mathbf{I} & & & & \\ -\Phi(t_1, t_0) & \mathbf{I} & & & \\ & \ddots & \ddots & & \\ & & -\Phi(t_N, t_{N-1}) & \mathbf{I} & \\ & & & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (4)$$

and

$$\mathbf{Q}^{-1} = \text{diag}(\mathbf{K}_0^{-1}, \mathbf{Q}_{0,1}^{-1}, \dots, \mathbf{Q}_{N-1,N}^{-1}, \mathbf{K}_N^{-1}). \quad (5)$$

Here Φ is the state transition matrix associated with $\mathbf{A}(t)$, \mathbf{Q} is a Grammian defined as $\mathbf{Q}_{i,i+1} = \int_{t_i}^{t_{i+1}} \Phi(t_{i+1}, s)\mathbf{F}(s)\mathbf{Q}_c\mathbf{F}(s)^T\Phi(t_{i+1}, s)^T ds$, and $\mathbf{K}_0, \mathbf{K}_N$ are desired covariances of the start and goal states.

We note that the likelihood probability $p(\mathbf{z}|\mathbf{x})$ describes in general the probability of the feasibility of the current trajectory candidate in (1). In this work we consider collision avoidance likelihood

$$p(\mathbf{z}|\mathbf{x}) \propto \exp(-\|\mathbf{h}(\mathbf{x})\|_{\Sigma_{obs}^{-1}}^2) \quad (6)$$

where $\|\mathbf{h}(\mathbf{x})\|_{\Sigma_{obs}^{-1}}^2$ is a penalty for the collision constraints. Clearly, the MAP problem (2) is equivalent to minimizing the cost function

$$\frac{1}{2}\|\mathbf{x} - \boldsymbol{\mu}\|_{\mathbf{K}^{-1}}^2 + \|\mathbf{h}(\mathbf{x})\|_{\Sigma_{obs}^{-1}}^2, \quad (7)$$

where $\|\cdot\|_{\mathbf{K}^{-1}}$ and $\|\cdot\|_{\Sigma_{obs}^{-1}}$ denote weighted 2-norm.

The prior in (2) can be decomposed into factors

$$p(\mathbf{x}) \propto f_0(x_0)f_N(x_N) \prod_{i=0}^{N-1} f_{gp}^i(x_i, x_{i+1}) \quad (8)$$

with

$$\begin{aligned} f_0(x_0) &= \exp\left(-\frac{1}{2}\|x_0 - \mu_0\|_{\mathbf{K}_0^{-1}}^2\right), \\ f_N(x_N) &= \exp\left(-\frac{1}{2}\|x_N - \mu_N\|_{\mathbf{K}_N^{-1}}^2\right), \\ f_{gp}^i(x_i, x_{i+1}) &= \exp\left(-\frac{1}{2}\|\Phi(t_{i+1}, t_i)(x_i - \mu_i) \right. \\ &\quad \left. - (x_{i+1} - \mu_{i+1})\|_{\mathbf{Q}_{i,i+1}^{-1}}^2\right), \end{aligned}$$

and the collision cost (6) can also be factorized into

$$\exp\left(-\frac{1}{2}\|\mathbf{h}(\mathbf{x})\|_{\Sigma_{obs}^{-1}}^2\right) = \prod_{i=0}^N f_i^{obs}(x_i), \quad (9)$$

where each factor

$$f_i^{obs}(x_i) \triangleq \exp \left(-\frac{1}{2} \|\mathbf{h}(x_i)\|_{\Sigma_{obs}^{-1}}^2 \right) \quad (10)$$

represents the collision cost evaluated at corresponding support states. The collision checking needs to be carried out at a very dense set of points along the trajectory. Gaussian process representation has the advantage that the intermediate collision-checking between the support states can be done through interpolation [8], which keeps the sparsity of the representation.

The assumptions in (8) and (9) together with the GP interpolation bring a sparse parameterization to our problem formulation and is greatly beneficial to the scalability of the proposed algorithm.

Finally, we remark that the MAP formulation (2) can be viewed as a discretization of the following trajectory optimization

$$\begin{aligned} \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot)} & \int_{t_0}^{t_N} \left[\frac{1}{2} \|\mathbf{u}(t)\|_{\mathbf{Q}_c^{-1}}^2 + \|\mathbf{h}(\mathbf{x}(t))\|_{\Sigma_{obs}^{-1}}^2 \right] dt \\ & + \frac{1}{2} \|\mathbf{x}(t_0) - \mu_0\|_{\mathbf{K}_0^{-1}} + \frac{1}{2} \|\mathbf{x}(t_N) - \mu_N\|_{\mathbf{K}_N^{-1}} \end{aligned} \quad (11a)$$

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{F}(t)\mathbf{u}(t) + b(t). \quad (11b)$$

To see this, note that, if we only evaluate $\mathbf{h}(\mathbf{x})$ at discretized time $\mathbf{t} = [t_0, \dots, t_N]$, then for a given $\mathbf{x} = [x_0, \dots, x_N]^T$, the optimization over $\mathbf{u}(\cdot)$ is a linear quadratic control problem for each time interval (t_i, t_{i+1}) and the corresponding closed-form minimum is exactly the exponent of f_{gp}^i .

B. Gaussian Variational Inference for Motion Planing

Though (2) is a probabilistic inference problem, the solution obtained in GPMP is still deterministic in the sense that it searches for a trajectory which maximizes the posterior probability. To better capture the uncertainties and risk presented in motion planning [11], we instead propose to approximate the full posterior distribution $p(\mathbf{x}|\mathbf{z})$ in (2). In particular, we propose the Gaussian variational inference approach to motion planning that seeks to minimize the distance between a Gaussian distribution and the true posterior, measured by KL divergence. It reads

$$\begin{aligned} q^* &= \arg \min_{q \in \mathcal{Q}} \text{KL}[q(\mathbf{x})||p(\mathbf{x}|\mathbf{z})] \\ &= \arg \min_{q \in \mathcal{Q}} \mathbb{E}_q[\log q(\mathbf{x}) - \log p(\mathbf{z}|\mathbf{x}) - \log p(\mathbf{x})] \\ &= \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q[\log p(\mathbf{z}|\mathbf{x}) - \text{KL}[q(\mathbf{x})||p(\mathbf{x})]] \end{aligned} \quad (12)$$

where \mathcal{Q} denotes the Gaussian distribution family. The expression $\mathbb{E}_q[\log p(\mathbf{z}|\mathbf{x}) - \text{KL}[q(\mathbf{x})||p(\mathbf{x})]]$ is known as the evidence lower bound (ELBO). The optimal distribution q^* encourages putting mass on the likelihood $p(\mathbf{z}|\mathbf{x})$ while minimizing its distance from the prior $p(\mathbf{x})$. It shows the trade-off between the smoothness and the collision avoidance.

An alternative form of (12) is

$$\begin{aligned} q^* &= \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z}) - \log q(\mathbf{x})] \\ &= \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})] + H(q) \end{aligned} \quad (13)$$

where $H(q) = -\mathbb{E}_q[\log(q)]$ is the entropy of the distribution. The objective can thus be interpreted as Gaussian process motion planning with an entropy regularization term.

To further balance the trade-off between the original prior-collision cost and the entropy cost, a temperature T can be introduced, pointing to

$$\begin{aligned} q^* &= \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})] + TH(q) \\ &= \arg \max_{q \in \mathcal{Q}} \frac{1}{T} \mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})] + H(q). \end{aligned} \quad (14)$$

When the temperature is low (small T), the optimization puts more weight on maintaining smoothness while avoiding obstacles. When the temperature is high, more weights are put on the system entropy cost to find solutions which have larger covariances so that they can tolerate larger uncertainties.

Remark: Formulation (14) shows an interpolation from the deterministic smooth-collision-avoiding objective (2) to an entropy regularized robust motion planning by changing the temperature T . To recover the deterministic solutions, as T approaches to 0, it can be shown [26] that obtained optimal value will tend to the minimal value for the original objective (2). Indeed, when $T \rightarrow 0$, the objective in (14) approaches $\mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})]$ with respect to $q \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. In this case, when $\boldsymbol{\Sigma}$ shrinks to 0, the objective function $\mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})]$ tends to $\log p(\boldsymbol{\mu}|\mathbf{z})$.

Finally, we note that the variational inference formulation (14) can be viewed as a time discretization of the following stochastic control problem

$$\begin{aligned} \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot)} & \mathbb{E} \left\{ \int_{t_0}^{t_N} \left[\frac{1}{2} \|\mathbf{u}(t)\|_{\mathbf{Q}_c^{-1}}^2 + \|\mathbf{h}(\mathbf{x}(t))\|_{\Sigma_{obs}^{-1}}^2 \right] dt \right. \\ & \left. + \frac{1}{2} \|\mathbf{x}(t_0) - \mu_0\|_{\mathbf{K}_0^{-1}} + \frac{1}{2} \|\mathbf{x}(t_N) - \mu_N\|_{\mathbf{K}_N^{-1}} \right\} \end{aligned} \quad (15a)$$

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{F}(t)(\mathbf{u}(t) + T\mathbf{w}(t)) + b(t). \quad (15b)$$

The proof is based on an equivalence relation between the quadratic control energy and the KL divergence $\text{KL}(q||p)$ [24]. The only difference between (15) and (11) is that the dynamics in (15) is disturbed by white noise $T\mathbf{w}(t)$. Thus, as T goes to zero, (15) should converge to (11).

IV. OPTIMIZATION SCHEME

GVI formulates the motion planning problem as an optimization over Gaussian distributions $q(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

Denote the concatenation of the mean and covariance in vector form as $\boldsymbol{\alpha} \triangleq (\boldsymbol{\mu}, \text{vec}(\boldsymbol{\Sigma}^{-1}))$. The inference objective then reads

$$V(q) = \text{KL}[q(\mathbf{x})||p(\mathbf{x}|\mathbf{z})] = \mathbb{E}_q[\log q(\mathbf{x}) - \log p(\mathbf{x}|\mathbf{z})]. \quad (16)$$

To solve this optimization, we utilize the natural gradient descent scheme. The factorized objective assumption which leads to a sparsity pattern of the problem is also leveraged to improve the scalability of our algorithm.

A. Natural Gradient Descent

For notation simplification, we denote $\psi(\mathbf{x}) = -\log p(\mathbf{x}|\mathbf{z})$. The derivatives w.r.t. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}^{-1}$ can be derived [14] explicitly

$$\frac{\partial V(q)}{\partial \boldsymbol{\mu}} = \boldsymbol{\Sigma}^{-1} \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})\psi(\mathbf{x})] \quad (17a)$$

$$\frac{\partial^2 V(q)}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}^T} = \boldsymbol{\Sigma}^{-1} \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu}^T)\psi(\mathbf{x})] \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1} \mathbb{E}[\psi(\mathbf{x})] \quad (17b)$$

$$\frac{\partial V(q)}{\partial \boldsymbol{\Sigma}^{-1}} = \frac{1}{2} \boldsymbol{\Sigma} \mathbb{E}[\psi(\mathbf{x})] - \frac{1}{2} \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \psi(\mathbf{x})] + \frac{1}{2} \boldsymbol{\Sigma}. \quad (17c)$$

All expectations are taken w.r.t. q . Comparing (17b) and (17c) we obtain

$$\frac{\partial^2 V(q)}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}^T} = \boldsymbol{\Sigma}^{-1} - 2 \boldsymbol{\Sigma}^{-1} \frac{\partial V(q)}{\partial \boldsymbol{\Sigma}^{-1}} \boldsymbol{\Sigma}^{-1}. \quad (18)$$

Having the relations in (17) and (18), for Gaussian distribution $q \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, a natural gradient descent update step w.r.t. objective function V can be calculated straightforward [27] as

$$\begin{bmatrix} \delta \boldsymbol{\mu} \\ \text{vec}(\delta \boldsymbol{\Sigma}^{-1}) \end{bmatrix} = - \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{0} \\ \mathbf{0} & 2(\boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{\Sigma}^{-1}) \end{bmatrix} \begin{bmatrix} \frac{\partial V}{\partial \boldsymbol{\mu}^T} \\ \text{vec}(\frac{\partial V}{\partial \boldsymbol{\Sigma}^{-1}}) \end{bmatrix}. \quad (19)$$

Using properties of the kronecker product and vectorizations of matrices, the update step in natural gradient is

$$\boldsymbol{\Sigma}^{-1} \delta \boldsymbol{\mu} = -\frac{\partial V}{\partial \boldsymbol{\mu}}, \quad \delta \boldsymbol{\Sigma}^{-1} = -2 \boldsymbol{\Sigma}^{-1} \frac{\partial V}{\partial \boldsymbol{\Sigma}^{-1}} \boldsymbol{\Sigma}^{-1}. \quad (20)$$

Notice that we write (20) in terms of $\boldsymbol{\Sigma}^{-1}$ to fully leverage its sparsity pattern. Comparing (17) and (20), we have

$$\delta \boldsymbol{\Sigma}^{-1} = \frac{\partial^2 V(q)}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}^T} - \boldsymbol{\Sigma}^{-1}. \quad (21)$$

Equations (20) and (21) tells that, to calculate the update $\delta \boldsymbol{\mu}, \delta \boldsymbol{\Sigma}^{-1}$, we only need to compute (17a) and (17b). The new variables are calculated using the updates, a step size $\gamma < 1$, and a constant R in a backtracking fashion as

$$\boldsymbol{\mu} \leftarrow \boldsymbol{\mu} + \gamma^R \times \delta \boldsymbol{\mu}, \quad \boldsymbol{\Sigma}^{-1} \leftarrow \boldsymbol{\Sigma}^{-1} + \gamma^R \times \delta \boldsymbol{\Sigma}^{-1}, \quad (22)$$

where $R > 1$ is increasing to shrink the step size for backtracking until the cost decreases. Line search algorithms [28] can also be deployed to obtain locally minimum solutions for this non-convex optimization.

B. Factorized Objectives

We next show that with factorized cost functions, the update step in the algorithm will preserve the sparsity pattern of $\boldsymbol{\Sigma}^{-1}$. Under the factorized assumptions (8) and (9), and denote

$\psi_k(\mathbf{x}_k) = -\log p(\mathbf{x}_k|\mathbf{z})$, (16) also factorizes

$$\begin{aligned} V(q) &= \mathbb{E}_q[\log q(\mathbf{x})] - \sum_{k=1}^K \mathbb{E}_{q_k}[\psi_k(\mathbf{x}_k)] \\ &= \mathbb{E}_q[\log q(\mathbf{x})] - \sum_{k=1}^K \mathbb{E}_{q_k}[\log p(\mathbf{x}_k) + \log p(\mathbf{z}|\mathbf{x}_k)] \\ &\triangleq \frac{1}{2} \log(|\boldsymbol{\Sigma}^{-1}|) + \sum_{k=1}^K V_k(q_k) \end{aligned} \quad (23)$$

where $V_k(q_k)$'s are factored costs and \mathbf{x}_k are the corresponding subsets of variables to the k th factor. We assume that \mathbf{x}_k can be transformed from \mathbf{x} using a linear mapping \mathbf{M}_k , i.e., $\mathbf{x}_k = \mathbf{M}_k \mathbf{x}$, and the marginal Gaussian $q_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. The relation between the joint and the factorized variables reads

$$\boldsymbol{\mu}_k = \mathbf{M}_k \boldsymbol{\mu}, \quad \boldsymbol{\Sigma}_k = \mathbf{M}_k \boldsymbol{\Sigma} \mathbf{M}_k^T. \quad (24)$$

In view of (17) and (20), to compute the updates $\delta \boldsymbol{\mu}$ and $\delta \boldsymbol{\Sigma}^{-1}$, we need to calculate the derivatives of the joint objective which also factorizes as

$$\frac{\partial V(q)}{\partial \boldsymbol{\mu}} = \sum_{k=1}^K \mathbf{M}_k^T \frac{\partial V_k(q_k)}{\partial \boldsymbol{\mu}_k}, \quad (25a)$$

$$\frac{\partial^2 V(q)}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}^T} = \sum_{k=1}^K \mathbf{M}_k^T \frac{\partial^2 V_k(q_k)}{\partial \boldsymbol{\mu}_k \partial \boldsymbol{\mu}_k^T} \mathbf{M}_k. \quad (25b)$$

The factorized derivatives $\frac{\partial V_k}{\partial \boldsymbol{\mu}_k}$ and $\frac{\partial^2 V_k}{\partial \boldsymbol{\mu}_k \partial \boldsymbol{\mu}_k^T}$ will have the same expressions as in (17) w.r.t. marginal distributions $q_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ and marginal factors $\psi(\mathbf{x}_k)$

$$\frac{\partial V_k}{\partial \boldsymbol{\mu}_k} = \boldsymbol{\Sigma}_k^{-1} \mathbb{E}_{q_k}[(\mathbf{x}_k - \boldsymbol{\mu}_k)\psi(\mathbf{x}_k)], \quad (26a)$$

$$\begin{aligned} \frac{\partial^2 V_k}{\partial \boldsymbol{\mu}_k \partial \boldsymbol{\mu}_k^T} &= \boldsymbol{\Sigma}_k^{-1} \mathbb{E}_{q_k}[(\mathbf{x}_k - \boldsymbol{\mu}_k)(\mathbf{x}_k - \boldsymbol{\mu}_k^T)\psi(\mathbf{x}_k)] \boldsymbol{\Sigma}_k^{-1} \\ &\quad - \boldsymbol{\Sigma}_k^{-1} \mathbb{E}_{q_k}[\psi(\mathbf{x}_k)]. \end{aligned} \quad (26b)$$

From (21), (25) and (26) we see that the sparsity pattern of the precision matrix $\boldsymbol{\Sigma}^{-1}$ is preserved after the transitions between the joint and factorized updates.

From (24) we know that a joint covariance matrix $\boldsymbol{\Sigma}$ is computed in each update step. Throughout the iterations $\boldsymbol{\Sigma}^{-1}$ remains sparse, but $\boldsymbol{\Sigma}$ need not to be. However, because of the consistent sparsity pattern, efficient methods [29] exist in sparse linear algebra literature to compute only the parts of $\boldsymbol{\Sigma}$ corresponding to the non-zero elements in $\boldsymbol{\Sigma}^{-1}$. Alternatively, Gaussian belief propagation [30][31] can also solve the marginal covariance efficiently. The expectations in (26) are approximately evaluated using Gauss-Hermite quadrature [32] in this work. We note that when the posterior $p(\mathbf{x}|\mathbf{z})$ is linear, then expectations in (17) have closed-form, which greatly accelerates the algorithm.

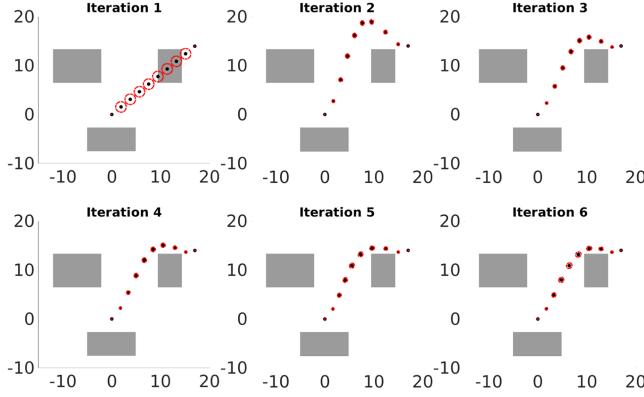


Fig. 1. Converging process with $T = 10$, $\mathbf{Q}_c = 0.8\mathbf{I}$, $\Sigma_{obs} = 0.004\mathbf{I}$, $\epsilon = 0.7$. Linear interpolated initialization for μ and $10\mathbf{I}$ for Σ^{-1} .

V. EXPERIMENTS

In all our experiments, we consider a constant-velocity model in (3). Let

$$\mathbf{A}(t) = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \mathbf{b}(t) = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \mathbf{F}(t) = \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix}. \quad (27)$$

The transition matrix Φ , matrices \mathbf{Q}_i , \mathbf{Q}_i^{-1} , \mathbf{Q} , and \mathbf{B} in (4) can be calculated explicitly [10]. The likelihood function is defined the same as in [8] [18] by

$$\mathbf{h}(\mathbf{x}) = \mathbf{c}_\epsilon(\mathbf{d}(FK(\mathbf{x}))) \quad (28)$$

where $FK(\cdot)$ is the forward kinematics, $\mathbf{d}(\cdot)$ is the signed distance function given a signed distance field (SDF), and $\mathbf{c}_\epsilon(\cdot)$ is the hinge loss function

$$\mathbf{c}_\epsilon(y) = \begin{cases} 0, & \text{if } y \geq \epsilon \\ \epsilon - y, & \text{if } y < \epsilon. \end{cases} \quad (29)$$

When evaluating the signed distance function $\mathbf{d}(\cdot)$, robots are modeled as balls with fixed radius r [8] at designated locations. The minimum distance from robots to obstacles is efficiently computed using the distance between centers of the balls to the obstacles and the ball radius. In this paper, to highlight the convergence of the algorithm, GP interpolation is not involved in any experiments.

A. 2 d Point Robot Collision Avoidance

The first experiment is conducted with a planar point robot, which better captures the covariances by plotting ellipsoids. Fig. 1 shows the convergence of the support states. Black dots represent μ , and the red ellipsoids draw the 0.997 confidence region contour. We initialize μ using a linear interpolation between the start and goal states, and initialize Σ^{-1} using isotropic matrices.

a) *Trade-off between motion planning and system entropy*: The cost function in (23) contains two parts: a motion planning including prior and collision costs, and a regularized entropy cost. Fig. 2 shows the evolution of different costs and the total cost, where the prior and collision costs are factorized, and

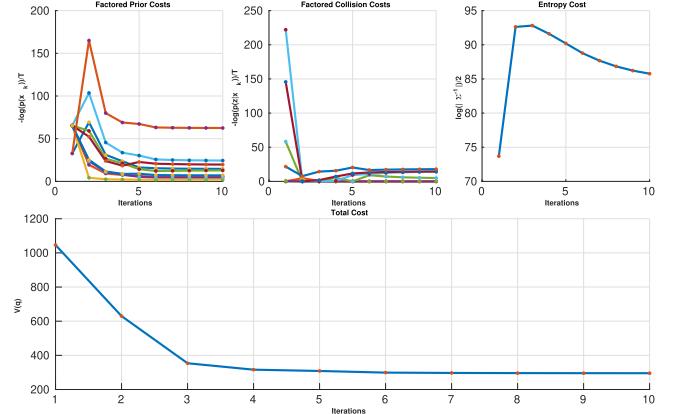


Fig. 2. Decomposed and total costs. Prior and collision costs on the factor level and the entropy cost on the joint level.

the cost on the entropy $\frac{1}{2} \log(|\Sigma^{-1}|)$ is computed on the joint level. As shown in Fig. 2, during the first several iterations the prior and collision costs on each factor decreases, meaning that the system gets rid of the obstacle while maximizing trajectory smoothness and system dynamics assumptions imposed by the prior. Meanwhile, the entropy costs increase. After the system is safe and smooth, the algorithm moves to the region where the entropy cost decreases. During the two phases, the total loss decreases. This trade-off process is also reflected in the Fig. 1. The covariance pivots shrink while the system is avoiding the obstacles, and increase after the system is safe and smooth.

b) *Planning with high temperature*: In (14), a temperature T is introduced to alter the weights between planning and entropy objectives. To achieve feasible trajectories, we use small T . However in low temperature regions, little changes on Σ will happen due to the low weight on the entropy cost. One motivation of the proposed formulation is to leverage the entropy in order to have wider-spread distributions in all areas, since the 3σ area measures the size of the safe regions in a probabilistic sense. Higher temperature promotes the system's entropy, but put less weights on the feasibility part. A compromise is to use a near-feasible initialization with high temperature. The initialization for the mean μ can either be the output of a lower temperature optimization as a re-planning, or from a higher level sampling based planner. Fig. 3 shows the converging process of the iterations for a high temperature re-planning. We note that the low temperature planning and the high temperature re-planning can be done in a consecutive manner in the optimization.

B. More Challenging Planning Problems

In the next set of experiments we show that by introducing entropy regularization to the deterministic formulation, we gain flexibility in solution searching as well as a risk-measuring metric. We illustrate using several experiments. In paragraph (a), to test the performance in hard tasks, we conduct long range planning in cluttered environments. In (b) we use a narrow gap environment to show that stochasticity brings flexibility in choosing collision-checking radius, compared with deterministic baseline; In (c) we show that stochasticity help explore

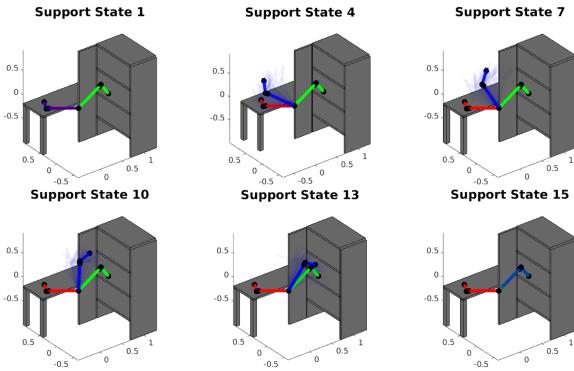


Fig. 12. Mean values and samples for selected support states.

samples distribute wider, representing higher entropy, and in the more constrained areas, there are less freedom.

b) *7-DOF WAM arm model*: Solving the optimization in the space of distributions brings additional computation complexities compared with the deterministic formulation. However, the factorized cost function (23) and partial update schemes (26) mitigate the problem. In addition, there exist more efficient methods in evaluating the integrals in (26), which can further accelerate the algorithm. We evaluated the proposed algorithm on a 7-DOF WAM Arm robot in a more realistic dataset, the optimized mean and samples are shown in Figs. 11 and 12.

VI. CONCLUSION

In this work we proposed a Gaussian variational inference framework to approach motion planning as a probability inference. On top of the Gaussian process representation of the trajectory, we calculate an optimal Gaussian distribution over the trajectories. Natural gradient descent scheme was deployed to solve the GVI. Factorized cost functions brings a sparsity pattern into the framework, and Gaussian assumption brings an explicit update scheme which converges quickly to locally minimum solutions. Alternatively, the proposed framework can be viewed as motion planning with entropy regularization. Experiments show that the proposed method achieves smooth collision-free trajectories, and also provides more robust solutions than deterministic baseline methods, especially in challenging environments. The limitation of the proposed algorithm is the computation complexity increased by introducing additional optimization variables, which is a trade-off for the additional distributional information gain. However, this issue can be mitigated by leveraging the problem's sparsity pattern and more advanced integration estimation techniques.

REFERENCES

- [1] J.-C. Latombe, *Robot Motion Planning*, vol. 124. Berlin, Germany:Springer, 2012.
- [2] R. Tedrake, I. R. Manchester, M. Tobenkin, and J. W. Roberts, "LQR-trees: Feedback motion planning via sums-of-squares verification," *Int. J. Robot. Res.*, vol. 29, no. 8, pp. 1038–1052, 2010.
- [3] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *Int. J. Robot. Res.*, vol. 36, no. 8, pp. 947–982, 2017.
- [4] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 4569–4574.
- [5] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *Int. J. Robot. Res.*, vol. 32, no. 9–10, pp. 1194–1227, 2013.
- [6] J. V. Den, S. B. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *Int. J. Robot. Res.*, vol. 31, no. 11, pp. 1263–1278, 2012.
- [7] T. Osa, "Multimodal trajectory optimization for motion planning," *Int. J. Robot. Res.*, vol. 39, no. 8, pp. 983–1001, 2020.
- [8] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuous-time Gaussian process motion planning via probabilistic inference," *Int. J. Robot. Res.*, vol. 37, no. 11, pp. 1319–1340, 2018.
- [9] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *J. Amer. Stat. Assoc.*, vol. 112, no. 518, pp. 859–877, 2017.
- [10] T. D. Barfoot, C. H. Tong, and S. Särkkä, "Batch continuous-time trajectory estimation as exactly sparse Gaussian process regression," in *Proc. Robot.: Sci. Syst.*, 2014.
- [11] A. Lambert and B. Boots, "Entropy regularized motion planning via stein variational inference," 2021, *arXiv:2107.05146*.
- [12] B. D. Ziebart et al., "Maximum entropy inverse reinforcement learning," in *Proc. AAAI Conf. Artif. Intell.*. Chicago, IL, USA, 2008, vol. 8, pp. 1433–1438.
- [13] B. Eysenbach and S. Levine, "Maximum entropy RL (provably) solves some robust RL problems," in *Proc. Int. Conf. Learn. Representations*, 2022.
- [14] M. Opper and C. Archambeau, "The variational Gaussian approximation revisited," *IEEE Neural Computation*, vol. 21, no. 3, pp. 786–792, Mar. 2009.
- [15] T. D. Barfoot, J. R. Forbes, and D. J. Yoon, "Exactly sparse Gaussian variational inference with application to derivative-free batch nonlinear state estimation," *Int. J. Robot. Res.*, vol. 39, no. 13, pp. 1473–1502, 2020.
- [16] S. M. LaValle et al., "Rapidly-exploring random trees: A new tool for path planning," *Annu. Res. Rep.*, 1998.
- [17] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Automat.*, vol. 12, no. 4, pp. 566–580, Aug. 1996.
- [18] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 489–494.
- [19] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Proc. Robot.: Sci. Syst. IX*, 2013.
- [20] J. Schulman et al., "Motion planning with sequential convex optimization and convex collision checking," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [21] M. Mukadam, X. Yan, and B. Boots, "Gaussian process motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 9–15.
- [22] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 1168–1175.
- [23] T. Osa, "Motion planning by learning the solution manifold in trajectory optimization," *Int. J. Robot. Res.*, vol. 41, no. 3, pp. 281–311, 2022.
- [24] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal transport over a linear dynamical system," *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2137–2152, May 2017.
- [25] Y. Chen, T. T. Georgiou, and M. Pavon, "On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint," *J. Optim. Theory Appl.*, vol. 169, no. 2, pp. 671–691, 2016.
- [26] T. Hazan and A. Shashua, "Norm-product belief propagation: Primal-dual message-passing for approximate inference," *IEEE Trans. Inf. Theory*, vol. 56, no. 12, pp. 6294–6316, Dec. 2010.
- [27] J. R. Magnus and H. Neudecker, *Matrix Differential Calculus With Applications in Statistics and Econometrics*. Hoboken, NJ, USA: Wiley, 2019.
- [28] S. Wright and J. Nocedal, "Numerical optimization," *Springer Sci.*, vol. 35, no. 67–68, 1999, Art. no. 7.
- [29] F. Broussolle, "State estimation in power systems: Detecting bad data through the sparse inverse matrix method," *IEEE Trans. Power App. Syst.*, vol. PAS-97, no. 3, pp. 678–682, May 1978.
- [30] D. Bickson, "Gaussian belief propagation: Theory and application," 2008, *arXiv:0811.2518*.
- [31] J. Ortiz, T. Evans, and A. J. Davison, "A visual introduction to Gaussian belief propagation," 2021, *arXiv:2107.02308*.
- [32] I. Arasaratnam, S. Haykin, and R. J. Elliott, "Discrete-time nonlinear filtering algorithms using Gauss–Hermite quadrature," *Proc. IEEE*, vol. 95, no. 5, pp. 953–977, May 2007.