

# Ergodic Exploration using Tensor Train: Applications in Insertion Tasks

Suhan Shetty<sup>1,2</sup>, João Silvério<sup>1</sup>, and Sylvain Calinon<sup>1,2</sup>

**Abstract**—In robotics, ergodic control extends the tracking principle by specifying a probability distribution over an area to cover instead of a trajectory to track. The original problem is formulated as a spectral multiscale coverage problem, typically requiring the spatial distribution to be decomposed as Fourier series. This approach does not scale well to control problems requiring exploration in search space of more than 2 dimensions. To address this issue, we propose the use of tensor trains, a recent low-rank tensor decomposition technique from the field of multilinear algebra. The proposed solution is efficient, both computationally and storage-wise, hence making it suitable for its online implementation in robotic systems. The approach is applied to a peg-in-hole insertion task requiring full 6D end-effector poses, implemented with a 7-axis Franka Emika Panda robot. In this experiment, ergodic exploration allows the task to be achieved without requiring the use of force/torque sensors.

**Index Terms**—Ergodic control, low-rank approximation, tensor methods, tensor train, tensor factorization, peg-in-hole insertion task, learning from demonstration.

## I. INTRODUCTION

**A**UTONOMOUS systems are often encountered with coverage tasks such as localization, tracking, and active learning. In such tasks, the agent might be required to explore a region of its state space, either due to the nature of the task at hand (e.g. surveillance) or due to uncertainties induced by sensory inaccuracies (e.g. peg-in-hole insertion). In such problems, the coverage task can be specified by a reference probability density function, which encodes the importance of exploration at any point of the state space. For such problems, a pattern-based coverage approach (e.g., a “lawnmower-type” strategy), as commonly used in low-dimensional state space, is not scalable, and hence not applicable to most of the applications encountered in practice [1]. Maximizing information gain, another popular approach to circumvent uncertainty, is not suitable for exploration since the coverage is likely to be concentrated in regions around information maxima disproportionately over the period of exploration [2].

Ergodic control provides an elegant solution to design control policies for such autonomous systems, in order to equip them with natural search behaviors. For a given reference probability density function over a domain of interest in the state space of the robot, a dynamical system is said to be ergodic if the fraction of time spent in a given region




Fig. 1: The 6D robot trajectory generated by ergodic controller offline for the insertion task.

is proportional to the probability mass of that region [3]. This is formalized in ergodic theory, where the goal is to characterize how ergodic a given dynamical system is. Ergodic control, on the other hand, aims to design a control policy for a given autonomous system so that the trajectory evolution of the resulting dynamical system is ergodic for the reference probability distribution. Systems engineered in such a way have already found applications in robotics [2], [4]. The original approach to ergodic control is called Spectral Multiscale Coverage (SMC) and involves spectral analysis of the dynamical system evolution [3]. This elegant method was proposed for point-mass systems having receding horizon control with infinitesimal control horizon. This original work paved the way for various extensions, with other types of dynamical systems and finite-horizon controllers [4], [2], [5]. The idea behind SMC [3] is to minimize a metric, called the ergodic metric, that quantifies the match between the Fourier coefficients of a reference distribution and those of the time-averaged statistics of the system trajectory. As we will see in Section III-B, this method unfortunately suffers from the curse of dimensionality, prohibiting its applications to search spaces with more than 2 or 3 dimensions, which are often encountered in robot manipulation problems (e.g. exploration in the task space of an end-effector is often a 6D problem).

In this paper, we propose a solution to overcome the challenges in SMC by using low-rank tensor approximation

<sup>1</sup> Idiap Research Institute, Martigny, Switzerland.

<sup>2</sup> École Polytechnique Fédérale de Lausanne (EPFL), Switzerland.

E-mails: name.surname@idiap.ch

This work was supported by the CoLLaboratE project (<https://collaborate-project.eu>), funded by the EU within H2020-DT-FOF-02-2018 under grant agreement 820767, and by the LEARN-REAL project (<https://learn-real.eu>), funded by the Swiss National Science Foundation (CHIST-ERA project).

techniques, in the form of a tensor train (TT), and hence expanding the domain of ergodic control to robot manipulation. Figure 1 shows an ergodic exploration behavior generated by the proposed method for the peg-in-hole task considered in this paper.

We showcase our approach in a 6D peg-in-hole insertion task using a robot manipulator, by considering the position and orientation of the end-effector<sup>1</sup>. In peg-in-hole scenarios, perception and modeling inaccuracies often compromise success, requiring the robot to leverage smart control strategies. Here, we propose to apply ergodic control to facilitate the insertion by letting the robot explore around the hole location in the 6D state space of the end-effector. In this application, we rely on human demonstrations to specify the distribution that the robot should use for an ergodic exploration.

The main contribution of this work is an algorithm for SMC to generate ergodic exploration in multi-dimensional spaces, which was previously considered to be an intractable problem. In particular, we improve the state of the art by proposing:

- Fast ways to compute the Fourier coefficients of multivariate functions, a well-known bottleneck in the ergodic control literature;
- The use of tensor train to exploit the inherent low-rank structure in the problem, which is used to overcome the curse of dimensionality in both real-time computation and storage requirements and facilitate implementation of ergodic control online on robotic systems.

The proposed ergodic control algorithm paves the way for two additional contributions in robotics. Particularly, we:

- Extend ergodic control to peg-in-hole tasks solved with an online policy, with a novel, principled and theoretically-grounded exploratory strategy for insertion tasks that does not rely on specialized sensors but only on human demonstrations;
- Provide a formal way to perform ergodic exploration in orientation by relying on the  $S^3$  Riemannian manifold.

To the best of our knowledge, this is the first time ergodic control is implemented online on a physical robot for exploration in dimension greater than 2. Note that the strategies to mitigate the curse of dimensionality introduced in this paper have the potential to be applied to many other applications in robotics to address real-time computation and limited storage requirements. Similarly, the proposed control strategy is not limited to manipulation applications, and can be extended to other robotics scenarios requiring high-dimensional coverage (e.g. 3D-object modeling, 6D surveillance).<sup>2</sup> Finally, even though we rely on human demonstrations to obtain the reference distribution, this reference can in practice be specified/learned in different ways (e.g. from sensor uncertainty models).

The paper is organized as follows: Section II gives a literature survey on ergodic control, tensor methods and control strategies for insertion tasks. Section III covers the necessary mathematical background to introduce our contribution. Particularly, Section III-A introduces tensor algebra.

<sup>1</sup>A video of the experiment is available at <https://sites.google.com/view/ergodic-exploration/>

<sup>2</sup>The proposed algorithm has been numerically evaluated for state space of up to 15 dimensions.

In Section III-B, the mathematical formulation of ergodic control is described, where the underlying challenges of the algorithm are outlined. Section III-C briefly introduces tensor decomposition techniques and Section III-D gives an overview of the tensor train decomposition, which is the main tool for low-rank approximation used in this paper. In Section IV, we propose low-rank approximation using the tensor train as a solution for multidimensional ergodic exploration. In Section V, we evaluate the proposed algorithm in simulation. Lastly, in Section VI we showcase the results of our approach in a peg-in-hole insertion task using a torque-controlled 7-axis Franka Emika Panda robot.

## II. MOTIVATION AND RELATED WORK

### A. Ergodic Control

A solution to ergodic control was originally proposed by [3] using Spectral Multiscale Coverage (SMC) in the form of a feedback control law designed for multi-agent systems, with an objective defined so that the agents trajectories cover a reference probability distribution. Here, the system considered is a point-mass system. The control policy is obtained by solving an optimization problem with an ergodic metric as the cost function (see Section III-B). The ergodic metric compares the Fourier series coefficients associated to the spatial reference distribution and the trajectory evolution of the system.

Although other possible choices of basis functions would be interesting to investigate (e.g. wavelets), the Fourier transform holds essential properties that are relevant to the considered problem. In [6], we discuss the use of Fourier series within ergodic control, including their links to cosine basis functions, and their properties for reference distributions in the form of mixtures of Gaussians.

The ergodic metric can be used as a starting point to design other forms of controllers. For example, ergodic controllers have been proposed using nonlinear dynamical systems and finite control horizons [2], using projection-based trajectory optimization [7], or using hybrid systems theory [5]. An overview of these methods with finite control horizon can be found in [8].

The main drawback of these methods is that they suffer from the curse of dimensionality (see Section III-B) when the dimension of the state space for ergodic exploration increases. This is due to the computational complexity and storage demanded in working with the ergodic metric and the control policy derived from it. For low dimensional exploration tasks (2D), Dressel and Kochenderfer used supervised learning to reduce the computational burden [9]. However, it does not scale to higher dimensional problems. Several authors deviated from the approach used in SMC to tackle this limitation. In [10] and [11], the authors relied on a different ergodic metric based on a Kullback-Leibler (KL) divergence measure for finite sensor footprint, where the control policy is obtained using sampling-based techniques. Here, the ergodic metric (KL-divergence) is approximated using the samples from the reference distribution. Sampling-based methods avoid the curse of dimensionality but they can still be computationally

expensive to address the real-time computational requirements of robotics systems. Moreover, the performance of the method is heavily dependent on the quality of the samples obtained, which is hard to assess. While most sampling-based methods generate the ergodic trajectory offline, its online implementation on real robots, which is the focus of the current paper, is still a challenging problem. Based on [10], Abraham *et al.* provides an online version of ergodic control with KL-divergence as ergodic metric [11], at the expense of potentially losing ergodicity in the exploration (e.g., by limiting the search to high density regions).

This paper keeps the original methodology (SMC) proposed by [3], which is the foundation of most literature on ergodic control, and which has the advantage of providing closed form solutions for many of the commonly used models of dynamical systems (kinematics-based) [4], while providing multi-scale coverage behaviour. To do so, we propose a solution based on a low-rank tensor approximation to overcome the curse of dimensionality. The proposed algorithm has intuitive hyper-parameters that can be adjusted to address the storage and computational constraints of the application.

### B. Tensor Methods

As we will see in Section III-B, the difficulty in ergodic control essentially arises from the storage and computation with multi-dimensional arrays involved in the algorithm. Tensor methods have recently gained popularity in signal processing, physics, applied mathematics, and machine learning communities due to their efficiency in storing and working with multidimensional arrays. These methods exploit the structure inherent in multidimensional arrays such as symmetry, parallel proportionality, and separability to represent them compactly and robustly. Furthermore, they allow performing algebraic operations efficiently in the compact format. Thus, the storage complexity and the algebraic operations complexity are significantly reduced. For a survey of classical tensor methods, we refer the readers to [12]. For applications of tensor methods in signal processing and machine learning, we refer to [13], [14]. In control, tensor methods have been used in [15] and [16] to solve multidimensional optimal control problems which were previously considered to be intractable.

### C. Insertion Tasks

Peg-in-hole insertion is a typical and important problem in robotics. Many strategies to solve this problem depend on expensive force and torque sensors [17], [18]. Sensorless strategies [19], [20], [21], on the other hand, rely only on the state of the end-effector and provide a low cost solution. However, most of the sensorless strategies depend either on a predetermined trajectory that the robot end-effector needs to follow [21], or a full modeling of the insertion behavior [19], [20]. In [21], insertion is treated as a 3D (2D position and 1D orientation of the end-effector) trajectory tracking problem, where the reference trajectory for the robot end-effector is generated offline by using a coverage strategy such as ergodic control. As we will show, this strategy fails for the peg-in-hole insertion task considered in this paper, as a trajectory generated

offline is often not possible to track due to obstructions from the surface of the hole and the peg. To address this challenge, our approach instead formulates the coverage problem in an online manner, with exploration simultaneously in the full 6D state space of the robot end-effector (3D position and 3D orientation).

In order to handle the exploration in orientation jointly with the position, we extend the control strategy to Riemannian manifolds by modeling the probability distribution of orientations, see [22], [23] for details. Subsequently, we use an online implementation of ergodic control as the solution for coverage. The algorithm proposed in this paper allows us to run the ergodic controller online on the robot for 6D insertion tasks.

TABLE I: Description of key notations and variables.

$d$	Dimension of the state space
$K$	Number of elementary basis functions along each dimension ( $K^d$ basis functions for the state space)
$N$	Degree of Gaussian quadrature rule ( $N \in \mathbb{Z}^+$ )
$\mathbf{b}$	1D array (vector)
$\mathbf{A}$	2D array (matrix)
$\mathcal{G}$	tensor of order $> 2$
$\mathcal{G}_k$	$k$ -th element of a tensor $\mathcal{G}$ (a scalar)
$\mathcal{G}^i$	$i$ -th core (a third order tensor) of $\mathcal{G}$ in TT-format
$\mathcal{G}_{:, :, k}^i$	$k$ -th frontal slice (a matrix) of $\mathcal{G}^i$
$\Omega$	Search space for ergodic exploration: $\Omega = [0, L]^d$
$P(\mathbf{x})$	Reference probability distribution defined on $\Omega$
$\mathcal{P}$	Tensor formed by discretizing $P(\mathbf{x})$
$\phi_k(x)$	Elementary Fourier basis function for $[0, L]$ , $\phi_k(x) = \cos(\frac{2\pi(k-1)x}{L})$
$\phi(x)$	A vector of elementary Fourier basis functions for $[0, L]$ , $\phi(x) = (\phi_1(x), \dots, \phi_K(x))$
$\Phi(x)$	The Fourier basis functions for $\Omega$ (a $d$ -th order tensor)
$C_t(\mathbf{x})$	Spatial statistics of the trajectory evolution $\mathbf{x}(t)$ of the dynamic system
$\mathcal{W}_t$	Fourier series coefficients of $C_t(\mathbf{x})$

## III. PROBLEM DEFINITION AND BACKGROUND

In this section we lay out the mathematical background of our contribution. The notation used is summarized in Table I.

### A. Tensors

A tensor is a multidimensional array<sup>3</sup>, whose order corresponds to the number of modes (or dimensions) of the array. A vector is a first order tensor and a matrix is a second order tensor. The  $k$ -th element of a  $d$ -th order tensor  $\mathcal{X} \in \mathbb{R}^{K_1 \times \dots \times K_d}$ , with indices  $\mathbf{k} = (k_1, \dots, k_d)$  and  $k_i \in \{1, \dots, K_i\}$ , is denoted by  $\mathcal{X}_{\mathbf{k}}$ . Here,  $K_i \in \mathbb{Z}^+$  represents the size of  $i$ -th mode, with  $i \in \{1, \dots, d\}$ .

Fibers are the higher-order analogs of matrix rows and columns. A fiber is obtained by fixing every index but one.

<sup>3</sup>The notion of tensors used in this paper is not to be confused with the one commonly used in physics as a multilinear map with specific properties. In this paper, a tensor is just a multidimensional array.

The inner product of two tensors  $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{K_1 \times \dots \times K_d}$  is given by

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{\mathbf{k} \in \mathcal{K}} \mathbf{x}_{\mathbf{k}} \mathbf{y}_{\mathbf{k}},$$

where  $\mathcal{K} = \{\mathbf{k} = (k_1, \dots, k_d) : k_i \in \{1, \dots, K_i\}\}$ . The Frobenius norm of a tensor  $\mathcal{X}$  is defined by

$$\|\mathcal{X}\| = \langle \mathcal{X}, \mathcal{X} \rangle^{\frac{1}{2}}.$$

The outer product of a tuple of one-dimensional arrays  $(\mathbf{a}^1, \dots, \mathbf{a}^d)$ , with  $\mathbf{a}^i = (a_{1, \dots, K_i}^i) \in \mathbb{R}^{K_i}$  is denoted by  $\mathcal{X} = \mathbf{a}^1 \circ \dots \circ \mathbf{a}^d \in \mathbb{R}^{K_1 \times \dots \times K_d}$ , whose  $\mathbf{k}$ -th element is given by

$$\mathbf{x}_{\mathbf{k}} = a_{k_1}^1 \dots a_{k_d}^d.$$

In a wide variety of applications, observed data could be represented naturally as tensors [12], [24]. In this article tensors arise either due to the discretization of an underlying multivariate function or from the tensor representation of the variables involved in the ergodic control algorithm.

### B. Ergodic Control

Ergodic control considers a point-mass dynamical system whose trajectory evolves such that its time-averaged statistics matches a desired reference probability distribution. In the method proposed originally in [3], the problem reduces to minimizing a cost function called ergodic metric, evaluating the distance between the Fourier coefficients of the reference distribution and that of the time-averaged statistics of the trajectory evolution of the dynamical system.

We assume a bounded  $d$ -dimensional rectangular domain:  $\Omega = [0, L_1] \times \dots \times [0, L_d]$  with  $L_i > 0, \forall i \in \{1, \dots, d\}$ . Without loss of generality, we will consider  $L_i = L, \forall i \in \{1, \dots, d\}$ .  $\mathbf{x}(t) \in \mathbb{R}^d$  represents the trajectory of the dynamical system in the domain. The spatial statistics of the trajectory  $\mathbf{x}(t)$  is defined as the fraction of time spent by the dynamical system at each point of the domain:

$$C_t(\mathbf{x}) = \frac{1}{t} \int_{\tau=0}^t \delta(\mathbf{x}(\tau) - \mathbf{x}) d\tau,$$

where  $\delta$  is the Dirac delta function, and  $\mathbf{x} = (x_1, \dots, x_d) \in \Omega$  is a point in the domain.

Let  $P(\mathbf{x})$  be the reference probability distribution for the exploration defined on  $\Omega$ . The goal of ergodic control is to match the time-averaged spatial statistic  $C_t(\mathbf{x})$  with the spatial distribution  $P(\mathbf{x})$ . The idea is to choose  $K \in \mathbb{Z}^+$  orthonormal Fourier basis functions<sup>4</sup> satisfying the Neumann boundary conditions on the boundary of  $\Omega$ :  $\phi_k, \forall k \in \{1, \dots, K\}$ , for each variable  $x$ , which is then organized as  $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_K(\mathbf{x})) \in \mathbb{R}^K$ . Although the results that follow apply to any such choice of basis function, we will use  $\phi_k(x) = \cos(\frac{2\pi(k-1)x}{L})$  for numerical evaluation, see [6] for details. Then, orthonormal Fourier basis functions for  $\Omega$  can be obtained by the elements of the  $d$ -th order tensor formed by the outer product of these vectors:  $\Phi(\mathbf{x}) = \phi(x_1) \circ \dots \circ \phi(x_d) \in$

<sup>4</sup>In general, we can choose a different number of basis functions  $K_i$  for each dimension  $i \in \{1, \dots, d\}$ . Without loss of generality, we will assume  $K_i = K, \forall i$ .

$\mathbb{R}^{K \times \dots \times K}$ . With respect to this basis, the Fourier coefficients (cosine transforms) of  $P(\mathbf{x})$  can be represented by a  $d$ -th order tensor  $\hat{\mathcal{W}}$ . For a given index  $\mathbf{k} = (k_1, \dots, k_d) \in \mathcal{K}$ , we have  $\Phi_{\mathbf{k}}(\mathbf{x}) = \phi_{k_1}(x_1) \times \dots \times \phi_{k_d}(x_d)$ , and  $\hat{\mathcal{W}}_{\mathbf{k}}$  represents the Fourier coefficient w.r.t. the basis  $\Phi_{\mathbf{k}}$ , namely

$$\hat{\mathcal{W}}_{\mathbf{k}} = \int_{x_1=0}^L \dots \int_{x_d=0}^L P(\mathbf{x}) \Phi_{\mathbf{k}}(\mathbf{x}) dx_1 \dots dx_d. \quad (1)$$

The ergodic metric is then defined as

$$\xi(t) = \sum_{\mathbf{k} \in \mathcal{K}} \Lambda_{\mathbf{k}} (\mathcal{W}_{\mathbf{k}}(t) - \hat{\mathcal{W}}_{\mathbf{k}})^2, \quad (2)$$

where  $\Lambda_{\mathbf{k}} = (1 + \|\mathbf{k}\|^2)^{-\frac{d+1}{2}}$  are the weights for different frequencies,  $\mathcal{K} = \{\mathbf{k} = (k_1, \dots, k_d) : k_i \in \{1, \dots, K\}\}$  and  $K$  is a sufficiently large positive integer. This way, higher priority is given to lower frequency contents of the reference distribution (i.e., exploration of large scale features), hence resulting in a multi-scale exploration behavior.  $\mathcal{W}(t)$  is the Fourier coefficients for the spatial statistics of the trajectory evolution  $\mathbf{x}(t)$  at time  $t$  (i.e., of  $C_t(\mathbf{x})$ ), which is given by

$$\mathcal{W}(t) = \frac{1}{t} \int_{\tau=0}^t \Phi(\mathbf{x}(\tau)) d\tau. \quad (3)$$

The ergodic control objective is  $\lim_{t \rightarrow \infty} \xi(t) = 0$ . For a dynamical system  $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$  where  $\mathbf{x} \in \mathbb{R}^d$ , we want the ergodic dynamics w.r.t. the evolution of the states  $\mathbf{x}(t) \in \Omega \subset \mathbb{R}^d$ . For infinitesimal control horizon, the solution for first-order systems is given as (see [3], [6] for details)

$$\dot{x}_i(t) = \alpha \frac{b_i(t)}{\|\mathbf{b}(t)\|},$$

$$\text{with } b_i(t) = \sum_{\mathbf{k} \in \mathcal{K}} \Lambda_{\mathbf{k}} (\mathcal{W}_{\mathbf{k}}(t) - \hat{\mathcal{W}}_{\mathbf{k}}) \nabla_i \Phi_{\mathbf{k}}(\mathbf{x}(t)),$$

$$\mathbf{b} = (b_1, \dots, b_d), \quad (4)$$

$$\nabla_i \Phi(\mathbf{x}(t)) = \phi(x_1) \circ \dots \circ \frac{\partial \phi(x_i)}{\partial x} \circ \dots \circ \phi(x_d), \quad (5)$$

where  $\alpha > 0$  is a small real number.

For a fully actuated system with point-mass dynamics  $\dot{\mathbf{x}} = \mathbf{u}$ , with  $\mathbf{u} = (u_1, \dots, u_d)$  and maximum velocity  $u_{\max}$ , the control commands  $u_i, i \in \{1, \dots, d\}$  that minimize the ergodic metric are given by

$$u_i(t) = u_{\max} \frac{b_i(t)}{\|\mathbf{b}(t)\|}.$$

Similar expressions for control policies are available for other types of dynamical systems, such as second-order point-mass systems (acceleration command), or first order and second order Dubin's car models, see e.g. [3], [4]. The controller can also be extended to other nonlinear dynamic systems, see e.g. [8], [7], [5]. We demonstrate our approach using the simple point-mass system, as it captures the key challenges in scaling the ergodic control algorithm to higher-dimensional exploration, and because it remains a classical choice for ergodic exploration [25], [4], [26].

<sup>5</sup>In the control loop, each of the variables  $\hat{\mathcal{W}}, \mathcal{W}(t), \Lambda, \nabla_i \Phi(\mathbf{x}(t))$  need  $\mathcal{O}(K^d)$  floating-point elements and each binary operation involving them has computational complexity  $\mathcal{O}(K^d)$ .

**Algorithm 1** Ergodic control algorithm

---

**Input:**  $d, L, K, u_{\max}, T$ , and  $P(\mathbf{x})$   
**Preprocessing:**  
 Compute  $\hat{\mathcal{W}}$  (evaluate  $K^d$  multivariate integrals with (1))  
 Compute  $\Lambda$  ( $K^d$  function evaluations)  
**Initialise:**  $t = 0$ ,  $dt$  (time step),  $\mathbf{x}(0)$ ,  $\mathcal{W}(0)$ ,  $\mathbf{u}(0)$   
 $\{\text{Control Loop}^5\}$   
**while**  $t < T$  **do**  
    $t \leftarrow t + dt$   
   Update  $\mathbf{x}(t)$   
   Update  $\mathcal{W}(t)$  (use numerical integration)  
   Compute  $\nabla_i \Phi(\mathbf{x}(t)), \forall i \in \{1, \dots, d\}$   
   **for**  $i=1, \dots, d$  **do**  
      $b_i(t) = \sum_{\mathbf{k} \in \mathcal{K}} \Lambda_{\mathbf{k}} (\mathcal{W}_{\mathbf{k}}(t) - \hat{\mathcal{W}}_{\mathbf{k}}) \nabla_i \Phi_{\mathbf{k}}(\mathbf{x}(t))$   
   **end for**  
   Update  $\mathbf{u}(t) = u_{\max} \frac{\mathbf{b}(t)}{\|\mathbf{b}(t)\|}$   
**end while**

---

The ergodic control algorithm (see Algorithm 1) is simple but it suffers from the curse of dimensionality when the dimension  $d$  of the exploration domain  $\Omega$  increases, which typically limits the use of the algorithm to problems of 2 or 3 dimensions. This is a drawback since many applications, in practice, are of dimensions  $d > 3$ . For example, the end-effector of a robot manipulator has 6 DOF (position and orientation). A naive implementation of the above algorithm for an exploration in the task space of a manipulator is then not feasible in practice.

The main challenges in the algorithm are listed below:

- 1) *Computation and storage of the Fourier series coefficients of the reference distribution  $\hat{\mathcal{W}}$ :*

This requires evaluation of the multidimensional integral in (1)  $K^d$  times to completely determine  $\hat{\mathcal{W}}$ . Although this is a preprocessing step, the computational complexity involved and the storage requirement make it infeasible for engineering applications.

- 2) *Computation and storage of  $\Lambda$ :*

This is a preprocessing step. The computation of each element  $\Lambda_{\mathbf{k}}$  is straightforward, but the number of function evaluations to find the complete tensor  $\Lambda$  and the required storage grow exponentially (i.e.,  $K^d$ ).

- 3) *Real-time implementation of the control loop:*

The control loop will be very slow as the algebraic operations (such as addition, element-wise product, summation, Frobenius norm of tensors) are more time-consuming as the order of the tensors involved in computing the control policy increases. So, a realtime implementation of the control loop may not be possible.

To give an example of the computation time involved in ergodic control, we used the Python software implementation of ergodic control described in [3]<sup>6</sup>. By using  $K = 10$  for the Fourier series coefficients and a spherical Gaussian of variance 0.01 at the center of  $\Omega$  with  $L = 1$  as the reference probability distribution, the preprocessing time to

compute the coefficients<sup>7</sup>  $\hat{\mathcal{W}}$  is approximately 16s in 2D, 5400s in 3D, and 2300000s in 4D ( $d = 4$ ) with the multivariate integration (1) evaluated using the Python package *scipy.integrate.nquad*.<sup>8</sup> Note that the number of elements to be stored in these cases is  $10^d$ , and for larger  $d$  (such as  $d > 6$ ) it is highly likely that memory/storage requirement for each of the multidimensional arrays involved exceeds the limit. Moreover, the average time taken per control loop in the above setting for  $d = 3, 4, 5, 6, 7$  are  $7 \times 10^{-4}$ s,  $1 \times 10^{-3}$ s,  $4 \times 10^{-3}$ s,  $6 \times 10^{-2}$ s, and  $8 \times 10^{-1}$ s correspondingly. For  $d = 2$ , [9] uses a supervised learning approach based on convolutional neural network for the fast computation of the Fourier coefficients. For higher dimensional problems, as we will see next, the above-mentioned challenges can be solved using the capability of a state-of-the-art tensor decomposition technique called tensor train (TT). Using the proposed solution, for the above reference distribution with  $d = 7$ ,  $\hat{\mathcal{W}}$  can be represented using approximately  $Kd$  parameters and it can be computed in less than  $2 \times 10^{-3}$ s and the average control loop takes less than  $1 \times 10^{-3}$ s. Also, for other reference distributions and larger  $d$ , the pre-processing step and the control loop can be processed very fast.

### C. Tensor Decomposition Techniques

Tensor decomposition (or tensor factorization) extends matrix decomposition techniques (e.g. matrix singular value decomposition) to multidimensional arrays. They allow tensors to be encoded compactly using a set of lower-order and/or lower-rank tensors (called factors). These factor elements are operated, depending on the decomposition technique, using various algebraic operations to represent elements of the given tensor. The accuracy of the representation is often controlled by the *rank* of the tensor decomposition. The rank of a tensor has different meanings for different tensor decomposition techniques. If the tensor has some low-rank structure (often due to separability arising from symmetry or smoothness or parallel proportionality), the number of elements required to represent the given tensor using the decomposition techniques will be far fewer than the original tensor.<sup>9</sup> In addition to compactly representing tensors, tensor decomposition allows efficient application of algebraic operations in the compressed format.<sup>10</sup>

Popular tensor decomposition techniques include CANDECOMP/PARAFAC (CP), also known as canonical polyadic, Tucker, hierarchical Tucker (HT), and tensor train (TT), see [27] for a survey. For  $d = 2$ , all these techniques reduce to the well known singular value decomposition (SVD) of matrices. They can be interpreted as higher-order extensions of

<sup>7</sup>Here, each coefficient is computed independently.

<sup>8</sup><https://docs.scipy.org/doc/scipy/reference/tutorial/integrate.html>

<sup>9</sup>For example, in matrix singular value decomposition, the low-rank corresponds to the case when many of the singular values are zero or negligible.

<sup>10</sup>This is analogous to algebraic operations on large matrices in their decomposition such as SVD, LU, etc. If we consider the right multiplication of a matrix  $\mathbf{A} \in \mathbb{R}^{m \times s}$  with  $\mathbf{B} \in \mathbb{R}^{s \times n}$ , and if we have a decomposition of  $\mathbf{A} = \mathbf{Q}\mathbf{R}$  with  $\mathbf{R} \in \mathbb{R}^{p \times n}$  and  $p \ll s$ , then computing  $\mathbf{C} = \mathbf{Q}\mathbf{R}\mathbf{B}$  using the decomposition of  $\mathbf{A}$  is more efficient than directly computing  $\mathbf{C} = \mathbf{A}\mathbf{B}$  using the raw  $\mathbf{A}$ . Similarly, for scalar multiplication of  $\mathbf{A}$ , computing  $\mathbf{C} = (k\mathbf{Q})\mathbf{R}$  is more efficient than  $\mathbf{C} = k\mathbf{A}$ .

<sup>6</sup>The configuration of computing system used is mentioned in Section V.



SVD, where these decomposition techniques differ for tensors of order  $d > 2$ . There exist many powerful algorithms to find the decomposition for each technique, most of them are based on the idea of alternating least squares (ALS). The CP decomposition of a tensor is formed by sum of rank-1 tensors. Although the number of parameters used by CP decomposition is small and scales well to higher order tensors, the space of CP tensors with a finite maximal rank is not closed. This may result in an ill-posed problem to find the best low-rank approximation. The Tucker decomposition overcomes this issue but it suffers from the curse of dimensionality, see [12] for details.

The TT decomposition [28] shares the good properties of the Tucker decomposition because the space of the tensors in TT format (with a fixed maximal rank) forms a smooth manifold, thus allowing robust algorithms to determine the TT decomposition of a tensor. It also shares good properties with CP, as it scales well to high-order tensors. We exploit the tensor train (TT) decomposition in our approach, as it gathers the good properties of both CP and Tucker decomposition. In the next section, a brief overview of TT decomposition is given.

#### D. Tensor Train Decomposition

A  $d$ -th order tensor  $\mathcal{G} \in \mathbb{R}^{K_1 \times \dots \times K_d}$  in TT format is represented using a tuple of  $d$  third-order tensors  $(\mathcal{G}^1, \dots, \mathcal{G}^d)$ .

Here,  $\mathcal{G}^i \in \mathbb{R}^{r_{i-1} \times r_i \times K_i}$ ,  $i \in \{2, \dots, d-1\}$ ,  $\mathcal{G}^1 \in \mathbb{R}^{1 \times r_1 \times K_1}$  and  $\mathcal{G}^d \in \mathbb{R}^{r_{d-1} \times 1 \times K_d}$ . As shown in Fig. 2, the  $\mathbf{k}$ -th element, with  $\mathbf{k} \in \mathcal{K} = \{(k_1, \dots, k_d) : k_i \in \{1, \dots, K_i\}, i \in \{1, \dots, d\}\}$ , is given by

$$\mathcal{G}_{\mathbf{k}} = \mathcal{G}_{:::,k_1}^1 \mathcal{G}_{:::,k_2}^2 \cdots \mathcal{G}_{:::,k_d}^d \quad (6)$$

where  $\mathcal{G}_{:::,k_i}^i \in \mathbb{R}^{r_{i-1} \times r_i}$  represents the  $k_i$ -th frontal slice (a matrix) of the third order tensor  $\mathcal{G}^i$ . The *TT-rank* of the tensor in TT representation is then defined as the tuple  $\mathbf{r} = (r_1, r_2, \dots, r_{d-1})$ . We call  $r = \max(r_1, \dots, r_{d-1})$  as the *maximal TT rank*. For any given tensor, there always exists a TT decomposition (6).

The TT decomposition of a tensor can be considered as a particular way of writing the elements of a given tensor as a finite sum of separable products [27]. To explain this intuitively, consider a multivariate function  $g: \Omega_g \subset \mathbb{R}^d \rightarrow \mathbb{R}$  such that  $g(\mathbf{x}) = g(x_1, \dots, x_d)$  which can be approximated as a finite sum of products of one dimensional functions ( $g_{\hat{r}_i}^i: \mathbb{R} \rightarrow \mathbb{R}$ ,  $i \in \{1, \dots, d\}$  and  $\hat{r}_i \in \{1, \dots, r_i\}$ ,  $r_i \in \mathbb{Z}^+$ ), namely

$$g(x_1, \dots, x_d) \approx \sum_{\hat{r}_1=1}^{r_1} \cdots \sum_{\hat{r}_d=1}^{r_d} g_{\hat{r}_1}^1(x_1) \cdots g_{\hat{r}_d}^d(x_d). \quad (7)$$

Suppose we have a discretization of the domain  $\Omega_g \subset \mathbb{R}^d$  given by the set  $\{\mathbf{y} = (y_{k_1}, \dots, y_{k_d}) : y_{k_i} \in \mathbb{R}, k_i \in \{1, \dots, K_i\}, K_i \in \mathbb{Z}^+\}$  with the corresponding index set  $\mathcal{K}$  as define above. Let  $\mathcal{G}$  be the tensor formed by evaluating the function  $g$  at these discretization points (i.e.,  $\mathcal{G}_{\mathbf{k}} = g(y_{k_1}, \dots, y_{k_d}), \forall \mathbf{k} = (k_1, \dots, k_d) \in \mathcal{K}$ ). Then, the TT




Fig. 2: An element of a tensor in TT-format can be accessed by multiplying the selected slices (matrices represented in red color) of the core tensors (factors). The figure depicts an example for a 4th order tensor  $\mathcal{G} \in \mathbb{R}^{5 \times 6 \times 7 \times 8}$  of rank  $\mathbf{r} = (2, 3, 4)$ .

decomposition of  $\mathcal{G}$  is analogous to the representation (7) of the underlying multivariate function  $g$ , and it is given by

$$\mathcal{G}_{\mathbf{k}} = \sum_{\hat{r}_1=1}^{r_1} \cdots \sum_{\hat{r}_d=1}^{r_d} \mathcal{G}_{1,\hat{r}_1,k_1}^1 \cdots \mathcal{G}_{\hat{r}_{d-1},1,k_d}^d, \quad \forall \mathbf{k} \in \mathcal{K},$$

which is compactly written in (6) using matrix multiplication.

Let  $r$  be the maximal TT-rank and  $K = \max(K_1, \dots, K_d)$ . Then, the number of elements in the TT representation is  $\mathcal{O}(Kdr^2)$  as compared to  $\mathcal{O}(K^d)$  elements in the original tensor. For small  $r$ , the representation is thus very efficient. In addition to storage efficiency, many numerical algebraic operations on tensors such as addition, scalar multiplication, Hadamard (or element wise) product, Frobenius norm can be directly and efficiently applied in TT representation<sup>11</sup>. These operations are explained in [28], [29]. Most algebraic operations on tensors in TT format have computational complexity linear in  $d$  and  $K$ , and polynomial (often quadratic or cubic) in  $r$ . As we will see in Section IV, the proposed algorithm for ergodic control in this paper will exploit this efficiency in algebraic operation offered by TT format to speed up the computation in the control loop.

As explained above, the existence of low-rank structure (i.e., a low maximal TT-rank) of a given tensor is closely related to the separability of the underlying multivariate function. Although separability of functions is not a very well understood problem, it is known that smoothness<sup>12</sup> and symmetry of functions often induces better<sup>13</sup> separability of the functions. This has been exploited to solve PDEs and problems involving high-dimensional integration [27], and it is used in this paper to solve the challenges in ergodic control.

<sup>11</sup>For example, to gain an intuition about the efficiency of algebraic operations in TT format, consider two  $d$ -th order tensors of same shape  $\mathcal{G}$  and  $\mathcal{H}$  with known TT representations. The TT cores of the tensor  $\mathcal{S} = \mathcal{G} + \mathcal{H}$  can be directly obtained as  $\mathcal{S}_{:::,k}^i = \text{diag}(\mathcal{G}_{:::,k}^i, \mathcal{H}_{:::,k}^i)$  for  $i \in \{2, \dots, (d-1)\}$  and  $\mathcal{S}_{:::,k}^j = \text{concatenate}(\mathcal{G}_{:::,k}^j, \mathcal{H}_{:::,k}^j)$  for  $j \in \{1, d\}$ , where *diag* is the block diagonal matrix construction and *concatenate* represents vector concatenation operation. Also, the TT cores of the tensor formed by scalar multiplication  $\mathcal{Q} = c \cdot \mathcal{G}$  can be obtained by multiplying the scalar only to the first core:  $\mathcal{Q}^1 = c \cdot \mathcal{G}^1$  and  $\mathcal{Q}^i = \mathcal{G}^i$  for  $i \in \{2, \dots, d\}$ .

<sup>12</sup>By smoothness, we mean the degree of variation of the function across its domain. For example, a probability density function in the form of Gaussian mixture model (GMM) is considered to become less smooth as the number of mixture components (i.e., multi-modality) increases or the variance of the component Gaussians decreases (i.e., sharper peaks). More formally, the degree of smoothness can be defined using the properties of higher-order derivatives.

<sup>13</sup>By *better*, we mean fewer low-dimensional functions in the sum of products representation.

### Finding TT Decompositions

There exist many algorithms to find the TT decomposition of a tensor. The popular methods are TT-SVD, TT-DMRG, and TT-cross. TT-SVD and TT-DMRG, like matrix SVD, require the full tensor in memory to find the decomposition, and hence they are infeasible for higher order tensors. TT cross approximation (TT-cross) [30][31] is an extension of the cross approximation technique (also called CUR or skeleton decomposition) in matrix theory [32], see Fig. 3. It is appealing for many practical problems as it approximates the given tensor with controlled accuracy, by using only a small number of its elements. When applied to tensors, the cross approximation method needs access to only certain fibers of the original tensor at a time, and hence works as a black box method.

Given a function (or a procedure)  $g$  that evaluates an element of a tensor  $\mathcal{G}$  given its index and an approximation accuracy  $\epsilon$  in the Frobenius norm, TT-cross returns an approximate tensor in TT format  $\hat{\mathcal{G}} = \text{TT-cross}(g, \epsilon)$  to the tensor  $\mathcal{G}$  by querying only a portion of its elements. The expected error in approximation is less than the accuracy  $\epsilon$  specified.

Thus, TT-cross avoids the need to store explicitly the original tensor, which may not be possible for higher order tensors. It can approximate a  $d$ -th order tensor, with maximal TT-rank  $r$  and size of each mode  $K$ , using only  $\mathcal{O}(Kdr^2)$  samples and  $\mathcal{O}(Kdr^3)$  operations (flops). This is very efficient if the TT-rank  $r$  of the tensor is low, which is typically the case in many engineering applications, including robotics. In this paper, we will use TT-cross to find TT decompositions.

### TT-rounding

TT-rounding [28] is an important operation on a tensor in TT format. Most binary operations on tensors in TT format, although efficient, result in an increase in the TT-rank of the resultant tensor, where the resultant tensor is often not in its optimal TT representation. A repeated application of binary operations to a given tensor may result in an explosion of its TT-rank, which would effect the efficiency of subsequent operations on the tensor. For example, the addition of two tensors, both with TT-rank  $\mathbf{r} = (r_1, \dots, r_d)$ , results in a tensor in TT format with rank  $\mathbf{r} = (2r_1, \dots, 2r_d)$ . TT-rounding is an operation applied to tensors already in TT format to compress it to optimal TT representation and hence reduces its TT-rank. For a  $d$ -th order tensor  $\mathcal{G}$  in TT-format with maximal TT-rank  $r$ , TT-rounding has computational complexity  $\mathcal{O}(Kdr^3)$ . The TT-rounding procedure returns a tensor  $\hat{\mathcal{G}} = \text{TT-round}(\mathcal{G}, \hat{r})$ , for a given  $\hat{r} < r$ , such that its maximal TT-rank is less than  $\hat{r}$  and the Frobenius norm of the residual  $\|\hat{\mathcal{G}} - \mathcal{G}\|$  is as small as possible. Alternatively, we can specify an approximation accuracy  $\epsilon$  to a tensor  $\mathcal{G}$  in TT format and the TT-rounding returns a tensor  $\hat{\mathcal{G}} = \text{TT-round}(\mathcal{G}, \epsilon)$  with optimal TT-rank and  $\|\hat{\mathcal{G}} - \mathcal{G}\| \leq \epsilon$ .

## IV. ALGORITHM DESCRIPTION

In this section, we give details of the solution proposed in this paper to overcome the challenges mentioned in Section III-B. Additionally, as part of our proposed strategy for 6D exploration in manipulation tasks, we propose a Riemannian




Fig. 3: TT-cross is an extension of skeleton decomposition in matrix theory. A rank- $r$  matrix  $A \in \mathbb{R}^{m \times n}$  can be exactly recovered if we know  $r$  independent columns of  $A$  indexed by  $j = (j_1, \dots, j_r)$ ,  $A_{:,j} \in \mathbb{R}^{m \times r}$  and  $r$  independent rows of  $A$  indexed by  $i = (i_1, \dots, i_r)$ ,  $A_{i,:} \in \mathbb{R}^{r \times n}$  of matrix  $A$ , with their intersection  $A_{i,j} \in \mathbb{R}^{r \times r}$  being nonsingular. Then, by skeleton decomposition we have  $A = A_{:,j} A_{i,j}^{-1} A_{i,:}$ .

manifold extension to allow ergodic exploration for orientation data represented as unit quaternions.

We use the TT representation for the variables involved in the algorithm, namely  $\hat{\mathcal{W}}, \mathcal{W}(t), \Lambda, \nabla_i \Phi(x(t))$ .<sup>14</sup> These variables, as described below, can be compactly represented in TT format, within a desired accuracy. Since all the operations to find the control policy, at each step of the control loop, are done on variables in TT format, the computational complexity is significantly reduced.

$\Lambda$  can be found using the TT-cross approximation. Because of the involved symmetry (see the definition of  $\Lambda$ ), the resultant tensor in TT format is of very low rank. Maximal TT rank of 2 accurately captures the tensor with error less than  $10^{-2}$  in Frobenius norm and it can be computed in a fraction of a second (for  $d < 15$ ). We also represent  $\mathcal{W}(t)$  in TT format and use time integration in TT format [33] to compute it efficiently at each iteration of the control loop. Due to integration, as the TT-rank of this tensor may increase in an unbounded manner over time, we specify an upper bound (a hyperparameter) on its TT-rank. If other integration schemes are used, one can periodically use TT-rounding to cut off the TT-rank.

Finding  $\hat{\mathcal{W}}$  is a preprocessing step for the algorithm, and it is the most challenging part as it requires to evaluate  $K^d$  times a  $d$ -dimensional integral given by (1), if implemented naively. By using the properties of the TT format, the Fourier Coefficients can be computed in the following ways:

*Method 1: By exploiting the properties of Gaussian mixture model (GMM) and the efficiency of TT-cross for reference distributions in the form of a GMM.*

By using TT-cross the above integral needs to be computed only at the query points (Fourier coefficients) of the TT-cross, thus we obtain a significant savings in the number of evaluations of the multidimensional integral. This will require the integral in (1) to be evaluated  $\mathcal{O}(Kdr^2)$  times and  $r$  is often small due to the structure of  $\hat{\mathcal{W}}$ . However, computation of each multidimensional integral is still time consuming.

<sup>14</sup>By its definition in (5),  $\nabla_i \Phi(x(t)), \forall i \in \{1, \dots, d\}$  can be represented as a rank-1 TT. Its TT cores are  $(\phi(x_1), \dots, \phi(x_{i-1}), \frac{\partial \phi(x_i)}{\partial x}, \phi(x_{i+1}), \dots, \phi(x_d))$ . Similarly,  $\Phi(x)$  by its definition can be represented as a rank-1 TT, with TT cores  $(\phi(x_1), \dots, \phi(x_d))$

For reference distributions in the form of a GMM, we have presented an analytical expression for the above integral in [6]. This avoids computation of the multi-dimensional integration using numerical schemes for distributions in the form of GMM for reasonably<sup>15</sup> large  $d$  as we use the analytical expression to compute the Fourier coefficients at the query points of TT-cross. Thus the method directly exploits the structure (i.e., smoothness) of the Fourier coefficients in  $\hat{\mathcal{W}}$ .

*Method 2: By exploiting the properties of TT decomposition for arbitrary reference distributions.*

We provide a solution to the Fourier coefficients  $\hat{\mathcal{W}}$  without having to perform any multidimensional integration directly. By exploiting the properties of Gaussian quadrature rule for the integration of multivariate functions, we derive an analytical expression for the Fourier coefficients  $\hat{\mathcal{W}}$  in the TT representation for arbitrary functions. In this method, we exploit the smoothness of the reference probability distribution  $P(\mathbf{x})$  to find  $\hat{\mathcal{W}}$  indirectly by evaluating  $P(\mathbf{x})$  at a few points in its domain. This will be the method we use in the rest of the paper for finding  $\hat{\mathcal{W}}$  unless otherwise stated. This method to find Fourier coefficients is very efficient for smooth high-dimensional functions, which will be explained in the next section.

Note that, in practice, the reference distributions are smooth and/or the Fourier coefficients vary smoothly across their indices. Hence, a low-rank TT representation of  $\hat{\mathcal{W}}$  is expected. For GMM with very sharp peaks, one may prefer Method 1 as it relies directly on the structure of  $\hat{\mathcal{W}}$  and hence it could be faster for such non-smooth reference distributions in the form of GMM. However, the computational complexity involved in evaluating the analytical expression for each Fourier series coefficient in Method 1 grows in proportion to  $2^d$  and linearly with the number of mixture components in the GMM [6]. Moreover, Method 1 requires the GMM to have negligible tail outside the domain of exploration  $\Omega$ . Since Method 2 applies to arbitrary reference distributions and it is observed in practice to be equally fast compared to Method 1 for smooth GMM as reference distribution, in the remaining sections we consider only this method to find the Fourier series coefficients  $\hat{\mathcal{W}}$ .

### A. Finding the Fourier Series Coefficients

In this section, we provide an analytical expression for the Fourier coefficients of a smooth but arbitrary distribution introduced in method 2 of Section IV. The proof is inspired by [34] where they have used separability structure in TT representation to find polynomial approximations of multivariate functions. A similar strategy has been used in [35] to evaluate high-dimensional integration of smooth functions. We use this strategy to find an analytical expression for the Fourier coefficients of functions directly in TT representation. The proof relies on quadrature rules (see Appendix A for the details) for numerical integration of multivariate functions. There are many possible options to choose the quadrature rule depending on the type of function  $P(\mathbf{x})$  to be integrated.

<sup>15</sup>Here, we consider  $d$  approximately up to 10 as the computation complexity of each Fourier coefficient using the analytical solution provided in [6] grows in proportion to  $2^d$ . For details, refer to [6].

The following result applies to any choice of quadrature rule, however, for simplicity, we use Gaussian quadrature rule (G-Q) in this paper.

The idea is to find the TT representation  $\mathcal{P}$  of the multivariate function  $P(\mathbf{x})$  evaluated at the discretization induced from the quadrature rule. Then, as we will see below, the TT representation of  $\hat{\mathcal{W}}$  can be obtained directly using  $\mathcal{P}$ .

Let  $x_j \in \mathbb{R}$  be the discretization points of the interval  $[0, L]$  and  $\alpha_j$  be the weights obtained from the quadrature rule, where  $j \in \{1, \dots, N\}$ , with  $N$  representing the specified degree of approximation of the function. Then, we can discretize the domain  $\Omega$  at  $\mathbf{x}_j = (x_{j_1}, \dots, x_{j_d})$ , with  $\mathbf{j} \in \mathcal{J}$ , where  $\mathcal{J} = \{\mathbf{j} = (j_1, \dots, j_d) : j_i \in \{1, \dots, N\}\}$  is the index set. Let  $\mathcal{P}$  be the tensor formed by evaluating the reference distribution  $P(\mathbf{x})$  at the discretization points, so that  $\mathcal{P}_{\mathbf{j}} = P(\mathbf{x}_{\mathbf{j}}), \forall \mathbf{j} \in \mathcal{J}$ .

Let  $(\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^d)$  be the TT cores of  $\mathcal{P}$  in its TT representation<sup>16</sup>, so that for  $\mathbf{j} = (j_1, \dots, j_d) \in \mathcal{J}$  we have,

$$\mathcal{P}_{\mathbf{j}} = \mathcal{P}_{:::,j_1}^1 \mathcal{P}_{:::,j_2}^2 \cdots \mathcal{P}_{:::,j_d}^d,$$

then the TT cores of  $\hat{\mathcal{W}}$  are (see Appendix A for the proof)

$$\hat{\mathcal{W}}_{:::,k}^i = \sum_{j=1}^N \alpha_j \mathcal{P}_{:::,j}^i \phi_k(x_j), \quad \forall k \in \{1, \dots, K\}, \quad \forall i \in \{1, \dots, d\}, \quad (8)$$

so that

$$\hat{\mathcal{W}}_{\mathbf{k}} = \hat{\mathcal{W}}_{:::,k_1}^1 \cdots \hat{\mathcal{W}}_{:::,k_d}^d, \quad \forall \mathbf{k} \in \mathcal{K}.$$

Thus, we can compute the Fourier coefficients  $\hat{\mathcal{W}}$  by only investing in computing the TT decomposition  $\mathcal{P}$  of the discretized reference distribution. This can be done in  $\mathcal{O}(Ndr^2)$  function evaluations of  $P(\mathbf{x})$  using the TT-cross algorithm. The TT-rank of the tensor  $\hat{\mathcal{W}}$  will be same as that of the TT-rank of  $\mathcal{P}$ . For a smooth reference distribution,  $\mathcal{P}$  will have low TT-rank. This is a tremendous saving in computing the Fourier coefficients  $\hat{\mathcal{W}}$ , and thus it overcomes the curse of dimensionality. The TT based algorithm for ergodic control is outlined in Algorithm 2.<sup>17</sup>

### B. Ergodic Control on Riemannian Manifolds

Most manipulation tasks concern the full robot end-effector pose, which includes both its position and orientation. Hence, when designing exploration strategies for manipulation it is desirable to consider both. In the case of position, the ergodic control formulation in Section IV-A can be directly applied. However, since orientation data do not lie on a Euclidean space, exploration in orientation requires a special mathematical treatment. In this section, we extend ergodic control to handle data on a Riemannian manifold [22], [23], particularly the orientation manifold  $\mathcal{S}^3$ .

<sup>16</sup>This can be obtained using TT-cross:  $\mathcal{P} = \text{TT-Cross}(P(\mathbf{x}), \epsilon)$ .

<sup>17</sup>Our implementation of Algorithm 2 in Python can be found at <https://sites.google.com/view/ergodic-exploration/>.

<sup>18</sup>In the control loop, the memory of each variable and the computational complexity of each algebraic operation has complexity grow linearly with  $d$ . Thus it avoids the curse of dimensionality. In particular, the computation of  $b_i(t)$  requires a subtraction, a Hadamard product and an inner product involving tensors in TT format, hence it can be computed efficiently.



---

**Algorithm 2** Ergodic Control using TT
 

---

**Input:**  $d, L, K, u_{\max}, T$ , and  $P(\mathbf{x})$   
**Pre-Processing:**  
 Compute  $\mathcal{P}$  using TT-cross  
 Find  $\hat{\mathcal{W}}$  using (8)  
 Compute  $\mathbf{\Lambda}$  using TT-cross  
 TT-rounding of  $\hat{\mathcal{W}}$  (remove low-energy contents)  
**Initialise:**  $dt$  (time step),  $\mathbf{x}(0)$ ,  $\mathcal{W}(0)$  in TT format,  $\mathbf{u}(0)$ , and  $t = 0$   
 $\{\text{Control Loop}^{18}\}$   
**while**  $t < T$  **do**  
    $t \leftarrow t + dt$   
   Update  $\mathbf{x}(t)$  (time integration)  
   Update  $\mathcal{W}(t)$  (Use TT time integration [33] with a fixed maximal TT-rank. Alternatively, use numerical integration such as Euler integration followed by TT-rounding)  
  
   Compute  $\nabla_i \Phi(\mathbf{x}(t)), i \in \{1, \dots, d\}$  (rank-1 TT)  
    $\{\text{Using algebraic operations in TT}\}$   
   **for**  $i=1, \dots, d$  **do**  
      $b_i(t) = \sum_{\mathbf{k} \in \mathcal{K}} \mathbf{\Lambda}_{\mathbf{k}}(\mathcal{W}_{\mathbf{k}}(t) - \hat{\mathcal{W}}_{\mathbf{k}}) \nabla_i \Phi_{\mathbf{k}}(\mathbf{x}(t))$   
   **end for**  
   Update  $\mathbf{u}(t) = u_{\max} \frac{\mathbf{b}(t)}{\|\mathbf{b}(t)\|}$   
**end while**

---

The orientation of the robot end-effector can be represented by a unit quaternion  $\mathbf{q} \in \mathcal{S}^3$ , comprised of a scalar part  $q_s \in \mathbb{R}$  and a vector part  $\mathbf{q}_v \in \mathbb{R}^3$  such that  $\mathbf{q} = [q_s \ \mathbf{q}_v^\top]^\top$ . For any point on the manifold  $\mathbf{g} \in \mathcal{S}^3$  there exists a tangent space  $\mathcal{T}_{\mathbf{g}}\mathcal{S}^3$  in which standard Euclidean methods can be applied to orientation. The function that maps a quaternion  $\mathbf{q}$  from the manifold to a tangent space is called the *logarithmic map* and is given by

$$\text{Log}(\mathbf{q}) = \begin{cases} \text{acos}_*(q_s) \frac{\mathbf{q}_v}{\|\mathbf{q}_v\|}, & q_s \neq 1 \\ [0, 0, 0]^\top, & q_s = 1 \end{cases}, \quad (9)$$

where  $\text{acos}_*(\cdot)$  is a modified arc-cosine function [22]. Equation (9) maps  $\mathbf{q}$  to the tangent space of the manifold origin. The mapping of  $\mathbf{q}$  to the tangent space of an arbitrary point  $\mathbf{g}$  is given by [22]

$$\text{Log}_{\mathbf{g}}(\mathbf{q}) = \text{Log}(\bar{\mathbf{g}} * \mathbf{q}), \quad (10)$$

where  $(\bar{\cdot})$  and  $*$  denote the quaternion conjugate and product, respectively. The logarithmic map represents a unit quaternion  $\mathbf{q}$  as a 3-dimensional Euclidean vector  $\mathbf{v} \in \mathbb{R}^3$ . Quaternions can be retrieved from the tangent space through the *exponential map*

$$\text{Exp}(\mathbf{v}) = \begin{cases} \left[ \cos(\|\mathbf{v}\|), \sin(\|\mathbf{v}\|) \frac{\mathbf{v}^\top}{\|\mathbf{v}\|} \right]^\top, & \|\mathbf{v}\| \neq 0 \\ [1, 0, 0, 0]^\top, & \|\mathbf{v}\| = 0 \end{cases}, \quad (11)$$

which, analogously to (10), maps from an arbitrary tangent space  $\mathcal{T}_{\mathbf{g}}\mathcal{S}^3$  back to the manifold through

$$\text{Exp}_{\mathbf{g}}(\mathbf{v}) = \mathbf{g} * \text{Exp}(\mathbf{v}). \quad (12)$$

Given a set of unit quaternions (e.g. obtained from demonstrations), we formulate orientation-ergodic control by modeling their distribution in the tangent space of their mean. For  $M$  end-effector orientations  $\{\mathbf{q}_i\}_{i=1}^M$ , the mean on the manifold  $\boldsymbol{\mu} \in \mathcal{S}^3$  is computed iteratively with (see [22], [23] for details)

$$\mathbf{v} = \frac{1}{M} \sum_{i=1}^M \text{Log}_{\boldsymbol{\mu}}(\mathbf{q}_i), \quad \boldsymbol{\mu} \leftarrow \text{Exp}_{\boldsymbol{\mu}}(\mathbf{v}). \quad (13)$$

All quaternions in the dataset can thus be mapped to the tangent space of the mean through  $\{\mathbf{v}_i\}_{i=1}^M = \{\text{Log}_{\boldsymbol{\mu}}(\mathbf{q}_i)\}_{i=1}^M$ , allowing the proposed ergodic exploration (Algorithm 2) to be performed in  $\mathbb{R}^3$ , even for orientation. As desired orientations are computed, in the tangent space, at each time step by  $\hat{\mathbf{v}}(t) = \mathbf{v}(t) + \mathbf{u}(t)dt$ , the exponential map generates a desired unit quaternion for the robot to track, using

$$\hat{\mathbf{q}}(t) = \text{Exp}_{\boldsymbol{\mu}}(\hat{\mathbf{v}}(t)). \quad (14)$$

In this way, ergodic control for end-effector poses is formulated as a 6D problem, where the first 3 dimensions correspond to position and the last 3 to orientation.

## V. NUMERICAL EVALUATION

In this section, we demonstrate the computational efficiency of the TT-based algorithm for ergodic control through simulations. We use Method 2 described in Section IV to compute the Fourier series coefficients. The simulations are performed on a Lenovo Thinkpad personal computer with Intel(R) Core(TM) i7-8565U CPU at 1.80GHz with 24GB RAM. We use ttpy, a Python-based toolbox for working with TT.<sup>19</sup>

A naive implementation without using tensor decomposition techniques would require  $K^d$  elements to store each of the tensors:  $\hat{\mathcal{W}}$ ,  $\mathbf{\Lambda}$ , and  $\nabla_i \Phi(\mathbf{x})$  ( $i \in \{1, \dots, d\}$ ). However, a TT representation requires less than  $4Kd$  elements for  $\mathbf{\Lambda}$  (with approximation error  $10^{-2}$  in the Frobenius norm) and only  $Kd$  elements to exactly represent  $\nabla_i \Phi(\mathbf{x})$ . As  $\nabla_i \Phi(\mathbf{x})$  is a rank-1 tensor with explicit analytical expressions for its TT cores, it can be computed very fast. Computing the weights  $\mathbf{\Lambda}$  can be done in a fraction of a second for  $d \leq 15$ .

The computation and storage of  $\hat{\mathcal{W}}$  depends on the smoothness of the reference probability distribution. For the evaluation, we define our reference distribution as an isotropic Gaussian distribution at the centre of the domain with variance 0.015, where we used  $L = 1$ ,  $K = 5$ ,  $N = 10$  and an approximation accuracy of  $10^{-2}$  in the TT representation of  $\hat{\mathcal{W}}$  and  $\mathbf{\Lambda}$ . As can be seen in Fig. 4, the time taken to compute  $\hat{\mathcal{W}}$  grows approximately linearly with  $d$ , and it is less than a second for the chosen reference distribution with  $d \leq 10$  and the average time taken per control loop increases almost linearly with the number of dimensions  $d$ , whereas with a naive implementation (without using TT) the time taken in the control loop grows exponentially with  $d$ . The trend remains the same for other reference distributions, see Table II.

The computation of  $\mathcal{W}(t)$  in the control loop requires some attention. At each iteration of the control loop, the rank of the tensor  $\mathcal{W}(t)$  may keep increasing due to the integration, see

<sup>19</sup><https://github.com/oseledets/ttpy>

		Gaussian mixture model						Uniform distribution	
		2 components		4 components		6 components		With TT	Without TT
		With TT	Without TT	With TT	Without TT	With TT	Without TT		
<b>5D</b>	$\nabla_i \Phi$	50	$10^5$	50	$10^5$	50	$10^5$	50	$10^5$
	# parameters $\Lambda$	160	$10^5$	160	$10^5$	160	$10^5$	160	$10^5$
	$\hat{\mathcal{W}}$	160	$10^5$	548	$10^5$	1032	$10^5$	50	$10^5$
	Average time per loop	$2 \times 10^{-3}$	$4 \times 10^{-3}$	$3 \times 10^{-3}$	$4 \times 10^{-3}$	$3.6 \times 10^{-3}$	$4 \times 10^{-3}$	$2 \times 10^{-3}$	$4 \times 10^{-3}$
	Pre-processing time	0.2	—	0.87	—	2.4	—	$33 \times 10^{-3}$	—
<b>6D</b>	$\nabla_i \Phi$	60	$10^6$	60	$10^6$	60	$10^6$	60	$10^6$
	# parameters $\Lambda$	200	$10^6$	200	$10^6$	200	$10^6$	200	$10^6$
	$\hat{\mathcal{W}}$	200	$10^6$	695	$10^6$	1431	$10^6$	60	$10^6$
	Average time per loop	$3.2 \times 10^{-3}$	$63 \times 10^{-3}$	$4.6 \times 10^{-3}$	$63 \times 10^{-3}$	$5.4 \times 10^{-3}$	$63 \times 10^{-3}$	$3 \times 10^{-3}$	$63 \times 10^{-3}$
	Pre-processing time	$30 \times 10^{-3}$	—	1.26	—	3.6	—	$40 \times 10^{-3}$	—
<b>7D</b>	$\nabla_i \Phi$	70	$10^7$	70	$10^7$	70	$10^7$	70	$10^7$
	# parameters $\Lambda$	240	$10^7$	240	$10^7$	240	$10^7$	240	$10^7$
	$\hat{\mathcal{W}}$	233	$10^7$	860	$10^7$	1801	$10^7$	70	$10^7$
	Average time per loop	$4.8 \times 10^{-3}$	0.8	$6.8 \times 10^{-3}$	0.8	$7.5 \times 10^{-3}$	0.8	$3.6 \times 10^{-3}$	0.8
	Pre-processing time	$35 \times 10^{-3}$	—	1.53	—	4.9	—	$43 \times 10^{-3}$	—

TABLE II: Computational speed and storage requirements in ergodic control for different reference probability distributions with  $K = 10$  and  $L = 1m$ . The components of GMM are spherical Gaussians with 0.005 variance. All the tensors in TT format are approximated with an accuracy of  $10^{-2}$  in the Frobenius norm. The preprocessing time for the naive approach (without using TT) is not given in the table as it is computationally infeasible in the computing device used for the experiment.




Fig. 4: Time taken (linear scale) to compute the Fourier coefficients  $\hat{\mathcal{W}}$  (left) and the average time taken (in log-10 scale) per control loop (right) of the ergodic control algorithm using the proposed technique for  $K = 5$ ,  $l = 1$ , and a reference distribution in the form of an isotropic Gaussian with variance 0.015. In the right figure, the exponential growth in the computational complexity in the control loop can be observed for the standard approach (without using tensor decomposition), whereas the proposed approach avoids the curse of dimensionality. In the left figure, it can be observed that the computational complexity in computing the Fourier series coefficients using the proposed approach tends to grow linearly with  $d$ .

(3). This could be a problem if the time period of ergodic exploration is very high. So, it is necessary to upper bound the TT-rank of this tensor using TT-rounding with a specified maximal TT-rank. Setting an excessively low value for the maximal rank may lead to convergence issues and specifying a large value for maximal rank slows down the speed of computation of each control loop. Thus, this hyperparameter must be chosen carefully. In the numerical evaluation, our experiments revealed that a maximal rank of  $d \cdot r$  for  $\mathcal{W}(t)$ , where  $r$  is maximal TT-rank of  $\hat{\mathcal{W}}$ , worked well for  $d < 15$ .

The TT representation allows compact representation of the tensors involved and the storage complexity also grows linearly with  $d$ . These properties allow our algorithm to be implemented for real-time applications on devices with small memory and limited computational power, which are often the case in robotic systems. Another important property of our algorithm is that the approximation of the tensors involved such as  $\hat{\mathcal{W}}$ ,  $\mathcal{W}(t)$  and  $\Lambda$  can be controlled precisely using TT-rounding. In practice, TT-rounding with accuracy

$10^{-2}$  is observed to be sufficient for all practical purposes considered here. Furthermore, doing these approximations in the spectral domain results in negligible impact on the time domain behaviour of the system, thanks to Parseval's theorem as the approximation can be considered as a noise filter in the ergodic motion. This also provides a convenient trade-off between accuracy of approximation in ergodic exploration and the speed of computation in the control loop.

## VI. EXPERIMENT: SENSORLESS PEG-IN-HOLE INSERTION USING ERGODIC EXPLORATION

We evaluate the proposed approach in an insertion task. We formulate the insertion task as a 6D exploration problem where we simultaneously address the uncertainty about the insertion pose in position (3D) and orientation (3D) in the robot task space. Our method is well suited for peg-in-hole insertion tasks where uncertainties may arise from several sources, including variable grasps of the peg, unprecise locations of the hole, and unmodeled manufacturing defects of the involved components (gripper fingers, pegs and holes). In the considered experiment, the reference probability distribution for exploration is found using information from human demonstrations. The human demonstrates the key regions for exploration in the state space of the end-effector and we use a Gaussian mixture model (GMM) to model the reference probability distribution based on the datapoints collected during the demonstrations. As a means to intuitively show the effectiveness of our approach, we begin by comparing it to three baselines of exploratory behaviors commonly used in insertion tasks. For this we use a toy example in 2D and 3D.

### A. Simulation experiments

In this section we provide the motivation for using ergodic control for exploration, and its significance for insertion tasks, using simulation of exploration behavior in 2D and 3D.

We use a GMM as the reference distribution in the space  $\Omega$  with  $L = 1\text{m}$ . The GMM is chosen such that it has 6 equally weighted mixture components with its centers placed randomly in the exploration space  $\Omega$  and each component is a spherical Gaussian with variance 0.01. As a first metric, we compute the average time taken to reach a spherical region with volume 0.5% of the volume of  $\Omega$ . The spherical region is representative of the target detection region during exploration. For the insertion task, this corresponds to the set of end-effector states at which the peg is inside the hole. For all the trials, we fix a maximum duration of 1000s for detection (i.e., reaching the target region) and the magnitude of the point-mass-system velocity is constant and fixed to  $u_{\max} = 0.1\text{ m/s}$ . For the analysis we choose 10 different GMMs as described above and for each choice of GMM, we conduct 10 trials. For each trial, the center of the target region is chosen by sampling in  $\Omega$  from the reference GMM and the point-mass system starts with the same initial state:  $(0.5, 0.5)$  for 2D and  $(0.5, 0.5, 0)$  for 3D.

We compare four different exploration strategies:

- 1) **Strategy 1:** *Ergodic exploration* (proposed approach),
- 2) **Strategy 2:** *Sampling-based exploration*
- 3) **Strategy 3:** *Cylindrical spiral*, as a representative of sweeping patterns,
- 4) **Strategy 4:** *Mixture of ellipsoidal spirals*, as a sweeping pattern customized for GMM.

Fig. 5 shows an example of exploration behaviors for these different strategies.

In strategy 2, we explore by tracking samples from the GMM sequentially. In this approach, the simulated system tracks GMM samples using a constant speed, with a new sample being computed every time the previous one is reached. Unlike ergodic exploration, such approaches based on sampling from the reference distribution are typically inefficient at handling distributed information [2].

In strategy 3, we use an Archimedean spiral for 2D and its cylindrical extension for 3D (see Fig. 5). These are conventionally used in robotics as heuristic search strategies for uniform coverage in 2D and 3D search spaces (i.e. for uniform reference probability distributions) [36] [37]. Strategy 4 is similar to strategy 3 but uses spherical/ellipsoidal spiral trajectories that are customized to sweep the GMM search space. In this approach, the Gaussians are swept in sequence with spherical/ellipsoidal spiral paths starting from the centers of the Gaussian and sweeping the area up to a given number of standard deviations before moving to the next component. These approaches suffer from the curse of dimensionality and perform poorly for  $d > 3$ . Moreover, they require careful tuning of hyperparameters to generate efficient spiral paths. Furthermore, the resulting trajectories are deterministic and do not consider the stochasticity of target detection. Namely, the trajectory generated by such sweeping pattern passes through a given point in the search space only once. If the measurement system fails to detect the target during its first pass, the strategy has no future possibilities for detection.

The first metric we use for comparison is the average time taken to reach the target region for the first time. Table III

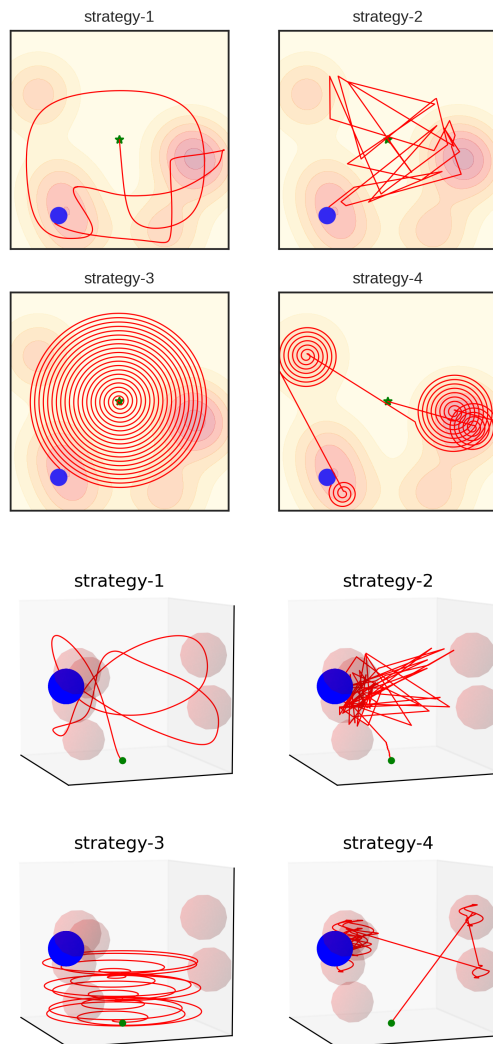


Fig. 5: Example trajectories (red) of the four different exploration strategies to reach a target region (blue sphere) within a reference probability distribution (in this case a GMM) for  $d = 2$  (top) and  $d = 3$  (bottom). The green point indicates the initial state of the point-mass system. The goal is to reach this target region the information about which is known to the search strategies only through the reference probability distribution.

shows the obtained results for 2D and 3D. We see that, on average, ergodic exploration reaches the target region faster due to its multiscale search behavior. Additionally, ergodic control has a 100% success rate. Despite being slower, spiral search is equally successful, but this success comes at the cost of optimally choosing the spiral parameters. This is often cumbersome in practice, especially in higher dimensions and considering that the tolerance of the detection region is often not known with high certainty.

A desirable property for exploration strategies is that the system takes into account the already visited regions to cover the unvisited regions more often. In order to evaluate this property, we consider a second metric: the cumulative average time to reach the target region over several successful attempts. We define this as  $\frac{T_c}{c}$ , where  $c$  is the number of successful attempts and  $T_c$  is the cumulative time until successful attempt




Fig. 6: Cumulative average time to reach a specified target region (spherical region with 1% of the volume of  $\Omega$  with  $d = 3$ ) with ergodic exploration (**left**), sampling-based strategy (**center**) and spiral movement (**right**). The x-axis represents the number of attempts to reach the target (reset count) and the y-axis represents the cumulative average of the time taken to reach the target at each number of attempts. Every time the target is reached, the system is re-initialized to the same initial state (starting a new attempt). The ergodic controller is aware of the lack of exploration inside the target region and tries to visit it more frequently. The cumulative average therefore converges to the time it takes to go from the initial state directly to the target at every re-initialization.

c. In this evaluation, as soon as the system reaches the boundary of the target region, we re-initialize it to the initial state and repeat this process for a fixed number of times. The results in Fig. 6 show that, for ergodic control, the cumulative average decreases with the number of successful attempts and converges to a fixed value. This suggests that the ergodic exploration tries to visit the unexplored regions more frequently as the number of attempts increases. In this case, the value that the cumulative average converges to is the time it takes to go from the initial state directly to the target at every re-initialization. This is an essential feature for insertion tasks as the exploration inside the hole (representing successful insertion of the peg) is not easy due to obstacles (e.g., uncertainties and collision of the surface of the hole against the peg) and we need the exploration strategy to drive exploration towards the unexplored region as the time evolves. This is an inherent property of ergodic exploration that the other approaches lack. This property plays a crucial role to cope with the stochasticity involved either in the measurement systems for detection, or the dynamics of the process (e.g., insertion task). To exploit this feature in real robot experiments, it is necessary that the ergodic controller is implemented *online* on the robot manipulator, i.e., that the controller knows the actual observed end-effector states. Our proposed algorithm for ergodic controller allows this online implementation on robot manipulators for  $d = 6$ , which is demonstrated in the next section.

In the insertion task, the target region for detection corresponds to a set of states of the peg that are necessary to be passed through for a successful insertion of the peg. This information is obtained from the reference probability distribution for exploration in the search space. Considering stochasticity is important for a search strategy to be useful in practice. In general, stochasticity may arise either from the measurement system (e.g., uncertainties in the location of the hole, grasp of the peg or manufacturing defects) and/or the dynamics of the system (e.g., stochasticity in the contact dynamics involved in the insertion). For example, during the insertion, the peg might be at the correct relative location to the hole according to the sensor system, but the insertion may still not be successful

TABLE III: Average time taken to reach a target region for different exploration strategies.

	Strategy	Success rate		Time taken (seconds)
		# Trials	# Success	
<b>2D</b>	Strategy 1	100	100	<b>66.9</b>
	Strategy 2	100	96	106.8
	Strategy 3	100	100	122.9
	Strategy 4	100	98	155.4
<b>3D</b>	Strategy 1	100	100	<b>84.7</b>
	Strategy 2	100	92	141.4
	Strategy 3	100	98	292.3
	Strategy 4	100	95	247.5

every time in that configuration due to the stochasticity of the process. We need the search strategy to explore more often in these target regions (correct configurations for insertion) to increase the likelihood of insertion. Ergodic control considers this stochasticity by driving the system to regions in the state space such that the amount of time it spends there is in proportion to the probability mass of that region. For more details on search strategies and their desired characteristics, see the seminal work of Koopman on the theory of search [38], [39], [40]. Ergodic exploration satisfies the standards for optimal search behavior set by Koopman. Although strategies 3 and 4 do not satisfy these properties, we included them in our evaluation for completeness.

### B. Experimental Setup for Peg-in-hole Task

We use a torque-controlled 7-axis Franka Emika Panda robot, with an insertion task based on the Siemens gear set benchmark (see Fig. 7)<sup>20</sup>, by using the 25.4mm-diameter peg and the 26.29mm-diameter receptacle, with the length of insertion of 47mm. We employed a Cartesian impedance controller to control the robot end-effector by computing a desired Cartesian wrench according to

$$\hat{\mathbf{f}} = \mathbf{K}_p \begin{bmatrix} \hat{\mathbf{p}} - \mathbf{p} \\ \text{Log}(\hat{\mathbf{q}} * \bar{\mathbf{q}}) \end{bmatrix} - \mathbf{K}_d \begin{bmatrix} \dot{\mathbf{p}} \\ \dot{\boldsymbol{\omega}} \end{bmatrix},$$

where  $\hat{\mathbf{p}} \in \mathbb{R}^3$ ,  $\hat{\mathbf{q}} \in \mathcal{S}^3$  are, respectively, the desired position and orientation of the end-effector (with  $\hat{\mathbf{q}}$  obtained from (14)),  $\mathbf{p}, \mathbf{q}, \dot{\mathbf{p}}, \dot{\boldsymbol{\omega}}$  are the end-effector position, orientation, linear and angular velocity and  $\mathbf{K}_p, \mathbf{K}_d$  are  $6 \times 6$  diagonal stiffness and damping gains, experimentally set to  $\mathbf{K}_p = \text{diag}(500, 500, 500, 160, 160, 160)$  and  $\mathbf{K}_d = \text{diag}(40, 40, 40, 10, 10, 10)$ . The symbol  $*$  denotes the unit quaternion product and  $\bar{\mathbf{q}}$  the conjugate of quaternion  $\mathbf{q}$ .  $\text{Log}(\cdot)$  is the logarithmic map defined in (9).

We obtain the desired robot joint torques with  $\hat{\boldsymbol{\tau}} = \mathbf{J}(\boldsymbol{\theta})^\top \hat{\mathbf{f}} + \mathbf{g}(\boldsymbol{\theta}) + \mathbf{h}(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}})$ , where  $\boldsymbol{\theta}, \dot{\boldsymbol{\theta}} \in \mathbb{R}^7$  are the robot joint positions and velocities and  $\mathbf{J} \in \mathbf{R}^{6 \times 7}$ ,  $\mathbf{g} \in \mathbb{R}^7$ ,  $\mathbf{h} \in \mathbb{R}^7$  are the Jacobian matrix of the end-effector, gravity and Coriolis terms. Note that the impedance gains were selected such that the robot remained compliant enough to safely interact with the environment during exploration, while still being able to track the ergodic trajectory.

<sup>20</sup><https://new.siemens.com/us/en/company/fairs-events/robot-learning.html>







Fig. 7: The Siemens gear benchmark used for evaluation of the peg-in-hole insertion task using ergodic control.



(a) A grasp with an offset in position



(b) A grasp with an offset in orientation

Fig. 8: Two instances of grasps typically appearing when testing the ergodic control for insertion. The demonstrations included different types of grasps to let the robot cope with this uncertainty during ergodic exploration.

We compare three different implementations of our approach. First, a *closed loop* version where the ergodic controller runs in real-time on the robot as an online coverage problem (Fig. 9). In that case, at every time step, we read the end-effector pose and feed it back to the ergodic controller, which computes the next pose to track based on the current state. Second, an *open loop* version where the controller tracks a reference ergodic trajectory computed offline (as used in [21] for a lower dimensional state space). Lastly, GMM-sampling-based exploration as presented in Section VI-A.

### C. Ergodic Controller Initialization

In our setup, the location of the hole is fixed (but unknown to the robot) and the main source of uncertainty comes from the different possible grasps of the peg by the end-effector, see Fig. 8. We model these uncertainties using a probability density function that indicates the importance of spending time




Fig. 9: Control diagram of *open* and *closed* loop implementations. In the latter, the ergodic controller takes the robot state (end-effector pose) into account while computing the desired pose. In this way, the history of observed states is kept and new desired poses are computed accordingly.




Fig. 10: Human demonstration of peg-in-hole insertion task. Datapoints are collected for different grasps of the peg through kinesthetic teaching to show the regions in the vicinity of the receptacle to be used by the ergodic controller.

in each region of the state space of the robot end-effector. We use ergodic control to insert the peg under such uncertainties.

To model the reference probability distribution for ergodic exploration of insertion, we collected  $M = 204$  data points using kinesthetic teaching, which corresponds to less than 2 minutes of recording, see Fig. 10. Each datapoint corresponds to the position and orientation of the end-effector holding the peg. The datapoints in the vicinity of the location of the hole were recorded for successful insertions with different orientation and position offsets inherent to the grasps of the peg. To give higher importance to insertion, almost half the data points were taken from the states corresponding to the peg within the hole. A variation of about 15mm in position of the grasp for each axis and about  $\sim 10^\circ$  variation in the orientation of the grasp were demonstrated (see accompanying video of the experiment).

Once data are collected, we preprocess the end-effector orientations as described in Section IV-B. Subsequently, we concatenate the position data with the obtained 3D orientation representation into a 6D vector. We then transform the data into the domain of ergodic control  $\Omega$  with  $L = 1$  (ergodic space) using a bijective linear transformation. We model the data points in this transformed space using a Gaussian mixture model (GMM) with full covariances. We empirically selected 8 mixture components with a minimal isotropic covariance prior of  $5 \times 10^{-3}$ . Figure 11 shows the obtained GMM for position and orientation (marginal distributions). The GMM is used as reference probability distribution  $P(\mathbf{x})$  for ergodic control in  $\Omega$ . With  $K = 10$  and  $N = 10$ , the computation of Fourier coefficients  $\hat{\mathcal{W}}$  for the reference probability distribution took less than 2 minutes. The closed-loop ergodic controller could be run at 100Hz ( $dt = 0.01$  in Algorithm 2) on the robot. The same settings are applied to the open-loop version of the insertion task. In this case, the trajectory is generated offline using ergodic control. It is then tracked using the above-mentioned settings of the impedance control. Figure 11 shows examples of generated ergodic trajectories. We allow a maximum of 40 seconds for insertion. Figure 1 shows trajectory generated from ergodic control in one of the trials of the experiment.



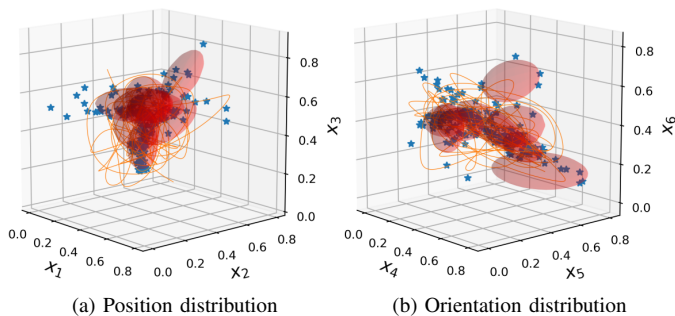


Fig. 11: Position and orientation marginals from the distribution used for full pose exploration (the figure does not show correlations between position and orientation). The 6D pose distribution is encoded as a GMM with 8 components (red ellipsoids) and full covariances, trained on a dataset of  $M = 204$  datapoints (blue points). The trajectory for the exploration (orange lines) is generated by ergodic control for the insertion task.

#### D. Experimental Results

The obtained results, summarized in Table IV, show that the *closed loop* approach clearly outperforms the other approaches that do not consider the history of observed states during the exploration. Indeed, while using the former, the robot was able to successfully insert the peg in the hole in 20 out of the 20 trials, where 10 were performed using the grasp in Fig. 8a and the other 10 using the one in Fig. 8b. On average, the insertion using the *closed loop* approach succeeded in 9.9s with a standard deviation of 8.5s. Figure 12 shows snapshots of the insertion using ergodic exploration in the *closed loop* setting. Note that the only required input was a set of demonstrations to show the robot the regions it should explore. This has proven important to deal with the uncertainty in the peg grasping pose. Note that interaction forces during exploration can also cause the grasp pose to change due to the limited gripping force of the robot (either from hardware and software limitations, or set on purpose to handle fragile objects). Our experimental results show that the closed loop ergodic strategy is also robust in these situations.

The open loop version succeeded in only 2 out of 10 trials and the naive approach of GMM-sampling-based exploration succeeded in none of the 5 trials. This shows the importance of exploiting the history of the real observed end-effector poses to compute control commands. This is particularly important for tasks requiring contacts with the environment, where the commands need to be re-evaluated according to the history of poses retrieved by the impedance controller. Notably, this allows the use of low gains to remain compliant and enable safe contacts. In the closed loop approach, the algorithm is aware of the locations that were previously effectively visited, which is exploited to fulfill the insertion task in an online and robust fashion.<sup>21</sup>

## VII. CONCLUSION AND FURTHER WORK

We proposed a solution to the multidimensional ergodic exploration problem with  $d > 3$ , which was previously considered to be intractable for applications in robotic manipulation.

<sup>21</sup>A video of the experiment is available at <https://sites.google.com/view/ergodic-exploration/>.

TABLE IV: Performance of the peg-in-hole task for different variations of the grasps.

Strategy	Success rate		Time taken (seconds)
	# Trials	Successful Trials	
Closed loop	20	20	$9.9 \pm 8.5$
Open loop	10	2	-
GMM-sampling	5	0	-




Fig. 12: Snapshots of an insertion using ergodic control in closed loop.

The proposed approach relies on tensor train decompositions and is evaluated in simulated examples of target detection and a real robot experiment using a peg-in-hole insertion task. The obtained results show that ergodic control has a 100% success rate in all tasks, succeeding faster than the competing approaches. In addition to the novelty of exploiting ergodic control in an online fashion for insertion tasks, the computational techniques we used demonstrate how algorithms in robotics, in general, can benefit from tensor methods to overcome the limitations in computational speed and storage requirements. To achieve this, tensor methods exploit the smoothness and symmetry of the functions underlying the problem. If the underlying problem does not have low rank structure, there may not be significant saving in the storage and computational cost using tensor methods. However, ergodic control problem as formulated in SMC has such low rank structure by design.

The main challenges of extending ergodic control to state space of more than 3 dimensions concern the computation of Fourier coefficients (preprocessing step) and the speed of the control loop (for online implementation). This article addressed both of these challenges using the properties of tensor train decomposition. We leveraged this improved ergodic control formulation to propose a sensorless strategy for peg-in-hole insertion tasks. We validated our approach with a compliant robot manipulator, where the 6D regions to explore are obtained from kinesthetic human demonstrations.

Our experimental results show that by using our approach, the robot is capable of achieving challenging insertion tasks without force/torque sensors.

By using the reproducible Siemens gears benchmark, insertion tasks with unknown gripping variations could be achieved in less than 10 seconds on average. This is, in part, due to the multi-scale exploration behavior that is inherent to the ergodic metric we employed. With such metric, the resulting controller first crudely explores the region of interest, and then progressively refines the search until insertion, efficiently exploiting information about the already covered regions. This is also due to the fast processing that we propose (through tensor train factorization), that allows ergodic control to be run in an online manner, even by considering distributions of full end-effector poses. Indeed, re-estimating the control commands on-the-fly allows the proposed ergodic control strategy to be used together with an impedance controller with low gains, which is important for tasks requiring contacts with the environment.

There is still room for improvement of the proposed approach. First, designing demonstrations for insertion tasks can be further optimized. In future work, we plan to study alternative demonstration strategies (including different distributions), which could speed up the insertion time using ergodic control. We will also apply the insertion tasks to different physical setups by using different sensory modalities. We also plan to study applications of our algorithm to other applications in robotics, either requiring exploration, active sensing, or, more generally, for applications requiring multidimensional basis functions with a low rank structure.

## APPENDIX A

### PROOF OF FOURIER COEFFICIENTS DECOMPOSITION

#### A one-dimensional integral

$$s = \int_{x=0}^l f(x) dx$$

can be computed numerically with the Gaussian-Quadrature (GQ) rule as

$$s = \sum_{j=1}^N \alpha_j f(x_j),$$

where  $N$  represents the degree of approximation (specified by the user),  $x_j$  represents the discretization points, and  $\alpha_j$  are the corresponding weights. For any polynomial function of degree less than  $2N - 1$ , the above summation gives exact result without any error in the integration. For a given  $N$ ,  $x_j$  and  $\alpha_j$  can be computed and are readily available using software packages for scientific computing.

We need to evaluate the multidimensional integral (1) to find  $\hat{\mathcal{W}}_k$ . Let  $x_j$  be the discretization points and  $\alpha_j$  the corresponding weights, with  $j \in \{1, \dots, N\}$ , obtained from the GQ rule for a given  $N$ . Let  $\mathcal{J} = \{\mathbf{j} = (j_1, \dots, j_d) : j_i \in \{1, \dots, N\}\}$  be the index set. We can discretize the domain  $\Omega$  at  $\mathbf{x}_j = (x_{j_1}, \dots, x_{j_d})$ , with  $\mathbf{j} \in \mathcal{J}$ . Let  $\mathcal{P}$  be the tensor formed by evaluating the reference distribution  $P(\mathbf{x})$  at the discretization points, i.e.,  $\mathcal{P}_j = P(\mathbf{x}_j)$ .

Then, we can evaluate (1) using GQ as

$$\begin{aligned} \hat{\mathcal{W}}_k &= \sum_{\mathbf{j} \in \mathcal{J}} \alpha_{j_1} \cdots \alpha_{j_d} P(\mathbf{x}_j) \Phi_k(\mathbf{x}_j) \\ &= \sum_{\mathbf{j} \in \mathcal{J}} \alpha_{j_1} \cdots \alpha_{j_d} P(\mathbf{x}_j) \phi_{i_1}(x_{j_1}) \cdots \phi_{i_d}(x_{j_d}), \quad \forall \mathbf{k} \in \mathcal{K}. \end{aligned} \quad (15)$$

Discretizing  $P(\mathbf{x})$  at the GQ points  $\mathbf{x}_j = (x_{j_1}, \dots, x_{j_d})$  with  $\mathbf{j} \in \mathcal{J}$ , we can get the tensor  $\mathcal{P}$ , with  $\mathcal{P}_j = P(\mathbf{x}_j)$ .

Consider a TT-representation of  $\mathcal{P}$  given by the TT-cores  $(\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^d)$ , so that for  $\mathbf{j} = (j_1, \dots, j_d) \in \mathcal{J}$  we have

$$\mathcal{P}_j = \mathcal{P}_{:::,j_1}^1 \mathcal{P}_{:::,j_2}^2 \cdots \mathcal{P}_{:::,j_d}^d.$$

Substituting the above expression in (15) yields

$$\begin{aligned} \hat{\mathcal{W}}_k &= \sum_{\mathbf{j} \in \mathcal{J}} \alpha_{j_1} \cdots \alpha_{j_d} \mathcal{P}_{:::,j_1}^1 \cdots \mathcal{P}_{:::,j_d}^d \phi_{k_1}(x_{j_1}) \cdots \phi_{k_d}(x_{j_d}) \\ &= \left( \sum_{j_1=1}^N \alpha_{j_1} \mathcal{P}_{:::,j_1}^1 \phi_{k_1}(x_{j_1}) \right) \cdots \left( \sum_{j_d=1}^N \alpha_{j_d} \mathcal{P}_{:::,j_d}^d \phi_{k_d}(x_{j_d}) \right). \end{aligned} \quad (16)$$

Also, we know that the TT-decomposition of  $\hat{\mathcal{W}}$  takes the form

$$\hat{\mathcal{W}}_k = \hat{\mathcal{W}}_{:::,k_1}^1 \cdots \hat{\mathcal{W}}_{:::,k_d}^d, \quad \forall \mathbf{k} \in \mathcal{K}. \quad (17)$$

Comparing (16) with (17), we obtain an expression for the TT-cores of the TT-decomposition of  $\hat{\mathcal{W}}$  as

$$\hat{\mathcal{W}}_{:::,k}^i = \sum_{j=1}^N \alpha_j \mathcal{P}_{:::,j}^i \phi_k(x_j), \quad \forall k \in (1, \dots, K), \quad \forall i \in (1, \dots, d). \quad (18)$$

## REFERENCES

- [1] A. Hubenko, V. A. Fonoberov, G. Mathew, and I. Mezić, "Multiscale adaptive search," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 41, no. 4, pp. 1076–1087, 2011.
- [2] L. M. Miller, Y. Silverman, M. A. MacIver, and T. D. Murphey, "Ergodic exploration of distributed information," *IEEE Transactions on Robotics*, vol. 32, pp. 36–52, 2016.
- [3] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4-5, pp. 432–442, 2011.
- [4] G. Mathew, S. Kannan, A. Surana, S. Bajekal, and K. R. Chevva, "Experimental implementation of spectral multiscale coverage and search algorithms for autonomous UAVs," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, 2013, p. 5182.
- [5] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey, "Real-time area coverage and target localization using receding-horizon ergodic exploration," *IEEE Transactions on Robotics*, vol. 34, pp. 62–80, 2018.
- [6] S. Calinon, "Mixture models for the analysis, edition, and synthesis of continuous time series," in *Mixture Models and Applications*, N. Bouguila and W. Fan, Eds. Springer, 2019, pp. 39–57.
- [7] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *American Control Conference*. IEEE, 2013, pp. 4196–4201.
- [8] L. Dressel and M. J. Kochenderfer, "Tutorial on the generation of ergodic trajectories with projection-based gradient descent," *IET Cyber-Phys. Syst.: Theory & Appl.*, vol. 4, pp. 89–100, 2019.
- [9] —, "Using neural networks to generate information maps for mobile sensors," in *IEEE Conference on Decision and Control (CDC)*, 2018, pp. 2555–2560.
- [10] E. Ayvali, H. Salman, and H. Choset, "Ergodic coverage in constrained environments using stochastic trajectory optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 5204–5210.

- [11] I. Abraham, A. Prabhakar, and T. D. Murphey, "An ergodic measure for active learning from equilibrium," *IEEE Transactions on Automation Science and Engineering*, 2021.
- [12] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM review*, vol. 51, no. 3, pp. 455–500, 2009.
- [13] A. Cichocki, D. P. Mandic, L. D. Lathauwer, G. Zhou, Q. Zhao, C. F. Caiafa, and A. H. Phan, "Tensor decompositions for signal processing applications: From two-way to multiway component analysis," *IEEE Signal Processing Magazine*, vol. 32, pp. 145–163, 2015.
- [14] S. Rabanser, O. Shchur, and S. Günnemann, "Introduction to tensor decompositions and their applications in machine learning," *ArXiv*, vol. 1711.10781, 2017.
- [15] M. B. Horowitz, A. Damle, and J. W. Burdick, "Linear Hamilton Jacobi Bellman equations in high dimensions," *53rd IEEE Conference on Decision and Control*, pp. 5880–5887, 2014.
- [16] A. A. Gorodetsky, S. Karaman, and Y. M. Marzouk, "Efficient high-dimensional stochastic optimal motion control using tensor-train decomposition," in *Robotics: Science and Systems*, 2015.
- [17] S.-k. Yun, "Compliant manipulation for peg-in-hole: Is passive compliance a key to learn contact motion?" in *IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1647–1652.
- [18] P. R. Giordano, A. Stemmer, K. Arbter, and A. Albu-Schaffer, "Robotic assembly of complex planar parts: An experimental evaluation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 3775–3782.
- [19] H. Park, J. Park, D.-H. Lee, J.-H. Park, M.-H. Baeg, and J.-H. Bae, "Compliance-based robotic peg-in-hole assembly strategy without force feedback," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 8, pp. 6299–6309, 2017.
- [20] M. P. Polverini, A. M. Zanchettin, S. Castello, and P. Rocco, "Sensorless and constraint based peg-in-hole task execution with a dual-arm robot," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 415–420.
- [21] D. Ehlers, M. Suomalainen, J. Lundell, and V. Kyrki, "Imitating human search strategies for assembly," in *International Conference on Robotics and Automation (ICRA)*, Montréal, Canada, May 2019, pp. 7821–7827.
- [22] M. J. A. Zeestraten, I. Havoutis, J. Silvério, S. Calinon, and D. G. Caldwell, "An approach for imitation learning on Riemannian manifolds," *IEEE Robotics and Automation Letters (RA-L)*, vol. 2, no. 3, pp. 1240–1247, June 2017.
- [23] S. Calinon, "Gaussians on Riemannian manifolds: Applications for robot learning and adaptive control," *IEEE Robotics and Automation Magazine (RAM)*, vol. 27, no. 2, pp. 33–45, June 2020.
- [24] P. M. Kroonenberg, *Applied multiway data analysis*. John Wiley & Sons, 2008, vol. 702.
- [25] T. Sahai, G. Mathew, and A. Surana, "A chaotic dynamical system that paints and samples," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 10760 – 10765, 2017, 20th IFAC World Congress.
- [26] G. Mathew, A. Surana, and I. Mezić, "Uniform coverage control of mobile sensor networks for dynamic target detection," *49th IEEE Conference on Decision and Control (CDC)*, pp. 7292–7299, 2010.
- [27] L. Grasedyck, D. Kressner, and C. Tobler, "A literature survey of low-rank tensor approximation techniques," *GAMM-Mitteilungen*, vol. 36(1), pp. 53–78, 2013.
- [28] I. V. Oseledets, "Tensor-train decomposition," *SIAM J. Scientific Computing*, vol. 33, pp. 2295–2317, 2011.
- [29] N. Lee and A. Cichocki, "Fundamental tensor operations for large-scale data analysis using tensor network formats," *Multidimensional Systems and Signal Processing*, vol. 29, no. 3, pp. 921–960, 2018.
- [30] D. V. Savostyanov and I. V. Oseledets, "Fast adaptive interpolation of multi-dimensional arrays in tensor train format," *The 2011 International Workshop on Multidimensional (nD) Systems*, pp. 1–8, 2011.
- [31] I. Oseledets and E. Tyrtshnikov, "TT-cross approximation for multidimensional arrays," *Linear Algebra and its Applications*, vol. 432, no. 1, pp. 70–88, 2010.
- [32] N. Kishore Kumar and J. Schneider, "Literature survey on low rank approximation of matrices," *Linear and Multilinear Algebra*, vol. 65, no. 11, pp. 2212–2244, 2017.
- [33] C. Lubich, I. V. Oseledets, and B. Vandereycken, "Time integration of tensor trains," *SIAM J. Numerical Analysis*, vol. 53, pp. 917–941, 2015.
- [34] D. Bigoni, A. P. Engsig-Karup, and Y. M. Marzouk, "Spectral tensor-train decomposition," *SIAM Journal on Scientific Computing*, vol. 38, no. 4, pp. A2405–A2439, 2016.
- [35] S. Dolgov and D. Savostyanov, "Parallel cross interpolation for high-precision calculation of high-dimensional integrals," *Computer Physics Communications*, vol. 246, p. 106869, 2020.
- [36] Hyeonjun Park, Ji-Hun Bae, Jae-Han Park, Moon-Hong Baeg, and Jaeheung Park, "Intuitive peg-in-hole assembly strategy with a compliant manipulator," in *IEEE ISR*, 2013, pp. 1–5.
- [37] J. C. Triyonoputro, W. Wan, and K. Harada, "Quickly inserting pegs into uncertain holes using multi-view images and deep network trained on synthetic data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 5792–5799.
- [38] B. O. Koopman, "The theory of search. I. kinematic bases," *Operations research*, vol. 4, no. 3, pp. 324–346, 1956.
- [39] —, "The theory of search. II. target detection," *Operations research*, vol. 4, no. 5, pp. 503–531, 1956.
- [40] —, "The theory of search: III. the optimum distribution of searching effort," *Operations research*, vol. 5, no. 5, pp. 613–626, 1957.



**Suhan Shetty** received the M.E. degree in mechanical engineering from the Indian Institute of Science, Bangalore, India. He is currently a Research Assistant with the Idiap Research Institute in Switzerland, working toward the Ph.D. degree in electrical engineering with the École Polytechnique Fédérale de Lausanne (EPFL). His Ph.D. thesis research is focusing on low-rank approximation techniques in robot learning.



**João Silvério** is a postdoctoral researcher at the Idiap Research Institute since July 2019. He received his M.Sc in Electrical and Computer Engineering (2011) from Instituto Superior Técnico (Lisbon, Portugal) and Ph.D. in Robotics (2017) from the University of Genoa (Genoa, Italy) and the Italian Institute of Technology, where he was also a postdoctoral researcher until May 2019. He is interested in machine learning for robotics, particularly imitation learning and control. Webpage: <http://joaosilverio.eu>.



**Sylvain Calinon** received the Ph.D. degree from the École Polytechnique Fédérale de Lausanne (EPFL) in 2007. He is a Senior Researcher at the Idiap Research Institute, and a Lecturer at the EPFL. From 2009 to 2014, he was a Team Leader at the Italian Institute of Technology. From 2007 to 2009, he was a Postdoc at EPFL. His research interests cover robot learning, human-robot collaboration, Riemannian geometry and optimal control. Website: <https://calinon.ch>.