

ASRO-DIO: Active Subspace Random Optimization based Depth Inertial Odometry

Jiazhao Zhang, Yijie Tang, He Wang, and Kai Xu

Abstract—High-dimensional nonlinear state estimation is at the heart of inertial-aided navigation systems (INS). Traditional methods usually rely on good initialization and find difficulty in handling large inter-frame transformations due to fast camera motion. We opt to tackle these challenges by solving the depth inertial odometry (DIO) problem with random optimization. To address the exponentially increased amount of candidate states sampled for the high-dimensional state space, we propose a highly efficient variant of random optimization based on the idea of active subspace. Our method identifies the active dimensions which contribute the most significantly to the decrease of the cost function in each iteration, and samples candidate states only within the corresponding subspace. This allows us to efficiently explore the 18D state space of DIO and achieve good optimality by sampling and evaluating only thousands of candidate states. Experiments show that our method attains highly robust and accurate DIO under fast camera motions and low light conditions, without needing a slow-motion warm-up for initialization.

Index Terms—Depth-inertial odometry, state estimation, evolution strategy, simultaneous localization and mapping.

I. INTRODUCTION

IN the fields of robotics and vision, extensive research has been devoted to inertial-aided navigation systems (INS) [1], [2] for motion estimation in GPS-denied environments such as indoor rooms. In the literature, INS is typically realized with either filtering-based [1], [3]–[9] or optimization-based approach [10]–[15]. The basic principle underlying both approaches is the linearization of the nonlinear problem of motion estimation. In particular, the iterated Kalman filter is proven to be equivalent to the Gaussian-Newton algorithm [2], [16]. The success of such linearization usually hinges on a good initialization and requires that the change of states in each time-step is small so that the nonlinearity is not high. This makes the existing methods sensitive to initialization and error-prone when handling fast camera motion.

To mitigate these issues, a promising approach is random optimization. Random optimization estimates states by sampling a population of candidates and evaluating them with a cost function. It excels at finding more global optima in highly nonlinear optimization problems [17]. Recently, it has also been adopted to online RGB-D reconstruction under fast camera motion, showing good performance [18]. However, existing state estimators based on random optimization are

mostly restricted to *low dimensional state space* (e.g., 6 DoFs of camera motion). To ensure a good exploration of state space, the amount of samples increases exponentially as the dimensionality grows [19]–[21]. This makes random optimization computationally prohibitive for INS where a much higher dimensional state space (e.g., 18 DoFs of IMU states) is involved.

In this work, we present *active subspace random optimization* for solving the high-dimensional nonlinear state estimation of INS. Our key observation is that only a fraction of the dimensions have a significant impact on the cost function and the subset of active dimensions changes over iteration steps. Similar findings have also been reported and exploited in alternative optimization frameworks [20]–[22]. During the iterations of random optimization, we dynamically identify active dimensions based on the sampling efficiency of a dimension. Given a dimension, its sampling efficiency is measured by the ratio of the amount of state update over the range of search along that dimension. The amount of state update is probed with a set of pre-sampled probing states.

With the active subspace, our method attains a good balance between the optimality and the computational cost of the optimization. As such, all states are estimated within the optimization framework; no specially designed initialization scheme is needed as in [15], [23], [24]. The main bottleneck of random optimization is the iterated sampling and the evaluation of candidate solutions (states). To accelerate the randomized search, we opt to pre-sample states in the solution space, forming a pre-sampled state template (PST), and then move and rescale the PST to explore and search in the solution space, similar to ROSEFusion [18]. Unlike ROSEFusion, where PST is constructed with uniform sampling, we take advantage of the prior knowledge of INS state space and work with different sampling spaces for different components (dimensions) of the states. In particular, orientation and gravity vector are sampled in $SO(3)$, velocity and position with uniform distribution, and IMU measurement error from Gaussian distribution. Furthermore, we devise different updating (moving, rescaling, and weighted averaging) mechanisms for the three sampling spaces for better efficiency.

Although our optimization approach is general and adaptable to vision-inertial odometry (VIO), we focus on depth-inertial odometry (DIO) in this work (see Fig. 1). While VIO has a large body of prior works, DIO has not been extensively studied, partly due to its special challenges (e.g., lack of reliably trackable features). However, DIO has several advantages such as robustness to poor lighting condition and resilience to

Jiazhao Zhang, Yijie Tang and Kai Xu are with the Department of Computer Science, National University of Defense Technology, Changsha, China (e-mail: zhngjzh@gmail.com; yjtang1024@gmail.com; kevin.kai.xu@gmail.com).

He Wang is with the Department of Computer Science, Peking University, Beijing, China (e-mail: hewang@pku.edu.cn).




Fig. 1. The tracking and mapping result of our depth inertial odometry for a sequence drawing letters “DIO”. The sequence is captured with a hand-held depth camera under challenging conditions including quick start (no slow-motion warm up; the red point and the red RGB frame), fast camera motion (blue, zoom in to see the motion blur), low light (orange). The tracked camera trajectory is depicted with green curve and the scene is reconstructed with TSDF-based depth fusion.

motion blur caused by fast motion, making it especially important to robotic applications. We adopt the TSDF (Truncated Signed Distance Field) based mapping framework [25] which allows for depth-based fusion. More importantly, TSDF map facilitates a purely depth-based cost function via measuring the depth-TSDF conformance [26]. The cost function is parallel-friendly and admits a GPU implementation, enabling fast evaluation of candidate solutions in random optimization.

We have evaluated our method on several benchmarks. On the ETH3D benchmark [27], our method achieves the state-of-the-art accuracy under fast camera motions (*i.e.*, shaking cameras). We also contribute a Robotic Arm Shaking (RAS) dataset of robot-operated fast-motion sequences with quality ground-truth trajectories. On RAS, our method shows high robustness and succeeds on all the sequences with low light conditions and high speed motions. Extensive quantitative and qualitative comparisons also demonstrate the significant advantage of our method over the state of the arts, under extreme conditions and without a slow-motion warm up for initialization. *All source code and dataset will be released.*

Our main contributions include:

- We present a practically robust tightly-coupled solution to DIO based on random optimization.
- We propose to utilize active subspace random optimization to tackle the high-dimensional state estimation problem in inertial-based navigation systems.
- We have implemented a navigation system based on our technique. It realizes robust and accurate tracking under challenging conditions such as low light and fast camera motions, without needing a slow-motion warm up for initialization.

II. RELATED WORK

a) Depth inertial navigation system: There is a large body of literature on inertial-aided navigation systems (INS). We decide to review only those which are highly related to our work. In [28], a tightly coupled EKF framework is proposed for pose estimation and IMU-RGBD extrinsic calibration.

Since then, most works focus on how to incorporate depth information into visual inertial odometry (VIO) frameworks. Representative approaches are formulated either as filters [29], [30] or with second-order optimization [23], [24], [31], [32].

In the context of depth inertial odometry (DIO) system, however, there are only a few existing works which leverage the ICP technique and develop loosely-coupled [33], [34] or tightly-coupled [31] navigation systems. Departing from these works, our work contributes the first random optimization framework for tightly coupled DIO, achieving highly robust and accurate tracking under various extreme conditions.

b) State estimation based on sampling-based optimization: In the robotics and computer vision community, perhaps the most commonly used sampling-based algorithm is particle filter, and more specifically, Rao-Blackwellized particle filter (RBPF) [35]. There has been extensive research on particle filter based SLAM [36]–[38] and object pose tracking [39]–[42]. Different from these works in which sampling-based algorithm is used for sequential state estimation across time steps, our method performs sampling-based random optimization to estimate the state *within a single time step*.

Another line of optimization approach is Evolutionary Algorithms (EAs), including Particle Swarm Optimization (PSO) [43], Evolution Strategies (ESs) [44], Random Optimization (RO) [45] and so on. These methods involve heuristically designed sampling and updating strategies, often requiring extremely large computing resources. EAs are limited to low-dimensional and time-insensitive scenarios. Recently, Zhang et. al [18] propose an online dense reconstruction method in which the core is 6-DoF camera tracking based on RO powered by pre-sampled state templates. We substantially extend this work to deal with the high dimensional solution space of DIO through leveraging active subspace.

c) Active subspace for high-dimensional optimization: High dimensional solution space induces very high computational and memory cost, and usually causes poor convergence. To tackle this difficulty, active subspace techniques and extensions [46] are becoming popular for dimensionality reduction on the fly during optimization. A classic method of obtaining active subspace is random projection [47], [48], which is simple but hard to be effective with a totally random selection. Recent works advance the active subspace idea [19], [21], [22] by evaluating the importance of different dimensions based on various heuristics or priors. Our method, sharing the same insights, identifies active subspace via measuring the sampling efficiency along different dimensions.

III. METHOD OVERVIEW

In this paper, we introduce a novel method for depth inertial odometry leveraging active subspace random optimization. Our system is based on a platform, on which we mount a depth sensor and a 6-axis inertial measurement unit (IMU). We assume that the transformation between the camera and the IMU is given by prior calibration. Our proposed method takes inputs a live stream of depth observations along with IMU measurements and outputs the 6 DoF pose of the platform, namely 3D orientation and 3D position. We formulate

DIO as an online optimization problem and construct a cost function based on depth-TSDF conformance [18] and IMU measurement residuals. Compared with ROSEFusion [18] that is depth-only, our method further leverages IMU measurement and thus incorporates the IMU related states into our state space, resulting in a high state space dimensionality. To achieve robust and real-time optimization for this high dimensional state estimation problem, we propose active subspace random optimization method. Note that our method doesn't rely on RGB images, which are shown in the paper only for visualization.

Our methodology is structured as follow: We formulate the optimization problem of DIO in Sec IV by first introducing the variables to be estimated (Sec IV-A and Sec IV-B), also called candidate states in random optimization. To evaluate the candidate states, we leverage the depth-TSDF conformance (Sec IV-C1) and IMU measurement residual (Sec IV-C2) to construct our tightly-coupled cost function (Sec IV-D). Based on the evaluated candidate states, we present the random optimization for state estimation (Sec V). After a brief review of the random optimization framework (Sec V-A), we then introduce how to extend the random optimization to high dimensional states estimation by the pre-defined state space (Sec V-B) and active subspace (Sec V-C).

IV. OPTIMIZATION FORMULATION

A. Notations and Definitions

In this paper, we deploy the following notations: superscripts $(\cdot)^W$, $(\cdot)^C$, and $(\cdot)^B$ denote variables in the world frame, camera frame and IMU body frame, respectively; subscripts $(\cdot)_C$ and $(\cdot)_B$ denote quantities associated with the camera and the IMU body, respectively; $(\hat{\cdot})$ denotes the values of measurements, in contrast with their true states (\cdot) ; subscripts $(\cdot)_t$ and $(\cdot)_k$ represent quantities at the time step t of the high frame-rate IMU measurements and quantities at the k -th camera frame, respectively.

Then, in the world frame, we can represent the position, velocity and orientation of the IMU as \mathbf{p}_B^W , \mathbf{v}_B^W , \mathbf{q}_B^W , respectively. As the IMU measurements are done in its body frames, we can represent the measurements as $\hat{\boldsymbol{\omega}}_t^B$ from the gyroscope and $\hat{\mathbf{a}}_t^B$ from the accelerometer and their corresponding measurements errors as $\mathbf{E}_{a_t}^B$ and $\mathbf{E}_{g_t}^B$, respectively. We further denote gravity vector as \mathbf{g}^W .

In this work, we are only interested to estimate the states at the frame rate of the depth camera. At the k -th camera frame, the state \mathbf{x}_k to be estimated is shown as the following:

$$\mathbf{x}_k = [\mathbf{p}_{B_k}^{W^T}, \mathbf{v}_{B_k}^{W^T}, \mathbf{q}_{B_k}^{W^T}, \mathbf{g}^{W^T}, \mathbf{E}_{a_k}^{B^T}, \mathbf{E}_{g_k}^{B^T}]^T. \quad (1)$$

Note that, our method only samples the imaginary part of the unit quaternion to represent rotation. Our actual state space is therefore 18-DoF. All the states, including gravity vector \mathbf{g}^W , are estimated by the random optimization framework without the need for specifically designed initialization like [15], [23].

B. IMU Kinematic Model

Following the existing works [1], [15], we adopt the IMU model as shown below:

$$\begin{aligned} \hat{\mathbf{a}}_t^B &= \mathbf{a}_t^B + \mathbf{q}_W^B \otimes \mathbf{g}^W + \mathbf{E}_{a_t}^B, \\ \hat{\boldsymbol{\omega}}_t^B &= \boldsymbol{\omega}_t^B + \mathbf{E}_{g_t}^B, \\ \mathbf{E}_{a_t}^B &= \mathbf{b}_{a_t}^B + \mathbf{n}_{a_t}^B, \\ \mathbf{E}_{g_t}^B &= \mathbf{b}_{g_t}^B + \mathbf{n}_{g_t}^B, \end{aligned} \quad (2)$$

where $\mathbf{q}_W^B = \mathbf{q}_B^{W^{-1}}$ is the corresponding quaternion for the rotation from the world frame to the IMU body frame, \otimes represents the Hamilton quaternion multiplication, the measurement errors \mathbf{E}_t^B (can either be $\mathbf{E}_{a_t}^B$ or $\mathbf{E}_{g_t}^B$) is assumed to be comprised of an bias term \mathbf{b}_t^B (i.e., $\mathbf{b}_{a_t}^B$ or $\mathbf{b}_{g_t}^B$) and an additive noise term \mathbf{n}_t^B (i.e., $\mathbf{n}_{a_t}^B$ or $\mathbf{n}_{g_t}^B$).

Here we assume the additive noise \mathbf{n}_t^B are white Gaussian noise, i.e., $\mathbf{n}_{a_t}^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_{n_a}^2)$, $\mathbf{n}_{g_t}^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_{n_g}^2)$, and the bias terms of gyroscope and accelerometer are modeled as random walk, i.e., $\dot{\mathbf{b}}_{g_t}^B = \mathbf{b}_{g_t}^B - \mathbf{b}_{g_{t-1}}^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_{b_g}^2)$, $\dot{\mathbf{b}}_{a_t}^B = \mathbf{b}_{a_t}^B - \mathbf{b}_{a_{t-1}}^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_{b_a}^2)$. Assuming noises from different sources and different time steps are independent, we can derive that the measurement errors \mathbf{E}_t^B are also random walks:

$$\begin{aligned} \dot{\mathbf{E}}_t^B &= \mathbf{E}_t^B - \mathbf{E}_{t-1}^B \\ &= \mathbf{b}_{t-1}^B + \dot{\mathbf{b}}_{t-1}^B + \mathbf{n}_t^B - \mathbf{b}_{t-1}^B - \mathbf{n}_{t-1}^B \\ &= \dot{\mathbf{b}}_{t-1}^B + \mathbf{n}_t^B - \mathbf{n}_{t-1}^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_{b^2} + 2\boldsymbol{\sigma}_n^2), \end{aligned} \quad (3)$$

We further assume, at time step 0, the initial measurement errors are from Gaussian distributions, i.e., $\mathbf{E}_0^B \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\sigma}_E^2)$.

C. Cost Function Terms

In this subsection, we would describe the cost function terms of the optimization problem for state estimation. To ensure the state estimation is robust to low lighting conditions or motion blur, we propose to incorporate the following depth-based cost terms in our cost function.

1) *Depth-TSDF Conformance*: Given the depth map I_k^d of frame k , and the so-far constructed TSDF $\psi : \mathbb{R}^3 \rightarrow \mathbb{R}$, the depth-TSDF conformance measures how well the 3D point cloud back-projected from I_k^d aligns to the zero-crossing surface of ψ , similar to [26] and [18]. The back-projection is based on the camera pose. Our goal is to find the camera pose $[\mathbf{q}_{C_k}^W | \mathbf{p}_{C_k}^W]$ that best fits the depth map I_k^d into the TSDF, where $\mathbf{q}_{C_k}^W$ and $\mathbf{p}_{C_k}^W$ are the orientation and position of camera in world coordinate, respectively.

For a pixel (i, j) and its corresponding depth value $I^d(i, j)$, we can obtain the corresponding 3D point \mathbf{p}_{ij}^C in camera coordinate system by back-projection. Then we can transform this 3D point to the world frame using $\mathbf{p}_{ij}^W = \mathbf{q}_C^W \otimes \mathbf{p}_{ij}^C + \mathbf{p}_C^W$. For a newly captured frame, we fit only the overlapping area observed by both the current and the previous frames to avoid over-estimation of fitness [18]. We compute the overlapping area O_k of the depth frame I_k^d as:

$$O_k = \{(i, j) \mid \pi_{k-1}(\pi_k)^{-1}[(i, j)] \in I_{k-1}^d\}, \quad (4)$$

where the π_k is the projection matrix of frame k . Note that, in Eq. (4) we use an approximate overlapping area between the

consecutive frames (π_{k-1}, π_k) . Under the online tracking, the transformation between consecutive time instants is relative small (tens of centimetres or degrees) even under fast motion meaning that overlapping area between consecutive frames is large. In our implementation, a rough π_k can be initialized with IMU propagation and then updated with camera pose after every iteration step during the optimization.

We then seek to find the camera pose $[\mathbf{q}_{C_k}^W | \mathbf{p}_{C_k}^W]$ such that the overlapping 3D points lie as close as possible to the zero crossing surface $\{\mathbf{x} | \psi(\mathbf{x}) = 0\}$ of TSDF. Following the same assumption in [26], the depth measurements contain Gaussian noise and that all pixels are independent and identically distributed. Therefore, the depth-TSDF conformance can be defined as:

$$\begin{aligned} \hat{\mathbf{r}}_{\mathcal{D}}(\mathbf{x}_k) &= \sum_{(i,j) \in O_k} \psi(\mathbf{q}_{C_k}^W \otimes \mathbf{p}_{ij}^C + \mathbf{p}_{C_k}^W)^2, \\ \mathbf{q}_{C_k}^W &= \mathbf{q}_{B_k}^W \otimes \mathbf{q}_B^C, \\ \mathbf{p}_{C_k}^W &= \mathbf{q}_{B_k}^W \otimes \mathbf{p}_B^C + \mathbf{p}_{B_k}^W, \end{aligned} \quad (5)$$

where the transformation between the camera and the IMU $[\mathbf{q}_B^C | \mathbf{p}_B^C]$ is fixed and known from prior calibration. Finally, we can normalize the depth-TSDF conformance by the number of valid points in the overlapping area:

$$\mathbf{r}_{\mathcal{D}}(\mathbf{x}_k) = \frac{\hat{\mathbf{r}}_{\mathcal{D}}(\mathbf{x}_k)}{|O_k|}. \quad (6)$$

The main advantage of the depth-TSDF conformance (Eq. (5)) over existing methods is correspondence-free, saving the need for descriptor matching. This also makes it robust to low lighting conditions or motion blur. Besides, the depth-TSDF conformance is parallel-friendly, which could be efficiently implemented with GPU for fast candidate states evaluation.

2) *IMU Measurement Residual*: Sharing the same insight with the primary INS works [12], [15], we perform IMU measurement residual minimization. At the frame k , we can obtain the IMU measurements $\hat{\mathbf{a}}_{t=t_{k-1}:t_k}^B$ and $\hat{\boldsymbol{\omega}}_{t=t_{k-1}:t_k}^B$ between two consecutive depth camera frames k and $k-1$, where t_{k-1} and t_k are the corresponding time steps of camera frame $k-1$ and k in the IMU rate. We can then propagate the IMU measurements to obtain the estimated rotation $\hat{\mathbf{q}}_{B_k}^W = Q(\mathbf{x}_{k-1}, \hat{\mathbf{a}}_{t=t_{k-1}:t_k}^B, \mathbf{E}_{a_k}^B)$ and position $\hat{\mathbf{p}}_{B_k}^W = P(\mathbf{x}_{k-1}, \hat{\mathbf{a}}_{t=t_{k-1}:t_k}^B, \hat{\boldsymbol{\omega}}_{t=t_{k-1}:t_k}^B, \mathbf{E}_{a_k}^B, \mathbf{E}_{g_k}^B)$, where $P(\cdot)$ and $Q(\cdot)$ are the mid-point propagation function introduced in [15], and $\mathbf{E}_{a_k}^B$ and $\mathbf{E}_{g_k}^B$ are included in our state \mathbf{x}_k and represent the average IMU measurement error from t_{k-1} to t_k . Assuming that the time interval between consecutive frames is nearly constant, it can be proved that both the two errors are approximative to random walk. (Eq. 3) [49]. Therefore we can define the IMU measurement residual as:

$$\begin{aligned} \mathbf{r}_{\mathcal{T}}(\mathbf{x}_k) &= \mathbf{p}_{B_k}^W - \hat{\mathbf{p}}_{B_k}^W, \\ \mathbf{r}_{\mathcal{R}}(\mathbf{x}_k) &= \theta(\mathbf{q}_{B_k}^W \otimes \hat{\mathbf{q}}_{B_k}^W), \end{aligned} \quad (7)$$

where the $\theta(\mathbf{q})$ extracts the rotation angle of a unit quaternion. For discrete-time implementation, we leverage the same mid-point algorithm as [15].

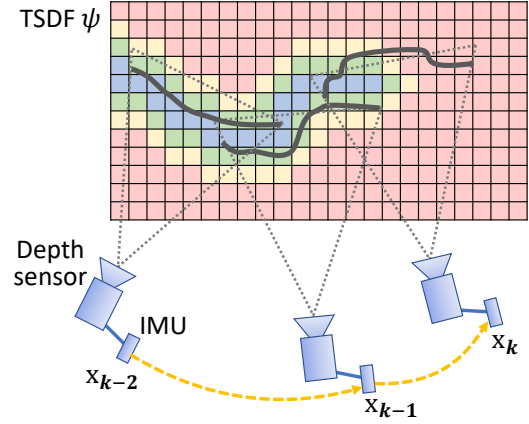


Fig. 2. Our cost function is composed of depth-TSDF conformance and the IMU measurement residual. Given a pose of the depth camera, a point cloud is back-projected from the depth map and embedded into the TSDF volume. The pose that causes the best alignment between the depth and the TSDF is the optimal solution. The transformation between the camera and the IMU is fixed.

D. The Formulation of Cost Function

The cost function (Fig. 2) is composed of the depth-TSDF conformance (Eq. (5)) and the IMU measurement residual (Eq. (7)):

$$\begin{aligned} \mathbf{x}_k^* &= \arg \min_{\mathbf{x}_k} \{w_1 \mathbf{r}_{\mathcal{D}}(\mathbf{x}_k) + w_2 \mathbf{r}_{\mathcal{R}}(\mathbf{x}_{k-1}, \mathbf{x}_k) + \\ &w_3 \|\mathbf{r}_{\mathcal{T}}(\mathbf{x}_{k-1}, \mathbf{x}_k)\|^2\}, \end{aligned} \quad (8)$$

where the w_1 , w_2 and w_3 are the weights of each cost term. We empirically set the weights as $w_1 = 1$, $w_2 = 1$ and $w_3 = 0.1$, considering that the double-integration term $\mathbf{r}_{\mathcal{T}}$ has a relatively large error, especially at the initialization stage. Note that, the cost function does not require RGB inputs, making our method naturally resilient to the low light conditions and motion blur.

V. ACTIVE SUBSPACE RANDOM OPTIMIZATION

To track the 6-DoF pose, we only need to solve the optimization problem in Eq. (8) at each camera frame. However, directly optimizing Eq. (8) brings three major challenges to state estimation algorithms. *First*, some states, like gravity vector and velocity, are unknown at initialization. Existing works require specially designed initialization scheme such as visual-inertial alignment [15]. *Second*, gradient can be undefined when the depth map is out of the valid scope of TSDF under very large camera motion. *Third*, high dimensional state estimation is highly computationally consuming, especially for sampling-based algorithms. To mitigate the challenges, we propose active subspace random optimization by sampling candidate states only within active subspace. We first briefly review random optimization, then introduce pre-sampled state templates and how to leverage them to replace random sampling, and finally describe our proposed active subspace random optimization method.

A. Random Optimization

Random optimization is a class of optimization methods, which does not require computing gradients of the objective function and thus is also known as derivative-free optimization. For a general optimization problem $\min_{\mathbf{x} \in \Omega} g(\mathbf{x})$, where \mathbf{x} represents optimization variables or the states being optimized in the state space Ω and g represents the objective function, the basic idea of random optimization is to iteratively sample around the current optimum \mathbf{x}^* and move the optimum to the best found position in the state space. More specifically, in each step of the optimization, a set of candidate states $\{\mathbf{x}^{(i)}\}_{i=1:N}$ are sampled from a pre-defined distribution, usually Gaussian distributions, centered round the so-far reached optimum, i.e., $\mathcal{N}(\mathbf{x}^*, \Sigma)$.

In this work, we adopt a fast and robust random optimization method introduced in [18], which we provide a brief recap below. The key idea is to pre-sample a set of random states $\{\mathbf{x}_{PST}^{(i)}\}_{i=1:N}$ around the origin of the state space Ω , which we refer as pre-sampled state template, or PST for short, denoted by Ω^{PST} , and use them to replace true random sampling operations used in each iteration of random optimization, for saving time in performing samplings and accelerating optimization.

At each iteration step j , [18] proposes to evolve Ω^{PST} in the following way: it rescales each pre-sampled states \mathbf{x}_{PST} from Ω^{PST} anisotropically using a scaling factor $\mathbf{r}_j \in \mathbb{R}^D$ ($D = 6$ for [18] but in our work $D = 18$) and then add each of them as a shift to the current optimum \mathbf{x}_{j-1}^* of step $j-1$, yielding a new set of candidate states Ω_j in replacement of true random sampled states around \mathbf{x}_{j-1}^* :

$$\Omega_j = \{\mathbf{x}_j^{(i)} \mid \mathbf{x}_j^{(i)} = \mathbf{r}_{j-1} \odot \mathbf{x}_{PST}^{(i)} + \mathbf{x}_{j-1}^*\}, \quad (9)$$

where \odot represents element-wise multiplication. We can then evaluate all the candidate states $\{\mathbf{x}_j^{(i)} \mid \mathbf{x}_j^{(i)} \in \Omega_j\}$ using the objective function $g(\cdot)$, followed by elitist selection which keeps only the ones that have lower costs than the previous optimum and forms Ω_j^e . Finally, we perform a weighted average among the selected candidate states, where the weight is proportional to the cost margin. This procedure can be formally described as follows:

$$\begin{aligned} \Omega_j^e &= \{\mathbf{x}_j^{(i)} \in \Omega_j \mid g(\mathbf{x}_j^{(i)}) < g(\mathbf{x}_{j-1}^*)\}, \\ \mathbf{w}^{(i)} &= g(\mathbf{x}_{j-1}^*) - g(\mathbf{x}_j^{(i)}), \\ \bar{\mathbf{w}}^{(i)} &= \mathbf{w}^{(i)} / \sum_{\mathbf{x}_j^{(i)} \in \Omega_j^e} \mathbf{w}^{(i)}, \\ \mathbf{x}_j^* &= \sum_{\mathbf{x}_j^{(i)} \in \Omega_j^e} \bar{\mathbf{w}}^{(i)} \mathbf{x}_j^{(i)}. \end{aligned} \quad (10)$$

Once the optimum is updated, we can easily adjust the scaling factor based on the dimensional deviation ($\mathbf{v} = \mathbf{x}_j^* - \mathbf{x}_{j-1}^*$) and objective/cost function $g(\cdot)$:

$$\hat{\mathbf{r}}_j = g(\mathbf{x}_j^*) \frac{\mathbf{v}}{\|\mathbf{v}\|}. \quad (11)$$

Here the $\frac{\mathbf{v}}{\|\mathbf{v}\|}$ drives the candidate states towards the so far best state, which can be scaled by the current best value of

objective function $g(\mathbf{x}_j^*)$. This means that the search range is adaptive to the performance of optimization, which helps the optimization converge more stably. In [18], $\mathbf{r}_j = \hat{\mathbf{r}}_j$. Our method proposes further adjustment for computing \mathbf{r}_j to control the active subspace (see Sec. V-C)

By leveraging this variant of random optimization, [18] shows impressive performance for low dimensional state estimation problem, e.g. 6-DoF camera tracking. However, we identify two issues of [18] when applying it to our problem: 1) to generate Ω_{PST} , [18] treats all state variables as Euclidean variables and always perform uniform sampling, which is problematic for rotations and other variables introduced in our state; 2) when the dimensionality of the state space increases, the amount of sampled candidate states to cover the state space needs to increase exponentially, making this method intractable for high dimensional state estimation problem, such as in INS and this work. We will tackle these two issues separately in Sec. V-B and Sec. V-C.

B. Sampling and Evolving Ω^{PST} for DIO

For depth inertial odometry, our state $\mathbf{x} = [\mathbf{p}_B^{WT}, \mathbf{v}_B^{WT}, \mathbf{q}_B^{WT}, \mathbf{g}^{WT}, \mathbf{E}_a^{BT}, \mathbf{E}_g^{BT}]^T$ contains both Euclidean and SO(3) variables that follow a variety of distributions. To pre-sample them and generate our Ω^{PST} , we want to make reasonable assumptions about the underlying distributions for each state variables, unlike [18] always uses uniform distribution. Our general guideline is that, since states in Ω_{PST} will be used as shift to the current state, their distributions need to resemble the distribution of the change in each state variable between two camera frames.

Here we can categorize the state variables into three types:

1) *Unit Quaternions*: The orientation \mathbf{q}_B^W is a quaternion representing a rotation in $SO(3)$. Gravity vector \mathbf{g}^W can also be seen as a rotation multiplies $[0, 0, g]^T$ ($g = 9.81$). Considering arbitrary and rapid change in camera orientation, we thus assume these quaternions follow uniform distribution in $SO(3)$. We then adopt the uniform rotation sampling algorithm proposed by [50]. And we use the positive-dot-product to check that all the quaternions are on the same hemisphere, which indicates the quaternions are continuously distributed. Note that, for these unit quaternions, we only use their imaginary part $[q_x, q_y, q_z]$ as state variables.

2) *Position and Velocity*: Due to the same reason, we assume the change in velocity \mathbf{v}_B^W and position \mathbf{p}_B^W may follow a uniform distribution, we thus sample the corresponding variables from uniform distribution $\mathcal{U}([-1, 1])$.

3) *IMU Measurement Errors*: According to Eq. (3), the change in IMU measurement errors $\mathbf{E}_a^B, \mathbf{E}_g^B$ follow Gaussian distribution, so we pre-sample them from $\mathcal{N}(0, \sigma_{\mathbf{E}_a} = 10^{-3})$ and $\mathcal{N}(0, \sigma_{\mathbf{E}_g} = 10^{-4})$, respectively. To ensure unbiased sampling, we use 3D Poisson disk sampling [51].

As $\mathbf{v}_B^W, \mathbf{p}_B^W, \mathbf{E}_a^B, \mathbf{E}_g^B$ are all Euclidean variables, we can use the previous way in [18] to evolve their states (Eq. (9)) and update state \mathbf{x} (Eq. (10)). Note that both uniform distribution and Gaussian distribution are closed under scaling and translation operations.




Fig. 3. Visualization of the rescaling of the state template (sphere) corresponding to rotations (represented with quaternions). The color map on the sphere depicts the Gauss map of the rotations. Initially, the rotations are perturbed with a random scaling (a). The dispersion of the rotations then converges over iteration steps (b). (c) shows anisotropic rescaling.

However, it is highly non-trivial to evolve and update unit quaternion variables, *i.e.*, $\mathbf{q}_B^W, \mathbf{g}^W$. To scale a unit quaternion $\mathbf{q} = q_w + q_x\mathbf{i} + q_y\mathbf{j} + q_z\mathbf{k}$, we devise a scale function $\phi_{\mathbf{q}}(\mathbf{q}, \mathbf{r}_q) : \mathcal{S}^3 \times \mathbb{R}^3 \rightarrow \mathcal{S}^3$, where \mathcal{S}^3 is the unit hypersphere that quaternions lie in. The imaginary part of $\mathbf{q}' = \phi_{\mathbf{q}, \mathbf{r}_q}$ output is simply the element-wise multiplication of the imaginary part of the quaternion and \mathbf{r}_q , and the real part is then computed using $q_w = \sqrt{1 - \mathbf{q}'_x{}^2 - \mathbf{q}'_y{}^2 - \mathbf{q}'_z{}^2}$. Finally, we can evolve the quaternion variables in the following way:

$$\mathbf{q}_j^{(i)} = \phi_{\mathbf{q}}(\mathbf{q}_{PST}^{(i)}, \mathbf{r}_{jq}) \otimes \mathbf{q}_{j-1}^*, \quad (12)$$

where $\mathbf{q}_{PST}^{(i)}$, \mathbf{q}_{j-1}^* , \mathbf{r}_{jq} are the corresponding portions in $\mathbf{x}_{PST}^{(i)}$, \mathbf{x}_{j-1}^* , \mathbf{r}_j . In fact, \mathbf{r}_q always stay between 0 and 1, leading to a decrease in the angle of rotation and attracting the rotation axis of \mathbf{q} to $\|\mathbf{r}_q\|$ (Fig. 3). Finally, to update a quaternion variable \mathbf{q}_j in \mathbf{x}_j^* , we need to perform weighted averaging over many quaternions, which we propose to do as the following:

$$\begin{aligned} \hat{\mathbf{q}}_j &= \sum_{\mathbf{x}_j^{(i)} \in \Omega_j^e} \bar{\mathbf{w}}^{(i)} \mathbf{q}_j^{(i)}, \\ \mathbf{q}_j^* &= \frac{\hat{\mathbf{q}}_j}{\|\hat{\mathbf{q}}_j\|}. \end{aligned} \quad (13)$$

There are other options for quaternions averaging, like Eigen-decomposition [52] or high dimensional rotation averaging [53]. However, for a relatively small rotation (tens of degrees) under the online tracking, the directly averaging (Eq. (13)) shows comparable performance with the fastest calculation (2.35x faster than [52] and 8.17x faster than [53]).

C. Active Subspace

To tackle the large dimensionality (18 DoFs) with random optimization, we propose the method of active subspace. More specially, we evaluate which dimensions in the state space are most crucial for the optimization and only optimize them.

We apply the sampling efficiency from sampling theory [54] and extend it to the dimensional sampling efficiency. The idea is that within the given sampling range/size, the favourable samples should lie as further as possible to the sampling centre, which means these samples introduces more posterior knowledge to the optimization. Given finite candidate




Fig. 4. The change of the sampling efficiency (Eq. (14)) of the 18 dimensions of the solution space over increasing iterations.

sampling states $(\mathbf{x}_d^{(i)})_{d=1:D}$ over given range \mathbf{r}_d , the dimensional sampling efficiency e_d of each dimensionality can be formulated as:

$$e_d = \frac{|\mathbf{x}_d^* - \bar{\mathbf{x}}_d|}{\mathbf{r}_d}, \quad (14)$$

where the $\bar{\mathbf{x}}_d$ is the sampling center of each dimensionality (local optimum at last step) and \mathbf{x}_d^* is new optimum calculated by Eq. (10).

Based on e_d , our method selects the top- k (in implementation, we set $k = 6$) dimensions with highest e_d as active subspace Ω^a . The active subspace represents the most crucial subspace that can be used to replace the whole state space in random optimization. But the dimensions in active subspace may not stay *high efficiency*, especially when close to a local optimal. Therefore, we still maintain a small search size for inactive dimensions and re-evaluate the e_d in each step for updating the active subspace. Here, we use the \mathbf{r}_d to control the active subspace and inactive subspace:

$$\mathbf{r}_d = \begin{cases} \hat{\mathbf{r}}_d + \epsilon, & \mathbf{x}_d \in \Omega^a \\ e_d^2 \hat{\mathbf{r}}_d + \epsilon, & \mathbf{x}_d \notin \Omega^a \end{cases}. \quad (15)$$

The square term of the e_d can significantly shorten the search size of inactive subspace and the $\hat{\mathbf{r}}$ is calculated from Eq. (11). We use a small number ϵ (10^{-3}) as a lower bound of search size to avoid degenerating dimensionality.

Fig. 4 visualizes the change of the sampling efficiency of different dimensions over increasing iteration steps. From the visualization, we can see that the dimensions corresponding to position contributes the most to the minimization of the cost function, while those corresponding to IMU measurement errors contributes the least. This means that the position part of the candidate states are initially far from the optimum. As the iteration proceeds, the sampling efficiency of all dimensions reduces to a low level, indicating convergence of optimization. In our implementation, we set the number of sampled candidate states to 3072 for a good trade-off between accuracy and speed.

VI. EXPERIMENTS

The experimental evaluation is designed and conducted to validate our key claims: first, the proposed active subspace random optimization can achieve good optimality for the high dimensional optimization problem of depth inertial odometry;

second, our depth inertial odometry is robust to low light conditions and fast camera motions without the need of warming up initialization.

In the video of our supplementary material, we show live demo of several challenging sequences. Especially on the camera shaking sequences, our method is able to achieve perhaps the fastest camera motion tracking ever seen for odometry in indoor environments.

A. Experimental Setup

We evaluate the methods on three indoor scene benchmark, including two public available benchmarks and one new benchmark, Robotic Arm Shaking (RAS) benchmark, contributed by this work. All benchmarks contain RGB, depth, and inertial measurements (though our method doesn't use RGB images).

Existing Benchmarks:

- The ETH3D benchmark [27] is captured by synchronized global shutter cameras together with an Asus Xtion Live Pro. We are especially interested in three challenging sequences, namely *camera_shake_1/2/3*, which contain severely fast motions without slow-motion warm up. For a comprehensive evaluation, we also test on 12 less aggressive sequences, named in four categories including *sofa*, *table einstein* and *mannequin_face*. They are explicitly marked as slow motion sequences in ETH3D.
- The VCU_RVI benchmark [55] contains a set of sequences captured by Structure Core under different environment conditions. We take three most challenging sequences, namely, lab-motion, lab-light and bumper, which are captured under fast camera motions, low light conditions and robot-scanning, respectively.

The Robotic Arm Shaking (RAS) benchmark. To thoroughly evaluate different odometry methods in the presences of fast motions, we curate a new benchmark, namely Robotic Arm Shaking (RAS). The RAS benchmark contains 24 sequences captured by Azure Kinect DK. We intentionally use a different depth sensor to the ones in the previous two datasets, which may help examine the methods against different depth sensor qualities. During capture, we fix the camera on a robot arm AUBO-i5 and move the arm following one pre-generated path which was pre-generated by human demonstration [56]. Unlike the existing datasets, we keep the path constant as a control variable but vary the lighting conditions (normal/low light conditions) and motion speeds (fast/slow with the maximum linear speeds above or below $2m/s$, respectively). The combination of two light conditions and two motion speeds yields four different settings, for each of which we capture six sequences to cover randomness in depth and IMU measurements. We obtain the ground truth trajectories using hand-eye calibration method [57].

Evaluation metrics. Following the existing works [24], [32], [55], we adopt the following evaluation metrics. 1) Absolute Trajectory Error (ATE): It measures the root mean squared distance between the ground truth and estimated trajectories. If a dataset contains many sequences, we use mean absolute

TABLE I

THE RESULTS OF ATE ON *camera_shake_1/2/3* SEQUENCES IN ETH3D. HERE CS_1/2/3 INDICATES WHICH SEQUENCE IN *camera_shake*, ATE IS MEASURED IN CM (THE LOWER, THE BETTER), AND – IN THE TABLE INDICATES FAILURE IN TRACKING. D-INERTIAL SLAM* INDICATES OUR OWN IMPLEMENTATION OF [31] AS AUTHOR'S CODE IS UNAVAILABLE.

	Inertial	Depth	RGB	CS_1	CS_2	CS_3
BAD SLAM [27]		✓	✓	–	–	–
DVO-SLAM [58]		✓	✓	9.40	–	–
ORB-SLAM2 [59]		✓	✓	–	6.89	–
ElasticFusion [60]		✓	✓	8.44	–	–
BundleFusion [61]		✓	✓	5.17	3.49	–
DROID-SLAM [62]		✓	✓	1.20	1.92	2.69
ROSEFusion [18]		✓		0.62	1.35	4.67
VINS-Mono [15]	✓		✓	–	–	–
VINS-RGBD [24]	✓	✓	✓	–	–	–
DUI-VIO [23]	✓	✓	✓	–	–	–
D-Inertial SLAM* [31]	✓	✓		–	–	–
Ours (ASRO-DIO)	✓	✓		0.59	0.98	2.37

trajectory error (mATE) to compute the mean ATE among all the successfully tracked trajectories. 2) Success Rate (SR): It is defined as the ratio of successfully tracked sequences to total sequences and thus indicates the robustness of navigation systems under challenging conditions. In both two metrics, a tracking is successful if it satisfies both the accuracy condition ($ATE < 5m$) and the completeness condition for at least half of the sequence. 3) Relative Pose Error (RPE): It evaluates the relative pose differences between the estimated and the ground-truth motion. This metric is well-suited for evaluating local trajectory accuracy.

B. Quantitative Comparisons

1) *Sequence camera_shake_1/2/3 in ETH3D:* The three *camera_shake* sequences mainly contain rapid shaking motions and exhibit increasing moving speeds from *camera_shake_1* to *camera_shake_3*. The shaking motions could cause significant difficulties in the initialization of inertial-aided methods. We compare our method with previous state-of-the-art (SOTA) RGB-D SLAM and inertial-aided methods based on traditional optimization methods (e.g., Gaussian-Newton and Levenberg-Marquardt). Besides, we also consider removing the dense photometric term of the tightly-coupled RGB-D-Inertial SLAM [31] to come up with a tightly-coupled DIO baseline. However, as the authors of [31] did not release their source code and the associated data for evaluation, we tried our best to faithfully implement their method and turn it into a depth-inertial version which we name as D-Inertial SLAM*. As shown in Table I, our method achieves successful tracking with the highest tracking accuracy on all three sequences.

Note that, all the other inertial-aided methods fail in all the three sequences. The reason is that the warm-up step of the existing inertial-aided methods are quite vulnerable to fast motion under which the input visual signals have been severely degraded by motion blur and the elevated nonlinearity in the optimization problem caused by large inter-frame rotations hampers their optimization. On the contrary, our method does not rely on RGB input and our random optimization is in

TABLE II

MEAN ATE (CM) ON 4 LESS AGGRESSIVE CATEGORIES INCLUDING *sofa*, *table*, *einstein* AND *mannequin_face* OF ETH3D. THE BEST AND THE SECOND BEST RESULTS FOR EACH CATEGORY ARE HIGHLIGHTED IN **blue** AND **green** COLORS, RESPECTIVELY.

	sofa	table	einstein	man._face
BAD SLAM	0.21	0.35	0.55	0.63
DROID-SLAM	0.61	0.90	0.54	0.39
ROSEFusion	1.39	2.57	1.19	2.90
DUI-VIO	7.89	3.72	1.84	2.43
D-Inertial SLAM*	-	5.77	6.13	5.63
Ours (ASRO-DIO)	0.57	0.83	0.91	1.59

principle more robust to high nonlinearity than second-order optimization methods.

Compared with ROSEFusion [18], which is the previous SOTA method on the three sequences and uses only depth, our method further incorporates IMU measurements into the optimization, taking the best of both worlds and thus leading to significant performance improvements, especially on the sequence *camera_shake_3* with the fastest camera motion.

2) *Slow camera motion sequences in ETH3D*: We have tested our method on 12 less aggressive sequences of ETH3D in Table II. They are explicitly marked as slow-motion sequences with no warm-up initialization. Here, we only show the average ATE of each category. The detailed results of the 12 sequences are provided in the supplemental material. Compared with the top two methods (BAD SLAM [27] and DROID-SLAM [62]), our method attains comparable performance. Note that our method does not include any post-processing of global pose optimization or loop closure which were employed by the two methods. Our method performs better than the inertial-based methods (DUI-VIO [23] and D-Inertial SLAM* [31]), clearly demonstrating its effectiveness in incorporating inertial information without needing a warm-up initialization.

3) *RAS*: We evaluate our method against the SOTA inertial-aided methods, including VINS-Mono, VINS-RGBD and DUI-VIO, on the 24 sequences of the RAS dataset under various combinations of acquisition conditions: slow motions (S), fast motions (F), normal lighting (L), and dark (low light) (D). “All” means the average results of all sequences. As shown in Table III, our method demonstrates strong robustness to low light conditions and fast motions by successfully tracking all the sequences while always maintaining the highest accuracy. Under low light conditions, our method is the only working method with perfect tracking, while all the other methods fail completely. The is because our method does not rely on RGB input by utilizing the depth-TSDF conformance as cost function, thus being naturally independent on lighting conditions. Under adequate lighting, we can see the switching from slow to fast motion leads to severe performance degradation for the alternative methods: VINS-Mono and VINS-RGBD both fail more frequently and their ATEs increase by 661.3% and 104.1%, respectively; the ATE of DUI-VIO also increases 195.7%. In contrast, our method incurs only a slightly decrease in accuracy and always maintains successful tracking (low ATEs). For the per-sequence results, please refer

TABLE III

COMPARING THE SUCCESS RATE (SR) AND ACCURACY (MEAN ATE (M)) OF FOUR METHODS ON THE ROBOTIC ARM SHAKING (RAS) BENCHMARK.

	VINS-Mono		VINS-RGBD		DUI-VIO		Ours (ASRO-DIO)	
	SR	ATE	SR	ATE	SR	ATE	SR	ATE
SL	83.3%	0.496	100%	0.073	100%	0.139	100%	0.044
FL	50.0%	3.28	83.3%	0.149	100%	0.411	100%	0.052
SD	0.00%		0.00%		0.00%		100%	0.042
FD	0.00%		0.00%		0.00%		100%	0.053
All	33.3%	0.154	45.83%	0.108	50%	0.275	100%	0.048

TABLE IV

COMPARISON THE SUCCESS RATE (SR) AND ACCURACY (MEAN ATE (M)) ON THE THREE CATEGORIES OF THE VCU_RVI DATASET: BUMPER, MOTION AND LIGHT.

		Bumper	Motion	Light	All
VINS-MONO	SR	80.0%	75.0%	75.0%	76.4%
	ATE	2.472	0.388	0.824	1.18
VINS-RGBD	SR	90.0%	33.3%	66.7%	61.7%
	ATE	0.514	0.247	1.301	0.763
DUI-VIO	SR	40.0%	91.7%	83.3%	76.4%
	ATE	0.336	0.154	0.921	0.477
DUI-VIO-NoBA	SR	40.0%	66.7%	75.0%	61.7%
	ATE	0.782	1.344	1.195	1.173
Ours (ASRO-DIO)	SR	100%	100%	100%	100%
	ATE	0.154	1.420	0.866	0.852

to the supplementary material.

4) *VCU_RVI sequences (lab-light, lab-motion and bumper)*: To avoid the bias induced by the quality of initialization, we skip the warm-up part and start from three uniformly sampled timestamps. The average results are reported. “All” means average results of all the three sequences. Note that ATE (m) is measured only for successfully tracked sequences. The results are demonstrated in Table IV.

Our method shows robustness over the three sequences, especially on the *bumper* sequences which contain severe camera shaking when the wheeled robot moves on a bumpy ground plane. However, on the sequences of *lab-light* and *lab-motion* whose trajectories contain multiple loops, the accuracy of our method is lower than the alternative methods. The reason is due to the lack of global pose optimization or loop closure in our method and these post-processing can be the following future works. To reveal the effect of these post-processing, we also compare to a downgraded version of DUI-VIO which turns off the loop closure and bundle adjustment, named as DUI-VIO-NoBA. This is achieved by setting a small number of keyframes ($n = 3$). Our method works significant better than DUI-VIO-NoBA in both robustness and accuracy. Furthermore, with two keyframes, the DUI-VIO fails on all the sequences.

5) *Comparison of fast initialization*: To test the ability of fast initialization, we evaluate our method and DUI-VIO on uniformly sampled short sequence of 10 seconds, to see if they can quickly start up. For each short sequence, we skip the warm-up part and start from a uniformly sampled timestamp. The results are reported in Table V. Note that RPE (m) is measured only for successfully tracked sequences. Here, DUI-VIO adopts a visual-inertial alignment algorithm for system initialization, which is widely used in existing optimization

TABLE V

COMPARISON OF SUCCESS RATE (SR) AND ACCURACY (MEAN RPE (M)) ON SHORT SEQUENCES SAMPLED FROM THREE VCU_RVI SEQUENCES. "NUM" IS THE NUMBER OF FRAGMENTS FOR EACH CATEGORY. THE AVERAGE RESULTS ARE REPORTED ON ALL FRAGMENTS. "ALL" MEANS AVERAGE RESULTS OVER ALL THREE CATEGORIES. RPE IS MEASURED ONLY FOR SUCCESSFULLY TRACKED SEQUENCES.

	Num	DUI-VIO		Ours (ASRO-DIO)	
		SR	RPE	SR	RPE
Bumper	15	93.3%	0.073	100%	0.020
Motion	18	77.8%	0.305	100%	0.056
Light	18	55.6%	0.068	100%	0.048
All	51	74.5%	0.153	100%	0.037




Fig. 5. Plots of cost value at different iteration steps over all frames for our ASRO and GN. For our method, we also show the range of cost value of all sampled states. The results are reported for both camera_shake_3 (fast motions) and sofa_3 (slow motions).

methods [15], [32]. Our method, on the other hand, does not require an explicit initialization step. The results in Table V show that our method achieves better performance on the short sequences with 100% success rate and thus faster initialization. More details can be found in the supplementary material.

6) Comparison with second-order optimization method:

To verify the superiority of our method (ASRO) over the second-order optimization methods, we conduct a comparison based on the same cost function proposed in this work (Eq. (8)). To do so, we build a Gauss-Newton (GN) based method where we use second-order derivatives of depth-to-TSDF based objective, similar to [26]. The two methods are tested on two sequences, *i.e.*, camera_shake_3 (fast motion) and sofa_3 (slow motion) of the ETH3D benchmark. We compare the performance of optimization by evaluating the converging cost (Eq. (8)) for each frame with the same initial (center) state. To make a fair comparison, we let both methods iterate sufficiently to converge.

As shown in Fig. 5 (left, fast motion), GN quickly converges to an average cost of 0.47, which is significantly higher than 0.217 of ours. The reason is that the least square optimization finds difficulty in handling highly nonlinear situations, *e.g.*, large rotation, as found by many previous works [18], [63]. Another reason is that, with the depth-TSDF based objective, only 68.3% of the points of a depth frame lie in the valid region of TSDF (within 20cm from the surface) to ensure a valid gradient evaluation. The limited gradient information makes GN easily get stuck in local optima. In Fig. 5 (right, slow motion), the cost values of GN are mostly larger than the medium half of ASRO (see the slim boxes). This demonstrates

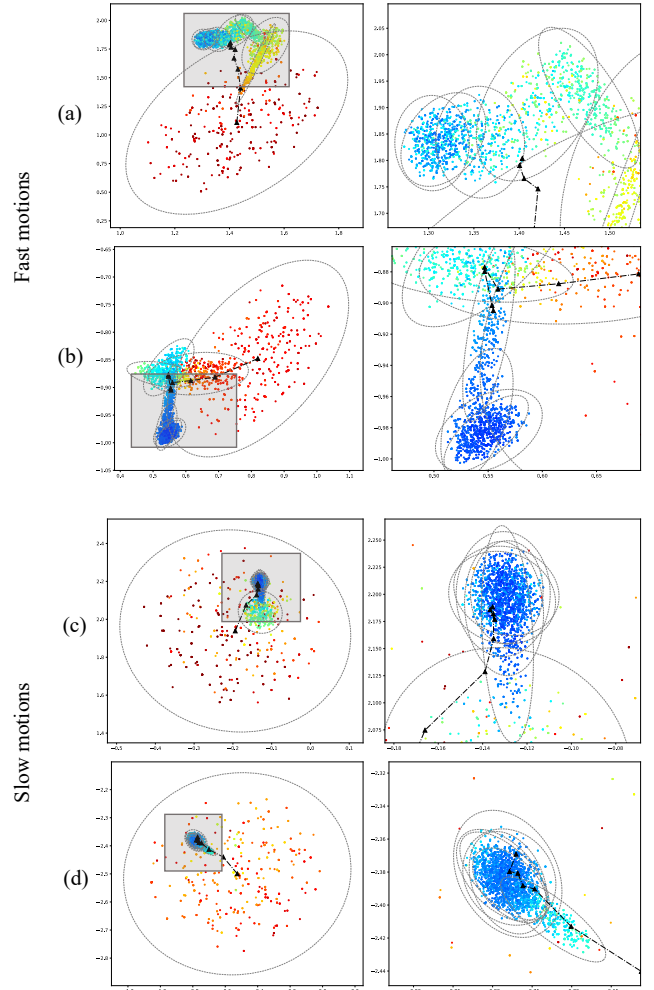


Fig. 6. 2D visualization of optimization process of our ASRO and GN for a given frame. For our method, we plot the cost value (color-coded) of sampled states at different iteration steps of optimization. The samples of each iteration are circled with a grey ellipse. For GN, we plot the states being optimized with black triangles (connected sequentially with black line segments). The right plot in each row is a zoom-in view of the grey box in the left plot.

that our method outperforms the GN with a significantly faster convergence on two separate sequences.

In the plots of Fig. 6, we provide a breakdown study of the performance of ASRO vs. GN. The plots visualize the progressive evolution of state estimation. The states are embedded in 2D using sparse random projection [64]. Starting with the same initial state, GN quickly gets stuck in a local optimum when initialized far from the global optimum. In contrast, ASRO realizes a good trade-off between exploration and exploitation and converges to much better local minima for both slow- and fast-motion sequences.

C. Qualitative Results

1) Real captured sequences under challenging conditions:

To demonstrate the superiority of our method under low light conditions and fast camera motions without warm up initialization, we further collect the following extremely challenging trajectories using man-hold Azure Kinect DK:

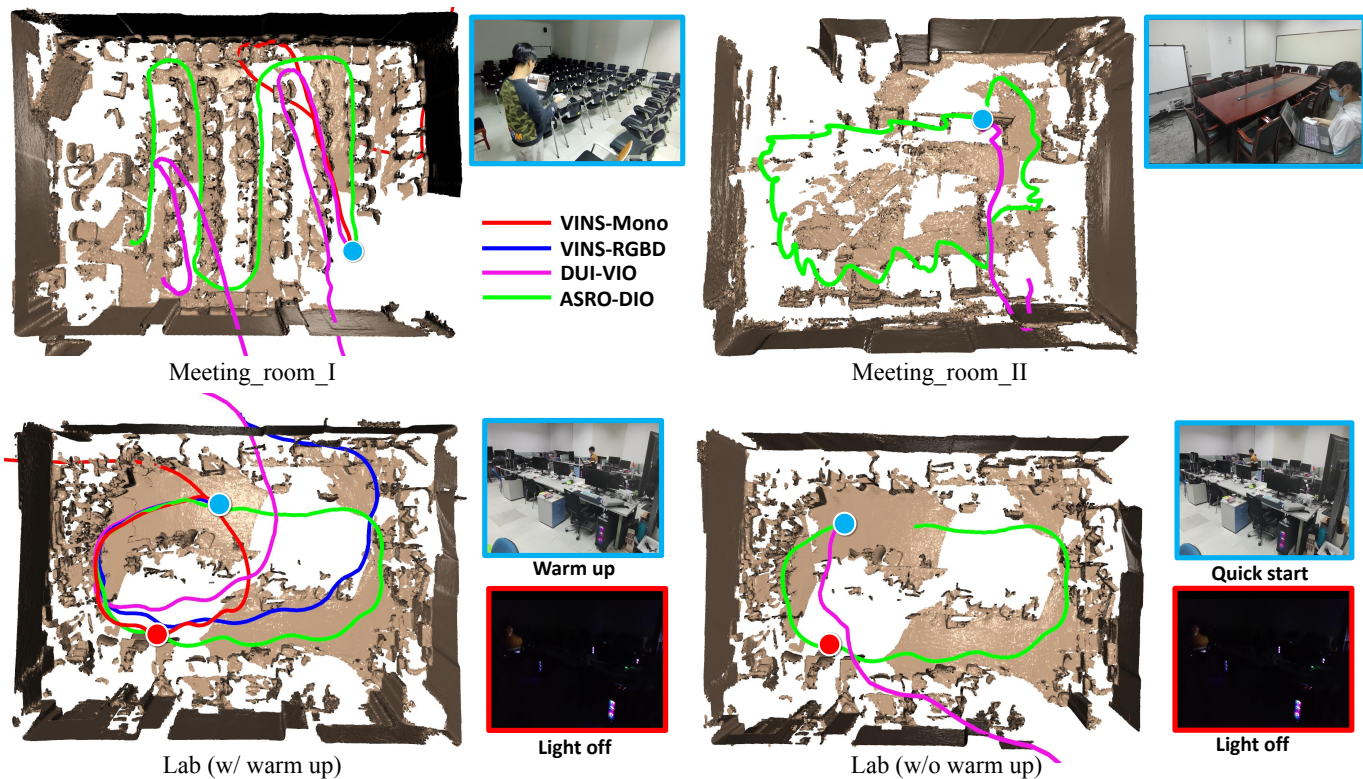


Fig. 7. Visual results of tracking and mapping on four challenging sequences. For each sequence, we compare the tracked trajectories of our method and three alternatives. The trajectory is not drawn when the method failed on a sequence completely. The 3D reconstruction is obtained with our method, via meshing the zero-crossing surface of the TSDF.

- *Lab*: One sequence was captured under fast camera motions but with slow motions at the beginning (blue point) for warm up initialization. At the half of the scanning (red point), the light was turned off (See Fig. 7 Lab (w/ warm up)). We also get rid of the slow motions from the beginning and start the methods directly under the fast motions (see Fig. 7 Lab (w/o warm up)).
- *Meeting_room*: two visual comparisons on challenging sequences captured under fast camera motions (Fig. 7 *Meeting_room_I* and *Meeting_room_II*).

Here we provide qualitative results (please refer to the video in supplementary material for more information). As shown in Fig. 7 Lab (w/ warm up)), DUI-VIO, VINS-RGBD and VINS-Mono all fail after the light is turned off, due to their failure to obtain informative visual features under low light conditions. Meanwhile, in Fig. 7 Lab (w/o warm up), they also suffer from erroneous initialization under the fast motions. And even after the initialization, the output trajectory of DUI-VIO still suffer from large drift under fast motions. In contrast, our method performs well under all the challenging conditions.

2) *ETH3D sequences*: For a more comprehensive study, we further provide the results of our method under both the fast-motion and slow-motion sequences on ETH3D benchmark. For fast-motion sequences, we select the camera_shake_1 and camera_shake_2 for evaluation. Specifically, We plot the 3D visual trajectories of our method against the ground-truth trajectories in Fig 8. Here our method obtains robust tracking results while the existing inertial-based methods failed




Fig. 8. 3D plots of trajectories of our method on camera_shake_1/2 [27]. The ground-truth is indicated as the grey dash line. And the colors on our trajectories encode the relative error against ground-truth (the closer to blue, the smaller).

(Table I). And a related breakdown study is conducted in Fig. 10, where we plot the position and orientation errors along the time. From Fig. 10, the well-aligned curves demonstrate that our method can perform robust tracking under vibratory shaking motions. For slow-motion sequences, we compare the trajectories of our approach with DUI-VIO. This experiment shares the same slow-motion sequences in Table II. The results are shown in Fig. 9. Our approach demonstrates more accurate tracking performance against DUI-VIO. Moreover, our trajectories are more complete because of the fast and robust initialization at the beginning of each given sequence. More results on ETH3D can be found in the supplemental

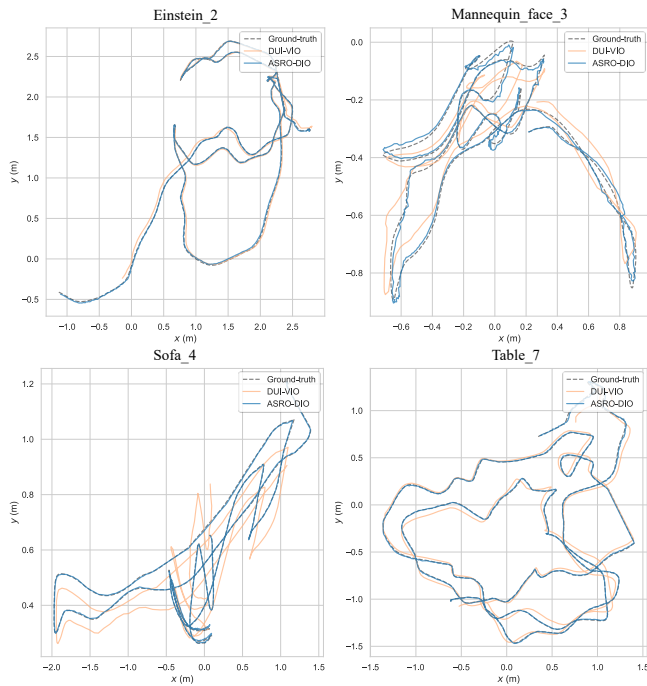


Fig. 9. 2D plots of trajectories of our method on slow-motion sequences of ETH3D [27]. The ground-truth is indicated as the grey dash line. And the colors of our approach and DUI-VIO are encoded in orange and blue, respectively.

material.

D. Ablation Studies on the RAS Benchmark

In this section, we conduct ablation studies on the fast-motion sequences of RAS to validate the necessity of the two key designs of our method, *i.e.*, active subspace and pre-defined sampling space (Sec. V-B). In Table VI, we add the designs to the basic random optimization framework (Sec V-A), including the use of gyroscope readings (Gyr) and accelerometer readings (Acc), and the utilization of active subspace (AS) and pre-defined space (PS), tested for different DoFs of search space. All methods share the same parameter setting, including the number of sampling states (3072), the initial scaling factor (following ROSEFusion [18]) and the maximum iteration times (20). M1, M2 and M3 share the same random optimization framework of ROSEFusion performing uniform sampling for all dimensions. Comparing M1 with M2, we notice that simply leveraging gyroscope measurements and thus incorporating \mathbf{E}_g^B into the state space, forming a 9-DoF optimization, could lead to higher accuracy. Here, the gyroscope provides a rough prior of the change in rotation between consecutive frames, which can benefit the optimization at the cost of a slight increase of the dimensionality of state space. However, when we further introduce accelerometer measurements, as in M3, the state space becomes 18D, immediately causing a significant performance drop, with an accuracy even worse than M1. By introducing active subspace in M4, the accuracy is greatly improved and is further boosted by utilizing pre-defined state space (M5).




Fig. 10. Plots of position and orientation errors of our method versus ground-truth on camera_shake_2.

TABLE VI
ABLATION STUDY OF VARIOUS DESIGN CHOICES, EXPERIMENT SETTINGS AND ALGORITHMIC COMPONENTS.

Method	DoFs	Gyr	Acc	AS	PS	Mean ATE
M1	6DoFs					5.27cm
M2	9DoFs	✓				4.94cm
M3	18DoFs	✓	✓			5.99cm
M4	18DoFs	✓	✓	✓		4.89cm
M5	18DoFs	✓	✓	✓	✓	4.32cm

In Fig. 11, we plot the cost function values (Eq. (IV-D)) of all frames and all sequences at different iteration steps with M3 and M5. For each method, we plot both the average and the range of cost values of all states (left) and the individual boxplot (middle and right). In accord with the result in Fig. 11, we notice that our method achieves fast convergence with higher confidence. And compared to the smooth convergence of M3, our method shows the *rugged* convergence which could be caused by the active subspace update.




Fig. 11. Plots of the average and the range of cost function values at different iteration steps for our method (blue) and random optimization in 18-DoF state space (orange).

E. Runtime

We have implemented our core algorithm based on C++ and CUDA. Both the main optimization pipeline and the volumetric fusion run on a workstation with an Intel® Core™ i7-5930K CPU @ 3.50GHz \times 12 with 32G RAM and an Nvidia GeForce RTX 2080 SUPER GPU with 8G memory. To enable flexible scanning, we implemented a front-end program running on a laptop for RGB-D-inertial measurements capturing, compressing, and streaming to the workstation via WiFi. Our pipeline runs with a framerate of 30Hz for all shown test sequences. The readers are welcome to watch our accompanying video.


VII. CONCLUSION

We have presented an efficient and robust depth inertial odometry method based on random optimization with active subspace. Through identifying active subspace and sampling candidate states only within it, our method is able to explore the 18D state space of DIO efficiently and achieve good optimality under extremely challenging conditions. Although we adopt a depth-TSDF-based cost function based on which fitness evaluation is parallel-friendly, our method is general and adaptable to other settings. In the future, we would like to investigate the integration of visual features especially when the camera motion slows down, to ensure high tracking accuracy at all time, thus enjoying the strength of both worlds. We would also like to exploit random optimization in back-end optimizations such as bundle adjustment and pose graph optimization. Finally, our method could be further improved by more efficient sampling strategies like CMA-ES [65].

REFERENCES

- [1] G. Huang, “Visual-inertial navigation: A concise review,” in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 9572–9582.
- [2] D. Scaramuzza and Z. Zhang, “Visual-inertial odometry of aerial robots,” *arXiv preprint arXiv:1906.03289*, 2019.
- [3] J. Kim and S. Sukkarieh, “Real-time implementation of airborne inertial-slam,” *Robotics and Autonomous Systems*, vol. 55, no. 1, pp. 62–71, 2007.
- [4] A. I. Mourikis and S. I. Roumeliotis, “A multi-state constraint kalman filter for vision-aided inertial navigation,” in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 2007, pp. 3565–3572.
- [5] M. Bryson and S. Sukkarieh, “Observability analysis and active control for airborne slam,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, no. 1, pp. 261–280, 2008.
- [6] Z. Huai and G. Huang, “Robocentric visual-inertial odometry,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6319–6326.
- [7] A. I. Mourikis, N. Trawny, S. I. Roumeliotis, A. E. Johnson, A. Ansar, and L. Matthies, “Vision-aided inertial navigation for spacecraft entry, descent, and landing,” *IEEE Transactions on Robotics*, vol. 25, no. 2, pp. 264–280, 2009.
- [8] M. Brossard, S. Bonnabel, and A. Barrau, “Invariant kalman filtering for visual inertial slam,” in *2018 21st International Conference on Information Fusion (FUSION)*. IEEE, 2018, pp. 2021–2028.
- [9] S. Ebcin and M. Veth, “Tightly-coupled image-aided inertial navigation using the unscented kalman filter,” in *Proceedings of the 20th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2007)*, 2007, pp. 1851–1860.
- [10] D. Strelow and S. Singh, “Motion estimation from image and inertial measurements,” *The International Journal of Robotics Research*, vol. 23, no. 12, pp. 1157–1195, 2004.
- [11] V. Indelman, S. Williams, M. Kaess, and F. Dellaert, “Information fusion in navigation systems via factor graph based incremental smoothing,” *Robotics and Autonomous Systems*, vol. 61, no. 8, pp. 721–738, 2013.
- [12] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, “On-manifold preintegration for real-time visual-inertial odometry,” *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 1–21, 2016.
- [13] K. Eickenhoff, P. Geneva, and G. Huang, “Closed-form preintegration methods for graph-based visual-inertial navigation,” *The International Journal of Robotics Research*, vol. 38, no. 5, pp. 563–586, 2019.
- [14] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, “Keyframe-based visual-inertial odometry using nonlinear optimization,” *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.
- [15] T. Qin, P. Li, and S. Shen, “Vins-mono: A robust and versatile monocular visual-inertial state estimator,” *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [16] B. M. Bell and F. W. Cathey, “The iterated kalman filter update as a gauss-newton method,” *IEEE Transactions on Automatic Control*, vol. 38, no. 2, pp. 294–297, 1993.
- [17] Y. Akimoto, A. Auger, and N. Hansen, “Cma-es and advanced adaptation mechanisms,” in *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion*, 2016, pp. 533–562.
- [18] J. Zhang, C. Zhu, L. Zheng, and K. Xu, “Rosefusion: random optimization for online dense reconstruction under fast camera motion,” *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, pp. 1–17, 2021.
- [19] Z. Wang, M. Zoghi, F. Hutter, D. Matheson, and N. De Freitas, “Bayesian optimization in high dimensions via random embeddings,” in *Twenty-Third international joint conference on artificial intelligence*, 2013.
- [20] R. M. Gower, D. Kovalev, F. Lieder, and P. Richtárik, “Rsn: randomized subspace newton,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 616–625.
- [21] J. Lacotte, M. Pilanci, and M. Pavone, “High-dimensional optimization in adaptive random subspaces,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 10 847–10 857.
- [22] K. M. Choromanski, A. Pacchiano, J. Parker-Holder, Y. Tang, and V. Sindhvani, “From complexity to simplicity: Adaptive es-active subspaces for blackbox optimization,” *Advances in Neural Information Processing Systems*, vol. 32, pp. 10 299–10 309, 2019.
- [23] H. Zhang and C. Ye, “Dui-vio: Depth uncertainty incorporated visual inertial odometry based on an rgb-d camera,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5002–5008.
- [24] Z. Shan, R. Li, and S. Schwertfeger, “Rgbd-inertial trajectory estimation and mapping for ground robots,” *Sensors*, vol. 19, no. 10, p. 2251, 2019.
- [25] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, “Kinectfusion: Real-time dense surface mapping and tracking,” in *2011 10th IEEE international symposium on mixed and augmented reality*. IEEE, 2011, pp. 127–136.
- [26] E. Bylow, J. Sturm, C. Kerl, F. Kahl, and D. Cremers, “Real-time camera tracking and 3d reconstruction using signed distance functions,” in *Robotics: Science and systems (RSS) conference 2013*. Robotics: Science and Systems, 2013.
- [27] T. Schops, T. Sattler, and M. Pollefeys, “Bad slam: Bundle adjusted direct rgb-d slam,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 134–144.
- [28] C. X. Guo and S. I. Roumeliotis, “Imu-rgbd camera 3d pose estimation and extrinsic calibration: Observability analysis and consistency

- improvement,” in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 2935–2942.
- [29] N. Brunetto, S. Salti, N. Fioraio, T. Cavallari, and L. Stefano, “Fusion of inertial and visual measurements for rgb-d slam on mobile devices,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2015, pp. 1–9.
- [30] Z. Zhu, F. Xu, C. Yan, X. Hao, X. Ji, Y. Zhang, and Q. Dai, “Real-time indoor scene reconstruction with rgb-d and inertial input,” in *2019 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2019, pp. 7–12.
- [31] T. Laidlow, M. Bloesch, W. Li, and S. Leutenegger, “Dense rgb-d-inertial slam with map deformations,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6741–6748.
- [32] A. Tyagi, Y. Liang, S. Wang, and D. Bai, “Dvio: Depth-aided visual inertial odometry for rgb-d sensors,” in *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2021, pp. 193–201.
- [33] M. Nießner, A. Dai, and M. Fisher, “Combining inertial navigation and icp for real-time 3d surface reconstruction,” in *Eurographics (Short Papers)*. Citeseer, 2014, pp. 13–16.
- [34] M. Camurri, S. Bazeille, D. G. Caldwell, and C. Semini, “Real-time depth and inertial fusion for local slam on dynamic legged robots,” in *2015 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2015, pp. 259–264.
- [35] C. Andrieu and A. Doucet, “Particle filtering for partially observed gaussian state space models,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 64, no. 4, pp. 827–836, 2002.
- [36] C. Choi and H. I. Christensen, “Robust 3d visual tracking using particle filtering on the special euclidean group: A combined approach of keypoint and edge features,” *The International Journal of Robotics Research*, vol. 31, no. 4, pp. 498–519, 2012.
- [37] A. Gil, Ó. Reinoso, M. Ballesta, and M. Juliá, “Multi-robot visual slam using a rao-blackwellized particle filter,” *Robotics and Autonomous Systems*, vol. 58, no. 1, pp. 68–80, 2010.
- [38] G. Grisetti, C. Stachniss, and W. Burgard, “Improved techniques for grid mapping with rao-blackwellized particle filters,” *IEEE transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.
- [39] E. Arnaud and E. Mémin, “An efficient rao-blackwellized particle filter for object tracking,” in *IEEE International Conference on Image Processing 2005*, vol. 2. IEEE, 2005, pp. II–426.
- [40] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, “Poserbpf: A rao-blackwellized particle filter for 6-d object pose tracking,” *IEEE Transactions on Robotics*, 2021.
- [41] M. Nieto, A. Cortés, O. Otaegui, J. Arróspide, and L. Salgado, “Real-time lane tracking using rao-blackwellized particle filter,” *Journal of Real-Time Image Processing*, vol. 11, no. 1, pp. 179–191, 2016.
- [42] G. Huang, “Particle filtering with analytically guided sampling,” *Advanced Robotics*, vol. 31, no. 17, pp. 932–945, 2017.
- [43] M. R. Bonyadi and Z. Michalewicz, “Particle swarm optimization for single objective continuous space problems: a review,” *Evolutionary computation*, vol. 25, no. 1, pp. 1–54, 2017.
- [44] H.-G. Beyer and H.-P. Schwefel, “Evolution strategies—a comprehensive introduction,” *Natural computing*, vol. 1, no. 1, pp. 3–52, 2002.
- [45] J. Matyas *et al.*, “Random optimization,” *Automation and Remote control*, vol. 26, no. 2, pp. 246–253, 1965.
- [46] P. G. Constantine, *Active subspaces: Emerging ideas for dimension reduction in parameter studies*. SIAM, 2015.
- [47] J. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization,” *Journal of machine learning research*, vol. 13, no. 2, 2012.
- [48] S. S. Vempala, *The random projection method*. American Mathematical Soc., 2005, vol. 65.
- [49] F. Spitzer, *Principles of random walk*. Springer Science & Business Media, 2001, vol. 34.
- [50] K. Shoemake, “Uniform random rotations,” in *Graphics Gems III (IBM Version)*. Elsevier, 1992, pp. 124–132.
- [51] R. Bridson, “Fast poisson disk sampling in arbitrary dimensions,” *SIGGRAPH sketches*, vol. 10, no. 1, 2007.
- [52] F. L. Markley, Y. Cheng, J. L. Crassidis, and Y. Oshman, “Averaging quaternions,” *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 4, pp. 1193–1197, 2007.
- [53] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, “On the continuity of rotation representations in neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [54] D. K. Nassiuma, “Survey sampling: Theory and methods,” 2001.
- [55] H. Zhang, L. Jin, and C. Ye, “The vcu-rvi benchmark: Evaluating visual inertial odometry for indoor navigation applications with an rgb-d camera,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6209–6214.
- [56] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, “Survey: Robot programming by demonstration,” Springer, Tech. Rep., 2008.
- [57] R. Horaud and F. Dornaika, “Hand-eye calibration,” *The international journal of robotics research*, vol. 14, no. 3, pp. 195–210, 1995.
- [58] C. Kerl, J. Sturm, and D. Cremers, “Dense visual slam for rgb-d cameras,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 2100–2106.
- [59] R. Mur-Artal and J. D. Tardós, “Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras,” *IEEE transactions on robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [60] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison, “Elasticfusion: Dense slam without a pose graph.” *Robotics: Science and Systems*, 2015.
- [61] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, “Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration,” *ACM Transactions on Graphics (ToG)*, vol. 36, no. 4, p. 1, 2017.
- [62] Z. Teed and J. Deng, “Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [63] J. Huang, X. Shi, X. Liu, K. Zhou, L.-Y. Wei, S.-H. Teng, H. Bao, B. Guo, and H.-Y. Shum, “Subspace gradient domain mesh deformation,” in *ACM SIGGRAPH 2006 Papers*, 2006, pp. 1126–1134.
- [64] P. Li, T. J. Hastie, and K. W. Church, “Very sparse random projections,” in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006, pp. 287–296.
- [65] N. Hansen, “The cma evolution strategy: A tutorial,” *arXiv preprint arXiv:1604.00772*, 2016.



Jiazhao Zhang received the B.Eng. degree in software engineering from Shandong University, Jinan, China, in 2019, and the M.S. degree in computer science from National University of Defense technology, Changsha, China in 2021. He is currently working toward the Ph.D. degree in computer science with Peking University, Beijing, China.

His current research interests include state estimation, robust odometry systems and dense reconstruction.



Yijie Tang received the B.Eng. degree in Software Engineering from Wuhan University, Wuhan, China, in 2019, and he is currently working toward the M.S. degree in computer science from the National University of Defense Technology, Changsha, China.

His research interests include online 3D reconstruction, Simultaneous Localization and Mapping.



He Wang is a tenure-track assistant professor in the Center on Frontiers of Computing Studies at Peking University, where he founds and leads Embodied Perception and InteraCtion (EPIC) Lab.

Prior to joining Peking University, he received his Ph.D. degree from Stanford University in 2021 under the advisory of Prof. Leonidas J. Guibas and his Bachelor's degree in 2014 from Tsinghua University.

His research interests span 3D vision, robotics, and machine learning, with a special focus on embodied AI. He has published more than 20 papers on top vision and learning conferences (CVPR/ICCV/ECCV/NeurIPS) with 8 of his works receiving CVPR/ICCV orals and one work receiving Eurographics 2019 best paper honorable mention. He serves as an area chair in CVPR 2022 and WACV 2022.



Kai Xu is a Professor at the School of Computer, National University of Defense Technology, where he received his Ph.D. in 2011. He is an adjunct professor of Simon Fraser University. During 2017-2018, he was a visiting research scientist at Princeton University.

His research interests include geometric modeling and shape analysis, especially on data-driven approaches to the problems in those directions, as well as 3D vision and robotic applications.

He has published over 100 research papers, including 27 SIGGRAPH/TOG papers. He serves on the editorial board of ACM Transactions on Graphics, Computer Graphics Forum, Computers Graphics, and The Visual Computer. He also served as program co-chair and PC member for several prestigious conferences. He has co-organized several SIGGRAPH courses, Eurographics STAR tutorials and CVPR workshops. His research work can be found in his personal website:www.kevinkaixu.net