

Debiasing 6-DOF IMU via Hierarchical Learning of Continuous Bias Dynamics

Ben Liu¹, Tzu-Yuan Lin², Wei Zhang^{1†}, and Maani Ghaffari²

[†]Corresponding Author

¹Southern University of Science and Technology

Email: liub2021@mail.sustech.edu.cn; zhangw3@sustech.edu.cn

²University of Michigan

Email: tzuyuan@umich.edu; maanigj@umich.edu

Abstract—This paper develops a deep learning approach to the online debiasing of IMU gyroscopes and accelerometers. Most existing methods rely on implicitly learning a bias term to compensate for raw IMU data. Explicit bias learning has recently shown its potential as a more interpretable and motion-independent alternative. However, it remains underexplored and faces challenges, particularly the need for ground truth bias data, which is rarely available. To address this, we propose a neural ordinary differential equation (NODE) framework that explicitly models continuous bias dynamics, requiring only pose ground truth, often available in datasets. This is achieved by extending the canonical NODE framework to the matrix Lie group for IMU kinematics with a hierarchical training strategy. The validation on two public datasets and one real-world experiment demonstrates significant accuracy improvements in IMU measurements, reducing errors in both pure IMU integration and visual-inertial odometry.

I. INTRODUCTION

Inertial Measurement Units (IMUs) are essential in robotic applications, providing angular velocity and acceleration measurements that support various state estimation tasks. A notable use of IMUs is in Visual-Inertial Odometry (VIO) [32, 33, 34, 15, 39], where IMU and camera data are fused to estimate a robot’s orientation, velocity, and position. However, low-cost IMUs are often prone to significant noise and bias [22], leading to inaccurate measurements that can compromise VIO performance. This becomes even more critical in scenarios where cameras fail due to adverse environmental conditions [6, 27], leaving the IMU as the sole source of odometry information. In such scenarios, the odometry output heavily depends on the accuracy of IMU data. Therefore, deriving accurate measurements from raw IMU data that is noisy and biased is critical for robust state estimation.

To address the inaccuracies inherent in raw IMU data, *calibration* methods are utilized to enhance measurement accuracy for a specific IMU device [9]. The typical method uses a linear model to model axis misalignment and scale factors, while the bias is often modeled as a constant or a Brownian motion process [35, 37]. However, the IMU model is complex and difficult to represent accurately, particularly because the bias can be time-varying and influenced by factors such as temperature and vibration. With the advancement of deep learning, researchers have begun leveraging neural networks to




Fig. 1: Training process for the explicit evolution of bias. The subscript notation $u_{n:k}$ represents $u_n, u_{n+1}, \dots, u_k$. The existing method [6] models bias evolution using a discrete approach, which requires ground-truth bias values during training. In contrast, our method employs a continuous model to capture bias dynamics and does not rely on ground-truth bias for training.

calibrate IMUs, aiming to capture ignored or simplified components in model-based methods. Brossard et al. [5] employs a convolutional neural network to directly learn a correction term for the gyroscope based on local windows of IMU data. Similar methods have been extended to both gyroscopes and accelerometers [29]. However, it has been noted that such methods may not effectively distinguish whether the network learns deviation characteristics from the motion pattern or from the IMU itself [6]. To address this issue, Buchanan et al. [6] proposed explicitly modeling the IMU bias evolution using a neural network to achieve motion-independent calibration. However, this method requires ground truth bias data from sensor fusion with other sensors like LiDAR or cameras. The accuracy of bias estimation thus depends heavily on the fusion algorithm, and the efforts to obtain the bias ground truth are non-trivial.

In this work, we propose a framework to model the continuous dynamics of the bias *without* the need for ground truth bias data. The contributions of this work can be summarized as follows:

1. We propose a novel loss formulation that bypasses the requirement for bias ground truth during training. Moreover,

its hierarchical structure makes network design and tuning more efficient.

2. We model the bias dynamics as a vector field in the Lie algebra with a learning approach, allowing us to utilize canonical neural ordinary differential equation tools effectively. The vector field formulation allows us to obtain continuous dynamics of the bias, resulting in lightweight networks that deliver superior performance.
3. We evaluate our method on two public datasets and one real-world experiment, demonstrating that the proposed method provides more accurate IMU measurements than existing methods and has generalization capability.
4. We offer an open-source project available at https://github.com/UMich-CURLY/Debias_IMU.git.

II. RELATED WORK

We review some learning-based methods in IMU estimation and their relationship with our method.

A. Learning Motion Pattern from IMU

One direction is to extract the motion pattern information from IMU data. Early work Yan et al. [41] utilized a combination of Support Vector Machine (SVM) and Support Vector Regression (SVR) to predict a person's motion displacement based on local windowed IMU data. Building on this approach, the same authors later employed a neural network to perform a similar task [20]. Both studies only focused specifically on 2D motion scenarios. Subsequently, Liu et al. [30] advanced this line of research by integrating the learned displacement into an Extended Kalman Filter (EKF), using the displacement estimates as updates with IMU-based predictions. These methods assume that motion within a given category exhibits finite patterns that can be effectively regressed using IMU data. However, these approaches require large datasets for training and are sensitive to the motion context [6]; for instance, a model trained on flat-ground motion may not generalize well to stair-climbing scenarios. Therefore, when the training data is not sufficient and motion contexts are complex, device-specific calibration is necessary. In addition, some of these methods use calibrated IMU data as input [20, 30], focusing on incorporating the motion information to improve estimation accuracies. On the contrary, our method focuses on denoising and debiasing the raw IMU measurements.

B. Learning Calibration Model for IMU

Another line of research assumes that neural networks can provide accurate IMU measurements from raw inaccurate data, typically restricted to a single device. These refined IMU measurements can be used in inertial-based odometry, such as VIO, or to learn motion patterns as in prior methods.

Esfahani et al. [13] introduced a Long Short-Term Memory (LSTM) framework to denoise gyroscope data and directly estimate orientation, where the orientation integration process is implicitly learned by the neural network. Similarly, Sun et al. [38] applied an LSTM to predict both orientation and position. To improve short-term accuracy, they incorporated an extra

EKF that uses raw IMU data to correct orientation, observing that the LSTM performs well over long durations but poorly in shorter timeframes. These methods combine denoising and integration into a single framework, which is hard to separate. In contrast, Brossard et al. [5] focused solely on gyroscope denoising by modeling bias and noise as a single term. By compensating for this term, the clean angular velocity can be obtained and can subsequently be numerically integrated into orientation. Their approach uses a dilated convolutional neural network that takes windowed IMU data as input. Building on this framework, Huang et al. [21] explored the use of a temporal convolutional network structure and demonstrated improved performance, while Liu et al. [29] extended the approach to handle both gyroscope and accelerometer data simultaneously. To ensure the generality, Zhang et al. [42] construct the loss function based on partial integration terms and employ a recurrent neural network to estimate accurate IMU measurements. All these methods implicitly learn the noise or bias of IMU data, regardless of whether the numerical integration is explicitly processed or implicit embedded in the neural network.

More recently, Buchanan et al. [6] proposed explicitly modeling bias evolution using a transformer or LSTM, which is embedded within a Maximum-a-Posteriori (MAP) framework alongside other sensor inputs. While this explicit bias modeling demonstrates generalizability across motion patterns, it relies on accurate bias ground truth derived from additional sensor fusion algorithms, which are challenging to obtain and sensitive to the fusion method.

In this work, we also focus on explicitly learning the bias evolution process but with a different approach: modeling the continuous dynamics of the bias. This framework reduces neural network complexity and eliminates the need for bias ground truth, making it lightweight and more practical.

III. PROBLEM STATEMENT

A. IMU Model

An IMU measures angular velocity, denoted as ω_t , and linear acceleration (include gravity), denoted as a_t , both expressed in the IMU frame. A commonly adopted measurement model for these quantities is presented in [4, 18, 25]:

$$\tilde{\omega}_t = \omega_t + b_t^g + n_t^g, \quad \tilde{a}_t = a_t + b_t^a + n_t^a, \quad (1)$$

where the superscript $(\cdot)$ indicates measured quantities. $\omega_t, a_t \in \mathbb{R}^3$ are the true angular velocity and true linear acceleration, respectively, both expressed in the IMU frame. n_t^g and n_t^a are the Gaussian noise. $b_t^g, b_t^a \in \mathbb{R}^3$ are the biases of the gyroscope and accelerometer, respectively. The bias is modeled as a Brownian motion as

$$\dot{b}_t^g = \eta_g, \quad \dot{b}_t^a = \eta_a, \quad (2)$$

where η_g and η_a follow the zero-mean Gaussian distribution. However, this simplified model can significantly deviate from the true behavior, resulting in inaccurate estimates of w_t, a_t .

The IMU measurements are often utilized with the following kinematics of a single rigid body [18]:

$$\dot{R}_t = R_t \omega_t^\times, \quad \dot{v}_t = R_t a_t + g, \quad \dot{p}_t = v_t, \quad (3)$$

where $R_t \in \text{SO}(3)$ is the orientation of the body frame. $\text{SO}(3)$ is the *special orthogonal group*. $v_t \in \mathbb{R}^3$ is the linear velocity of the body frame expressed in the world frame, $p_t \in \mathbb{R}^3$ is the position of the body frame. $(\cdot)^\times : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ is a cross-product operator that satisfies $a^\times b = a \times b$, $\forall a, b \in \mathbb{R}^3$, where $\mathfrak{so}(3)$ is the vector space of 3-by-3 *skew-symmetric matrices* which is also the *Lie algebra* of $\text{SO}(3)$. $g \in \mathbb{R}^3$ is the gravity expressed in the world frame. In this work, we assume the IMU frame is the body frame.

Accurate pose (represented as a triple (R, v, p) in the paper) is critical in robotics applications, which often involve integrating IMU measurements. However, direct integration of the measured $\tilde{\omega}_t$ and $\tilde{a}_t$ in (3) leads to rapidly accumulating errors in pose over time due to noise and biases. Therefore, to get accurate integration results, it is necessary to develop a “filter” to give accurate or clean angular velocity and acceleration from raw IMU data.

For simplicity of notation, we will use the subscript x_k to represent the value of x at time t_k , indicating the discretization of the continuous variable x_t throughout the rest of the paper.

B. Problem Formulation

To describe the problem of getting the accurate IMU data from raw measurements, we can formulate it as follows:

Problem 1: Consider true ω_t and a_t as deterministic parameters. Given the raw discrete measurement $\tilde{\omega}_{0:k}$ and $\tilde{a}_{0:k}$ from the initial time t_0 to current time t_k , find a suitable estimator for ω_k and a_k :

$$(\hat{\omega}_k, \hat{a}_k) = \mathcal{F}(\tilde{\omega}_{0:k}, \tilde{a}_{0:k}), \quad (4)$$

where $\mathcal{F}$ represents the estimator.

We aim to develop a casual estimator that can operate online using only the *past* IMU measurements during inference time. This is an open problem that remains to be explored and is also challenging due to several factors. First, even small errors in w_t and a_t can accumulate into significant pose errors due to the integration process, which requires the estimator to meet strict accuracy standards. Second, the estimator uses only IMU information rather than fusing it with other sensors. The information resource is limited, requiring us to carefully design a suitable IMU measurement model for the estimator. Third, obtaining ground truth for ω_t and a_t is often impractical, making it challenging to establish a suitable evaluation metric for the estimator. In addition, this makes data-driven approaches particularly difficult, as they depend on suitable loss and reliable ground truth for training.

IV. LEARNING BIAS DYNAMICS

In this section, we propose an estimator for problem 1 by using a neural ordinary differential equation framework [10] for bias dynamics modeling.

A. IMU Measurement Model

Adopting (1), we model the IMU measurements as corrupted by some biases and additive Gaussian noise. Instead of modeling the bias dynamics as Brownian motions, we consider the bias b_t^g and b_t^a to be some deterministic variables with nonlinear dynamics as:

$$\dot{b}_t^g = f_g(b_t^g, \tilde{u}_t), \quad \dot{b}_t^a = f_a(b_t^a, \tilde{u}_t), \quad (5)$$

where $\tilde{u}_t := (\tilde{\omega}_t, \tilde{a}_t)$ denotes the raw IMU measurements. Since u_t serves as the control input to the bias dynamics, (5) can be viewed as input-dependent vector fields. As a result, with known initial conditions, one can simply integrate the deterministic dynamics to obtain the bias at any time.

With the known bias from (5), combining model (1), we can obtain an estimate of w_t, a_t as

$$\hat{\omega}_t = \tilde{\omega}_t - b_t^g, \quad \hat{a}_t = \tilde{a}_t - b_t^a. \quad (6)$$

This estimation can be regarded as that the noise is negligible. Such an estimator is reasonable because the estimation in (6) corresponds to a Least-Squares Estimation (LSE) applied to the model $\tilde{\omega}_t - b_t^g = w_t + n_t^g$ using only a single measurement. A more comprehensive treatment of the noise term can be achieved by integrating additional sensor data within a multi-sensor fusion framework such as [40, 28]. In this work, however, we focus on the estimation method described in (6), leaving noise handling as the future work.

According to (5) and (6), estimating the true w_t, a_t at time t only requires IMU measurements and the initial condition. The initial bias condition can always be determined from stationary IMU data. Therefore, this provides a solution to problem 1.

Remark 1: Compared to the *stochastic differential equation* (SDE) model for bias in conventional assumption, our method can be regarded as only considering the *drift* of the SDE, which reduces to an ordinary differential equation.

B. Neural Ordinary Differential Equations

We model the vector fields in (5) with a *neural ordinary differential equation* (NODE) architecture introduced in [10]. The NODE seeks to model the dynamics of a state as $\dot{z}_t = f(z_t; \theta)$ using a neural network, where θ is the parameters of the network. Instead of directly fitting $\dot{z}_t$, the loss $L(\cdot)$ of the network is defined based on z_t through an ODE solver:

$$L(z_T) = L(z_0 + \int_{t_0}^{t_T} f(z_t; \theta) dt), \quad (7)$$

where the integral can be solved by an ODE solver with methods such as the Euler method, RK45, etc. Therefore, the NODE framework takes the initial condition z_0 and time t_T , and outputs the state z_T after integration. To train this network, the gradient can be calculated by numerical backpropagation or memory-efficient *adjoint* method [10, 24]. Such a NODE framework allows us to use the sequence of observations of the state $\tilde{z}_0, \tilde{z}_1, \dots, \tilde{z}_T$ to fit the dynamics $\dot{z}_t$.

The ODE in (5) depends on additional control inputs u_t , which can not be handled by the canonical NODE framework. Moreover, the control inputs are measured in a discrete manner

rather than a continuous formulation. To address this issue, we utilize the idea of *neural controlled differential equations* [23]. The discrete control input $u_i \in \mathbb{R}^n$ can be interpolated using a continuous spline $\mathcal{S} : [t_0, t_T] \rightarrow \mathbb{R}^n$ such that $\mathcal{S}(t_i) = u_i$, where $t_0 \leq t_i \leq t_N$. Therefore, the control input can be modeled as a function of t given discrete $\tilde{u}_i$. To maintain consistency with the formulation in (7), where the dynamics do not explicitly depend on t , the time t can be augmented to the original state as $z_{\text{aug}} = [z_t, t]$.

Define $b_t := (b_t^g, b_t^a)$ and apply a continuous spline described above to handle the discrete control input $\tilde{u}_i$, we can reformulate the underlying bias dynamics (5) in the NODE framework as

$$\begin{bmatrix} \dot{t} \\ \dot{b}_t \end{bmatrix} = \begin{bmatrix} 1 \\ f(b_t, \mathcal{S}_{\tilde{u}}(t); \theta) \end{bmatrix}. \quad (8)$$

Using this ODE, the bias at time t_T is given by

$$b_T = b_0 + \int_{t_0}^{t_T} f(b_t, \mathcal{S}_{\tilde{u}}(t); \theta) dt. \quad (9)$$

t can be ignored in the loss function since it does not depend on network parameters. After defining suitable networks and the loss function of b_t , we can use the NODE framework to learn the bias dynamics.

V. NETWORK STRUCTURE AND IMPLEMENTATION

We have outlined the general concept of our proposed method; in this section, we present the detailed learning framework. The overall process is illustrated in Fig. 2.

A. Neural Network Structure and Input Spline

To model the bias vector field, we utilize a simple multilayer perceptron with residual connections [19] for both gyroscope and accelerometer, which can be represented as

$$\dot{b}_t^g = f_g(b_t^g, \tilde{\omega}_t, \dot{\tilde{\omega}}_t; \theta), \quad \dot{b}_t^a = f_a(b_t^a, \tilde{a}_t, \dot{\tilde{a}}_t; \theta), \quad (10)$$

where θ represents the parameters of neural networks and is omitted in the rest of this paper for simplicity. The explicit modeling of the bias vector field allows the network to remain lightweight, making it significantly simpler than most existing methods. This reduced model complexity lowers computational resource requirements. In our experiments, we find the residual connections can accelerate the model convergence.

In (10), we include additional derivatives $\dot{\tilde{\omega}}_t$ and $\dot{\tilde{a}}_t$ as the input, which is uncommon in the literature. This additional information allows us to model the components of bias that depend on ω_t and a_t . For example, if the measurement is linear in the true value as $\tilde{\omega}_t = A\omega_t$, where $A \in \mathbb{R}^{3 \times 3}$ is invertible, the bias dynamics should be

$$\dot{b}_t^g = \frac{d}{dt}(\tilde{\omega}_t - \omega_t) = (I - A^{-1})\dot{\tilde{\omega}}_t, \quad (11)$$

which is a function of $\dot{\tilde{\omega}}_t$. This demonstrates the importance of including $\dot{\tilde{\omega}}_t$ and $\dot{\tilde{a}}_t$, as the linear model is often used to model axis misalignment in an IMU [35, 5].

To handle the input into a continuous function of time t as (9), we utilize a *cubic Hermite spline* with the rule of backward

differences [31] to interpolate the discrete input. This is a cubic spline $\mathcal{S}_{\tilde{u}}$ satisfying the following:

$$\mathcal{S}_t(t_i) = \tilde{u}_i, \quad \dot{\mathcal{S}}_t(t_i) = (\tilde{u}_i - \tilde{u}_{i-1})/(t_i - t_{i-1}), \quad (12)$$

where $(t_i, \tilde{u}_i)$ is the discrete data. Therefore, the continuous spline for IMU measurements is given by $(\tilde{\omega}_t, \tilde{a}_t) = \mathcal{S}_{\tilde{u}}(t)$. Such a spline has continuous derivatives, which can benefit the convergence of the NODE [31]. More importantly, this spline is causal, which means the spline up to t_i will not depend on $\tilde{u}_{i+1}$. This enables the spline to be generated in real-time as new measurements are received, allowing the network to perform online inference.

B. Loss Construction

To train the NODE (9), a straightforward loss could be defined as a scalar function of estimated bias and the true bias $L(\hat{b}_t, b_t)$. However, obtaining the true bias is challenging. One can use multi-sensor fusion to estimate bias [6], but this requires extra resources and depends on the accuracy of the estimation algorithm. We proposed directly using the pose ground truth to construct the loss for learning bias dynamics.

Combining the bias dynamics (5) and the IMU kinematics (3), we can get the ODE as

$$\begin{aligned} \dot{R}_t &= R_t(\tilde{\omega}_t - b_t^g)^\times \\ \dot{b}_t^g &= f_g(b_t^g, \tilde{\omega}_t, \dot{\tilde{\omega}}_t) \\ \dot{v}_t &= R_t(\tilde{a}_t - b_t^a) + g \\ \dot{p}_t &= v_t \\ \dot{b}_t^a &= f_a(b_t^a, \tilde{a}_t, \dot{\tilde{a}}_t). \end{aligned} \quad (13)$$

By solving the entire ODE, we can construct the loss function using the pose ground truth. For orientation, we apply a mean-squared error (MSE) loss defined as

$$L_R = \frac{1}{N} \sum_{k=1}^N \|\text{Log}(\hat{R}_k R_k^T)\|_2^2, \quad (14)$$

where $\text{Log}(\cdot) : \text{SO}(3) \rightarrow \mathbb{R}^3$ is the vectorized *logarithm* of $\text{SO}(3)$ [14], $\hat{R}_k$ is the solution by solving (13), R_k is the ground truth. For velocity and position, we also use an MSE loss:

$$L_{v,p} = \frac{1}{N} \sum_{k=1}^N (\|\hat{p}_k - p_k\|_2^2 + \|\hat{v}_k - v_k\|_2^2). \quad (15)$$

where $\hat{p}_k, \hat{v}_k$ are the solution from (13) and p_k, v_k are the ground truth. Therefore, the total loss for training NODE (13) is given by

$$L = L_R + L_{v,p}. \quad (16)$$

By coupling the bias dynamics and the IMU kinematics in (13), the neural networks' parameters of bias dynamics can be optimized using the loss (16) that only depend on pose ground truth.




Fig. 2: The framework for learning bias dynamics. The bias dynamics are modeled by NODE, trained in a hierarchical manner. We first train the gyroscope component, followed by the accelerometer component. At each stage, the IMU raw data is represented as a continuous spline and serves as control input to the bias dynamics. Given initial conditions, the pose and bias along the trajectory are obtained through integration, with only the pose contributing to the loss function. The ODE on the manifold $\text{SO}(3)$ is reformulated as a Lie algebra ODE, enabling an efficient solution via canonical NODE.

C. Solving ODE on Lie Group via Lie Algebra

Embedding the ODE (13) within the standard NODE framework presents a challenge since it involves the manifold $\text{SO}(3)$, whereas the standard NODE framework is designed for Euclidean spaces. To address this, we utilize a Lie group method [12], reformulating the ODE on the *Lie algebra*, a linear space naturally compatible with the NODE framework. In this section, we focus on addressing the first equation in (13), the non-Euclidean part, which can be easily integrated into the original ODE (see Sec. V-D). Consider the following differential equation:

$$\dot{R}_t = R_t \omega_t^\times, \quad (17)$$

where ω_t is a function that only depends on t . The following theorem holds:

Theorem 1: The solution of the differential equation (17) with initial condition $R(0) = R_0$ is given by $R_t = R_0 \text{Exp}(\xi_t)$, where $\xi_t \in \mathbb{R}^3$ is the solution of

$$\dot{\xi}_t = J_r^{-1}(\xi_t)\omega_t, \quad \xi_t(0) = 0, \quad (18)$$

as long as $\|\xi_t\| < 2\pi$. $J_r^{-1} \in \mathbb{R}^{3 \times 3}$ is the inverse of the right Jacobian of $\text{SO}(3)$ [1, 11], which is well-defined under the condition $\|\xi_t\| < 2\pi$. $\square$

The proof is given in the Appendix. The proof generally follows [12, Theorem 7.1], however, which gives a left derivative version for the general Lie groups. Here we provide the right derivative version with matrix formulation, which is more compatible with IMU kinematics.

The theorem indicates we can obtain the local solution of ODE (17) by solving an ODE in Euclidean space. To solve the ODE on $\text{SO}(3)$ over time, we need to change the *chart* by resetting R_n as the initial condition and ξ_t to zeros before

the Jacobian approaches the singularity, i.e., $\|\xi_t\| \rightarrow 2\pi$. The general idea can be summarized as follows: parameterize the point R_0 on $\text{SO}(3)$ into $\mathbb{R}^3$ using the local chart around R_0 , solving the corresponding ODE in $\mathbb{R}^3$ to determine the next desired point, and then map the point in $\mathbb{R}^3$ back to $\text{SO}(3)$ to get R_1 . Repeat this process then we can get the solution R_t , see Algorithm 1.

Algorithm 1 Solving ODE for rotation kinematics

- 1: **Input:** Initial condition R_0 , input ω_t , time $[t_0, t_T]$
- 2: Initialize $R_0 \leftarrow R_0$, $\xi_t = 0$, $t \leftarrow t_0$
- 3: **while** $t \leq t_T$ **do**
- 4: $\xi_t = \text{ODESolver}\{\dot{\xi}_t = J_r^{-1}(\xi_t)\omega_t\}$
- 5: $R_t = R_0 \text{Exp}(\xi_t)$
- 6: **if** $\|\xi_t\| > \pi$ **then** $R_0 \leftarrow R_t$, $\xi_t \leftarrow 0$
- 7: Increase t (adaptive or fixed step, $t = t_0, t_1, \dots, t_T$)
- 8: **end while**
- 9: **Output:** Solution R_t , where $t = t_0, t_1, \dots, t_T$.

Compared to the canonical Euclidean ODE, our approach introduces only an additional chart transition as ξ_t approaches the singularity. We present a PyTorch-based ODE solver that extends the standard NODE framework to accommodate $\text{SO}(3)$ dynamics. The gradient of this NODE is computed using numerical backpropagation. The corresponding adjoint method for calculating gradient is beyond the scope of this work, which can be the future work.

Remark 2: The singularity issue in Theorem 1 typically requires changing the chart at each integration step to ensure well-definedness. However, in our application, the angular velocity will not jump too much. Therefore, we relax this condition and only change the chart when $\|\xi\|$ closes to π to

improve the computation efficiency. If singularities still occur, the threshold can be reduced further until chart updates are triggered at each step.

D. Hierarchical Training Framework

The NODE framework for learning bias dynamics can be regarded as a continuous recurrent neural network (RNN) [36], which is hard to train. To address this, we propose a hierarchical strategy to make neural networks converge faster.

Under our hypothesis, the biases of the gyroscope b_t^g and accelerometer b_t^a , are decoupled. This means their dynamics are independent of each other. We can adopt a two-stage training process by leveraging the IMU kinematics in (3), which indicates that velocity and position are dependent on rotation but not vice versa. First, we train $\dot{b}_t^g$ for the rotational component, and once complete, we fix $\dot{b}_t^g$ to train $\dot{b}_t^a$.

For the rotational component, we use the method described in the previous section and solve the following ODE:

$$\begin{aligned}\dot{\xi}_t &= J_r^{-1}(\xi_t)(\tilde{\omega}_t - b_t^g) \\ \dot{b}_t^g &= f_g(b_t^g, \tilde{\omega}_t, \dot{\tilde{\omega}}_t),\end{aligned}\quad (19)$$

where R_t can be recovered by $R_t = R_0 \text{Exp}(\xi_t)$. The loss function L_R is defined in (14). Once this stage is complete, the network parameters of f_g are fixed, and we proceed by combining (19) with the following equations to handle the accelerometer component:

$$\begin{aligned}\dot{v}_t &= R_0 \text{Exp}(\xi_t)(\tilde{a}_t - b_t^a) + g \\ \dot{p}_t &= v_t \\ \dot{b}_t^a &= f_a(b_t^a, \tilde{a}_t, \dot{\tilde{a}}_t),\end{aligned}\quad (20)$$

where $R_t = R_0 \text{Exp}(\xi_t)$ has been substituted. The corresponding loss function $L_{v,p}$ is defined as (15). The entire training process is illustrated in Fig. 2.

In each training stage, the most straightforward approach would be to train the neural ODE by integrating over the entire sequence, calculating the loss, and optimizing the parameters. However, this method is computationally expensive and memory-intensive, making it impractical for implementation. To address these limitations, we adopt the approach proposed in [10], where training is performed using short segments of the sequence. Instead of integrating over the entire sequence, we compute the vector field by integrating over multiple short time intervals, significantly reducing computational cost and memory usage. This approach can be expressed as:

$$(\hat{R}, \hat{v}, \hat{p})_{s:s+N} = \text{ODEINT}_{\text{SO}(3)}(R_s, b_s^g, v_s, p_s, b_s^a), \quad (21)$$

where s represents the initial time step of short time intervals and N represents the length of the intervals. In the implementation, we set $N = 16$, making a trade-off between being sufficiently long to effectively learn the vector field and short enough to avoid excessive training time. This training strategy requires the initial condition for each interval. Usually, we can get the pose ground truth but hard to the biased ground truth.

We utilize an approximation of the initial bias condition as

$$\begin{aligned}b_k^g &= \tilde{\omega}_k - \text{Log}(R_k^T R_{k+1}) / (t_{k+1} - t_k) \\ b_k^a &= \tilde{a}_k - R_k^T ((v_{k+1} - v_k) / (t_{k+1} - t_k) - g),\end{aligned}\quad (22)$$

where R_k, v_k is assumed known as the ground truth. The initial bias for each time interval, determined in this manner, is affected by noise from IMU measurements and pose estimates. However, as long as the noise remains within a reasonable range, the proposed method can effectively manage this uncertainty. In this context, the noise in the bias initialization can be regarded as data augmentation.

VI. EXPERIMENTAL RESULTS

In this section, we demonstrate the proposed method's accuracy in angular velocity and acceleration. We compare pure IMU integration and its application in visual-inertial odometry against other methods on two public datasets. We also evaluate its generalization capability on a real-world experiment.

A. Public Datasets and Training Platform

1) *EUROC* [7]: This is a VIO dataset for a micro aerial vehicle (MAV). It comprises 11 trajectories with lengths ranging from approximately 36.5 to 130.9 meters and durations between 99 to 182 seconds. The IMU used is the ADIS16448, operating at 200 Hz, providing the angular velocity and the linear acceleration. The dataset includes stereo camera data at 20 Hz. Ground truth poses are provided by a motion capture system and a laser tracker.

2) *TUM-VI* [37]: This is a handheld VIO dataset. It contains 28 sequences, covering a total distance of approximately 20 km, with both indoor and outdoor scenes. The IMU used is the Bosch BMI160, which also provides the angular velocity and the linear acceleration at a frequency of 200 Hz. The dataset contains the stereo camera data at 20 Hz. Ground truth poses are obtained from a motion capture system but are available only for the indoor environment, resulting in some sequences with incomplete pose ground truth. For our experiments, we select the 6 room sequences, which contain the longest ground truth data, each lasting 2–3 minutes.

3) *Data Arrangement and Platform*: We divide the data into training, validation, and test sets for both the EUROC and TUM-VI datasets. For training and validation, we use the same sequences, with the first 80% designated for training and the last 20% for validation. The test set consists of entirely unseen sequences. In the EUROC dataset, following [5], we use 6 sequences for training and the remaining 5 for testing. For the TUM-VI dataset, 3 sequences are used for training and 3 for testing.

For both datasets, the data is assumed to be synthesized, with raw IMU data serving as the network's input. Different from EUROC, TUM-VI also provided model-based calibrated IMU data. The training is processed on a laptop with RTX4060, 8G memory. We choose the Adam-optimizer [26] with a StepLR scheduler for the learning rate. The epoch is chosen 1800 for all trainings. For the EUROC dataset, the learning rate is set as 0.005 for both the gyroscope and




Fig. 3: Absolute error and relative error.

accelerometer. We set nearly the same parameters for TUM-VI dataset. The integration method chosen is the Euler method for fast training.

B. Evaluation Metric

We utilize the *absolute error* and *relative error* [43] as the metric, which are the common criteria in VIO systems. More specifically, the following metric in toolbox *evo* [16] is used.

1) *Absolute Error*: After the alignment of the estimated trajectory and ground trajectory, the *Absolute Orientation Error* (AOE) is defined as

$$\text{AOE} = \sqrt{\frac{1}{N} \sum_{k=1}^N \|\text{Log}(\hat{R}_k^T R_k)\|_2^2}, \quad (23)$$

where $\hat{R}_k$ is the estimated orientation and R_k is the ground truth. The *Absolute Position Error* (APE) is defined as

$$\text{APE} = \sqrt{\frac{1}{N} \sum_{k=1}^N \|p_k - \hat{p}_k\|_2^2}, \quad (24)$$

where $\hat{p}_k$ is the estimated position and p_k is the ground truth.

2) *Relative Error*: The idea of relative error is to compare multiple sub-trajectories of two trajectories rather than the entire ones. To define the relative error, all sub-trajectories whose length is d m will be collected using the ground truth, denote $X_{s_0:s_1}$, where $X \in SE(3)$ made of R and p . The increment for each sub-trajectory will be calculated using $\Delta X_s = X_{s_0}^{-1} X_{s_1}$. Then the difference between the estimation and ground truth for each sub-trajectory is defined as $X_{e,s} = \Delta X_s^{-1} \Delta \hat{X}_s$. Extracting the orientation and position components, $X_{e,s} = (R_{e,s}, p_{e,s})$, we obtain the *Relative Orientation Error* (ROE) and *Relative Position Error* (RPE) as follows:

$$\text{ROE} = \frac{1}{D} \sum_{s=1}^D \|\text{Log}(R_{e,s})\|_2 \quad \text{RPE} = \frac{1}{D} \sum_{s=1}^D \|p_{e,s}\|_2. \quad (25)$$

Fig. 3 shows the difference between absolute errors and relative errors. Since absolute error is sensitive to the time that errors occur [43, 15], relative error provides more informative insights into the comparison of two trajectories.

C. Compared Methods

In each experiment, the IMU data from the following sources are used.

- 1) **Raw IMU**: Raw uncalibrated data.
- 2) **Linear**: A simple linear model with constant biases, $\hat{\omega} = A_g \tilde{\omega} + b_g$, $\hat{a} = A_a \tilde{a} + b_a$. The parameters A and b are

modeled using a single linear layer network, which is trained by considering $\hat{b} \equiv 0$ in the proposed framework.

- 3) **M.B.** [5]: A CNN-based neural network, which implicitly inference the bias using a window of IMU raw data. However, it only provides the Gyro. calibration. The parameters for their method are set to the defaults provided in their open-source project.
- 4) **Proposed**: The calibrated IMU data by the proposed method.

D. Debiasing Results on EUROC and TUM-VI

To demonstrate that the proposed method provides more accurate angular velocity and acceleration estimates, we present two types of results. First, we integrate the IMU measurements alone and evaluate the orientation and position errors. Second, we substitute the debiased IMU data into a VIO algorithm and examine the resulting orientation and position errors. The brief IMU grades are shown in Tab. III for a more intuitive understanding of the improvements of the proposed method.

The absolute and relative pose errors from pure IMU integration are shown in Table I and II, respectively. Since the method in [5] only provides angular velocity estimates, raw acceleration data is used. This scenario reflects cases where camera data is unavailable, only IMU provides odometry information for a VIO. It can be found that raw IMU data is unreliable, implying it is necessary to have a calibration for IMU. The proposed method outperforms other approaches in both absolute and relative error metrics. Notably, orientation errors are significantly smaller than position errors. Our findings show that pure integration with debiased IMU data yields accurate orientation along the test trajectories, see Fig. 4. The comparison of different method for orientation is shown in Fig. 5a and 5b, which only shows the error of estimations for better visualization. Calibrated IMU generally gives good results and the proposed method gives better results than other compared methods. However, position estimates remain unreliable due to the inherent challenges of double integration and sensitivity to orientation estimation. Instead, we present the velocity estimation in Fig. 6a and 6b. The result from the M.B. method [5] is omitted, as it does not account for accelerometer debiasing. The velocity estimation is reliable only over short intervals, while long-term integration leads to divergence. The significant velocity error is partly due to its dependence not only on accelerometer measurements but also on the accuracy of orientation estimation. The velocity along the z-axis (yaw axis) exhibits the smallest error. A possible reason is that z-axis velocity is influenced by pitch and roll, which tend to be less aggressive than yaw during most motions. Although pure integration over long trajectories is rare in practice, these results highlight the strengths of the proposed method.

We also employ our network to provide clean IMU measurements as input to OpenVINS [15], a VIO algorithm, instead of using raw IMU data. This approach has greater practical relevance. The absolute errors and relative errors are presented in Table I and Tab. II. Fig. 7 and 8 further plot the

TABLE I: Absolute Orientation Error (AOE) and Absolute Position Error (APE) of pure IMU integration and a VIO for EUROC and TUM-VI. Less is better. The pure integration results in position from raw IMU data have been removed due to significant errors.

Dataset	Seq.	AOE / APE of IMU Integration (deg / m)				AOE / APE of VIO (deg / m)			
		Raw IMU	Linear	M.B. [5]	Proposed	Raw	Linear	M.B. [5]	Proposed
EUROC	MH_02	123.05/-	5.01/2204.16	4.54/1647.48	3.12/249.79	1.73/0.30	0.86/0.19	1.41/0.21	1.01/0.22
	MH_04	130.33/-	4.08/636.25	1.47/433.94	1.06/145.40	2.07/0.43	0.95/0.82	1.52/1.16	0.65/0.31
	V1_01	113.47/-	4.89/2031.47	1.93/1862.44	1.27/1712.75	1.83/0.12	1.83/0.11	1.20/0.10	1.44/0.11
	V1_03	120.13/-	2.52/693.59	1.58/407.72	2.26/614.36	2.82/0.11	1.16/0.11	1.67/0.16	0.85/0.12
	V2_02	116.92/-	4.62/1210.81	4.85/494.58	4.27/1111.75	1.87/0.15	2.26/0.16	2.16/0.16	2.60/0.16
	Average	120.78/-	4.22/1355.26	2.87/969.23	2.40/766.81	2.07/0.22	1.41/0.28	1.59/0.36	1.31/0.18
TUM-VI	Room2	32.82/-	15.02/3355.75	15.60/12268.32	1.69/1216.82	failed	10.04/0.39	9.80/0.32	5.08/0.35
	Room4	70.19/-	4.07/1615.27	17.90/7782.74	1.71/672.13	failed	2.56/0.22	9.58/0.23	2.52/0.11
	Room6	65.15/-	8.47/2333.58	19.10/11669.17	1.97/1139.77	failed	1.55/0.10	8.92/0.15	3.04/0.05
	Average	56.06/-	9.19/2434.87	17.53/10573.41	1.79/1009.58	failed	4.72/0.24	9.43/0.23	3.55/0.17

TABLE II: Relative Orientation Error (ROE) and Relative Position Error (RPE) of pure IMU integration and a VIO for EUROC and TUM-VI. Less is better. Len. represents the given distance for finding evaluation pairs.

Dataset	Len. (m)	ROE / RPE of IMU Integration (deg / m)				ROE / RPE of VIO (deg / m)			
		Raw IMU	Linear	M.B. [5]	Proposed	Raw	Linear	M.B. [5]	Proposed
EUROC	5	38.51/4823.72	0.99/241.61	0.89/167.74	0.71/133.79	0.71/0.09	0.67/0.11	0.73/0.11	0.65/0.08
	10	62.11/9290.19	1.44/456.60	1.25/321.93	0.98/256.06	0.93/0.12	0.93/0.16	0.96/0.15	0.88/0.11
	15	79.74/14029.04	1.81/690.52	1.47/487.35	1.21/390.69	1.01/0.14	1.04/0.19	1.04/0.19	0.97/0.14
	20	93.17/18735.49	2.18/919.75	1.71/654.11	1.44/517.72	1.08/0.16	1.16/0.22	1.14/0.23	1.06/0.15
TUM-VI	5	12.12/2546.67	1.58/356.56	4.47/1604.39	0.85/156.37	failed	1.07/0.07	1.54/0.06	0.86/0.03
	10	17.98/4829.26	2.45/685.14	6.26/3083.10	1.00/299.48	failed	1.29/0.09	1.92/0.07	0.95/0.03
	15	23.44/7125.78	3.31/1018.76	8.00/4583.68	1.16/444.15	failed	1.66/0.12	2.32/0.09	1.08/0.04
	20	27.21/9281.92	4.08/1338.19	8.68/6019.29	1.23/581.83	failed	2.01/0.13	2.64/0.09	1.23/0.04




Fig. 4: The Euler angles obtained by only integrating the IMU data. The results are very close to the ground truth, which implies that the debiased gyroscope yields good performance.

orientation errors and positions of different methods in VIO settings. The performance improvements with our method are not as significant compared to pure IMU integration since the extra vision sensor can also provide pose estimation itself. It is worth noting that using raw IMU data from the TUM-VI dataset causes the VIO algorithm to fail under our experimental settings, which gives significant errors. Although the raw IMU data performance is bad in pure IMU integration, it performs well in the Euroc dataset, primarily due to the bias estimation capabilities of the VIO algorithm. However, this bias estimation is not always reliable, as demonstrated by its failure with raw IMU data in the TUM-VI dataset. From the tables, the proposed method yields better results in both absolute and relative errors within the VIO framework,

where the IMU is coupled with other vision sensors. This demonstrates that the accurate IMU measurements provided by our debiasing method can also contribute to improved performance in inertial-based odometry.


TABLE III: IMU Grades

Dataset	IMU Model	Grade
EUROC	ADIS16448	Industrial/tactical
TUM-VI	Bosch BMI160	Consumer/low-cost
FETCH	Not declared	Consumer/low-cost


E. Debiasing Results on Real-world Experiment

To further demonstrate the generalization capability of the proposed method, we also apply it to a real-world dataset, the FETCH dataset. This dataset is collected from an indoor mobile robot, including six sequences with simple motion patterns, each lasting 30-40 seconds, and one longer trajectory with random motion lasting approximately 90 seconds. We use the longest trajectory for testing and the remaining sequences for training. The experiment platform is shown in Fig. 9.


To demonstrate the practicality of our method, we used the same network parameters on FETCH as for the EUROC dataset. Results in Table IV show that our method also works well for real-world dataset, consistent with its performance on public datasets. The orientation and velocity estimation is shown in Fig. 5c and 6c, respectively. It is surprisingly found that the linear model performs comparably to the proposed method and even achieves better positional accuracy. This can be attributed to the simplicity of motion patterns in



(a) EUROCC: MH_04. Pure integration orientation errors.




(b) TUM: Room4. Pure integration orientation errors.




(c) FETCH: 05_random. Pure integration orientation errors.

Fig. 5: The coordinates of the errors, $\text{Log}(R_{gt}^T \hat{R})$. The estimates are obtained by only integrating the IMU data. The unit is converted to a degree. Closer to zero yields better results.


the FETCH dataset, as we argue that learning-based methods would inevitably capture and leverage such patterns. To support this interpretation, we analyze the performance of the linear model on the more diverse EUROCC and TUM-VI datasets, which contain a wider range of motion patterns. In these cases, the linear model performs worse than the learning-based method in both cases, with a particularly poor performance on the TUM-VI dataset, which involves human motion patterns, further validating our hypothesis. Based on



(a) EUROCC: MH_04. Pure integration velocity estimation.



(b) TUM: Room4. Pure integration velocity estimation.




(c) FETCH: 05_random. Pure integration velocity estimation.

Fig. 6: The Velocity estimation obtained by only integrating the IMU data. Closer to the ground truth yields better results.


these findings, we recommend first attempting a linear model for IMU calibration to assess whether it meets the desired accuracy requirements. If not, a learning-based approach can then be implemented to achieve better performance.

F. Discussion on Implementation

Integration Method: There is no significant performance difference between the Euler method and more advanced integration methods. This is likely because the high frequency of IMU data and the coupling of integration errors with



(a) EUROC: MH_04. VIO orientation errors.



(b) TUM: Room4. VIO orientation errors.

Fig. 7: The coordinates of the errors, $\text{Log}(R_{gt}^T \hat{R})$. The estimates are obtained by VIO with debiased IMU data. The unit is converted to a degree. Closer to zero yields better results.

TABLE IV: Absolute Orientation Error (AOE) and Absolute Position Error (APE) of pure IMU integration for FETCH. Less is better.


Seq.	AOE / APE of IMU Integration (deg / m)			
	Raw IMU	Linear	M.B. [5]	Proposed
05_random	73.28/1139.07	2.42/119.70	2.78/146.45	1.96/132.57

unknown noise in IMU measurements diminish the impact of the integration method used.


Integration Interval Length: While longer integration windows generally produce smoother bias estimates, simply increasing the window length does not improve accuracy, see Tab. V.

Data Pre-processing: Pre-processing plays a critical role. We observed that the ground truth in the TUM-VI dataset is noisier than in the EUROC dataset, resulting in noisier initial bias conditions. Applying an appropriate smoother to refine the initial bias conditions leads to better training outcomes.

Method Limitations: Although the proposed method is a device-specific calibration approach with some robustness across varying motion patterns, it still learns motion pattern characteristics during training. As a result, when the available motion pattern information is limited, the method may perform similarly to a simple linear model, losing the advantages in accuracy. Another limitation is that the training time and memory



(a) EUROC: MH_04. VIO position errors.



(b) TUM: Room4. VIO position errors.

Fig. 8: The position error, $\hat{p} - p_{gt}$. The estimates are obtained by VIO with debiased IMU data. Closer to the ground truth yields better results.





Fig. 9: Fetch experiment platform. Fetch is a mobile manipulation, the base equips a low-cost IMU. A motion capture collects the pose ground truth.

consumption will increase as the integration interval lengthens, which is a common issue in RNN-like networks. The memory consumption issue can be mitigated by developing the adjoint method in the future.

Additionally, we provide an ablation study on integration length, with results shown in Tab. V on the EUROC dataset. The accuracy is poor with short integration lengths mainly due to noises. The accuracy can improve as the integration length increases, however, the accuracy saturates around $N = 10$. Since both training time and GPU memory consumption grow

TABLE V: Ablation Study on Integration Length. We present results on the EUROC dataset showing how performance varies with different integration lengths. AOE and APE are computed from pure integration. Since position drifts quickly, we report APE on velocity instead.

Integration Length N	4 (0.2 s)	8 (0.4 s)	12 (0.6 s)	16 (0.8 s)	32 (0.16 s)	48 (0.24 s)	64 (0.32 s)	96 (0.48 s)
AOE (deg)	91.59	12.53	2.50	2.71	2.62	2.45	3.48	2.70
APE (velocity) (m/s)	257.89	32.74	19.94	17.75	20.82	18.02	16.36	18.58
Training Time (s)	135.62	300.83	460.36	573.63	1223.78	1823.72	2483.10	3610.63
GPU Memory (MB)	590.73	418.34	476.53	533.20	762.34	988.30	1217.71	1671.22



roughly linearly with the integration length, we set $N = 16$ in our experiments to balance accuracy and resource usage. Our findings are consistent with [10], where GPU memory consumption scales as $\mathcal{O}(N)$ when the adjoint method is not used.

VII. CONCLUSION

In this work, we propose a learning-based method to obtain accurate angular velocity and acceleration using only raw IMU measurements. The proposed method explicitly learns the bias dynamics of the IMU using a NODE framework without the need for the bias ground truth. With the known bias dynamics, we can compensate for the bias and obtain a clean IMU measurement. Experiments on two popular datasets and one real-world experiment demonstrate that the proposed method outperforms the existing approach and has practical meaning.

In the future, this work can be extended to incorporate covariance propagation through the learned dynamics, allowing integration with any existing IMU-based odometry systems. Although the current work does not fully address bias propagation with covariance, one promising application is to use the unbiased IMU data within the *invariant extended Kalman filter* framework, which effectively eliminates the nonlinearities in IMU kinematics [3, 2, 8]. Additionally, the backpropagation of the NODE on Lie algebra formulation can be extended into more efficient adjoint method.

ACKNOWLEDGMENT

Ben Liu would like to thank Hongbo Zhu for his invaluable assistance in this work. Wei Zhang was supported by the Guangdong Science and Technology Program under Grant No. 2024B1212010002 and 2019QN01X793 for this work. M. Ghaffari was supported by AFOSR MURI FA9550-23-1-0400.

APPENDIX

A. Proof of Theorem 1

We need to show $\dot{R}_t$ meets the differential equation and $R_t(0) = R_0$. First, we need the derivative of the exponential

map for $\text{SO}(3)$ [17, Theorem 5.4]:

$$\frac{d \exp(\xi^\times)}{dt} = \exp(\xi^\times) \left(\sum_{k=0}^{\infty} (-1)^k \frac{\text{ad}_{\xi^\times}^k}{(k+1)!} \cdot \left(\frac{d\xi^\times}{dt} \right) \right) \quad (26)$$

where $\text{ad}_\xi \in \mathbb{R}^{3 \times 3}$ is the *adjoint map*. This can be expressed in further simplified formulation [1, eq. 7.77a]:

$$\frac{d \text{Exp}(\xi)}{dt} = \text{Exp}(\xi) (J_r(\xi) \dot{\xi})^\times \quad (27)$$

which can be substituted into $R_t = R_0 \text{Exp}(\xi_t)$, then

$$\dot{R}_t = R_0 \text{Exp}(\xi_t) (J_r(\xi_t) J_r^{-1}(\xi_t) \omega_t)^\times = R_t \omega^\times \quad (28)$$

with the initial condition: $R_t(0) = R_0 \text{Exp}(\mathbf{0}) = R_0$. $\square$

REFERENCES

- [1] Timothy D Barfoot. *State estimation for robotics*. Cambridge University Press, 2017.
- [2] Axel Barrau and Silvere Bonnabel. Invariant kalman filtering. *Annual Review of Control, Robotics, and Autonomous Systems*, 1:237–257, 2018.
- [3] Axel Barrau and Silvère Bonnabel. The invariant extended kalman filter as a stable observer. *IEEE Transactions on Automatic Control*, 62(4):1797–1812, 2017. doi: 10.1109/TAC.2016.2594085.
- [4] Michael Bloesch, Marco Hutter, Mark A Hoepflinger, Stefan Leutenegger, Christian Gehring, C David Remy, and Roland Siegwart. State estimation for legged robots-consistent fusion of leg kinematics and imu. *Robotics*, 17:17–24, 2013.
- [5] Martin Brossard, Silvere Bonnabel, and Axel Barrau. Denoising imu gyroscopes with deep learning for open-loop attitude estimation. *IEEE Robotics and Automation Letters*, 5(3):4796–4803, 2020.
- [6] Russell Buchanan, Varun Agrawal, Marco Camurri, Frank Dellaert, and Maurice Fallon. Deep imu bias inference for robust visual-inertial odometry with factor graphs. *IEEE Robotics and Automation Letters*, 8(1):41–48, 2022.
- [7] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The euroc micro aerial vehicle datasets. *The International Journal of Robotics Research*, 35(10):1157–1163, 2016.
- [8] Paul Chauchat, Axel Barrau, and Silvere Bonnabel. Invariant smoothing on lie groups. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1703–1710, 2018. doi: 10.1109/IROS.2018.8594068.

- [9] Changhao Chen and Xianfei Pan. Deep learning for inertial positioning: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 2024.
- [10] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [11] Gregory S. Chirikjian. *Stochastic Models, Information Theory, and Lie Groups, Volume 2*. Birkhäuser Boston, MA, 2012.
- [12] Christian Lubich Ernst Hairer, Gerhard Wanner. *Geometric Numerical Integration*. Springer Berlin, Heidelberg, 2006.
- [13] Mahdi Abolfazli Esfahani, Han Wang, Keyu Wu, and Shenghai Yuan. Orinet: Robust 3-d orientation estimation with a single particular imu. *IEEE Robotics and Automation Letters*, 5(2): 399–406, 2019.
- [14] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Transactions on Robotics*, 33(1):1–21, 2017. doi: 10.1109/TRO.2016.2597321.
- [15] Patrick Geneva, Kevin Eickenhoff, Woosik Lee, Yulin Yang, and Guoquan Huang. Openvins: A research platform for visual-inertial estimation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4666–4672. IEEE, 2020.
- [16] Michael Grupp. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>, 2017.
- [17] Brian C. Hall. *Lie Groups, Lie Algebras, and Representations*. Springer Cham, 2015.
- [18] Ross Hartley, Maani Ghaffari, Ryan M Eustice, and Jessy W Grizzle. Contact-aided invariant extended kalman filtering for robot state estimation. *The International Journal of Robotics Research*, 39(4):402–430, 2020.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [20] Sachini Herath, Hang Yan, and Yasutaka Furukawa. Ronin: Robust neural inertial navigation in the wild: Benchmark, evaluations, and new methods. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3146–3152, 2020. doi: 10.1109/ICRA40945.2020.9196860.
- [21] Fengrong Huang, Zhen Wang, Luran Xing, and Chunyan Gao. A mems imu gyroscope calibration method based on deep learning. *IEEE Transactions on Instrumentation and Measurement*, 71:1–9, 2022.
- [22] NovAtel Inc. Apn064 bulletin. <https://hexagondownloads.blob.core.windows.net/public/Novatel/assets/Documents/Bulletins/APN064/APN064.pdf>, 2025. Accessed: 2025-03-23.
- [23] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *Advances in Neural Information Processing Systems*, 33:6696–6707, 2020.
- [24] Patrick Kidger, Ricky TQ Chen, and Terry J Lyons. ” hey, that’s not an ode”: Faster ode adjoints via seminorms. In *ICML*, pages 5443–5452, 2021.
- [25] Joon-Ha Kim, Seungwoo Hong, Gwanghyeon Ji, Seunghun Jeon, Jemin Hwangbo, Jun-Ho Oh, and Hae-Won Park. Legged robot state estimation with dynamic contact event information. *IEEE Robotics and Automation Letters*, 6(4):6733–6740, 2021. doi: 10.1109/LRA.2021.3093876.
- [26] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] Tzu-Yuan Lin, Ray Zhang, Justin Yu, and Maani Ghaffari. Legged robot state estimation using invariant kalman filtering and learned contact events. In *5th Annual Conference on Robot Learning*, 2021. URL <https://openreview.net/forum?id=yt3tDB67lc5>.
- [28] Tzu-Yuan Lin, Tingjun Li, Wenzhe Tong, and Maani Ghaffari. Proprioceptive invariant robot state estimation. *arXiv preprint arXiv:2311.04320*, 2023.
- [29] Huakun Liu, Xin Wei, Monica Perusquía-Hernández, Naoya Isoyama, Hideaki Uchiyama, and Kiyoshi Kiyokawa. Duet: Improving inertial-based odometry via deep imu online calibration. *IEEE Transactions on Instrumentation and Measurement*, 2023.
- [30] Wenxin Liu, David Caruso, Eddy Ilg, Jing Dong, Anastasios I Mourikis, Kostas Daniilidis, Vijay Kumar, and Jakob Engel. Tlio: Tight learned inertial odometry. *IEEE Robotics and Automation Letters*, 5(4):5653–5660, 2020.
- [31] James Morrill, Patrick Kidger, Lingyi Yang, and Terry Lyons. On the choice of interpolation scheme for neural cdes. *Transactions on Machine Learning Research*, 2022(9), 2022.
- [32] Anastasios I Mourikis and Stergios I Roumeliotis. A multi-state constraint kalman filter for vision-aided inertial navigation. In *Proceedings 2007 IEEE international conference on robotics and automation*, pages 3565–3572. IEEE, 2007.
- [33] Tong Qin, Peiliang Li, and Shaojie Shen. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE transactions on robotics*, 34(4):1004–1020, 2018.
- [34] Tong Qin, Shaozu Cao, Jie Pan, and Shaojie Shen. A general optimization-based framework for global pose estimation with multiple sensors. *arXiv preprint arXiv:1901.03642*, 2019.
- [35] Joern Rehder, Janosch Nikolic, Thomas Schneider, Timo Hinzmann, and Roland Siegwart. Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4304–4311. IEEE, 2016.
- [36] Yulia Rubanova, Ricky TQ Chen, and David K Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*, 32, 2019.
- [37] David Schubert, Thore Goll, Nikolaus Demmel, Vladyslav Usenko, Jörg Stückler, and Daniel Cremers. The tum vi benchmark for evaluating visual-inertial odometry. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1680–1687. IEEE, 2018.
- [38] Scott Sun, Dennis Melamed, and Kris Kitani. Idol: Inertial deep orientation-estimation and localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6128–6137, 2021.
- [39] Pieter van Goor and Robert Mahony. Eqvio: An equivariant filter for visual-inertial odometry. *IEEE Transactions on Robotics*, 39(5):3567–3585, 2023. doi: 10.1109/TRO.2023.3289587.
- [40] David Wisth, Marco Camurri, and Maurice Fallon. Vilens: Visual, inertial, lidar, and leg odometry for all-terrain legged robots. *IEEE Transactions on Robotics*, 39(1):309–326, 2023. doi: 10.1109/TRO.2022.3193788.
- [41] Hang Yan, Qi Shan, and Yasutaka Furukawa. Ridi: Robust imu double integration. In *Proceedings of the European conference on computer vision (ECCV)*, pages 621–636, 2018.
- [42] Ming Zhang, Mingming Zhang, Yiming Chen, and Mingyang Li. Imu data processing for inertial aided navigation: A recurrent neural network based approach. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3992–3998. IEEE, 2021.
- [43] Zichao Zhang and Davide Scaramuzza. A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7244–7251. IEEE, 2018.