

# TLIO: Tight Learned Inertial Odometry

Wenxin Liu<sup>1,2</sup>, David Caruso<sup>2</sup>, Eddy Ilg<sup>2</sup>, Jing Dong<sup>2</sup>,  
Anastasios I. Mourikis<sup>2</sup>, Kostas Daniilidis<sup>1</sup>, Vijay Kumar<sup>1</sup>, and Jakob Engel<sup>2</sup>

**Abstract**—In this work we propose a tightly-coupled Extended Kalman Filter framework for IMU-only state estimation. Strap-down IMU measurements provide relative state estimates based on IMU kinematic motion model. However the integration of measurements is sensitive to sensor bias and noise, causing significant drift within seconds. Recent research by Yan et al. (RoNIN) and Chen et al. (IONet) showed the capability of using trained neural networks to obtain accurate 2D displacement estimates from segments of IMU data and obtained good position estimates from concatenating them. This paper demonstrates a network that regresses 3D displacement estimates and its uncertainty, giving us the ability to tightly fuse the relative state measurement into a stochastic cloning EKF to solve for pose, velocity and sensor biases. We show that our network, trained with pedestrian data from a headset, can produce statistically consistent measurement and uncertainty to be used as the update step in the filter, and the tightly-coupled system outperforms velocity integration approaches in position estimates, and AHRS attitude filter in orientation estimates.

Video materials and code can be found on our project page:

[cathias.github.io/TLIO](https://cathias.github.io/TLIO)

## I. INTRODUCTION

Visual-Inertial Navigation Systems (VINS) in recent years have seen tremendous success, enabling a wide range of applications from mobile devices to autonomous systems. Using only light-weight cameras and Inertial Measurement Units (IMUs), VINS achieve high accuracy in tracking at low cost and provide one of the best solutions for localization and navigation on constrained platforms. With the unique combination of these advantages, VINS have been the de facto standard for demanding applications such as Virtual/Augmented Reality (VR/AR) on mobile phones or headsets.

Despite the impressive performance of state-of-the-art VINS algorithms, demands from consumer AR/VR products are posing new challenges on state estimation, pushing the research frontier. Visual-Inertial Odometry (VIO) largely relies on consistent image tracking, leading to failure cases under extreme lighting or camera blockage conditions, such as in a dark room or inside a pocket. High frequency image processing makes power consumption a bottleneck for sustainable long-term operations. In addition, widespread

Wenxin Liu, Kostas Daniilidis, and Vijay Kumar are with GRASP Lab, University of Pennsylvania, Philadelphia, PA 19104 USA. [wenxinl@seas.upenn.edu](mailto:wenxinl@seas.upenn.edu); [kostas@cis.upenn.edu](mailto:kostas@cis.upenn.edu); [kumar@cis.upenn.edu](mailto:kumar@cis.upenn.edu)

David Caruso, Eddy Ilg, Jing Dong, Anastasios I. Mourikis, and Jakob Engel are with Facebook Reality Labs, Redmond, WA 98052 USA. [dcaruso@fb.com](mailto:dcaruso@fb.com); [eddyilg@fb.com](mailto:eddyilg@fb.com); [jingdong@fb.com](mailto:jingdong@fb.com); [mourikis@fb.com](mailto:mourikis@fb.com); [jakob.engel@fb.com](mailto:jakob.engel@fb.com)




Fig. 1: An example 3D trajectory estimated by TLIO from a single MEMS IMU, showing the user repeatedly walking up and down a staircase. Our method estimates full orientation and translation in 3D: On the left, we show the estimated  $x, y, z$  position (blue) as well as associated uncertainties ( $\pm 3\sigma$  in red). On the right, we show the resulting 3D trajectory. For both plots, we show the output of state-of-the-art visual-inertial odometry (green) for comparison.

camera usage carries privacy implications. Targeting an alternative to the state-of-the-art VINS algorithms for pedestrian applications, this paper focuses on consumer grade IMU-only state estimation problem known as strap-down Inertial Navigation System (INS) or Dead Reckoning [1].

Inertial Pedestrian Dead Reckoning (PDR) has gained increasing interest with the price decline and widespread usage of MEMS sensors in the past two decades [2]. Strap-down pedestrian IMU navigation faces the challenge of the accumulation of sensor errors since the IMU kinematic model provides only relative state estimates. To compensate for the errors and reduce drift without the aid of other sensors or floor maps, the existing approaches rely on the prior knowledge of human walking motion, in particular, steps. One way to make use of steps is Zero-velocity UPdaTe (ZUPT) [3]. It detects when the foot touches the ground to generate pseudo-measurement updates in an Extended Kalman Filter (EKF) framework to calibrate IMU biases. However, it only works well when the sensor is attached to the foot where step detection is obvious. Another category is step counting [4] which does not require the sensor to be attached to the foot. Such systems consist of multiple submodules: the identification and classification of steps, data segmentation and step length prediction, all of which require heavy tuning of hand-designed heuristics or machine learning.

In parallel, recent research advances (IONet [5] and RoNIN [6]) have shown that integrating average velocity estimates from a trained neural network results in highly accurate 2D trajectory reconstruction using only IMU data from pedestrian hand-carried devices. These results showed

the ability of networks to learn a translation motion model from pedestrian data. In this work, we draw inspiration from these new findings to build an IMU-only dead reckoning system trained with data collected from a head-mounted device, similar to a VR headset. This paper has two major contributions:

- 1) We propose a network design to regress both the 3D displacement and the corresponding covariance, and show the network’s capability of providing both accurate and statistically consistent outputs trained on our pedestrian dataset.
- 2) We propose a complete state estimation system combining the neural network with an EKF in a tightly-coupled formulation that jointly estimates position, orientation, velocity, and IMU biases with only pedestrian IMU data.

Our approach presents significant advantages over other IMU-only state estimation approaches. Comparing to traditional PDR methods, it avoids the limitations and complexities of step-counting/stride or gait detection by learning a displacement motion model valid for any data segments, simplifying the process and improving accuracy and robustness. Comparing to common deep learning approaches like RoNIN, it elevates the problem onto 3D domain and does not require an external ZUPT-based orientation estimator. This tight fusion approach reduces average yaw and position drift by 27% and 33% respectively on our test dataset comparing to the best performing RoNIN velocity concatenation baseline approach. Fig. 1 shows an example of the estimated trajectory. We name our method TIGHT LEARNED INERTIAL ODOMETRY (TLIO).

## II. RELATED WORKS

**VIO.** Visual-Inertial Odometry has been well studied in the literature [7]. Without power and privacy issues and when static visual features are clearly trackable, VIO is the standard method for state estimation.

**PDR.** Solutions to Pedestrian Dead Reckoning problems can be divided into two categories: Bayesian filtering [8] and step counting [4]. With IMU as the only source of information for both propagation and update steps of an EKF, the measurements need to be carefully constructed. Different types of pseudo measurement approaches include Zero-velocity UPdaTe (ZUPT) [3] and Zero Angular Rate Update (ZARU) [9] which detects when the system is static, and heuristic heading reduction (HDR) [10] which identifies walking in a straight line to calibrate gyroscope bias. Foot-mounted IMU sensor enables reliable detection of contact from which accurate velocity constraints can be derived [11], however such information is not present for hand carried devices, and various signal processing and machine learning algorithms such as peak detection in time/frequency domain and activity classification have been explored [4]. These algorithms are complex and involve lots of tuning to accurately estimate a trajectory. It is also an active research direction as deep learning approaches solving for subproblems




Fig. 2: System block diagram. The IMU buffer provides segments of gravity-aligned IMU measurements to the network, using rotation from the filter state. The network outputs displacement  $\hat{d}$  and uncertainty  $\hat{u}$  used as measurement update to the filter. The filter estimates rotation, velocity, position and IMU biases at IMU rate.

such as gait classification [12] and stride detection [13] are investigated.

**Deep Learning.** The emergence of deep learning provides new possibilities to extract information from IMU data. Estimating average velocity from IMU data segments has seen great success in position estimation. IONet [5] first proposed an LSTM structure to output relative displacement in 2D polar coordinates and concatenate to obtain position. RIDI [14] and RoNIN [6] both assume orientation is known from an Android device to rotate IMU data into a gravity-aligned frame. While RIDI regresses velocity to optimize bias but still uses double integration from the corrected IMU data for position, RoNIN directly integrates the regressed velocity and showed better results. This shows that a statistical IMU motion model can outperform the traditional kinematic IMU model in scenarios that can be captured with training data.

In addition to using networks alone for pose estimates, Backprop KF [15] proposes an end-to-end differentiable Kalman filter framework. Because the loss function is on the accuracy of the filter outputs, the noise parameters are trained to produce the best state estimate, and do not necessarily best capture the measurement error model. AI-IMU [16] uses this approach to estimate IMU noise parameters and measurement uncertainties for car applications. In this work, we also combine deep learning with a Kalman filter. However, instead of training an end-to-end system, we leverage a state-of-the-art visual-inertial fusion algorithm as ground-truth for supervised learning of the measurement function itself. We use a likelihood loss function to jointly train for 3D displacement and uncertainty, which are directly used as input to the EKF in the measurement update.

## III. SYSTEM DESIGN

Our estimation system takes IMU data as input and has two major components: the network and the filter. Figure 2 shows a high-level diagram of the system.

The first component is a convolutional neural network trained to regress 3D relative displacement and uncertainty between two time instants given the segment of IMU data in between. The network is required to infer positional displacement over a short timespan from acceleration and angular velocity measurements alone, without access to the initial velocity. This is intentional: It forces the network

to learn only a prior expressing "typical human motion" and leaves model-based state propagation, i.e. acceleration integration to propagate velocity, and respective uncertainty propagation, to the second system component, the EKF. In order for this to work well, we found that it is important to include the inferred prior's uncertainty to the network output, allowing the network to encode how much motion model prior it obtained from the input measurements.

The EKF estimates the current state: 3D position, velocity, orientation and IMU biases, and a sparse set of past poses. The EKF propagates with raw IMU samples and uses network outputs for measurement updates. We define the measurement in a local *gravity-aligned frame* to decouple global yaw information from the relative state measurement (see Sec. V-D). The propagation from raw IMU data provides a model-based kinematic motion model, and the neural network provides a statistical motion model. The filter tightly couples these two sources of information.

At runtime, the raw IMU samples are interpolated to the network input frequency and rotated to a local gravity-aligned frame using rotations estimated from the filter state and gyroscope data. A gravity-aligned frame always has gravity pointing downward. Placing data in this frame implicitly gives the gravity direction information to the network.

Note that in the proposed approach, IMU data is being used twice, both as direct input for state propagation and indirectly through the network as the measurement. This violates the independence assumptions the EKF is based on: the errors in the IMU network input (initial orientation, sensor bias, and sensor noise) would propagate to its output which is used to correct the propagation results polluted with the same noise, which could lead to estimation inconsistency. We address this issue with training techniques (see Sec. IV) to reduce this error propagation by adding random perturbations to IMU bias and gravity direction during training. We show in Section VII that these techniques successfully improved the robustness of the network to sensor bias and rotation inaccuracies.

## IV. NEURAL STATISTICAL MOTION MODEL

### A. Architecture and Loss Function Design

Our network uses a 1D version of ResNet18 architecture proposed in [17]. The input dimension to the network is  $N \times 6$ , consisting of  $N$  IMU samples in the gravity-aligned frame. The output of the network contains two 3D vectors: the displacement estimates  $\hat{\mathbf{d}}$  and their uncertainties  $\hat{\mathbf{u}}$  which parametrize the diagonal entries of the covariance. The two vectors have independent fully-connected blocks extending the main ResNet architecture.

We make use of two different loss functions during training: the Mean Square Error (MSE) and the Gaussian Maximum Likelihood loss. The MSE loss on the trained dataset is defined as:

$$\mathcal{L}_{\text{MSE}}(\mathbf{d}, \hat{\mathbf{d}}) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{d}_i - \hat{\mathbf{d}}_i\|^2 \quad (1)$$

where  $\hat{\mathbf{d}} = \{\hat{\mathbf{d}}_i\}_{i \leq n}$  are the 3D displacement output of the network and  $\mathbf{d} = \{\mathbf{d}_i\}_{i \leq n}$  are the ground truth displacement.  $n$  is the number of data in the training set.

We define the Maximum Likelihood loss as the negative log-likelihood of the displacement according to the regressed Gaussian distribution:

$$\begin{aligned} \mathcal{L}_{\text{ML}}(\mathbf{d}, \hat{\mathbf{d}}, \hat{\mathbf{\Sigma}}) &= \frac{1}{n} \sum_{i=1}^n -\log \left( \frac{1}{\sqrt{8\pi^3 \det(\hat{\mathbf{\Sigma}}_i)}} e^{-\frac{1}{2} \|\mathbf{d}_i - \hat{\mathbf{d}}_i\|_{\hat{\mathbf{\Sigma}}_i}^2} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \left( \frac{1}{2} \log \det(\hat{\mathbf{\Sigma}}_i) + \frac{1}{2} \|\mathbf{d}_i - \hat{\mathbf{d}}_i\|_{\hat{\mathbf{\Sigma}}_i}^2 \right) + \text{Cst} \quad (2) \end{aligned}$$

where  $\hat{\mathbf{\Sigma}} = \{\hat{\mathbf{\Sigma}}_i\}_{i \leq n}$  are the  $3 \times 3$  covariance matrices for  $i^{\text{th}}$  data as a function of the network uncertainty output vector  $\hat{\mathbf{u}}_i$ .  $\hat{\mathbf{\Sigma}}_i$  has 6 degrees of freedom, and there are various covariance parametrizations for neural network uncertainty estimation [18]. In this paper, we simply assume a diagonal covariance output, parametrized by 3 coefficients written as:

$$\hat{\mathbf{\Sigma}}_i(\hat{\mathbf{u}}_i) = \text{diag}(e^{2\hat{u}_{xi}}, e^{2\hat{u}_{yi}}, e^{2\hat{u}_{zi}}) \quad (3)$$

The diagonal assumption decouples each axis, while regressing the logarithm of the standard deviations removes the singularity around zero in the loss function, adding numerical stabilization and helping the convergence in the optimization process. This choice constrains the principal axis of the uncertainty ellipses to be along the gravity-aligned frame axis.

### B. Data Collection and Implementation Details

We use a dataset collected by a custom rig where an IMU (Bosch BMI055) is mounted on a headset rigidly attached to the cameras. The full dataset contains more than 400 sequences totaling 60 hours of pedestrian data that pictures a variety of activities including walking, standing still, organizing the kitchen, playing pool, going up and down the stairs etc. It was captured with multiple different physical devices by more than 5 people to depict a wide range of individual motion patterns and IMU systematic errors. A state-of-the-art visual-inertial filter based on [19] provides position estimates at 1000 Hz on the entire dataset. We use these results both as supervision data in the training set and as ground truth in the test set. The dataset is split into 80% training, 10% validation and 10% test subsets randomly.

For network training, we use an overlapping sliding window on each sequence to collect input samples. Each window contains  $N$  IMU samples of total size  $N \times 6$ . In our final system we choose  $N = 200$  for 200 Hz IMU data. We want the network to capture a motion model with respect to the gravity-aligned IMU frame, therefore the IMU samples in each window are rotated from the IMU frame to a gravity-aligned frame built from the orientation at the beginning of the window. We use visual-inertial ground-truth rotation for that purpose. The supervision data for the network output is computed as the difference of the ground-truth position between two instants expressed in the same headset-centered, gravity-aligned frame.

During training, because we assume the headset can

be worn at an arbitrary heading angle with respect to the walking direction, we augment the input data for the network to be yaw invariant by giving a random horizontal rotation to each sample following RoNIN [6]. In our final estimator, the network is fed with potentially inaccurate input from the filter, especially at the initialization stage. We simulate this at training time by random perturbations on the sensor bias and the gravity direction to reduce network sensitivity to these input errors. To simulate bias, we generate additive bias vectors with each component independently sampled from uniform distribution in  $[-0.2, 0.2] m/s^2$  or  $[-0.05, 0.05] rad/s$  for each input sample. Gravity direction is perturbed by rotating those samples along a random horizontal rotation axis with magnitude sampled from  $[0, 5^\circ]$ .

Optimization is done through the Adam optimizer. We used an initial learning rate of 0.0001, zero weight decay, and dropouts with a probability of 0.5 for the fully connected layers. We observe that training for  $\mathcal{L}_{ML}$  directly does not converge. Therefore we first train with  $\mathcal{L}_{MSE}$  for 10 epochs until the network stabilizes, then switch to  $\mathcal{L}_{ML}$  until the network fully converges. It takes around another 10 epochs to converge and a total of 4 hours of training time is needed on an NVIDIA DGX computer.

## V. STOCHASTIC CLONING EXTENDED KALMAN FILTER

The EKF in our system tightly integrates the displacements predicted by the network with a statistical IMU model as used in other inertial navigation systems. As the displacement estimates from the network express constraints on pairs of past states, we adopt a stochastic cloning framework [20]. Similar to [19], we maintain a sliding window of  $m$  poses in the filter state. In contrast to [19], however, we only apply constraints between *pairs* of poses, and these constraints are derived from the network described in the preceding section, rather than camera data.

### A. State Definition

At each instant, the full state of the EKF is defined as:

$$\mathbf{X} = (\xi_1, \dots, \xi_m, \mathbf{s})$$

where  $\xi_i$ ,  $i = 1, \dots, m$  are past (cloned) states, and  $\mathbf{s}$  is the current state. More specifically,

$$\xi_i = ({}^w\mathbf{R}_i, {}^w\mathbf{p}_i), \quad \mathbf{s} = ({}^w\mathbf{R}, {}^w\mathbf{v}, {}^w\mathbf{p}, \mathbf{b}_g, \mathbf{b}_a)$$

We express  ${}^w\mathbf{R}$  as the rotation matrix that transforms a point from the IMU frame to the world frame, and  ${}^w\mathbf{v}$  and  ${}^w\mathbf{p}$  are respectively the velocity and position expressed in the world frame. In the following, we will drop the explicit superscript for conciseness.  $\mathbf{b}_g, \mathbf{b}_a$  are the IMU gyroscope and accelerometer biases.

As commonly done in such setups, we apply the error-based filtering methodology to linearize locally on the manifold of the minimal parametrization of the rotation. More specifically, the filter covariance is defined as the covariance of the following error-state:

$$\tilde{\xi}_i = (\tilde{\theta}_i, \delta\tilde{\mathbf{p}}_i), \quad \tilde{\mathbf{s}} = (\tilde{\theta}, \delta\tilde{\mathbf{v}}, \delta\tilde{\mathbf{p}}, \delta\tilde{\mathbf{b}}_g, \delta\tilde{\mathbf{b}}_a)$$

where *tilde* indicates errors for every state as the difference between the estimate and the real value, except for rotation error which is defined as  $\tilde{\theta} = \log_{SO3}(\hat{\mathbf{R}}\mathbf{R}^{-1}) \in \mathfrak{so}(3)$  where  $\log_{SO3}$  denotes the logarithm map of rotation. The full error-state is of dimension  $6m + 15$ , where  $m$  is the total number of past states, and  $\tilde{\mathbf{s}}$  has dimension of 15.

### B. IMU model

We assume that the IMU sensor provides direct measurements of the non-gravitational acceleration  $\mathbf{a}$  and angular velocity  $\boldsymbol{\omega}$ , expressed in the IMU frame. We assume as it is common for this class of sensor that the measurements are polluted with noise  $\mathbf{n}_{g;a}$  and bias  $\mathbf{b}_{g;a}$ .

$$\boldsymbol{\omega} = \boldsymbol{\omega}_{\text{true}} + \mathbf{b}_g + \mathbf{n}_g, \quad \mathbf{a} = \mathbf{a}_{\text{true}} + \mathbf{b}_a + \mathbf{n}_a.$$

$\mathbf{n}_g$  and  $\mathbf{n}_a$  are random noise variables following a zero-centered Gaussian distribution. Evolution of biases is modeled as a random walk process with discrete parameters  $\boldsymbol{\eta}_{gd}$  and  $\boldsymbol{\eta}_{ad}$  over the IMU sampling period  $\Delta t$ .

### C. State Propagation and Augmentation

The filter propagates the current state  $\mathbf{s}$  with IMU data using a kinematic motion model. If the current timestamp is associated to a measurement update, stochastic cloning is performed together with propagation in one step. During cloning, a new state  $\xi$  is appended to the past state.

1) *Propagation Model*: We use the strapdown inertial kinematic equation assuming a uniform gravity field  ${}^w\mathbf{g}$  and ignoring Coriolis forces and the earth's curvature:

$${}^w\hat{\mathbf{R}}_{k+1} = {}^w\hat{\mathbf{R}}_k \exp_{SO3}((\boldsymbol{\omega}_k - \hat{\mathbf{b}}_{gk})\Delta t) \quad (4)$$

$${}^w\hat{\mathbf{v}}_{k+1} = {}^w\hat{\mathbf{v}}_k + {}^w\mathbf{g}\Delta t + {}^w\hat{\mathbf{R}}_k(\mathbf{a}_k - \hat{\mathbf{b}}_{ak})\Delta t \quad (5)$$

$${}^w\hat{\mathbf{p}}_{k+1} = {}^w\hat{\mathbf{p}}_k + {}^w\hat{\mathbf{v}}_k\Delta t + \frac{1}{2}\Delta t^2({}^w\mathbf{g} + {}^w\hat{\mathbf{R}}_k(\mathbf{a}_k - \hat{\mathbf{b}}_{ak})) \quad (6)$$

$$\hat{\mathbf{b}}_{g(k+1)} = \hat{\mathbf{b}}_{gk} + \boldsymbol{\eta}_{gdk} \quad (7)$$

$$\hat{\mathbf{b}}_{a(k+1)} = \hat{\mathbf{b}}_{ak} + \boldsymbol{\eta}_{adk}. \quad (8)$$

Here,  $\exp_{SO3}$  denotes the SO(3) exponential map, the inverse function of  $\log_{SO3}$ . The discrete-time noise  $\boldsymbol{\eta}_{gdk}, \boldsymbol{\eta}_{adk}$  have been defined above and follow a Gaussian distribution.

The linearized error propagation can be written as:

$$\tilde{\mathbf{s}}_{k+1} = \mathbf{A}_k^s \tilde{\mathbf{s}}_k + \mathbf{B}_k^s \mathbf{n}_k \quad (9)$$

where  $\mathbf{n}_k = [\mathbf{n}_{\omega k}, \mathbf{n}_{ak}, \boldsymbol{\eta}_{gdk}, \boldsymbol{\eta}_{adk}]^T$  is a vector containing all random input. The superscript  $s$  indicates that these matrices correspond to the current state  $\mathbf{s}$ , and the subscript brackets indicate matrix dimensions. The corresponding propagation of the state covariance  $\mathbf{P}$  is:

$$\mathbf{P}_{k+1} = \mathbf{A}_k \mathbf{P}_k \mathbf{A}_k^T + \mathbf{B}_k \mathbf{W} \mathbf{B}_k^T, \quad (10)$$

$$\mathbf{A}_k = \begin{bmatrix} \mathbf{I}_{6m} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_k^s \end{bmatrix}, \mathbf{B}_k = \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_k^s \end{bmatrix} \quad (11)$$

with  $\mathbf{W}_{(12,12)}$  the covariance matrices of sensor noise and bias random walk noise. The state covariance  $\mathbf{P}$  is of the dimension  $(6m + 15)$  of the full state  $\mathbf{X}$ . In the implementation, multiple propagation steps can be combined together to reduce the computational cost [20].

2) *State Augmentation*: In our system, state augmentations are performed at the measurement update frequency. During an augmentation step, the state dimension is incremented through propagation with cloning:

$$\mathbf{P}_{k+1} = \bar{\mathbf{A}}_k \mathbf{P}_k \bar{\mathbf{A}}_k^T + \bar{\mathbf{B}}_k \mathbf{W} \bar{\mathbf{B}}_k^T, \quad (12)$$

$$\bar{\mathbf{A}}_k = \begin{bmatrix} \mathbf{I}_{6m} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_k^\xi \\ \mathbf{0} & \mathbf{A}_k^s \end{bmatrix}, \bar{\mathbf{B}}_k = \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_k^\xi \\ \mathbf{B}_k^s \end{bmatrix} \quad (13)$$

$\bar{\mathbf{A}}_k$  is now a copy operation plus the augmented and current state propagation operations.  $\mathbf{A}_k^\xi$  and  $\mathbf{B}_k^\xi$  are partial propagation matrices for rotation and position only. After the increment, the dimension of the state vector increases by 6. Old past states will be pruned in the marginalization step, see Sec.V-E.

#### D. Measurement Update

It would be natural to define the measurement function as the displacement  $\mathbf{d}$  expressed in the world frame, however such measurement function would imply heading observability at the filter level. Heading is theoretically unobservable as the IMU propagation model and the learned prior are invariant to the change of yaw angle. In order to prevent injecting spurious information into the filter, we carefully define the measurement function  $h$  as the 3D displacement in a *local gravity-aligned* frame. This frame is anchored to a clone state  $i$ . Vectors in this frame can be obtained by rotating the corresponding world frame vectors by the yaw rotation matrix  $\mathbf{R}_\gamma$  of the state rotation matrix  $\mathbf{R}_i$ . We decompose  $\mathbf{R}_i$  using extrinsic "XYZ" Euler angle convention:  $\mathbf{R}_i = \mathbf{R}_\gamma \mathbf{R}_\beta \mathbf{R}_\alpha$ , where  $\alpha, \beta, \gamma$  correspond to roll, pitch and yaw respectively.  $h$  then writes:

$$h(\mathbf{X}) = \mathbf{R}_\gamma^T (\mathbf{w} \mathbf{p}_j - \mathbf{w} \mathbf{p}_i) = \hat{\mathbf{d}}_{ij} + \boldsymbol{\eta}_{d_{ij}}. \quad (14)$$

$\hat{\mathbf{d}}_{ij}$  is the network output displacement between state  $i$  and  $j$ , using IMU samples rotated to the local gravity-aligned frame anchored at pose  $i$  as input.  $\boldsymbol{\eta}_{d_{ij}}$  is a random variable that, we assume, follows the normal distribution  $\mathcal{N}(0, \hat{\boldsymbol{\Sigma}}_{ij})$  given by the network.

The linearization of the measurement function  $h$  with respect to the error state yields the linearized measurement matrix  $\mathbf{H}_{(3,6m+15)}$ . It has only non zero values in the  $3 \times 3$  blocks corresponding to  $\tilde{\boldsymbol{\theta}}_i, \tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_i$  and  $\tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_j$ .

$$\mathbf{H}_{\tilde{\boldsymbol{\theta}}_i} = \frac{\partial h(\mathbf{X})}{\partial \tilde{\boldsymbol{\theta}}_i} = \hat{\mathbf{R}}_\gamma^T [\mathbf{w} \hat{\mathbf{p}}_j - \mathbf{w} \hat{\mathbf{p}}_i]_\times \mathbf{H}_z \quad (15)$$

$$\mathbf{H}_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_i} = \frac{\partial h(\mathbf{X})}{\partial \tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_i} = -\hat{\mathbf{R}}_\gamma^T \quad (16)$$

$$\mathbf{H}_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_j} = \frac{\partial h(\mathbf{X})}{\partial \tilde{\boldsymbol{\delta}}\tilde{\mathbf{p}}_j} = \hat{\mathbf{R}}_\gamma^T \quad (17)$$

$$\text{where } \mathbf{H}_z = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \cos \gamma \tan \beta & \sin \gamma \tan \beta & 1 \end{bmatrix} \quad (18)$$

In the above expression,  $[\mathbf{x}]_\times$  is the skew-symmetric matrix built from vector  $\mathbf{x}$ . The singularity  $\cos \beta = 0$  occurs when the person wearing the headset is looking straight down or

straight up: we simply discard the update for these cases. Note that these situations are unlikely and never occur in our dataset.

Finally,  $\mathbf{H}$  and  $\hat{\boldsymbol{\Sigma}}_{ij}$  are used to compute the Kalman gain to update the state  $\mathbf{X}$  and the covariance  $\mathbf{P}$  as follows:

$$\mathbf{K} = \mathbf{P} \mathbf{H}^T (\mathbf{H} \mathbf{P} \mathbf{H}^T + \hat{\boldsymbol{\Sigma}}_{ij})^{-1} \quad (19)$$

$$\mathbf{X} \leftarrow \mathbf{X} \oplus \left( \mathbf{K} (h(\mathbf{X}) - \hat{\mathbf{d}}_{ij}) \right) \quad (20)$$

$$\mathbf{P} \leftarrow (\mathbf{I} - \mathbf{K} \mathbf{H}) \mathbf{P} (\mathbf{I} - \mathbf{K} \mathbf{H})^T + \mathbf{K} \hat{\boldsymbol{\Sigma}}_{ij} \mathbf{K}^T \quad (21)$$

Operator  $\oplus$  denotes the regular addition operation except for rotation where the update operation writes  $\mathbf{R} \leftarrow \text{Exp}(\boldsymbol{\theta}) \mathbf{R}$ .

We use a  $\chi_2$  test to protect the filter estimate against occasional wrong network output: we discard the update when the normalized innovation error  $\|\mathbf{d}_{ij} - \hat{\mathbf{d}}_{ij}\|_{\mathbf{H} \mathbf{P} \mathbf{H}^T + \hat{\boldsymbol{\Sigma}}_{ij}}$  is greater than a threshold. We choose here the threshold value of 11.345 which corresponds to the 99<sup>th</sup> percentile of the  $\chi_2$  distribution with 3 degrees of freedom.

In practice, when using the measurement covariance  $\hat{\boldsymbol{\Sigma}}_{ij}$  from the network, we scale the covariance by a factor of 10 to compensate for the temporal correlation of measurements as noted in Sec. VII-A.1.

#### E. State size, Marginalization and Initialization

As soon as an update is processed, all states prior to this update window are marginalized. This is performed by removing the corresponding  $\boldsymbol{\xi}$  from the state and the corresponding covariance entries. The number of states kept in the filter depends on two parameters: (i) the displacement duration window for which the network is trained and (ii) the update frequency. For instance, for an update frequency of 20 Hz and a window of 1 s, there will be at most 21 past states in the filter. The two parameters can be set independently: if the period of update is smaller than the displacement window, the consecutive measurements will be generated from the overlapping windows of IMU data, while if the update period is too long, some IMU measurements would not be given to the network. We observe that our system has better performance with a higher update frequency (see Sec. VII-B.2), and we use 20 Hz in our final system.

Our EKF needs a good initialization to converge. In this work we assumed the initial state is given by an external procedure. In our experiments we initialize the speed, roll, and pitch with the ground-truth estimate assigning a large covariance, while the yaw and the position are initialized with a strong prior in order to fix the gauge. Biases are initialized as with the values of an initial factory calibration, and we refine the biases online to account for its evolution due to turn-on change, temperature effect, or aging.

In practice, we use the following values to initialize the state covariance:  $\sigma_{\mathbf{w}} = 0.1$  m/s,  $\sigma_{\mathbf{b}_a} = 0.2$  m/s<sup>2</sup>,  $\sigma_{\mathbf{b}_g} = 1.10^{-4}$  rad/s,  $\sigma_\theta = \text{diag}(10, 10, 0.1)$  deg. We also compensate all the input IMU samples for non-orthogonality scale factor and g-sensitivity from the factory calibration.

## VI. EXPERIMENTAL RESULTS

We evaluate our system on the test split of our dataset. Each of these 37 trajectories contains between 3 to 7 minutes

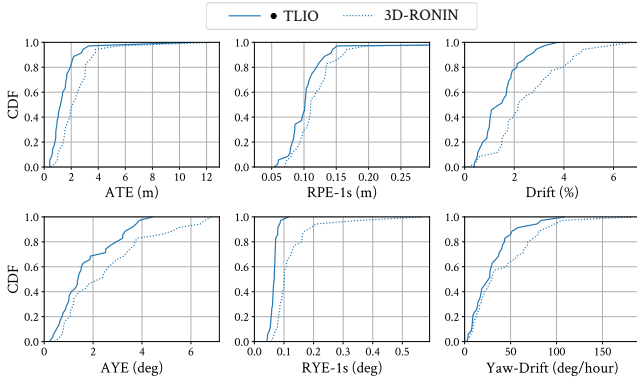


Fig. 3: Comparing TLIO to 3D-RONIN on the error metrics with respect to the ground-truth. Each plot shows the cumulative density function of the chosen metric on the entire test set. The steeper the plots are the better the performance. TLIO shows significantly lower value than 3D-RONIN for all presented metrics.

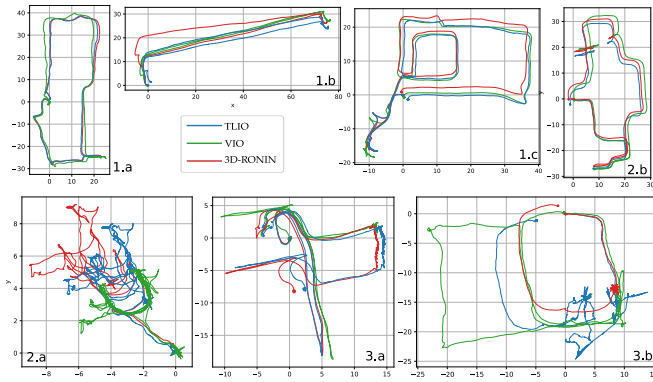


Fig. 4: Selection of trajectories. 1.x and 2.x are the good and hard cases and 3.x show examples of failures. 2.a. shows a case with very wrong network outputs. 2.b shows a case of bad initialization in the filter leading to yaw drift. In 3.a and 3.b, unusual motions not present in the training set are performed: side-stepping repeatedly for 3 minutes and rolling on a chair.

of human activities. We compare the performance of two pose estimators against a state-of-the-art VIO implementation based on [19] and considered as ground-truth (GT):

- TLIO: the tightly-coupled EKF with displacement update from the trained network presented in this work.
- 3D-RONIN: an estimator where the displacements from the same trained network are concatenated in the direction given by an engineered AHRS attitude filter, resembling what smartphones have. We borrow the name from [6], but we extended the method to 3D and we trained the network on our dataset to get a fair comparison. Note that our dataset contains trajectories involving climbing staircases, sitting motions, and walking outdoors on uneven terrain, where the 2D PDR method of [6] is not directly applicable.

#### A. Metrics definitions

In order to assess our system performance, for each dataset of length  $n$ , we define the following metrics:

- **ATE (m)**:  $\sqrt{\frac{1}{n} \sum_i \|\mathbf{w} \mathbf{p}_i - \mathbf{w} \hat{\mathbf{p}}_i\|^2}$

The Absolute Translation Error indicates the spatial closeness of position estimate to the GT over the whole trajectory, computed as root-mean square error (RMSE) between these sets of values.

- **RTE- $\Delta t$  (m)**:

$$\sqrt{\frac{1}{n} \sum_i \|\mathbf{w} \mathbf{p}_{i+\Delta t} - \mathbf{w} \mathbf{p}_i - \mathbf{R}_\gamma \hat{\mathbf{R}}_\gamma^T (\mathbf{w} \hat{\mathbf{p}}_{i+\Delta t} - \mathbf{w} \hat{\mathbf{p}}_i)\|^2}$$

The Relative Translation Error indicates the local spatial closeness of position estimate to the GT over a window of duration  $\Delta t$ . We use 1s in this analysis. We remove the yaw drift at the beginning of the window so that the relative measure is not affected by accumulated errors.

- **DR (%)**:  $(\|\mathbf{w} \mathbf{p}_n - \mathbf{w} \hat{\mathbf{p}}_n\|) / (\text{trajectory-length})$

The final translation drift over the distance traveled.

We compute similar metrics for yaw angle  $\gamma$ , measuring the quality of the unobservable orientation around gravity:

- **AYE ( $^\circ$ )**:  $\sqrt{(\frac{1}{n} \sum_i \|\gamma_i - \hat{\gamma}_i\|^2)}$  is the Absolute Yaw Error.
- **RYE- $\Delta t$  ( $^\circ$ )**:  $\sqrt{\frac{1}{n} \sum_i \|\gamma_{i+\Delta t} - \gamma_i - (\hat{\gamma}_{i+\Delta t} - \hat{\gamma}_i)\|^2}$  is the Relative Yaw Error
- **Yaw-DR ( $^\circ/\text{hour}$ )**:  $(\gamma_n - \hat{\gamma}_n) / (\text{sequence-duration})$  is the Yaw Drift over time.

#### B. System Performance

Figure 3 shows the distribution of the metrics across the entire test set. It demonstrates that our system consistently performs better than the decoupled orientation and position estimator 3D-RONIN on all metrics.

On 3D position, TLIO performs better than 3D-RONIN which integrates average velocities. Integration approach has the advantage of being robust to outliers since the measurements at each instant are decoupled. The result shows not only the benefit of integrating IMU kinematic model, but also the overall robustness of the filter, which comes from the quality of the covariance output from the network and the effectiveness of outlier rejection with  $\chi_2$  test.

Our system also has a smaller yaw drift than 3D-RONIN, indicating that even without any hand engineered heuristics or step detection, using displacement estimates from a trained statistical model outperforms a smartphone AHRS attitude filter. This shows that the EKF can accurately estimate not only position and velocity, but also gyroscope biases.

Figure 4 shows a hand-picked collection of 7 trajectories containing the best, typical and worst cases of TLIO and 3D-RONIN, while Figure 1 shows a 3D visualization of one additional sequence. See the corresponding captions for more details about the trajectories and failure cases.

### VII. COMPONENTS AND VARIATION STUDIES

We investigate different variations of the network settings in three aspects: input IMU frequency, time interval  $\Delta t_{ij}$  over which displacement is computed, and the total time interval included in the network input. We do not consider future data because we aim for a causal filter. The name convention we use is built from these aspects. For example, 200hz-05s-3s is a model that takes 3s of 200 Hz input data and regresses for the displacement over the last 0.5s.




Fig. 5: Network variations with different past and total window sizes evaluated on the test set. Upper left: ATE of 3D-ROIN concatenation result comparing to GT trajectories. Upper right: average MSE loss on the test set. Adding past data to the network input reduces average MSE loss over samples, but it does not imply lower ATE over trajectories.

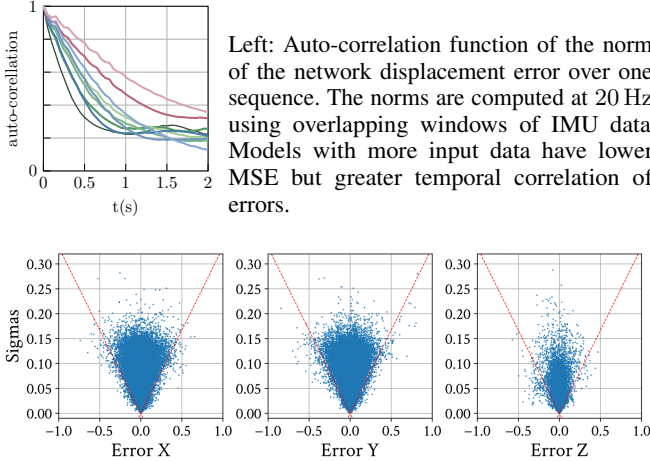


Fig. 6: Uncertainty returned by the network (standard-deviation  $\sigma$ ) plotted against errors (m) on three axes  $x, y, z$  respectively. Points with larger errors have larger variance outputs, and over 99% of the points fall inside the  $3\sigma$  cone region indicated by the dashed red line. We have 0.70% for  $x, y$  and 0.47% for  $z$  axis of error points outside  $3\sigma$  bounds respectively.

To easily keep track of the different models, in all following figures, we prefix with “•” the version of our algorithm whose results were presented in Sec. VI: 200hz-1s-1s.

### A. Network Statistical Model

In this section, we evaluate the deep learning component of our system in isolation.

1) *Does using more data help?*: Figure 5 shows a comparison of several variations of the window duration, keeping the IMU frequency fixed to 200 Hz. It is surprising that lower MSE loss over the same displacement window does not imply lower ATE. This is because MSE does not say anything about the sum of the errors: the network can make more accurate predictions per sample by seeing more data, while at the same time produces more correlated results as the inputs overlap over a longer period of time. This introduces a trade-off when integrating displacement, and we observe similar trajectory ATEs in different network time window variants. For further experiments we will use 200hz-1s-1s except noted otherwise.

2) *Consistency of learned covariance*: We collected around 60 000 samples at 5 Hz on the entire test set for this analysis and the sensitivity analysis section below. Fig. 6 plots the




Fig. 7: Sensitivity of the network prediction on bias and gravity direction noise. At test time the input IMU samples are perturbed for accelerometer bias, gyroscope bias, and gravity direction as described in Sec. IV-B but with various ranges. The three models under comparison are trained with different data augmentation procedures: no perturbation (noptrb), perturbing only the bias (bias), and perturbing both the bias and the gravity direction (bias-grav). We observe that training with perturbation significantly improves robustness on accelerometer bias and gravity direction.

sigma outputs of the network against the errors on the displacement. We observe over 99% of the error points fall inside the  $3\sigma$  region, showing the network outputs are consistent statistically. In addition, we compute the Mahalanobis distance of the output displacement errors scaled by the output covariances in 3D. We found that only 0.30% of the samples were beyond the 99 percentile critical value of the  $\chi^2$  distribution, all of which are from the failure case shown in Fig. 4.3.b). From this analysis, we conclude that the uncertainty output of the network grows as expected with its own error albeit being often slightly conservative. We observe an asymmetry between horizontal and vertical errors distribution; along vertical axis errors and sigmas concentrate on lower values. This is not surprising: people usually either walk on a flat ground or on stairs which may make expressing a prior easier on the  $z$  axis.

3) *Sensitivity Analysis*: Figure 7 shows the network output robustness to input perturbation on bias and gravity direction. We observe that the models trained with bias perturbation are more robust to IMU bias, in particular accelerometer bias. The model trained with gravity perturbation is significantly more robust to gravity direction errors. These network training techniques effectively help reduce the propagation of errors from input to the output, protecting the independence assumption required in a Kalman Filter framework. Therefore we choose 200hz-1s-1s trained with bias and gravity perturbation as our system model.

### B. EKF System

1) *Ablation Study*: We compare system variants using a hand-tuned constant covariance and networks trained without likelihood loss to show the benefit of using the regressed uncertainty for TLIO. To find the best constant covariance parameters, we use the standard deviation of the measurement errors on the test set -  $\text{diag}(0.051, 0.051, 0.013)\text{m}$  - multiplied by a factor yielding the best median ATE tuned by grid-search. Fig. 8 shows the Cumulative Density Function of ATE, Drift and AYE over the test set of various system configurations. We observe that using 3D-ROIN, training the network with only MSE loss gives a more accurate estimate. However, using a fixed covariance, such network variants in a filter system achieve similar performance. What makes




Fig. 8: Performance comparison showing the effectiveness of using a learned covariance. TLIO-mse and 3D-ROIN-mse use a network trained with only MSE loss. TLIO-fixcov and TLIO-mse use a hand-tuned constant covariance for all the measurements. 3D-ROIN and 3D-ROIN-mse concatenate displacement estimates directly without considering uncertainty. TLIO achieves the best performance with the covariance regressed from the network. Constant covariance approaches do not reach 100% for ATE and drift in this illustration due to a failure case.

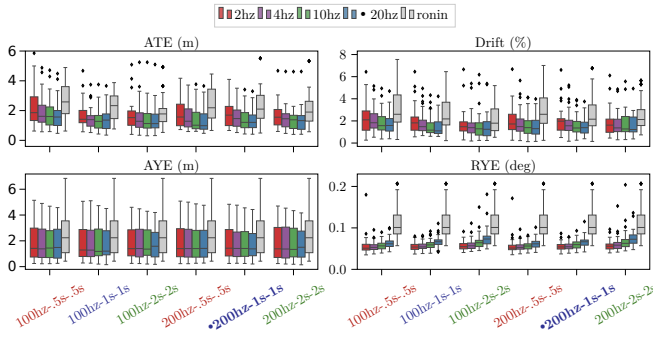


Fig. 9: Performance of different system configurations. Each group corresponds to a network model, as indicated by the x axis labels. Within each group, the colored boxplots differ by update frequency, and the grey ones show 3D-ROIN as baseline. The filtering approach constantly outperforms 3D-ROIN on all the experimented models. High frequency update yields the best ATE and drift in spite of temporal correlation of the measurements. RYE-1s increases with update frequency, indicating more jitter on the yaw estimates.

a difference is the regressed covariance from the network, which significantly improves the filter in ATE and Drift. We also notice a dataset where using a fixed covariance loses track due to initialization failure, showing the adaptive property of the regressed uncertainty for different inputs, improving system robustness. Comparing to 3D-ROIN-mse, TLIO has an improvement of 27% and 31% on the average yaw and position drift respectively.

2) *Timing parameters*: Fig. 9 shows the performance comparison between variations on network IMU frequency,  $\Delta t_{ij}$ , and filter measurement-update frequency. ATE and drift are reduced when using high frequency measurements. Once again, we observe minor differences between network parameter variations on the monitored metrics. This shows that our entire pipeline is not sensitive to these choices of parameters, and TLIO outperforms 3D-ROIN in all cases.

## VIII. CONCLUSION

In this paper we propose a tightly-coupled inertial odometry algorithm introducing a learned component in an EKF framework. We present the network regressing 3D displacement and uncertainty, and the EKF fusing the displacement measure-

ment to estimate the full state. We train a network that learns a prior on the displacement distributions given IMU data from statistical motion patterns. We show through experimental results and variation studies that the network outputs are statistically consistent, and the filter outperforms the state-of-the-art trajectory estimator using velocity integration on position estimates, as well as a model-based AHRS attitude filter on orientation. This demonstrates that with a learned prior, an IMU sensor alone can provide enough information to do low drift pose estimation and calibration for pedestrian dead-reckoning. As common with learning approaches, the system is limited by the scope of training data. Unusual motions cause system failure as discussed with examples. Whether similar approach can be generalized to wider use cases such as legged robots is still an unexplored but promising field of research.

## REFERENCES

- [1] P. D. Groves, "Principles of GNSS, inertial, and multisensor integrated navigation systems, [book review]," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 30, no. 2, pp. 26–27, Feb. 2015.
- [2] R. Harle, "A survey of indoor inertial positioning systems for pedestrians," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1281–1293, 2013.
- [3] E. Foxlin, "Pedestrian tracking with shoe-mounted inertial sensors," *IEEE CG&A*, vol. 25, no. 6, pp. 38–46, 2005.
- [4] A. Brajdic and R. Harle, "Walk detection and step counting on unconstrained smartphones," in *UbiComp*, 2013, pp. 225–234.
- [5] C. Chen, X. Lu, A. Markham, and N. Trigoni, "IONet: Learning to cure the curse of drift in inertial odometry," *AAAI*, 2018, pp. 6468–6476.
- [6] H. Yan, S. Herath, and Y. Furukawa, "RoNIN: Robust neural inertial navigation in the wild: Benchmark, evaluations, and new methods," *ICRA*, pp. 3146–3152, 2020.
- [7] C. Chen, H. Zhu, M. Li, and S. You, "A review of visual-inertial simultaneous localization and mapping from filtering-based and optimization-based perspectives," *Robotics*, vol. 7, no. 3, p. 45, 2018.
- [8] A. R. Jimenez, F. Seco, J. C. Prieto, and J. Guevara, "Indoor pedestrian navigation using an INS/EKF framework for yaw drift reduction and a foot-mounted IMU," in *7th WPNC*. IEEE, 2010.
- [9] S. Rajagopal, "Personal dead reckoning system with shoe mounted inertial sensors," Master's Degree Project, Stockholm, Sweden, 2008.
- [10] J. Borenstein, L. Ojeda, and S. Kwanmuang, "Heuristic reduction of gyro drift in IMU-based personnel tracking systems," in *Optics and Photonics in Global Homeland Security V and Biometric Technology for Human Identification VI*, vol. 7306, 2009, p. 73061H.
- [11] H. Ju, S. Y. Park, and C. G. Park, "Foot mounted inertial navigation system using estimated velocity during the contact phase," *IPIN*, 2016.
- [12] O. Dehzangi, M. Taherisadr, and R. ChagalVala, "IMU-based gait recognition using convolutional neural networks and multi-sensor fusion," *Sensors*, vol. 17, no. 12, p. 2735, 2017.
- [13] B. Beaufile, F. Chazal, M. Grelet, and B. Michel, "Robust stride detector from ankle-mounted inertial sensors for pedestrian navigation and activity recognition with machine learning approaches," *Sensors*, vol. 19, no. 30, p. 4491, 2019.
- [14] H. Yan, Q. Shan, and Y. Furukawa, "RIDI: Robust IMU double integration," *ECCV*, 2018, pp. 621–636.
- [15] T. Haarnoja, A. Ajay, S. Levine, and P. Abbeel, "Backprop KF" learning discriminative deterministic state estimators," *NIPS*, 2016, pp. 4376–4384.
- [16] M. Brossard, A. Barrau, and S. Bonnabel, "AI-IMU dead-reckoning," *IEEE Transactions on Intelligent Vehicles*, 2019.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CVPR*, 2016, pp. 770–778.
- [18] R. L. Russell and C. Reale, "Multivariate uncertainty in deep learning," *arXiv preprint arXiv:1910.14215*, 2019.
- [19] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint kalman filter for vision-aided inertial navigation," *ICRA*, 2007, pp. 3565–3572.
- [20] S. I. Roumeliotis and J. W. Burdick, "Stochastic cloning: A generalized framework for processing relative state measurements," *ICRA*, vol. 2, 2002, pp. 1788–1795.