

Denoising IMU Gyroscopes with Deep Learning for Open-Loop Attitude Estimation

Martin BROSSARD*, Silvère BONNABEL*[†] and Axel BARRAU*[‡]

*Centre for Robotics, MINES ParisTech, PSL Research University, 60 Boulevard Saint-Michel, 75006 Paris, France

[†]University of New Caledonia, ISEA, 98851, Noumea Cedex, New Caledonia

[‡]Safran Tech, Groupe Safran, Rue des Jeunes Bois-Châteaufort, 78772, Magny Les Hameaux Cedex, France

Abstract—This paper proposes a learning method for denoising gyroscopes of Inertial Measurement Units (IMUs) using ground truth data, and estimating in real time the orientation (attitude) of a robot in dead reckoning. The obtained algorithm outperforms the state-of-the-art on the (unseen) test sequences. The obtained performances are achieved thanks to a well-chosen model, a proper loss function for orientation increments, and through the identification of key points when training with high-frequency inertial data. Our approach builds upon a neural network based on dilated convolutions, without requiring any recurrent neural network. We demonstrate how efficient our strategy is for 3D attitude estimation on the EuRoC and TUM-VI datasets. Interestingly, we observe our dead reckoning algorithm manages to beat top-ranked visual-inertial odometry systems in terms of attitude estimation although it does not use vision sensors. We believe this paper offers new perspectives for visual-inertial localization and constitutes a step toward more efficient learning methods involving IMUs. Our open-source implementation is available at <https://github.com/mbrossard/denoise-imu-gyro>.

Index Terms—localization, odometry, deep learning in robotics and automation, autonomous systems navigation, drones

I. INTRODUCTION

Inertial Measurement Units (IMUs) consist of gyroscopes (gyros) that measure angular velocities, i.e., the rate of change of the sensor's orientation, and accelerometers (acceleros) that measure accelerations [1]. IMUs allow estimating a robot's trajectory relative to its starting position, a task called odometry [2] or dead reckoning. Small and cheap IMUs are ubiquitous in smartphones, industrial and robotics applications, but suffer from difficulties in estimating various errors, such as axis-misalignment, scale factors and time-varying offsets [3,4]. Hence, not only are IMU signals noisy, but they are biased. In the present paper, we propose to leverage deep learning for denoising the gyroscopes of an IMU, that is, reduce noise and biases, based on gyros and acceleros' signals. As a byproduct, we obtain accurate attitude (i.e. orientation) estimates simply through open-loop integration of the obtained noise-free gyro measurements.

A. Links and Differences with Existing Literature

IMUs are generally combined with complementary sensors to obtain robust pose estimates in sensor-fusion systems [5], using cameras in Visual-Inertial Odometry (VIO) [6]–[8], LiDAR, GNSS, or side information about the model [9]–[12]. To obtain accurate pose estimates, a proper IMU calibration is




Fig. 1. Schematic illustration of our method. The convolutional neural network computes gyro corrections (based on past IMU measurements) that filter undesirable errors in the raw IMU signals. We then perform *open-loop* time integration on the noise-free measurements for regressing low frequency errors between ground truth and estimated orientation increments.

required, see e.g. the widely used *Kalibr* library [3,13], which computes offline the underlying IMU intrinsic calibration parameters, and extrinsic parameters between the camera and IMU. Our approach, recapped in Figure 1, applies to any system equipped with an IMU. It estimates offline the IMU calibration parameters and extends methods such as [3,13] to time-varying and difficult to model signal corrections.

Machine learning (more specifically deep learning) has been recently leveraged to perform LiDAR, visual-inertial, and purely inertial localization, where methods are divided into supervised [14]–[18] and unsupervised [19] ones. Most works extract relevant features in the sensors' signals which are propagated in time through recurrent neural networks, whereas [20] proposes convolutional neural networks for pedestrian inertial navigation. A related approach [21] applies reinforcement learning to guide the user in calibrating visual-inertial rigs. Our method is supervised (we require ground truth poses for training), leverages convolutions rather than recurrent architectures, and outperforms the latter approach. We obtain excellent results while requiring considerably fewer data and less time. Finally, the reference [17] estimates orientation with an IMU and recurrent neural networks, but our approach proves simpler.

B. Contributions

Our main contributions are as follows:

- detailed modelling of the problem of learning orientation increments from low-cost IMUs;
- the convolutional neural network which regresses gyro corrections and whose features are carefully selected;
- the training procedure which involves a trackable loss function for estimating relative orientation increments;

- the approach evaluation on datasets acquired by a drone and a hand-held device [22,23], where our method outperforms [17] and competes with VIO methods [6,7] although it does not use vision;
- perspectives towards more efficient VIO and IMU based learning methods;
- a publicly available open-source code, where training is done in 5 minutes per dataset.

II. KINEMATIC & LOW-COST IMU MODELS

We detail in this section our model.

A. Kinematic Model based on Orientation Increments

The 3D orientation of a rigid platform is obtained by integrating orientation increments, that is, gyro outputs of an IMU, through

$$\mathbf{R}_n = \mathbf{R}_{n-1} \exp(\boldsymbol{\omega}_n dt), \quad (1)$$

where the rotation matrix $\mathbf{R}_n \in SO(3)$ at timestamp n maps the IMU frame to the global frame, the angular velocity $\boldsymbol{\omega}_n \in \mathbb{R}^3$ is averaged during dt , and with $\exp(\cdot)$ the $SO(3)$ exponential map. The model (1) successively integrates in open-loop $\boldsymbol{\omega}_n$ and propagates estimation errors. Indeed, if $\hat{\mathbf{R}}_n$ denotes an estimate of \mathbf{R}_n , we see the error present in $\hat{\mathbf{R}}_{n-1}$ is propagated in $\hat{\mathbf{R}}_n$ through (1).

B. Low-Cost Inertial Measurement Unit (IMU) Sensor Model

The IMU provides noisy and biased measurements of angular rate $\boldsymbol{\omega}_n$ and specific acceleration \mathbf{a}_n at high frequency (200 Hz in our experiments) as, see [3,4],

$$\mathbf{u}_n^{\text{IMU}} = \begin{bmatrix} \boldsymbol{\omega}_n^{\text{IMU}} \\ \mathbf{a}_n^{\text{IMU}} \end{bmatrix} = \mathbf{C} \begin{bmatrix} \boldsymbol{\omega}_n \\ \mathbf{a}_n \end{bmatrix} + \mathbf{b}_n + \boldsymbol{\eta}_n, \quad (2)$$

where $\mathbf{b}_n \in \mathbb{R}^6$ are quasi-constant biases, $\boldsymbol{\eta}_n \in \mathbb{R}^6$ are commonly assumed zero-mean white Gaussian noises, and

$$\mathbf{a}_n = \mathbf{R}_{n-1}^T ((\mathbf{v}_n - \mathbf{v}_{n-1})/dt - \mathbf{g}) \in \mathbb{R}^3 \quad (3)$$

is the acceleration in the IMU frame without the effects of gravity \mathbf{g} , with $\mathbf{v}_n \in \mathbb{R}^3$ the IMU velocity expressed in the global frame. The intrinsic *calibration* matrix

$$\mathbf{C} = \begin{bmatrix} \mathbf{S}_\omega \mathbf{M}_\omega & \mathbf{A} \\ \mathbf{0}_{3 \times 3} & \mathbf{S}_a \mathbf{M}_a \end{bmatrix} \approx \mathbf{I}_6 \quad (4)$$

contains the information for correcting signals: axis misalignments (matrices $\mathbf{M}_\omega \approx \mathbf{I}_3$ and $\mathbf{M}_a \approx \mathbf{I}_3$); scale factors (diagonal matrices $\mathbf{S}_\omega \approx \mathbf{I}_3$ and $\mathbf{S}_a \approx \mathbf{I}_3$); and effect of linear accelerations on gyro measurements, a.k.a. g-sensitivity (matrix $\mathbf{A} \approx \mathbf{0}_{3 \times 3}$). Remaining intrinsic parameters, e.g. level-arm between gyro and accelerometer, are found in [3,4].

We now make three remarks regarding (1)-(4):

- 1) equations (2)-(4) represent a model that *approximates* reality. Indeed, calibration parameters \mathbf{C} and biases \mathbf{b}_n should both depend on time as they vary with temperature and stress [1,4], but they are difficult to estimate in real-time. Then, vibrations and platform excitations due to, e.g., rotors, make Gaussian noise $\boldsymbol{\eta}_n$ colored in practice [24], albeit commonly assumed white;

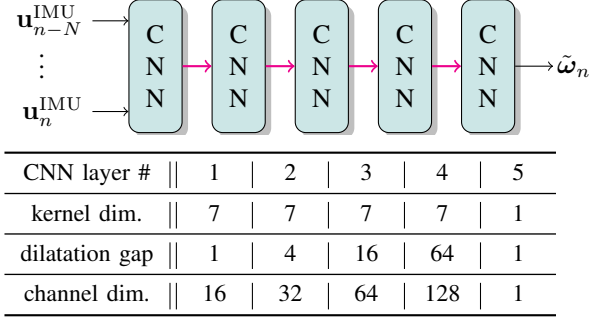


Fig. 2. Proposed neural network structure which computes gyro correction $\tilde{\omega}_n$ in (5) from the N past IMU measurements. We set the Convolutional Neural Networks (CNNs) as indicated in the table, and define between two convolutional layers a batchnorm layer [25] and a smooth GELU activation function [26] (magenta arrows). The adopted structure defines the window size $N = \max(\text{kernel dim.} \times \text{dilatation gap}) = 7 \times 64 = 448$, corresponding to 2.24 s of past information, and has been found after trial-and-error on datasets [22,23].

- 2) substituting actual measurements $\boldsymbol{\omega}_n^{\text{IMU}}$ in place of true value $\boldsymbol{\omega}_n$ in (1) leads generally to quick drift (in a few seconds) and poor orientation estimates;
- 3) in terms of method evaluation, one should always compare the learning method with respect to results obtained with a properly calibrated IMU as a proper estimation of the parameters \mathbf{C} and \mathbf{b}_n in (2) actually leads to surprisingly precise results, see Section IV.

III. LEARNING METHOD FOR DENOISING THE IMU

We describe our approach for regression of noise-free gyro increments $\boldsymbol{\omega}_n$ in (2) in order to obtain accurate orientation estimates by integration (1). Our goal thus boils down to estimating \mathbf{b}_n , $\boldsymbol{\eta}_n$, and correcting the unknown \mathbf{C} , using IMU signals $\mathbf{u}_n^{\text{IMU}}$ and some ground truth information.

A. Proposed Gyro Correction Model

Leveraging the analysis of Section II, we compute the noise-free increments as

$$\hat{\omega}_n = \hat{\mathbf{C}}_\omega \boldsymbol{\omega}_n^{\text{IMU}} + \tilde{\omega}_n, \quad (5)$$

with $\hat{\mathbf{C}}_\omega = \hat{\mathbf{S}}_\omega \hat{\mathbf{M}}_\omega \in \mathbb{R}^{3 \times 3}$ the intrinsic parameters that account for gyro axis-misalignment and scale factors, and where the gyro bias is included in the gyro correction $\tilde{\omega}_n = \hat{\mathbf{c}}_n + \hat{\mathbf{b}}$, where $\hat{\mathbf{c}}_n$ is time-varying and $\hat{\mathbf{b}}$ is the static bias. Explicitly considering the small accelerometer influence \mathbf{A} , see (2)-(4), does not affect the results, so it is ignored.

We search to compute $\tilde{\omega}_n$ and $\hat{\mathbf{C}}_\omega$. The neural network described in Section III-B below computes $\tilde{\omega}_n$ using information present in a past local moving window of size N around $\boldsymbol{\omega}_n^{\text{IMU}}$. By contrast, we let $\hat{\mathbf{C}}_\omega$ be static parameters initialized at \mathbf{I}_3 , and optimized during training since each considered dataset uses *one* IMU. The learning problem involving a time varying $\hat{\mathbf{C}}_\omega$ and/or multiple IMUs is left for future work.

The consequences of opting for the model (5) and the proposed network structure are as follows. First, the corrected gyro is initialized on the original gyro $\hat{\omega}_n \approx \boldsymbol{\omega}_n^{\text{IMU}}$, with $\hat{\mathbf{C}}_\omega = \mathbf{I}_3$ and $\tilde{\omega}_n \approx \mathbf{0}_3$ before training. This way, the method improves the estimates from the first training epoch. Then,




Fig. 3. Time efficient computation of the loss (9) by viewing the orientation integration (8) as a tree of matrix multiplications. Computation for length j requires $\log_2(j)$ “batch” operations as denoted by the blue vertical arrow on the left. We see we need 3 batches of parallel operations for $j = 8$ on the chart above. In the same way, we only need 5 operations for $j = 32$.

our method is intrinsically robust to overfitting as measurements outside the local window, i.e., whose timestamps are less than $n - N$ or greater than n , see Figure 2, do not contribute to inferring $\hat{\omega}_n$. This allows training the method with 8 or less minutes of data, see Section IV-A.

B. Dilated Convolutional Neural Network Structure

We define here the neural network structure which infers the gyro correction as

$$\hat{\omega}_n = f(\mathbf{u}_{n-N}^{\text{IMU}}, \dots, \mathbf{u}_n^{\text{IMU}}), \quad (6)$$

where $f(\cdot)$ is the function defined by the neural network. The network should extract information at temporal multi-scales and compute smooth corrections. Note that, the input of the network consists of IMU data, that is, gyros, but also accelerometers’ signals. Indeed, from (3), if the velocity varies slowly between successive increments we have

$$\begin{aligned} \mathbf{a}_{n+1} - \mathbf{a}_n &\approx -(\mathbf{R}_n - \mathbf{R}_{n-1})^T \mathbf{g} \\ &\approx -(\exp(-\boldsymbol{\omega}_n dt) - \mathbf{I}_3) \mathbf{R}_{n-1}^T \mathbf{g}, \end{aligned} \quad (7)$$

which also provides information about angular velocity. This illustrates intuitively how \mathbf{a}_n can be used, although the neural network does not assume small velocity variation and works actually well in the presence of large accelerations.

We leverage dilated convolutions that infer a correction based on a local window of $N = 448$ previous measurements, which represents 2.24 s of information before timestamp n in our experiments. Dilated convolutions are based on convolutions applied to an input with defined dilatation gap, see [27], which: *i*) support exponential expansion of the receptive field, i.e., N , without loss of resolution or coverage; *ii*) are computationally efficient with few memory consumption; and *iii*) maintain the temporal ordering of data. We thus expect the network to detect and correct various features such as rotor vibrations that are not modeled in (2). Our configuration displayed in Figure 2 requires learning 77 052 parameters, which is *extremely* low and contrasts with recent (visual-) inertial learning methods, see e.g. [19] Figure 2, where IMU processing only requires more than 2 600 000 parameters.

C. Loss Function based on Integrated Gyro Increments

Defining a loss function directly based on errors $\omega_n - \hat{\omega}_n$ requires having a ground truth ω_n at IMU frequency




Fig. 4. Horizontal ground truth trajectory for the sequence `room 1` of [23]. Ground truth is periodically absent due to occlusions of the hand-held device from the motion capture system, see the color lines on the right of the figure.

(200 Hz), which is not feasible in practice as the best tracking systems are accurate at 20-120 Hz. In place, we suggest defining a loss based on the following integrated increments

$$\delta \mathbf{R}_{i,i+j} = \mathbf{R}_i^T \mathbf{R}_{i+j} = \prod_{k=i}^{i+j-1} \exp(\boldsymbol{\omega}_k), \quad (8)$$

i.e., where the IMU frequency is reduced by a factor j . We then compute the loss for a given j as

$$\mathcal{L}_j = \sum_i \rho \left(\log \left(\delta \mathbf{R}_{i,i+j} \delta \hat{\mathbf{R}}_{i,i+j}^T \right) \right), \quad (9)$$

where $\log(\cdot)$ is the $SO(3)$ logarithm map, and $\rho(\cdot)$ is the Huber loss function. We set in our experiments the Huber loss parameter to 0.005, and define our loss function as

$$\mathcal{L} = \mathcal{L}_{16} + \mathcal{L}_{32}. \quad (10)$$

The motivations for (9)-(10) are as follows:

- the choice of Huber loss $\rho(\cdot)$ yields robustness to ground truth outliers;
- (8) is invariant to rotations, which suits well IMUs, whose gyro and accelerometer measurements are respectively invariant to rotations and yaw changes [2,28], i.e., left shifting $\mathbf{R}_n \leftarrow \delta \mathbf{R} \mathbf{R}_n$ and $\hat{\mathbf{R}}_n \leftarrow \delta \mathbf{R} \hat{\mathbf{R}}_n$ with $\delta \mathbf{R}_n \in SO(3)$ leaves (9) unchanged;
- the choice of (10) corresponds to error increments at $200/16 \approx 12$ Hz and $200/32 \approx 6$ Hz, which is barely slower than ground truth. Setting too high a j , or in the extreme case using a loss based on the overall orientation error $\mathbf{R}_n^T \hat{\mathbf{R}}_n$, would make the algorithm prone to overfitting, and would make the method too sensitive to specific trajectory patterns of training data.

D. Efficient Computation of (8)-(10)

First, note that thanks to parallelization applying e.g. $\exp(\cdot)$, to one or parallelly to many $\hat{\omega}_n$ takes similar execution time on a GPU (we found experimentally that one operation takes 5 ms whereas 10 million operations in parallel take 70 ms, that is, the time per operation drops to 7 ns). We call an operation that is parallelly applied to many instances a

dataset	sequence	VINS-Mono [7]	VINS-Mono (loop-closure)	Open-VINS [6]	Open-VINS (proposed)	zero motion	raw IMU	OriNet [17]	calibrated IMU (proposed)	proposed IMU
EuRoC [22]	MH 02 easy	1.34/1.32	0.57/0.50	1.11/1.05	1.21/1.12	44.4/43.7	146/130	5.75/0.51	7.09/1.49	1.39/0.85
	MH 04 difficult	1.44/1.40	1.06/1.00	1.60/1.16	1.40/0.89	42.3/41.9	130/77.9	8.85/7.27	5.64/2.53	1.40/ 0.25
	V1 01 easy	0.97/0.90	0.57/0.44	0.80/0.67	0.80/0.67	114/76	71.3/71.2	6.36/2.09	6.65/3.95	1.13/0.49
	V1 03 difficult	4.72/4.68	4.06/4.00	2.32/2.27	2.25/2.20	81.4/80.5	119/84.9	14.7/11.5	3.56/2.04	2.70/ 0.96
	V2 02 medium	2.58/2.41	1.83/1.61	1.85/1.61	1.81/1.57	93.9/93.5	117/86	11.7/6.03	4.63/2.30	3.85/2.25
	average	2.21/2.14	1.62/1.52	1.55/1.37	1.50/1.30	66.1/66.1	125/89.0	9.46/5.48	5.51/2.46	2.10/ 0.96
TUM-VI [23]	room 2	0.60/0.45	0.69/0.50	2.47/2.36	1.95/1.84	91.8/90.4	118/88.1	—/—	10.6/10.5	1.31/1.18
	room 4	0.76/0.63	0.66/0.51	0.97/0.88	0.93/0.83	107/103	74.1/48.2	—/—	2.43/2.30	1.48/0.85
	room 6	0.58/0.38	0.54/0.33	0.63/0.51	0.60/0.51	138/131	94.0/76.1	—/—	4.39/4.31	1.04/0.57
	average	0.66/0.49	0.63/0.45	1.33/1.25	1.12/1.05	112/108	95.7/70.8	—/—	5.82/5.72	1.28/0.82

Table 1. Absolute Orientation Error (AOE) in terms of **3D orientation/yaw**, in degree, on the *test* sequences. We see our approach competes with VIO (while entirely based on IMU signals) and outperforms other inertial methods. Results from OriNet are unavailable on the TUM-VI dataset.

batch operation. That said, an apparent drawback of (8) is to require j matrix multiplications, i.e. j operations. However, first, we may compute ground truth $\delta \mathbf{R}_{i,i+j}$ only once, store it, and then we only need to compute $\delta \hat{\mathbf{R}}_{i,i+j}$ multiple times. Second, by viewing (8) as a tree of matrix multiplications, see Figure 3, we reduce the computation to $\log_2(j)$ batch GPU operations only. We finally apply sub-sampling and take one i every j timestamps to avoid counting multiple times the same increment. Training speed is thus increased by a factor $32/\log_2(32) \approx 6$.

E. Training with Data Augmentation

Data augmentation is a way to significantly increase the diversity of data available for training without actually collecting new data, to avoid overfitting. This may be done for the IMU model of Section II by adding Gaussian noise η_n , adding static bias \mathbf{b}_n , uncalibrating the IMU, and shifting the orientation of the IMU in the accelerometer measurement. The two first points were noted in [17], whereas the two latter are to the best of our knowledge novel.

Although each point may increase the diversity of data, we found they do not necessarily improve the results. We opted for addition of a Gaussian noise (only), during each training epoch, whose standard deviation is the half the standard deviation that the dataset provides (0.01 deg/s).

IV. EXPERIMENTS

We evaluate the method in terms of 3D orientation, and yaw estimates, that are pivotal for long-run odometry [2,29].

A. Dataset Descriptions

We divide data into training, validation, and test sets, see Chapter I.5.3 of [30]. We optimize the neural network and calibration parameters on the training set. The validation set provides a *biased* evaluation, as it is used for training: data are seen, although never used for “learning”. The test set is then the gold standard to provide an *unbiased* evaluation. It is only used once training (using the training and validation sets) is finished. The datasets we use are as follows.

1) *EuRoC*: the dataset [22] contains image and inertial data at 200 Hz from a *micro aerial vehicle* divided into 11 flight trajectories of 2-3 minutes in two environments. The ADIS16448 IMU is *uncalibrated* and we note ground truth from laser tracker and motion capture system is accurately time-synchronized with the IMU, although dynamic motions deteriorate the measurement accuracy. As already noticed in [6], ground truth for the sequence V1 01 *easy* needs to be recomputed. The sequences MH and V1 were acquired in two days, and the sequences V2 more than three months later. Thus the network should adapt to small varying bias.

We define the train set as the first 50 s of the six sequences MH{01, 03, 05}, V1{02}, V2{01, 03}, the validation set as the remaining parts of these sequences, and the test set consists of the five remaining sequences. We show in Section IV-E that using only 8 minutes of accurate data for training - the beginning and end of each trajectory are the most accurately measured - is sufficient to obtain relevant results.

2) *TUM-VI*: the recent dataset [23] consists of visual-inertial sequences in different scenes from an *hand-held* device. The cheap BMI160 IMU logs data at 200 Hz and was properly *calibrated*. Ground truth is accurately time-synchronized with the IMU, although each sequence contains periodic instants of duration 0.2 s where ground truth is *unavailable* as the acquisition platform was hidden from the motion capture system, see Figure 4. We take the 6 room sequences, which are the sequences having longest ground truth (2-3 minutes each).

We define the training set as the first 50 s of the sequences room 1, room 3, and room 5, the validation set as the remaining ending parts of these sequences, and we set the test set as the 3 other room sequences. This corresponds to 45 000 training data points (90 000 for EuRoC) which is in the same order as the number of optimized parameters, 77 052, and requires regularization techniques such as weight decay and dropout during training.

B. Method Implementation & Training

Our open-source method is implemented based on PyTorch 1.5, where we configure the training hyperparameters as follows. We set weight decay with parameter 0.1, and dropout

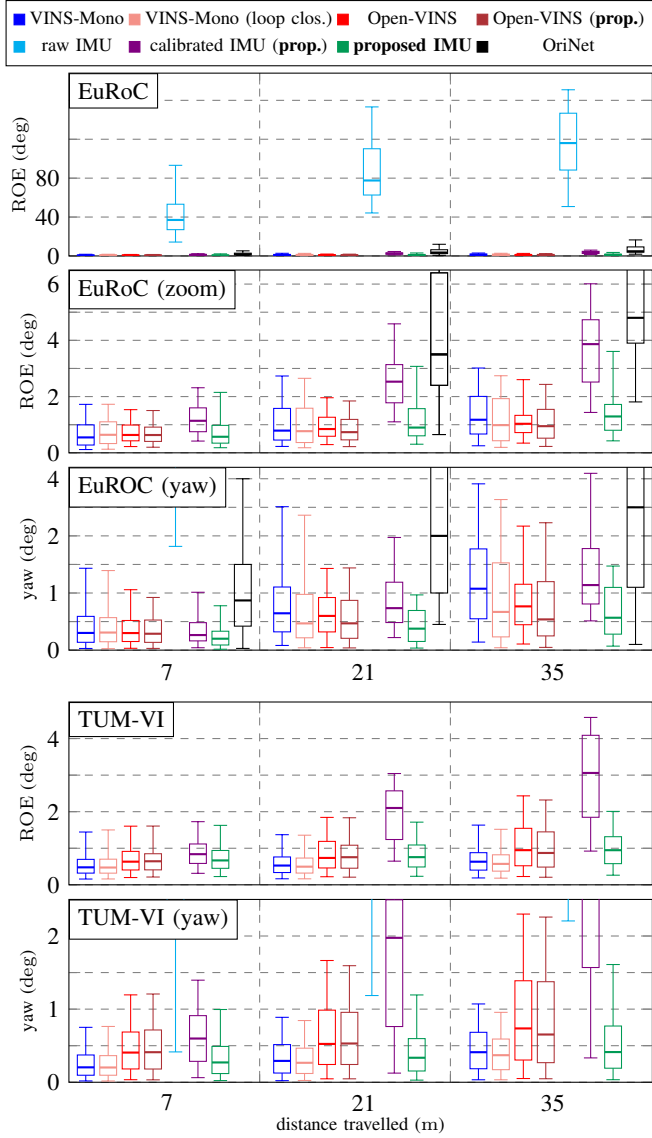


Fig. 5. Relative Orientation Error (ROE) in terms of 3D orientation and yaw errors on the test sequences. Our method outperforms calibrated IMU and competes with VIO methods albeit based only on IMU signals. Raw IMU is way off and results from OriNet on TUM-VI are unavailable.

with 0.1, which is the probability of an element to be set equal to zero. Both techniques reduce overfitting. We choose the ADAM optimizer [31] with cosines warning restart scheduler [32] where learning rate is initialized at 0.01. We train for 1800 epochs, which is very fast as it takes less than 5 minutes for each dataset with a GTX 1080 GPU device.

C. Compared Methods

We compare methods based on IMU, and/or camera.

1) *Methods Based on the IMU Only*: we compare the following methods

- **raw IMU**, that is, uncalibrated IMU. It refers also to the proposed method once initialized but not trained;
- **OriNet** [17], which is based on *recurrent neural networks*, and whose training and test sets correspond to ours;

- **calibrated IMU**, that is, our method where the 12 parameters $\hat{\mathbf{C}}_\omega$ and $\hat{\omega}_n$ are *constant*, nonzero, and optimized. This may be achieved by enforcing null input, i.e., setting the IMU gyros and accelerometers in (6) as $\mathbf{u}_n = \mathbf{0}_6$, both during training and evaluation;
- **proposed IMU**, which is our learning based method described in Section III.

2) *Methods Based on Camera and the IMU*: we run each of the following method in the same setting ten times, and then average the results, on a Dell Precision Tower 7910 workstation desktop, to avoid deterioration due to computational limitations [29]. We compare:

- **VINS-Mono** [7], a monocular VIO framework with notable performances on the EuRoC benchmark [29];
- **VINS-Mono (loop closure)**, which is the original VINS-Mono [7] reinforced with loop-closure ability;
- **Open-VINS** [6], a versatile filter-based visual-inertial estimator for which we choose the stereo configuration, and that is top-ranked on the drone dataset of [33];
- **Open-VINS (proposed)**, which is Open-VINS of [6] but where gyro inputs are the proposed corrected gyro measurements (5) output by our method (trained on sequences that are of course different from those used for evaluation).

D. Evaluation Metrics

We evaluate the above methods using the following metrics that we compute with the toolbox of [6].

1) *Absolute Orientation Error (AOE)*: which computes the mean square error between the ground truth and estimates for a given sequence as

$$\text{AOE} = \sqrt{\sum_{n=1}^M \frac{1}{M} \|\log(\mathbf{R}_n^T \hat{\mathbf{R}}_n)\|_2^2}, \quad (11)$$

with M the sequence length, $\log(\cdot)$ the $SO(3)$ logarithm map, and where the estimated trajectory has been aligned on the ground truth at the first instant $n = 0$.

2) *Relative Orientation Error (ROE)*: which is computed as [28]

$$\text{ROE} = \|\log(\delta \mathbf{R}_{n,g(n)}^T \delta \hat{\mathbf{R}}_{n,g(n)})\|_2, \quad (12)$$

for each pair of timestamps $(n, g(n))$ representing an IMU displacement of 7, 21 or 35 meters. Collecting the error (12) for all the pairs of sub-trajectories generates a collection of errors where informative statistics such as the median and percentiles are calculated. As [6,28,29], we strongly recommend ROE for comparing odometry estimation methods since AOE is highly sensitive to the time when the estimation error occurs. We finally use slight variants of (11)-(12) when considering *yaw* (only) errors, and note errors of visual methods generally scale as travelled distance whereas errors of inertial only methods scale as time. We provide in the present paper errors w.r.t. distance travelled to favor comparison with benchmarks such as [29], and same conclusions hold when computing ROE for different times.




Fig. 6. Orientation estimates on the test sequence MH 04 difficult of [22] (left), and room 4 of [23] (right). Our method removes errors of the calibrated IMU and competes with VIO algorithms.

E. Results

Results are given in terms of AOE and ROE respectively in Table 1 and Figure 5. Figure 6 illustrates roll, pitch and yaw estimates for a test sequence of each dataset, and Figure 7 shows orientation errors. We note that:

1) *Uncalibrated IMU is Unreliable*: raw IMU estimates deviate from ground truth in less than 10s, see Figure 6.

2) *Calibrated IMU Outperforms OriNet*: calibrating the IMU via our optimization method leads to surprisingly accurate results, see e.g., Figure 6 (right) where the estimate is difficult to distinguish from ground truth. This evidences cheap sensors can provide very accurate information once they are correctly calibrated.

3) *The Proposed Method Outperforms Inertial Methods*: OriNet [17] is outperformed. Moreover, the proposed method improves accurately calibrated IMU by a factor 2 to 4. It obtains as low as a median error of 1.34 deg/min and 0.68 deg/min on respectively EuRoC and TUM-VI datasets.

4) *The Proposed Method Competes with VIO*: our IMU only method is accurate even on the high motion dynamics present in both datasets, see Figure 6, and competes with VINS-Mono and Open-VINS, although trained with only a few minutes of data.

Finally, as the performance of each method depends on the dataset and the algorithm setting, see Figure 5, it is difficult to conclude which VIO algorithm is the best.

F. Further Results and Comments

We provide a few more comments, supported by further experimental results.

1) *Small Corrections Might Lead to Large Improvement*: the calibrated and corrected gyro signals are visually undistinguishable: differences between them rely in corrections $\tilde{\omega}_n$ of few deg/s, see Figure 8. However, they bring significant improvement in the estimates. This confirms the interest of leveraging neural networks for model correction (2)-(4).

2) *The Proposed Method is Well Suited to Yaw Estimation*: according to Table 1 and Figure 5, we see yaw estimates are particularly accurate. Indeed, VIO methods are able to recover at any time roll and pitch thanks to accelerometers, but the yaw estimates drift over time. In contrast our dead-reckoning method never has access to information allowing it to recover roll and pitch during testing, and nor does it use “future” information such as VINS-Mono with loop-closure ability. We finally note that accurate yaw estimates could be fruitful for yaw-independent VIO methods such as [8].

3) *Correcting Gyro Slightly Improves Open-VINS* [6]: both methods based on Open-VINS perform similarly, which is not surprising as camera alone already provides accurate orientation estimates and the gyros assist stereo cameras.

4) *Our Method Requires few Computational Ressources*: each VIO method performs here at its best while resorting to high computational requirements, and we expect our method - once trained - is very attractive when running onboard with




Fig. 7. Orientation errors on the sequence `room 4` of [23]. Our method removes errors of the calibrated IMU and competes with VIO algorithms.

limited resources. Note for instance that the proposed method performs 3 times better in terms of yaw estimates than a slightly restricted VINS-Mono, see Figure 3 of [29].

V. DISCUSSION

We now provide the community with feedback regarding the method and its implementation. Notably, we emphasize a few points that seem key to a successful implementation when working with a low-cost high frequency IMU.

A. Key Points Regarding the Dataset

One should be careful regarding quality of data, especially when IMU are sampled at high-frequency. This concerns:

- 1) *IMU Signal*: the IMU signal acquisition should be correct with constant sampling time.
- 2) *Ground Truth Pose Accuracy*: we note the EuRoC ground truth accuracy is better at the beginning of the trajectory. Thus, training with only the first 50s of the training sequences proves sufficient, and best.
- 3) *Ground Truth Time-Alignment*: the time alignment between ground truth and IMU is key for success, otherwise the method is prone to learn a time delay.

We admit that our approach requires a proper dataset, which is what constitutes its main *limitation*.

B. Key Points Regarding the Neural Network

Our conclusions about the neural network are as follows.




Fig. 8. Gyro correction $\tilde{\omega}_n$ on the sequence `room 4` of [23]. We see we manage to divide orientation errors by a factor at least 2 w.r.t. calibrated IMU applying corrections whose amplitude is as low as 1 deg/s (max).

1) *Activation Function*: the GELU and other smooth activation functions [26], such as ELU, perform well, whereas ReLU based network is more prone to overfit. We believe ReLU activation function favors sharp corrections which do not make sense when dealing with physical signals.

2) *Neural Network Hyperparameters*: increasing the depth, channel and/or kernel sizes of the network, see Figure 2, does not systematically lead to better results. We tuned these hyperparameters with random search, although more sophisticated methods such as [34] exist.

3) *Normalization Layer*: batchnorm layer improves both training speed and accuracy [25], and is highly recommended.

4) *Accelerometers are Relevant*: we tested our approach without accelerometers, a.k.a. ablation study experiment. Then the neural network obtains 50 % more errors than proposed IMU in terms of ROE. It indicates the utility of accelerometers and confirms the insights of Section III-B.

5) *Recurrent Neural Network obtains Higher Errors*: we also compare our method to an identical approach but where LSTM replaces the dilated convolutions in the network structure. This type of neural network obtains around 40 % more errors than proposed IMU in terms of ROE, while requiring also more time to be trained.

C. Key Points Regarding Training

As in any machine learning application, the neural network architecture is a key component among others [30]. Our comments regarding training are as follows:

- 1) *Optimizer*: the ADAM optimizer [31] performs well.
- 2) *Learning Rate Scheduler*: adopting a learning rate policy with cosinus warming restart [32] leads to substantial improvement and helps to find a correct learning rate.
- 3) *Regularization*: dropout and weight decay hyperparameters are key to avoid overfitting. Each has a range of ideal values which is quickly tuned with basic grid-search.

D. Remaining Key Points

We finally outline two points that may prove useful to the practitioner:

1) *Orientation Implementation*: we did not find any difference between rotation matrix or quaternion loss function implementation once numerical issues are solved, e.g., by enforcing quaternion unit norm. Both implementations result in similar accuracy performance and execution time.

2) *Generalization and Transfer Learning*: it may prove useful to assess to what extent a learning method is generalizable. The extension of the method, trained on one dataset, to another device or to the same device on another platform is considered as challenging, though, and left for future work.

VI. CONCLUSION

This paper proposes a deep-learning method for denoising IMU gyroscopes and obtains remarkable accurate attitude estimates with only a low-cost IMU, and outperforms state-of-the-art [17]. The core of the approach is based on a careful design and feature selection of a dilated convolutional network, and an appropriate loss function leveraged for training on orientation increment at the ground truth frequency. This leads to a method that is robust to overfitting, efficient and fast to train, which serves as offline IMU calibration and may enhance it. As a remarkable byproduct, the method competes with state-of-the-art visual-inertial methods in terms of attitude estimates on a drone and hand-held device datasets, where we simply integrate noise-free gyro measurements.

We believe the present paper offers new perspectives for (visual-)inertial learning methods. Future work will address new challenges in three directions: learning from multiple IMUs (in the current method one IMU serves for training and testing); learning from moderately accurate ground truth that can be output by visual-inertial localization systems; and denoising accelerometers based on relative increments obtained from preintegration theory [5,35].

ACKNOWLEDGEMENTS

The authors wish to thank Jeffrey Delmerico and Mahdi Abolfazli Esfahani for sharing respectively the results of the VIO benchmark [29] and the results of OriNet [17].

REFERENCES

- [1] M. Kok, J. D. Hol, and T. B. Schön, "Using Inertial Sensors for Position and Orientation Estimation," *Foundations and Trends® in Signal Processing*, vol. 11, no. 1-2, pp. 1–153, 2017.
- [2] D. Scaramuzza and Z. Zhang, "Visual-Inertial Odometry of Aerial Robots," *Encyclopedia of Robotics*, 2019.
- [3] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending Kalibr: Calibrating the Extrinsic of Multiple IMUs and of Individual Axes," in *ICRA*. IEEE, 2016, pp. 4304–4311.
- [4] J. Rohac, M. Sipos, and J. Simanek, "Calibration of Low-cost Triaxial Inertial Sensors," *IEEE Instrumentation & Measurement Magazine*, vol. 18, no. 6, pp. 32–38, 2015.
- [5] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-Manifold Preintegration for Real-Time Visual-Inertial Odometry," *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 1–21, 2017.
- [6] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVINS: A Research Platform for Visual-Inertial Estimation," *IROS Workshop on Visual-Inertial Navigation: Challenges and Applications*, 2019.
- [7] T. Qin, P. Li, and S. Shen, "VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [8] J. Svacha, G. Loianno, and V. Kumar, "Inertial Yaw-Independent Velocity and Attitude Estimation for High-Speed Quadrotor Flight," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1109–1116, 2019.
- [9] M. Brossard, A. Barrau, and S. Bonnabel, "RINS-W: Robust Inertial Navigation System on Wheels," in *IROS*. IEEE, 2019.
- [10] —, "AI-IMU Dead-Reckoning," *IEEE Transactions on Intelligent Vehicles*, 2020.
- [11] S. Madgwick, A. Harrison, and R. Vaidyanathan, "Estimation of IMU and MARG Orientation using a Gradient Descent Algorithm," in *International Conference on Rehabilitation Robotics*. IEEE, 2011, pp. 1–7.
- [12] A. Solin, S. Cortes, E. Rahtu, and J. Kannala, "Inertial Odometry on Handheld Smartphones," in *FUSION*, 2018.
- [13] P. Furgale, J. Rehder, and R. Siegwart, "Unified Temporal and Spatial Calibration for Multi-Sensor Systems," in *IROS*. IEEE, 2013, pp. 1280–1286.
- [14] C. Chen, X. Lu, A. Markham, and N. Trigoni, "IONet: Learning to Cure the Curse of Drift in Inertial Odometry," in *Conference on Artificial Intelligence (AAAI)*, 2018.
- [15] R. Clark, S. Wang, H. Wen, A. Markham, and N. Trigoni, "VINet: Visual-Inertial Odometry as a Sequence-to-Sequence Learning Problem," *Conference on Artificial Intelligence (AAAI)*, 2017.
- [16] H. Yan, Q. Shan, and Y. Furukawa, "RIDI: Robust IMU Double Integration," in *ECCV*, 2018.
- [17] M. A. Esfahani, H. Wang, K. Wu, and S. Yuan, "OriNet: Robust 3-D Orientation Estimation With a Single Particular IMU," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 399–406, 2020.
- [18] M. A. Esfahani, H. Wang, K. Wu, and S. Yuan, "AbolDeepIO: A Novel Deep Inertial Odometry Network for Autonomous Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 5, pp. 1941–1950, 2020.
- [19] Y. Almalioglu, M. Turan, A. E. Sari, M. R. U. Saputra, P. P. B. de Gusmão, A. Markham, and N. Trigoni, "SelfVIO: Self-Supervised Deep Monocular Visual-Inertial Odometry and Depth Estimation," *arXiv*, 2019.
- [20] H. Yan, S. Herath, and Y. Furukawa, "RoNIN: Robust Neural Inertial Navigation in the Wild: Benchmark, Evaluations, and New Methods," *arXiv*, 2019.
- [21] F. Nobre and C. Heckman, "Learning to Calibrate: Reinforcement Learning for Guided Calibration of Visual-Inertial Rigs," *The International Journal of Robotics Research*, vol. 38, no. 12-13, pp. 1388–1402, 2019.
- [22] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The EuRoC Micro Aerial Vehicle Datasets," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.
- [23] D. Schubert, T. Goll, N. Demmel, V. Usenko, J. Stuckler, and D. Cremers, "The TUM VI Benchmark for Evaluating Visual-Inertial Odometry," in *IROS*. IEEE, 2018, pp. 1680–1687.
- [24] G. Lu and F. Zhang, "IMU-Based Attitude Estimation in the Presence of Narrow-Band Noise," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 2, pp. 841–852, 2019.
- [25] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in *ICLR*, vol. 37, 2015.
- [26] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for Activation Functions," *ICLR*, 2018.
- [27] F. Yu and V. Koltun, "Multi-Scale Context Aggregation by Dilated Convolutions," *ICLR*, 2016.
- [28] Z. Zhang and D. Scaramuzza, "A Tutorial on Quantitative Trajectory Evaluation for Visual-(Inertial) Odometry," in *International Conference on Intelligent Robots and Systems*. IEEE, 2018, pp. 7244–7251.
- [29] J. Delmerico and D. Scaramuzza, "A Benchmark Comparison of Monocular Visual-Inertial Odometry Algorithms for Flying Robots," in *ICRA*. IEEE, 2018, pp. 2502–2509.
- [30] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. The MIT press, 2016.
- [31] D. P. Kingma and J. Ba, "ADAM: A Method for Stochastic Optimization," in *ICLR*, 2014.
- [32] I. Loshchilov and F. Hutter, "SGDR: Stochastic Gradient Descent with Warm Restarts," *ICLR*, 2016.
- [33] J. Delmerico, T. Cieslewski, H. Rebecq, M. Faessler, and D. Scaramuzza, "Are We Ready for Autonomous Drone Racing? The UZH-FPV Drone Racing Dataset," in *ICRA*. IEEE, 2019, pp. 6713–6719.
- [34] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6765–6816, 2017.
- [35] A. Barrau and S. Bonnabel, "A Mathematical Framework for IMU Error Propagation with Applications to Preintegration," in *ICRA*. IEEE, 2020.