Learning Inertial Odometry for Dynamic Legged Robot State Estimation

Anonymous Author(s)
Affiliation
Address
email

Abstract: This paper introduces a novel proprioceptive state estimator for legged robots based on a learned displacement measurement from IMU data. Recent research in pedestrian tracking has shown that motion can be inferred from inertial data with linear drift using convolutional neural networks. A learned inertial displacement measurement can improve state estimation in challenging scenarios where leg odometry is unreliable, such as slipping and compressible terrains. Our proposed work adapts this idea and learns to estimate a displacement measurement from IMU data which is then fused with traditional leg odometry. Our approach greatly reduces the drift of proprioceptive state estimation, which is critical for legged robots deployed in vision and lidar denied environments such as foggy sewers or dusty mines. We present results of the inertial relative measurement within an EKF and an incremental fixed-lag factor graph estimator. We compared the two approaches using data from several real robot experiments crossing challenging terrains. Our results show a reduction of relative pose error by $37\,\%$ in challenging scenarios when compared to traditional kinematic-inertial estimators. We also demonstrate a $22\,\%$ reduction in error when used with vision systems in visually degraded environments such as an underground mine.

Keywords: Legged Robots, Inertial Navigation, Deep Neural Networks

1 Introduction

Recent advances in legged robot locomotion have motivated the use of quadruped robots in dull and dirty industrial applications, such as inspection and monitoring. To perform these tasks autonomously, accurate state estimation with limited drift over long periods of time is fundamental.

When operating in challenging environments such as mines and sewers, dust particles and vapor can degrade camera and lidar systems up to the point of unusability [1] (Fig. 1). Proprioceptive methods can work in these environments, as they fuse information only from joint kinematics and Inertial Measurement Unit (IMU) data [2, 3, 4]. However, slip events or deformation of the robot's foot or terrain introduce non-Gaussian error that accumulates over time and leads to unbounded drift.

Separately, recent advances in pedestrian tracking using only IMU data with machine learning have shown promising results [5, 6]. The key insight of these works is that it is possible to learn a *motion prior* from inertial data which could be robust to changing biases.

In this paper, we propose a deep neural network capable of learning a motion prior on the locomotion of a legged robot. The network takes a buffer of IMU data and outputs a displacement measurement and its associated uncertainty. Our learning method is similar to [5], however, instead of using only IMU measurements, we also fuse them with kinematic information from joint encoders. We demonstrate the application of this measurement in both a filtering context and as part of a factor-graph optimization problem to track the motion of an ANYbotics ANYmal quadrupedal robot. We present an extensive experimental validation on real robot data in challenging conditions such as slipping, uneven, and soft terrain.

The contributions in this paper can be listed as follows: 1) First kinematics-inertial estimator based on learned IMU displacement measurements; 2) Integration with both a filter-based estimator and

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

Figure 1: **Left:** The ANYmal C100 quadrupedal robot exploring a limestone mine. **Right:** Images from the on-board camera which demonstrate the vision challenges in such an environment. The top shows dust particles while the bottom shows reflections off walls influencing the camera exposure.

a factor graph, merging concepts from machine learning and model based methods; 3) Extensive testing on the ANYmal quadruped platform, including application to field experiments of the robot navigating a mine.

2 Related Work

Here we briefly summarize relevant work on proprioceptive state estimation for legged robots (Section 2.1). We also detail related research in learned inertial state estimation (Section 2.2).

2.1 Kinematic-Inertial State Estimation in Legged Robots

Lin et al. [7] presented the first state estimator for legged robots which fused information from IMU and kinematics in an Extended Kalman Filter (EKF). Their method was highly dependent on the morphology of their hexapedal robot assuming three feet always in contact. Bloesch et al. [2] introduced a new filtering approach called the Two-State Implicit Filter (TSIF). This method employs residual-based modeling of the available information, eliminating the need for an explicit process model. TSIF is the default ANYmal robot state estimator which is used in our experiments.

More recently, Pronto [8] has been presented as an EKF-based state estimator with the capacity to fuse pose corrections from exteroceptive sensing. This is done by maintaining a history of states, measurements and covariances, and applying corrections to the past trajectory asynchronously.

Factor graphs have also been used for legged robot state estimation. Hartley et al. [4] proposed a *contact factor* to link robot states through kinematics. As the robot takes a step, the foot in contact with the ground provides a constraint on the robot's pose. The authors improved their method in [9] with a *hybrid preintegrated contact factor* so that multiple foot steps or measurements can be combined into a single factor, reducing the number of nodes in the graph. Wisth et al. [10] used factor graphs to fuse visual information with inertial and kinematics. In [11], the same authors improved on the factor graph formulation by adding feature tracking-based lidar measurements.

2.2 Learned Inertial Navigation

A recent research direction in inertial navigation systems (INS) is the application of machine learning to infer motion directly from IMU data — typically focused on tracking a handheld mobile phone. These methods have shown impressive performance when compared to traditional inertial-only pedestrian tracking methods, however these methods are susceptible to failure if the trained model is applied outside of the training domain (e.g., a network trained with pedestrian motion would not generalize to a person traveling by bicycle).

Chen et al. [6] introduced IoNet, the first data driven INS which estimates odometry using deep recurrent neural networks. The network was trained to relate buffered data from an IMU into 2D displacements and orientation changes. Their method outperformed both naive integration and existing heuristic pedestrian tracking methods. Yan et al. [12] proposed RoNIN, a network architecture which infers 2D velocity and orientation from raw IMU data; they claim to outperform IoNet. Both approaches assume planar motion.

The first approach estimating 3D motion was proposed by Liu et al. [5], who presented the TLIO filtering framework including learned motion estimates from batches of IMU signals. An EKF was

used to filter the full 6 DoF state and process updates were performed by traditional IMU integration. They trained a network with $40\,\mathrm{h}$ of pedestrian data to output displacement and covariance from the buffered IMU data. These measurements were provided to the EKF as relative position measurements and allowed the filter to correct for biases online. Their approach showed impressive results for IMU only navigation but was unable to integrate other sensors such as kinematics or vision.

Currently, the application of inertial learning to robotics is limited. Brossard et al. [13] demonstrated IMU-only state estimation specifically for wheeled vehicles, heavily exploiting the constrained motion model. Instead of learning a generic displacement measurement, they use a neural network to provide the uncertainty parameters on the motion for the EKF. Whereas, Zhang et al. [14] trained a network to learn the noise parameters of the IMU as opposed to the motion of the subject. Their RNN generates preintegrated rotation, velocity and acceleration terms with bias removed, however this makes integration with other frameworks such as filtering more difficult.

3 Problem Statement

Our objective is to estimate the state of a legged robot equipped with an IMU, joint sensors (encoders, torque sensors) and, optionally, exteroceptive sensors. Following standard conventions, we define a world frame $\mathtt{W}$, fixed to earth and aligned with gravity, and a series of frames attached to the robot, namely: the base frame $\mathtt{B}$, the IMU frame $\mathtt{I}$ and a camera frame $\mathtt{C}$. The transformation between $\mathtt{B}$,$\mathtt{I}$, and $\mathtt{C}$ are known by design. The robot's state at the current time k is defined as follows:

$$\boldsymbol{x}_k = [\mathbf{R}_k, \mathbf{p}_k, \mathbf{v}_k, \mathbf{b}^a_k, \mathbf{b}^g_k] \tag{1}$$

and includes the rotation $\mathbf{R}_{\mathtt{WB}} \in \mathrm{SO}(3)$ in world coordinates, the position ${}_{\mathtt{W}}\mathbf{p}_{\mathtt{WB}} \in \mathbb{R}^3$, the linear velocity ${}_{\mathtt{B}}\mathbf{v}_{\mathtt{WB}}$, and the gyro and accelerometer IMU biases ${}_{\mathtt{I}}\mathbf{b}^g$, ${}_{\mathtt{I}}\mathbf{b}^a$. The IMU measures rotational velocity ${}_{\mathtt{I}}\tilde{\boldsymbol{\omega}}_{\mathtt{WI}} \in \mathbb{R}^3$ and the proper acceleration ${}_{\mathtt{I}}\tilde{\mathbf{a}} \in \mathbb{R}^3$. These are corrupted by zero-mean Gaussian noise $\boldsymbol{\eta}^g, \boldsymbol{\eta}^a$ and by slowly changing biases:

$$\tilde{\boldsymbol{\omega}}(t) = \boldsymbol{\omega}(t) + \mathbf{b}^g(t) + \boldsymbol{\eta}^g(t) \qquad \tilde{\mathbf{a}}(t) = \mathbf{R}^{\mathsf{T}}({}_{\mathtt{W}}\mathbf{a}(t) - {}_{\mathtt{W}}\mathbf{g}) + \mathbf{b}^a(t) + \boldsymbol{\eta}^a(t) \tag{2}$$

where ${}_{\mathtt{W}}\mathbf{g}$ is the gravity vector in world frame. Without constraining measurements, the biases are not observable and simple IMU integration will quickly lead to runaway error.

4 Learned Displacement Measurement

In this section, we describe our learning method including network architecture and loss functions. Let $\mathcal{I}_{ij}$ be a sequence of N consecutive IMU measurements collected from time t_i to t_j and gravity aligned. From $\mathcal{I}_{ij}$, we want to predict the robot's linear position increment $\hat{\mathbf{d}}_i \in \mathbb{R}^3$. To this end, we define a nonlinear function $\mathcal{F}_\theta$ that will be approximated by a neural network:

$$\mathcal{F}_\theta : (\mathcal{I}_{ij}) \mapsto (\hat{\mathbf{d}}_i, \hat{\mathbf{u}}_i) \tag{3}$$

where $\hat{\mathbf{u}}_i = (\hat{u}^x_i, \hat{u}^y_i, \hat{u}^x_i)$ contains the logarithms of standard deviations from which the covariance of the displacement $\hat{\boldsymbol{\Sigma}}^d_i \in \mathbb{R}^{3\times 3}$ can be computed:

$$\hat{\boldsymbol{\Sigma}}^d_i = \mathrm{diag}\left(\exp(2\hat{u}^x_i), \exp(2\hat{u}^y_i), \exp(2\hat{u}^z_i)\right) \tag{4}$$

Following an approach similar to TLIO [5], we approximate $\mathcal{F}_\theta$ with a 1D version of the ResNet18 architecture [15]. We also make use of two loss functions, but in contrast to [5] we use a robust Huber loss kernel to account for outliers in the training set.

4.1 Loss Function

For the first 10 epochs, the loss function is defined by the Mean Square Error (MSE):

$$\mathcal{L}(\mathbf{d}, \hat{\mathbf{d}}) = \frac{1}{n}\sum_{k=1}^{n} \left\| \mathbf{d}_k - \hat{\mathbf{d}}_k \right\|^2, \tag{5}$$

where n is the number of training samples, $\hat{\mathbf{d}}$ is the network output, and $\mathbf{d}$ is the ground truth displacement from motion capture. Once the network has begun to converge, the loss function is

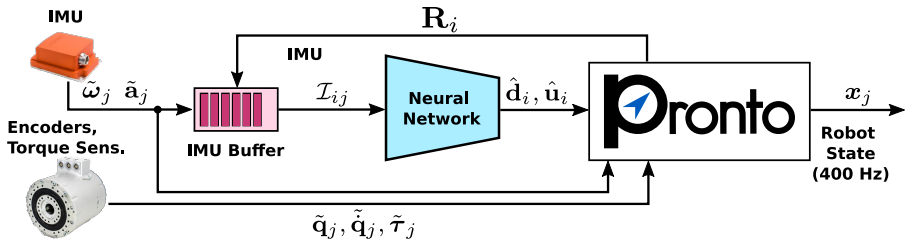

Figure 2: Block diagram of the full EKF-IKD pipeline. IMU and kinematic data are processed by Pronto which provides an estimate of the robot's orientation. This allows a buffer of IMU data to be aligned with gravity and provided as input to the neural network. Displacement measurements and their covariance can then be fused back into Pronto to improve the state estimate.

changed to a robust modification of the Gaussian Maximum Likelihood (GML) which allows the network to learn uncertainty as in [16] without being over influenced by outliers in the training data. The negative log-loss of the GML is given by:

$$\begin{aligned}\mathcal{L}(\mathbf{d}, \hat{\boldsymbol{\Sigma}}^d, \hat{\mathbf{d}}) &= \frac{1}{n}\sum_{k=1}^{n} -\log\left(\frac{1}{\sqrt{8\pi\det(\hat{\boldsymbol{\Sigma}}_k^d)}}\exp\left(\mathrm{H}(\mathbf{d}_k - \hat{\mathbf{d}}_k, \hat{\boldsymbol{\Sigma}}_k^d)\right)\right) \\ &= \frac{1}{n}\sum_{k=1}^{n}\left(\frac{1}{2}\log\det(\hat{\boldsymbol{\Sigma}}_k^d) + \mathrm{H}(\mathbf{d}_k - \hat{\mathbf{d}}_k, \hat{\boldsymbol{\Sigma}}_k^d) + \text{const.}\right) \end{aligned} \tag{6}$$

where $\mathrm{H}(\boldsymbol{\epsilon}_k, \hat{\boldsymbol{\Sigma}}_k^d)$ is the sum of the Huber loss $\mathrm{h}(\cdot,\cdot)$ applied to each axis of $\boldsymbol{\epsilon}_k = \mathbf{d}_k - \hat{\mathbf{d}}_k$:

$$\mathrm{H}(\boldsymbol{\epsilon}_k, \boldsymbol{\Sigma}_k^d) = \mathrm{h}(\epsilon_k^x, \sigma_k^x) + \mathrm{h}(\epsilon_k^y, \sigma_k^y) + \mathrm{h}(\epsilon_k^z, \sigma_k^z) \qquad \mathrm{h}(\epsilon, \sigma) = \begin{cases} \frac{1}{2}\epsilon^2(\sigma^2)^{-1} & \text{if } |\epsilon| < \delta, \\ \delta|\boldsymbol{\epsilon}| - \frac{1}{2}\delta^2 & \text{otherwise} \end{cases} \tag{7}$$

5 State Estimation

The network outputs a relative position measurement and its uncertainty, which is suitable for incorporation into both filtering and smoothing frameworks, as described in the following sections.

5.1 Filtering Estimator

We incorporate the learned displacement measurement into Pronto, an open-source modular legged robot estimation framework[8]. As Pronto maintains a history of states, measurements and covariances, it can easily incorporate relative measurements. The full system architecture is shown in Fig. 2: the IMU measurements from time t_i to t_j are accumulated and rotated to be gravity aligned at time t_i using the filter state $\mathbf{R}_i$. Then, they are passed to the network to estimate displacement and uncertainty. These are then fused into the filter as a relative position measurement:

$$\mathbf{z}_j = \mathbf{p}_i + \hat{\mathbf{d}}_i \tag{8}$$

where the robot's position $\mathbf{p}_i$ at time i is taken from the filter's history, while the displacement is from the network. The covariance associated to the measurement is the same as in Eq. 4.

We call our proposed filtering method EKF-IKD (EKF with inertial, kinematic, and learned displacement). For all training and experiments we used a buffer size of $N = 400$ (1 s of 400 Hz IMU data). Displacement measurements were estimated at 20 Hz, which are overlapping[1].

5.2 Factor Graph Estimator

In this section, we describe a novel factor graph architecture that incorporates the learned inertial factor. In contrast to filters, factor graph methods re-optimize the full or partial history of states, given all the measurements connected to them, allowing for more accurate and smoother estimates.

[1] We recognize that this overlap compromises independence of measurements but have not included analysis of this effect in this paper due to space limits.

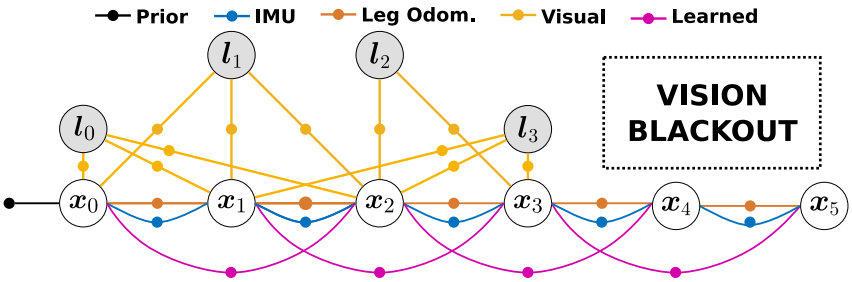

Figure 3: Proposed factor graph framework with IMU factor from [18], kinematics factor from [17] and our proposed learned factor which specially helps in visual blackouts.

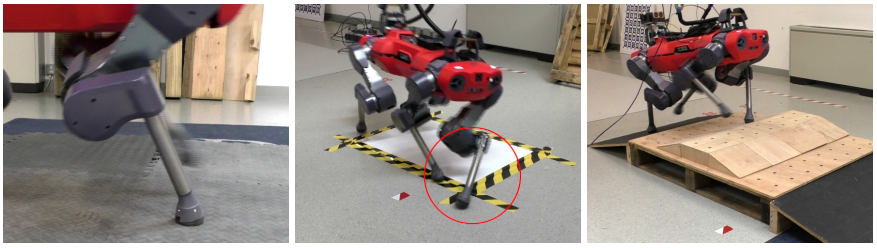

Figure 4: Lab experiments. **Left:** Robot foot compressing into soft terrain. **Middle:** Slip event while walking on slippery terrain. **Right:** Robot climbing over an obstacle.

Fig. 3 shows the factor graph architecture proposed in this work. When a new camera keyframe is available, a new node (white circle) is created. Similarly to [17], we connect consecutive nodes of the graph $\boldsymbol{x}_k$ and $\boldsymbol{x}_{k+1}$ with preintegrated IMU and leg odometry factors (blue and orange factors), while the visual features are modeled as individual landmarks (gray circles) and factors (yellow factors). In visually degraded environments, the visual landmarks might not be observable for a long period of time (e.g., from node $\boldsymbol{x}_3$ in the example), so it is important to have robust proprioceptive odometry to fall back to. To this end, we introduce the novel learned displacement factor to further constrain the robot motion (magenta line). We call the factor graph implementation of our method FG-VIKD (factor graph with vision, IMU, kinematics and learned displacement).

The learned displacement measurement $\mathbf{d}_i$ estimates the position of the robot between the start and end times of the buffer. The learned factor residual $\mathbf{r_d}$ can then be computed using $\hat{\boldsymbol{\Sigma}}_i^d$ from Eq. 4:

$$\mathbf{r_d} = \left\| \mathbf{R}_i^\mathsf{T}(\mathbf{p}_j - \mathbf{p}_i) - \hat{\mathbf{d}}_i \right\|^2_{\hat{\boldsymbol{\Sigma}}_i^d}, \tag{9}$$

Note that while the network output is at the same frequency of the keyframes, the history of measurements stored into the IMU buffer is longer than the period between two keyframes, so the nodes connected by the factor are not consecutive as for the IMU and leg odometry factor. In our case, the IMU buffer is $1\,\mathrm{s}$ long, while the keyframe period is $0.1\,\mathrm{s}$.

6 Lab Experiments

We performed several experiments with our ANYbotics ANYmal C100 robot as shown in Fig. 4. The robot was commanded to walk on different terrains in the lab: flat ground, soft ground, a slippery surface and an obstacle. Ground truth poses were measured using a Vicon motion capture system.

Flat Ground As a baseline example we show performance on rigid, flat terrain. In these conditions, kinematic state estimate is very accurate, however over long trajectories there is upward Z drift.

Soft Ground Deformable terrain poses a major challenge to legged robot state estimation as the specific moment of contact is ill-defined — the foot continues to move and compress the terrain after a contact is detected. This additional downward motion results in a hallucinated upward Z drift. Some realistic examples of deformable terrain include sand, mud or tall grass. To reproduce this effect we placed several layers of foam padding on the floor.

	Absolute Pose Error (APE) [m]					5 m Relative Pose Error (RPE) [m]				
Data	**TSIF**	**EKF-IK**	**EKF-IKD**	**FG-IK**	**FG-IKD**	**TSIF**	**EKF-IK**	**EKF-IKD**	**FG-IK**	**FG-IKD**
FLAT	**0.14**	0.42	0.20	0.36	**0.22**	**0.09**	0.19	0.16	0.25	**0.24**
SLIP.	0.21	0.38	**0.20**	0.23	**0.20**	0.13	0.20	**0.12**	0.18	**0.17**
SOFT	0.35	0.53	**0.24**	0.44	**0.39**	**0.15**	0.24	0.16	0.28	**0.23**
OBST.	**0.19**	0.39	0.27	0.23	**0.21**	**0.12**	0.22	0.22	**0.18**	**0.18**

Table 1: Experimental Results for Lab Experiments.

Slipping When a robot slips, such as on wet or loose terrain, certain kinematic assumptions are violated as the foot is in contact with the terrain but does not have zero velocity. This results in large spikes in odometry error. To create slip events in the lab, we placed a common office white board on the floor and covered it with lubricant (hand sanitizer).

Obstacle Terrain Blind climbing over an uneven obstacle course with ramps. This experiment was designed to show that our method can generalize to motions on non-planar surfaces.

6.1 Training

For the lab experiments we collected $2\,\mathrm{h}16\,\mathrm{min}$ of data of the robot walking in each of the four terrains. Four sequences were selected as test experiments and the remaining were split 75:25 into training and validation datasets. We used the Adam optimizer with an initial learning rate of 10^{-4} and a Huber delta of 1.35. Training lasted 200 epochs with the model minimizing validation error used in all lab experiments. Training took 6 hours on a laptop with an Nvidia Quadro P2000 GPU.

To collect training data we used the Vicon system for training reference. During the training step we remove the biases from the raw IMU signal and inject back random bias values. In order to estimate the true IMU biases, we create a factor graph where each node has a prior pose factor from a Vicon measurement. IMU factors connect the nodes allowing the biases to be estimated in batch optimization for all Vicon measurements.

6.2 Filtering Results

We present here results comparing three **filtering-based** proprioceptive estimators: TSIF [2], which is provided with the ANYmal robot and uses IMU and kinematics, Pronto (which we hereafter refer to as EKF-IK) which is open-source and also uses IMU and kinematics, and our proposed method which incorporates IMU, kinematics and the learned displacement measurement EKF-IKD.

We evaluated performance using two metrics: absolute pose error (APE) and relative pose error (RPE) – as defined in [19]. APE measures the amount of drift over an entire experiment. RPE measures the drift every X m traveled and is a more fine grained measure – we used $5\,\mathrm{m}$ for lab experiments and $10\,\mathrm{m}$ for the mine experiment. In all of the experiments, adding the learned measurement improved the state estimation significantly. As shown in Table 1, in comparison with EKF-IK, the APE was reduced by $46\,\%$ on average ($51\,\%$ on the most challenging terrain: slipping and soft) and the RPE was reduced by $22\,\%$ on average and $37\,\%$ on the challenging terrains.

Flat Ground The results for the filtering methods on flat ground are shown on the left of Fig. 5. There is a clear upward Z drift for TSIF and EKF-IK which is significantly reduced using the learned measurement in EKF-IKD .

Soft Ground The results for the filtering methods on soft ground are shown on the right of Fig. 5. The periods where the robot walked on the soft terrain are highlighted in gray. The upward Z drift for TSIF and EKF-IK is more pronounced but this is reduced in EKF-IKD.

Slipping As shown in Fig. 6, periods when the robot is slipping have high error. The addition of the learned measurement in EKF-IKD significantly reduces spikes in error, particularly at $60\,\mathrm{s}$.

Obstacle Terrain The obstacle terrain experiment demonstrates that our method can generalize to locomotion on non-planar terrain. The upper right plot of Fig. 6 shows that EKF-IKD reduces Z drift while preserving the true changes in elevation as the robot climbs over the terrain. In initial experiments, we found that if the training dataset did not include enough examples of the robot

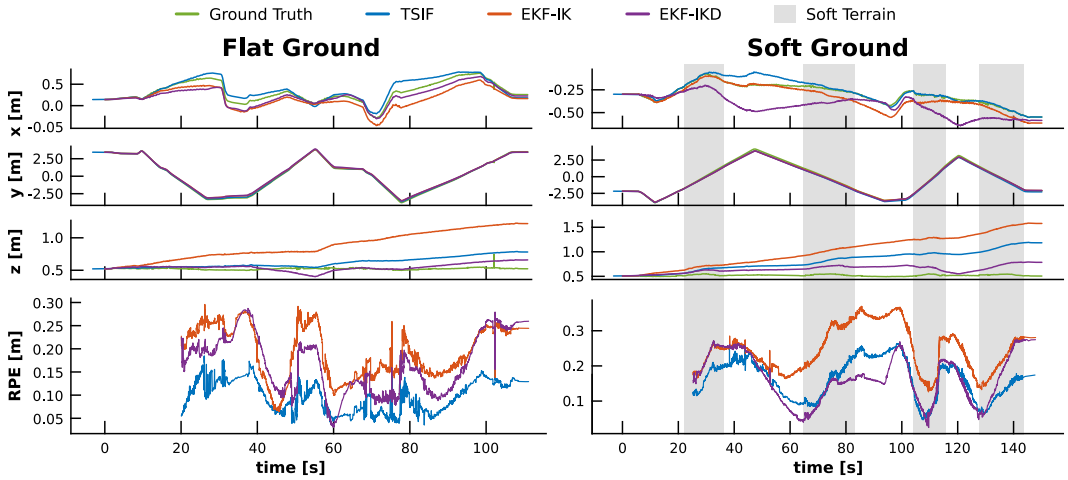

Figure 5: **Top:** Position estimate comparison for flat and soft ground. Note the upward z drift for TSIF and EKF-IK on the soft terrain (gray area). **Bottom:** $5\,\mathrm{m}$ Relative Pose Error.

	Absolute Pose Error (APE) [m]		**10 m Relative Pose Error (RPE)[m]**	
Data	**FG-VIK**	**FG-VIKD**	**FG-VIK**	**FG-VIKD**
MINE	4.63 (2.43)	**4.09** (2.21)	0.45 (0.32)	**0.35** (0.20)

Table 2: Experimental results for the Mine Experiment.

walking over elevated terrain, the network assumed constant planar elevation. The resulting state estimator would then ignore the elevation gain and loss from climbing on terrain.

6.3 Factor Graph Results

We present results comparing two **factor graph-based** proprioceptive estimators: FG-IKD which is our proposed factor graph implementation without vision data (simulating camera failure) and FG-IK which is the same factor graph without vision or the learned displacement. Since there is no vision in these cases nodes are created at a 40 Hz fixed frequency synced with the IMU measurements.

We demonstrate that by adding the learned factor to this factor graph without any vision, APE was reduced by $18\,\%$ and RPE by $7\,\%$. The error in these proprioceptive factor graph methods is higher than the filtering methods because the robot state is intentionally estimated at a lower rate to allow visual data to be incorporated. Additionally, without landmarks, FG-IK cannot take advantage of the graph structure. Nevertheless FG-IKD shows an improvement in proprioceptive estimation implying if a robot's vision system were to fail, our method would reduce drift.

7 Field Experiment: Underground Mine

We tested our method on a dataset collected in a completely dark underground mine as shown in Fig. 1. Whereas previous experiments evaluated improvements to proprioceptive state estimators, this experiment demonstrates application to vision systems in visually degraded environments. We teleoperated our ANYmal robot through a decommissioned limestone mine with large rooms, long corridors and narrow passageways.

A second training dataset totaling $1\,\mathrm{h}23\,\mathrm{min}$ was collected for the mine experiment. In the mine the robot walks with a different gait and controller requiring re-training. Training data for the mine was entirely generated in the lab and the resulting model was applied successfully to the field experiment.

7.1 Mine Experiment Results

We compared the factor graph-based state estimator VILENS [10]; which uses vision, IMU and kinematics (and hereafter we refer to as FG-VIK); to our method FG-VIKD which uses the same

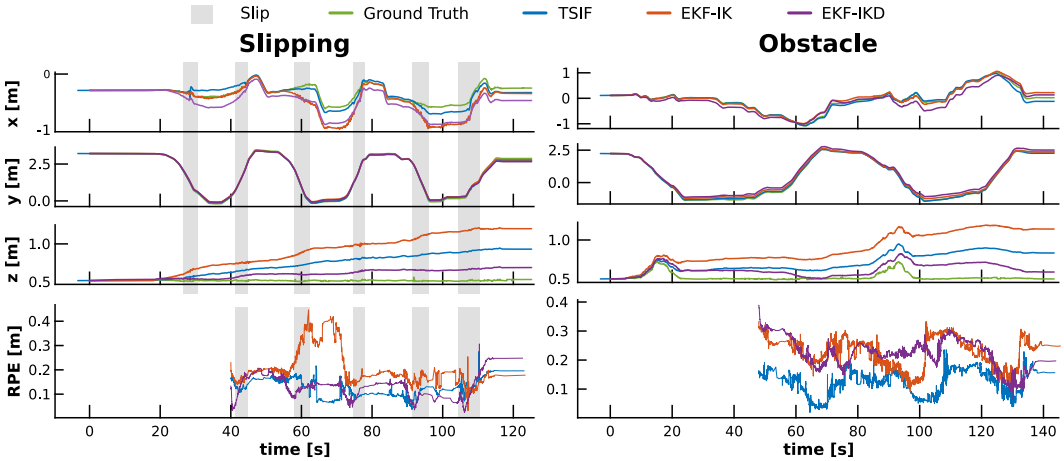

Figure 6: **Top:** Position estimate comparison for slipping and obstacle terrain. Note the upward z drift for TSIF and EKF-IK when slipping (gray areas). **Bottom:** $5\,\mathrm{m}$ Relative Pose Error.

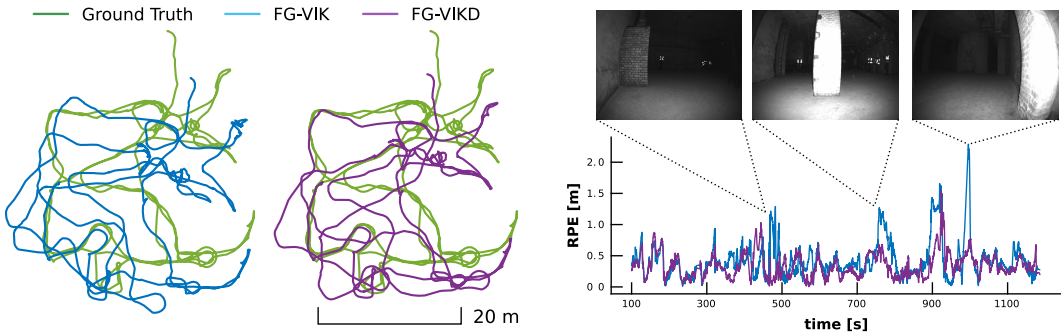

Figure 7: Results from the mine experiment. **Left:** Top-down perspective of FG-VIK vision-based factor graph and our method FG-VIKD. **Right:** $10\,\mathrm{m}$ RPE for both trajectories. We show images from three times where the error spikes in the baseline FG-VIK due to poor visibility.

sensors as well as the learned displacement factor. In the mine, vision sensing is severely affected, forcing FG-VIK to rely on its proprioceptive sensing and, in these situations, the learned factor can help keep error low. Our experiment lasted 20 minutes and over $472\,\mathrm{m}$ of walking.

The ground truth trajectory was estimated using the robot's lidar to localize in a high quality, prior point cloud scan of the area. Table 2 shows that RPE was reduced by $22\,\%$ by adding the learned factor. The trajectory and error over time are shown in Fig. 7 along with images from the front facing camera at times where the drift rate spiked. These spikes coincide with periods of poor feature illumination and visual tracking failure in the dark environment. The robot lights reflected off nearby surfaces triggering the camera to reduce its exposure making the image darker.

8 Conclusion

In this paper we introduced a learned displacement measurement derived from inertial sensing which can greatly reduce state estimation drift for legged robots in challenging situations. We show how this measurement can be used in both filtering and factor graph frameworks. We see the greatest benefit when the robot slips or the terrain is deformable. We show in real experiments with the ANYmal quadruped an reduction of $37\,\%$ in relative pose error for filter-based state estimation. We also show how the measurement can reduce error in a vision driven factor graph framework by $22\,\%$ in a real underground mine. A video demonstrating these experiments is provided as supplementary material. Future work in this area should make the learning more generalizable across different walking gaits as well as more tightly integrate the measurements with the estimator optimization.

References

[1] H. Kolvenbach, D. Wisth, R. Buchanan, G. Valsecchi, R. Grandia, M. Fallon, and M. Hutter. Towards autonomous inspection of concrete deterioration in sewers with legged robots. *Journal of Field Robotics*, 37(8):1314–1327, 2020. doi:https://doi.org/10.1002/rob.21964.

[2] M. Bloesch, M. Burri, H. Sommer, R. Siegwart, and M. Hutter. The two-state implicit filter recursive estimation for mobile robots. *IEEE Robotics and Automation Letters*, 3(1):573–580, 2018. doi:10.1109/LRA.2017.2776340.

[3] M. Camurri, M. Fallon, S. Bazeille, A. Radulescu, V. Barasuol, D. G. Caldwell, and C. Semini. Probabilistic Contact Estimation and Impact Detection for State Estimation of Quadruped Robots. In *IEEE Robotics and Automation Letters*, 2017.

[4] R. Hartley, J. Mangelson, L. Gan, M. G. Jadidi, J. M. Walls, R. M. Eustice, and J. W. Grizzle. Legged robot state-estimation through combined forward kinematic and preintegrated contact factors. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4422–4429, 2018. doi:10.1109/ICRA.2018.8460748.

[5] W. Liu, D. Caruso, E. Ilg, J. Dong, A. I. Mourikis, K. Daniilidis, V. Kumar, and J. Engel. TLIO: Tight learned inertial odometry. *IEEE Robotics and Automation Letters*, 5(4):5653–5660, 2020. doi:10.1109/LRA.2020.3007421.

[6] C. Chen, X. Lu, A. Markham, and N. Trigoni. IoNet: Learning to cure the curse of drift in inertial odometry. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 6468–6476, 2018.

[7] P. Lin, H. Komsuo, and D. E. Koditschek. Sensor Data Fusion for Body State Estimation in a Hexapod Robot With Dynamical Gaits. *IEEE Transactions on Robotics*, 22(5):932–943, 2006.

[8] M. Camurri, M. Ramezani, S. Nobili, and M. Fallon. Pronto: A Multi-Sensor State Estimator for Legged Robots in Real-World Scenarios. *Frontiers in Robotics and AI*, 7, 2020. doi:10.3389/frobt.2020.00068.

[9] R. Hartley, M. G. Jadidi, L. Gan, J. K. Huang, J. W. Grizzle, and R. M. Eustice. Hybrid Contact Preintegration for Visual-Inertial-Contact State Estimation Using Factor Graphs. *IEEE International Conference on Intelligent Robots and Systems*, pages 3783–3790, 2018. doi:10.1109/IROS.2018.8593801.

[10] D. Wisth, M. Camurri, and M. Fallon. Robust Legged Robot State Estimation Using Factor Graph Optimization. *IEEE Robotics and Automation Letters*, 4(4):4507–4514, 2019. doi:10.1109/LRA.2019.2933768.

[11] D. Wisth, M. Camurri, S. Das, and M. Fallon. Unified multi-modal landmark tracking for tightly coupled lidar-visual-inertial odometry. *IEEE Robotics and Automation Letters*, 6(2):1004–1011, Apr 2021. doi:10.1109/lra.2021.3056380.

[12] S. Herath, H. Yan, and Y. Furukawa. Ronin: Robust neural inertial navigation in the wild: Benchmark, evaluations, new methods. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3146–3152, 2020. doi:10.1109/ICRA40945.2020.9196860.

[13] M. Brossard, A. Barrau, and S. Bonnabel. AI-IMU dead-reckoning. *IEEE Transactions on Intelligent Vehicles*, 5(4):585–595, 2020. doi:10.1109/TIV.2020.2980758.

[14] M. Zhang, M. Zhang, Y. Chen, and M. Li. IMU data processing for inertial aided navigation: A recurrent neural network based approach. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi:10.1109/CVPR.2016.90.

[16] R. L. Russell and C. Reale. Multivariate uncertainty in deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

[17] D. Wisth, M. Camurri, and M. Fallon. Preintegrated velocity bias estimation to overcome contact nonlinearities in legged robot odometry. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 392–398, 2020. doi:10.1109/ICRA40945.2020.9197214.

[18] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. IMU preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation. *Robotics: Science and Systems*, 11, 2015. doi:10.15607/RSS.2015.XI.006.

[19] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of RGB-D SLAM systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 573–580, 2012. doi:10.1109/IROS.2012.6385773.