Deep Multi-Modal Contact Estimation for Invariant Observer Design on Quadruped Robots

Anonymous Author(s)
Affiliation
Address
email

Abstract: This work reports on developing a deep learning-based contact estimator for legged robots that bypasses the need for physical contact sensors and takes multi-modal proprioceptive sensory data from joint encoders, kinematics, and an inertial measurement unit as input. Unlike vision-based state estimators, proprioceptive state estimators are agnostic to perceptually degraded situations such as dark or foggy scenes. For legged robots, reliable kinematics and contact data are necessary to develop a proprioceptive state estimator. While some robots are equipped with dedicated contact sensors or springs to detect contact, some robots do not have dedicated contact sensors, and the addition of such sensors is non-trivial without redesigning the hardware. The trained deep network can accurately estimate contacts on different terrains and robot gaits and is deployed along a contact-aided invariant extended Kalman filter to generate odometry trajectories. The filter performs comparably to a state-of-the-art visual SLAM system.

Keywords: State Estimation, Deep Learning, Legged Robot

1 Introduction

Legged robots can traverse uneven terrains. This unique capability gives legged robots the potential to conduct scientific exploration in extreme environments, execute rescue missions in hazardous scenes, and can help humans in everyday tasks [1]. To accomplish these goals, knowledge of the robot's current state is necessary. While joint angles can be directly measured using encoders, the robot's pose and velocity require additional measurements and a mathematically sound data fusion method. Methods for state estimation rely on visual and proprioceptive information. Vision-based methods take advantage of the rich information provided by cameras and LiDARs. However, these methods might suffer from failure due to illumination changes or extreme environments such as snowstorms or foggy scenes [2]. On the other hand, proprioceptive state estimators take only internal measurements and are thus agnostic to perceptually degraded situations. Moreover, unlike vision sensors that typically operate with low frequency (e.g., $10 - 60\,\mathrm{Hz}$), proprioceptive sensors usually operate at a high update rate (e.g., $100-2000\,\mathrm{Hz}$). This high-frequency information provides odometry estimates for localization and mapping tasks [3, 4, 5] and it is also crucial for controllers and planners to maintain stability and execute planned policies [6]. The former application is the focus of this work.

Proprioceptive state estimators often fuse the measurements from an Inertial Measurement Unit (IMU) with leg odometry. Leg odometry, in particular, uses kinematics and contact data to update the state. Reliable measurements of kinematic and contact information become critical in this case. Not all legged robots are equipped with dedicated contact sensors or springs to detect contact [7, 8]. The addition of dedicated contact sensors is non-trivial and often leads to a redesign of the hardware.

In this paper, we develop a deep learning-based contact estimator that does not require dedicated sensors and instead uses joint encoders, kinematics, and IMU data. We create contact data sets using an MIT Mini Cheetah robot [8] on eight different terrains. The trained network can correctly estimate contact events across different terrains and robot gaits. We further deploy the contact estimator along with a contact-aided invariant extended Kalman filter (InEKF) [9] and show that the resulting odometry trajectory is comparable to a state-of-the-art visual SLAM algorithm (used as a proxy for

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

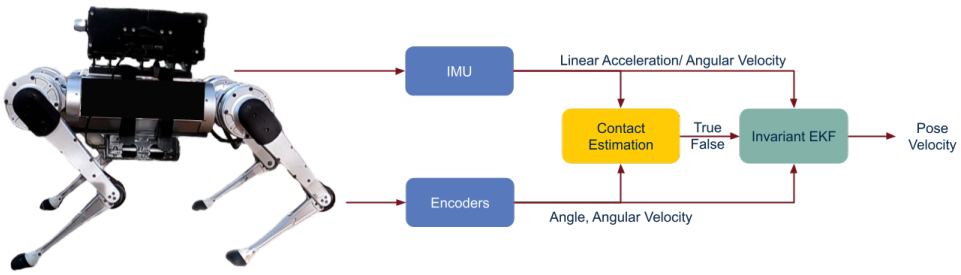

Figure 1: Proposed state estimator. The deep contact estimation network (Yellow) takes joint encoders data, IMU, and kinematics as input and classifies the correct contact state for the quadruped robot. The estimated contacts and the IMU and joint encoders data are fused inside a contact-aided invariant extended Kalman filter for state estimation. We tested the proposed pipeline on a MIT Mini Cheetah robot (Left).

ground truth). The contributions of this work are as follows: 1) Open-source quadruped contact data sets recorded using an MIT Mini Cheetah, 2) A lightweight multi-modal deep learning-based contact estimator that mitigates the need for physical contact sensors and gives reliable estimations on different terrains with different robot gaits, 3) A quadruped version of the contact-aided invariant EKF compatible with Lightweight Communications and Marshalling (LCM) [10, 11] interface. The data sets and software will be available for download after receiving the final decision, and 4) Experimental results for contact-aided state estimation using an MIT Mini Cheetah on eight different terrains.

2 Related Work

A thorough literature review on robot state estimation, including legged robots and contact-aided approaches, is available by Hartley et al. [9]. In this section, we focus on model-based and data-driven methods for contact estimation.

Model-Based Contact Estimation. Model-based approaches utilize physics models of the robot to estimate foot contacts. Traditionally, researchers segment the touchdown event of robot legs or prosthetic legs by thresholding on the estimated Ground Reaction Force (GRF) from the general equation of motion [12, 13, 14]. Although this method can detect touchdown events, the estimated GRF is often noisy and unreliable. De Luca et al. [15], Haddadin et al. [16] proposed a Generalized Momentum (GM) method for detecting contact events on robot manipulators. This GM-based method is, in fact, a filtered version of the work of Focchi et al. [12]. Although GM-based methods mitigate the noise problem in GRF estimation, an empirical threshold on the cut-off frequency is still required.

Hwangbo et al. [17] first introduces a probabilistic representation of the contact state and uses a method similar to a Hidden Markov Model (HMM) to fuse the dynamics and kinematics for contact estimation. They adopt a Monte-Carlo sampling algorithm to compute the transition model and verify their method against GM-based methods. Jenelten et al. [18] expands the HMM method and focuses on slippage detection. They demonstrate ANYmal [19], a quadruped robot, walking stably on slippery ground with their approach. The above two methods aim to detect contact as early as possible for the controller to maintain stability; however, we aim to detect reliable contacts for state estimation on various terrains. Since the mentioned works do not evaluate using a state estimator, it is unclear how their methods assist state estimation.

Bledt et al. [20] leverages both the GM-based methods and the probabilistic representation of contact states. They use a Kalman filter to fuse the gait phase scheduler information from the controller with the GRF estimated using the GM method. They demonstrate that this Kalman filtering approach can correctly estimate contacts and help the controller reduce the bouncing event upon touchdown. However, this method assumes the leg phase to be periodic as it uses the gait scheduler information in the prediction step of the Kalman filter. It could experience a loss in performance when the phase is heavily violated.

Data-Driven Contact Estimation. Data-driven approaches take advantage of the rapid development of recent machine learning techniques. Camurri et al. [21] uses logistic regression to learn the GRF threshold for contact detection. This work compares against heuristic-based thresholding on

GRF using a base state estimator. The result shows that the logistic regression classifier can double the performance of the state estimator. However, compared to deep learning methods, the performance of a logistic regression classifier gets saturated as the number of data increases [22, 23]. This method requires a specific training procedure for different gait, loading conditions of the robot, and individual terrain properties.

Rotella et al. [24] uses a fuzzy C-means clustering for the probability of contacts in all six end-effector degrees of freedom. They integrate the contact estimator with a base state estimator and show their approach performs considerably better than implementations that are purely based on measured normal force. However, this method assumes contact wrench sensors and additional IMU are available at each end-effector. Furthermore, this method was only tested in simulation. Its performance on real robots remains unknown. Piperakis et al. [25] proposes an unsupervised learning method for humanoid gait phase estimation. The authors employ Gaussian Mixture Models (GMMs) for clustering and show the accuracy by comparing to the ground-truth data and leg odometry. However, this work also assumes the availability of wrench/force sensors at each end-effector, and the clustering result is affected by the gait and data density.

The above methods either assume the availability of wrench/force sensors or are restricted by the nature of simple regression and are thus unable to generalize to different scenarios. In contrast, our work proposes a multi-modal deep learning-based contact estimator that does not require contact sensors and can generalize well to different gaits and terrain properties. Moreover, as more data becomes available, the network performance can be improved.

3 Background on Contact-Aided Invariant Extended Kalman Filtering

We adopt the contact-aided invariant Kalman filtering approach [9] that is well-suited for legged robot state estimation problems.

State Representation. We wish to estimate the orientation, velocity, and position of the body (IMU) frame in the world frame at any time stamp t, which are represented as $R_t \in \mathrm{SO}(3)$, $v_t \in \mathbb{R}^3$, and $p_t \in \mathbb{R}^3$, respectively. We define the state $X_t \in \mathrm{SE}_{L+2}(3)$ as $X_t := (R_t, v_t, p_t, d_{lt})$, where $d_{lt} \in \mathbb{R}^3$ is the position of a contact foot in the world frame. When a contact event is detected for leg l, d_{lt} is calculated via the forward kinematics and the current body state, and then augmented into the state to enforce the no-slip (zero-velocity) condition. Similarly, if the contact constraint breaks, the associated contact position d_{lt} will be marginalized (removed) from the state to allow foot movement. The matrix Lie group, $\mathrm{SE}_{L+2}(3)$ is an extension of $\mathrm{SE}(3)$ and it was introduced by Barrau [26].

When a new contact event is detected, the position of the corresponding foot is computed using $\bar{d}_t = \bar{p}_t + \bar{R}_t h_p(\tilde{\alpha}_t)$, where $h_p(\tilde{\alpha}_t)$ is the foot position relative to the body frame computed by the forward kinematics (the bar notation denotes estimated variables). The number L indicates the number of legs that is currently having contact with the ground (here $L \leq 4$).

Process Model. We assume the IMU measurements, i.e., angular velocity and linear acceleration, are corrupted by white Gaussian noise:

$$\tilde{\omega}_t = \omega_t + w_t^g, \quad w_t^g \sim \mathcal{GP}(0_{3,1}, \Sigma^g \delta(t-t')),\ \tilde{a}_t = a_t + w_t^a, \quad w_t^a \sim \mathcal{GP}(0_{3,1}, \Sigma^a \delta(t-t')),$$

where $\mathcal{GP}$ represents a Gaussian process and $\delta(t-t')$ is the Dirac delta function. To handle small foot slippage, we model the contact velocity noise in the body frame via a white Gaussian noise [27] $w_t^v \sim \mathcal{GP}(0_{3,1}, \Sigma^v \delta(t-t'))$. The process model of the individual term in the state becomes:

$$\frac{d}{dt}R_t = R_t(\tilde{\omega}_t - w_t^g)_\times,\ \frac{d}{dt}v_t = R_t\tilde{a}_t - w_t^a + g,\ \frac{d}{dt}p_t = v_t,\ \frac{d}{dt}d_t = R_t h_R(\tilde{\alpha}_t)(-w_t^v),$$

where $(\cdot)_\times$ denotes a 3×3 skew symmetric matrix, g is the gravity vector, $\tilde{\alpha}_t$ is the encoder measurements, and $h_R(\tilde{\alpha}_t)$ is the orientation of the contact frame in IMU (body) frame calculated by forward kinematics.

Measurement Model. In the filter update step, we use the contact constraint calculated earlier along with forward kinematics from current encoder measurements to correct the state. We assume the measurements from joint encoders are corrupted by white Gaussian noise: $\tilde{\alpha}_t = \alpha + w_t^\alpha, \quad w_t^\alpha \sim \mathcal{N}(0_{3,1}, \Sigma^\alpha \delta(t-t'))$. The assumption that the contact point is fixed in the world frame during a

contact event leads to $h_p(\tilde{\alpha}_t) = R_t^\mathsf{T}(d_t - p_t) + J_p(\tilde{\alpha}_t)w_t^\alpha$, where $J_p(\tilde{\alpha}_t)$ is the analytical Jacobian of the forward kinematics function. Given this measurement model, we use the right-invariant correction model described by Hartley et al. [9].

Implementation on Mini Cheetah. The implementation by Hartley et al. [9] is generalizable and open-sourced. In this work, we modify the open-sourced software for the MIT Mini Cheetah [8] by integrating the forward kinematics functions generated using FROST library [28]. An LCM interface is created for the filter to communicate with the robot. Furthermore, because the Mini Cheetah is not equipped with contact sensors, we develop a contact estimation framework to detect the contact events for the robot, as detailed in the next section.

4 Contact Estimation

In this section, we will discuss our deep learning approach for contact estimation. We define our contact state for each leg $l \in \{RF, LF, RH, LH\}$ as $C = [c_{RF} \quad c_{LF} \quad c_{RH} \quad c_{LH}]$ and $c_l \in \{0, 1\}$. Here, 0 denotes no contact, and 1 indicates a firm contact with the ground. Depending on the robot's motion, the contact state vector C can have totally 16 different states. (From all feet are in the air to all feet are in contact with the ground). We formulate it as a classification problem in our deep neural network and aim at estimating the correct contact state C given the input data. In order to follow typical classification pipelines, we map our contact state vector C to 16 different states $S \in \{0, 1, \ldots, 15\}$ by treating C as a binary value with 4 digits and using binary to decimal conversion as the function. For example, contact state $C_i = [0 \quad 1 \quad 1 \quad 0]$ is mapped to $S_i = 6$.

4.1 Preprocessing of the Input Data

The contact estimation network takes sensor measurements from joint encoders, IMU, and kinematics as input. For a synchronized time n, the sensor measurements are concatenated as a 54×1 array $z_n = [q_n \quad \dot{q}_n \quad a_n \quad \omega_n \quad p_{fn} \quad v_{fn}]$, where

$$
\begin{aligned}
q_n &= [q_{RF1n} \quad q_{RF2n} \quad q_{RF3n} \quad q_{LF1n} \quad \cdots \quad q_{LH3n}] \\
\dot{q}_n &= [\dot{q}_{RF1n} \quad \dot{q}_{RF2n} \quad \dot{q}_{RF3n} \quad \dot{q}_{LF1n} \quad \cdots \quad \dot{q}_{LH3n}] \\
a_n &= [a_{xn} \quad a_{yn} \quad a_{zn}]\,,\ \omega_n = [\omega_{xn} \quad \omega_{yn} \quad \omega_{zn}] \\
p_{fn} &= [p_{RFxn} \quad p_{RFyn} \quad p_{RFzn} \quad p_{LFxn} \quad \cdots \quad p_{LHzn}] \\
v_{fn} &= [v_{RFxn} \quad v_{RFyn} \quad v_{RFzn} \quad v_{LFxn} \quad \cdots \quad v_{LHzn}]\,.
\end{aligned}
$$

Here, q_n^T is a 12×1 vector containing all the joint encoder measurements (rad) at time n, $\dot{q}_n^\mathsf{T}$ is a 12×1 vector with joint angular velocity $(\mathrm{rad/sec})$, a_n holds the linear accelerations $(\mathrm{m/sec})$ from the IMU in the IMU frame, ω_n contains the angular velocity $(\mathrm{rad/sec})$ in the IMU frame, p_{fn}^T is a 12×1 vector with foot positions calculated from forward kinematics, and v_{fn}^T is a 12×1 vector that carries the linear velocities of each foot. It is worth noticing that both p_{fn} and v_{fn} are represented in the robot's hip frame. To infer the relationship between data across the time domain, for data point at time n, we create a window with size w and append previous measurements within this window into a 2D array $D_n = \left[z_{n-w}^\mathsf{T} \quad z_{n-w+1}^\mathsf{T} \quad \ldots \quad z_n^\mathsf{T}\right]^\mathsf{T}$. D_n is a $w \times 54$ array. (In our case, we use $w = 150$.) Each time, the network takes D_n as input and estimates the contact state S_n as output.

4.2 Network Architecture

The contact estimation network consists of 2 blocks of convolutions and 3 fully connected layers, as shown in Figure 2. Each block contains 2 1-dimensional convolution layers and a 1-dimensional max pooling. The convolution layers are designed to extract deep features from the input data. We choose a 1-dimensional kernel to increase computational efficiency in terms of memory usage and run time. The kernel moves along the time domain, and the data is padded to preserve dimensions. ReLU is employed as the nonlinear activation function for the convolution layer. The second convolution layer is applied with a dropout mechanism to prevent overfitting. At the end of each block, a 1-dimensional max-pooling layer is added to downsample the data.

The second block is connected to 3 fully connected layers that convert the deep features into the 16 classes we defined earlier. We also employ the dropout mechanisms in the first 2 fully connected

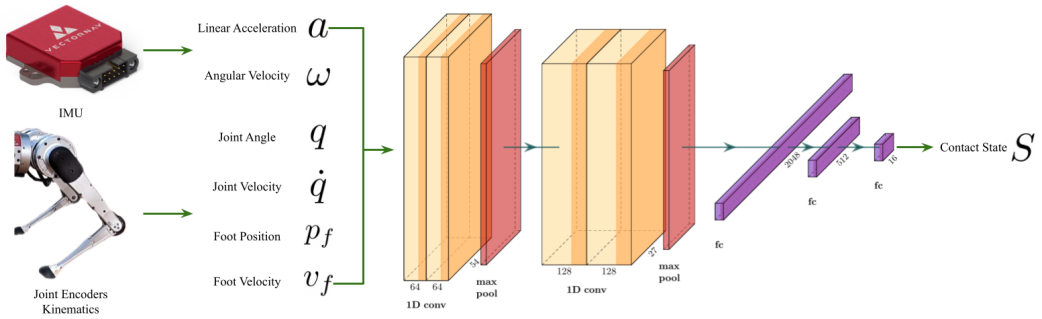

Figure 2: The architecture of the contact estimation network. The network takes in measurements from encoders and an IMU and outputs the contact state of the robot. The structure consists of 2 blocks and 3 fully connected layers. In each block, there are 2 1-dimensional convolution layers and 1 1-dimensional max pooling. The second convolution layer in each block is applied with a dropout mechanism to prevent overfitting.

Table 1: Number of data of each terrain in the contact data set.

				Terrain Type						
overall	air walking gait	air jumping gait	asphalt road	concrete	forest	grass	middle pebble	small pebble	rock road	sidewalk
1,013,441	44,386	48,972	94,615	465,144	72,144	103,392	44,442	52,669	45,819	58,115

layers to prevent overfitting. Finally, we formulate a classification problem using the cross entropy loss as $L(P_i) = -\log \frac{\exp(P_i)}{\sum_j \exp(P_j)}$. P_j is the probability output from the network of state j, and P_i is the probability of the ground truth state. The detailed network architecture is listed in the Appendix.

4.3 Contact Data Set

We create open-source contact data sets using an MIT Mini Cheetah robot. We record all the proprioceptive sensor measurements from the robot as LCM logs. Recorded measurements include joint encoders data, foot positions and velocities, IMU measurements, and estimated joint torques from an MIT controller [29]. The IMU measurements are received at $1000\,\text{Hz}$, while other data are recorded at $500\,\text{Hz}$. All data are upsampled to match the IMU frequency after being recorded. Around 1,000,000 data points are collected from various terrains, including asphalt road, concrete, forest, grass, middle pebbles, small pebbles, rock road, and sidewalk. To give the network negative examples of non-contact scenarios, we also record several sequences of the robot walking in the air by holding the robot up and giving the same controller command. (i.e., not having contact with the ground while operating the same gait.)

The grass data sets are collected inside an outdoor facility equipped with a motion capture system, which we utilize as a proxy for ground truth trajectory. Markers are attached to each foot and the robot body for the motion capture system to record the absolute foot positions and body pose in the world frame. For the rest of the data sets, in addition to the proprioceptive sensors, we also record RGB-D images with an Intel D455 camera mounted on top of the robot. We use these images in a state-of-the-art visual SLAM system, ORB SLAM2 [30], to generate the approximate ground truth trajectories. Figure 3b shows the picture of each terrain. The concrete data sets are recorded in a lab environment with polished concrete. The forest data sets are collected in a forest with random leaves, woods, and plants scattered around the ground. The middle pebbles data sets are recorded on pebbles of size ranging from 3 to 7 cm. The small pebbles data sets are recorded on pebbles ranging from 0.5 to 2 cm. Rock road consists of dirt and rocks of random sizes. It is worth mentioning that the contact data sets contain data from 3 different gaits and different robot loading.

To label ground truth contact points, we run an offline pre-processing algorithm with manual refinements. As shown in the Appendix, the pre-processing algorithm computes the local minimum and maximum value using the data points both in the future and past around current time stamp. Thus, it is unable to run online on the robot. On top of that, we also notice that this process creates false positives. Specifically, after inspecting slow-motion videos of Mini Cheetah's walking patterns, we observe a bouncing effect on the robot's foot upon touch down. This results in a sudden change of foot height. Because of this, we manually remove labels where a sudden change of the foot height take place. Figure 3c shows an example of labeled ground truth contacts overlapped with foot height. As seen in the figure, the initial "jump" in the signal is removed from the ground truth contact list.

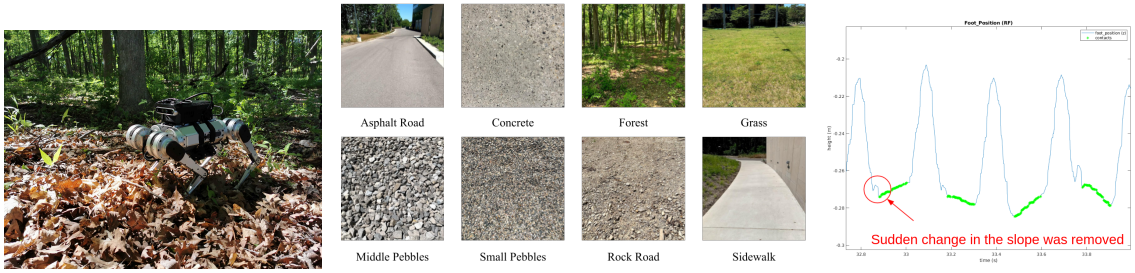

(a) Robot Setup for data collection. (b) Different ground types in the contact data set. (c) Example of the ground truth contacts.

Figure 3: (a) Robot configuration for all data sets except grass. Additional RGB-D images are recorded using an Intel D455 camera mounted on the robot. The images are used in ORB SLAM2 to generate ground truth trajectories. (b) Different ground types in the contact data set. (c) An example of labeled ground truth contacts (in green) overlapped with foot height measurements (in blue) in the hip frame. As seen in the figure, the initial "jump" in the signal is removed from contact ground truth labels.

Table 2: Accuracy comparison against baselines. The proposed method achieves the highest accuracy on both sequences. Although the gait cycle method has an accuracy closer to the proposed method, it does not remove false positives when gait cycle is violated.

Sequence	**Method**	**% Accuracy**				
		Leg RF	Leg LF	Leg RH	Leg LH	Leg Avg
Concrete Short Loop	GRF Thresholding	79.46	76.90	77.15	77.39	77.73
	Gait Cycle	88.64	87.34	86.93	87.72	87.66
	Proposed Method	95.78	93.84	95.07	95.49	95.04
Grass Test Sequence	GRF Thresholding	85.12	83.10	86.78	84.51	84.88
	Gait Cycle	95.38	94.56	91.17	90.67	92.95
	Proposed Method	98.82	98.14	98.42	97.85	98.31

5 Experimental Results

We train and evaluate our contact network on the contact data sets discussed in the previous section. Some sequences are reserved for testing. The rest of the data are split into validation, testing, and training sets with the ratio of 15%, 15%, and 70%, respectively. Table 1 lists the number of data being used from each terrain. A window size of 150 is used to allow the network to infer from the time domain. Each data point is normalized along the time domain to prevent scaling issues. We shuffled the training data to ensure generalizability in the model and reduce overfitting. We set the batch size to 30 and use 10^{-4} as the learning rate. The training process took 1.5 hours on an NVIDIA RTX 3090 GPU for 30 epochs. We note that after about seven epochs, the network can converge.

5.1 Contact Estimation

We now present the evaluation of the trained contact network. We evaluate the performance of the network in two ways. First, 16 class denotes the accuracy in terms of 16 contact states S we defined earlier. This case is a harsher way to evaluate the network since it requires all four legs to be correct simultaneously. We also compute the accuracy of the individual leg by comparing the estimated contact state of each leg with ground truth contacts individually. On average, the trained network achieves 87.54% of accuracy in terms of 16 contact states and 96.21% in each leg individually. Detailed accuracy across different terrain is listed in the Appendix.

5.2 Comparison Against Other Contact Detection Methods

In this experiment, we compare the accuracy of the proposed method with other contact detection methods with the Mini Cheetah robot. We implement a model-based approach as proposed in [12, 13, 14]. In this method, the estimated ground reaction force is computed via the general equation of motion with a low-pass filter. A fixed threshold is then set to detect the contact event. Because we don't have direct access to the motor current on the Mini Cheetah robot, we use the torque command from the controller as an approximation to the actual torque on the actuators. We also obtain the gait cycle command from the MIT controller [29] to serve as one contact segmentation method.

Table 2 lists the accuracy of above methods on two test sequence. We can see the proposed method achieves the highest accuracy on both sequences and the GRF thresholding has the worst perfor-

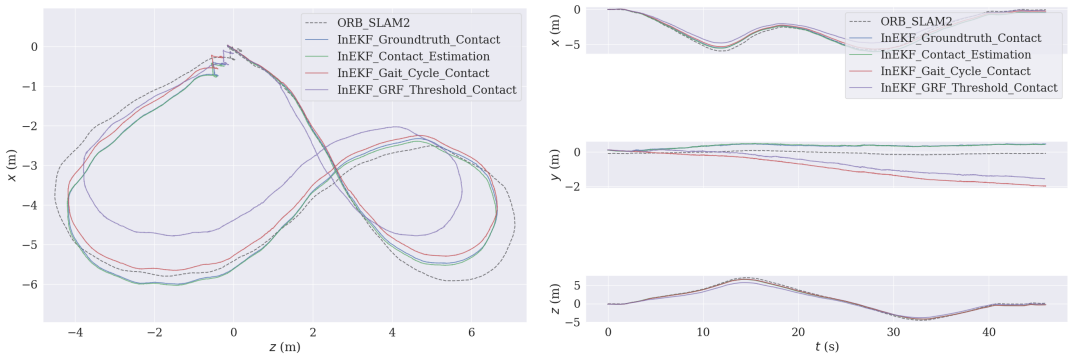

(a) Bird-eye View.

(b) XYZ Section View.

Figure 4: Concrete short loop test sequence.(a) The bird-eye view of the trajectories. The estimated trajectory is mapped to the camera frame. (Y pointing downward, and Z pointing forward.) With the estimated contacts, the InEKF can reproduce a similar trajectory to that uses ground truth contacts. (b) This figure shows the gait cycle method and GRF thresholding method creates a significant height drift in the InEKF state estimator.

mance. The gait cycle method does not remove the false positive signal when gait cycle is violated.

5.3 Contact-Aided Invariant Extended Kalman Filter

The estimated contacts are used with the contact-aided invariant extended Kalman filter described in Section 3 on both grass and concrete test sets. The estimated contacts are converted into LCM messages along with IMU and encoder measurements. The InEKF takes in the LCM messages and estimates the pose of the robot accordingly. We also run the filter with the ground truth contact data to serve as a reference.

Figure 4a and 4b show the trajectories from the InEKF with different contact sources on a concrete loop sequence. Qualitatively compared to the baseline contact detectors, the resulting trajectory with the proposed contact estimation has smaller drifts from the trajectory with ground truth contacts, especially in the height (y) axis.

5.4 Cassie-Series Robot Experiment

We also perform an experiment on a Cassie-series bipedal robot, developed by Agility Robotics, to show the generalizability of the proposed framework in different robots and training label sources. The robot is equipped a spring on each leg, which serves as a contact detection sensor. The contact state is defined as $C = [c_L \quad c_R]$. The decimal contact states becomes $S \in \{0, 1, \ldots, 3\}$. Similar to MiniCheetah, we concatenate the input data as $z_n = [q_n \quad \dot{q}_n \quad a_n \quad \omega_n \quad p_{fn} \quad v_{fn}]$. As the data from Cassie is recorded at a higher frequency, a window size of 600 is used in this experiment.
The data is recorded on polished concrete with a controller from [31]. Around $1,300,000$ data points are recorded. Two sequences are reserved as test sets, and two are used in training. We obtained two sources of ground truth contacts in this experiment - one from spring deflections; the other from kinematics. Spring contact labels are obtained by thresholding on the GRF computed from the spring deflection. We assume this is a more reliable contact source as it comes from a physical sensor. The kinematics contacts are generated following the same procedure described in Section 4.3. We train the network on two ground truth contacts separately and evaluate both networks against spring contacts.
Table 3 lists the accuracy of the two networks. We can see that when trained with ground truth contacts from spring deflection, the network can achieve above 98% of accuracy. When it is trained with the labels generated with kinematics, although a little bit of accuracy is lost, it still maintain above 95% accuracy on both leg. This indicates training with the labels generated from the process described in Section 4.3 only sacrifice a small amount of network accuracy. This experiment also shows the generalizability of the proposed framework. We demonstrate the proposed framework can achieve above 95% of accuracy on two distinct robot as long as one collect enough data for training.

Table 3: The accuracy of the network trained with data from a Cassie-series robot. The network is trained with labels generated from spring deflections and labels from kinematics separately. The results indicates training with the labels generated from the process described in Section 4.3 only sacrifice a small amount of network accuracy. This experiment also shows the generalizability of the proposed framework. We demonstrate the proposed framework can achieve above 95% of accuracy on two distinct robot as long as one collect enough data for training.

Contact Labels	**% Accuracy**			
	4 class	Leg L	Leg R	Leg Avg
Spring Deflection	98.03	99.82	98.19	99.01
Kinematics	93.32	95.86	97.45	96.65

Table 4: The average performance of different network structure on the test sets. Each "block" consists of 2 convolution layers followed by ReLU activation function and 1 max pooling, as described in Figure 2.

Network Structure	**% Accuracy**					
	16 class	Leg RF	Leg LF	Leg RH	Leg LH	Leg Avg
2 Blocks	87.54	96.73	95.69	95.94	96.50	96.21
1 Block	87.48	96.68	95.62	95.89	96.51	96.18
4 Blocks	85.87	95.61	95.54	94.83	95.39	95.34
Conv-Pool-Conv-Pool	86.45	96.16	95.36	94.90	96.00	95.60

5.5 Ablation Study and Runtime

We perform an ablation study on different network architectures to study how the performance changes. Table 4 lists the average performance of different network structure on the test sets. Each "block" consists of 2 convolution layers followed by ReLU activation function and 1 max pooling, as described in Figure 2. In the table, 2 Blocks is the proposed network detailed in Section 4. 1 Block refers to the network with only the first block and the fully connected layers of the proposed network. As for 4 Blocks, we added two additional blocks with 256 and 512 channels before the fully connected layers. Conv-Pool-Conv-Pool consists of only 2 convolution layers of sizes 64 and 128 and the fully connected layers. Each convolution layer is followed by a ReLU function and one max pooling. From the table, we can see that 2 Blocks has the best performance among all other structures. However, 1 Block maintains a smaller network while showing a slightly lower accuracy.

The inference speed on an NVIDIA RTX 3090 GPU is approximately 1100 Hz. The inference speed on an NVIDIA Jetson AGX Xavier, which we equipped on the robot, is around 830 Hz after TensorRT optimization. The frequency of the encoder measurements on Mini Cheetah is 500 Hz. This means the proposed contact estimator run at real-time without delay on the robot. The InEKF runs at 2000 Hz on an Intel i7-8750H CPU.

6 Discussion and Limitations

The developed contact estimator can capture contact events for a quadruped robot without force/contact sensors. However, the contacts are modeled as binary values, and the “quality” of contacts is not modeled in our contact estimation network. As a result, a perfect contact is always assumed when a contact event is detected. Furthermore, the covariance matrix of the contact measurements is set heuristically in the InEKF, which along with the uncertainties in the kinematic models, could be the potential sources of estimation error in the InEKF. In the future, we wish to expand the contact detection network to have a covariance estimation of the current contact state. This approach will allow the state estimator to treat contact information as a sensor measurement with an online uncertainty update and potentially decrease the drift in the state estimation.

We formulated the problem as a classification problem with 16 different contact states. This model can potentially be further improved by employing a multi-task learning approach [32, 33]. This approach allows to independently segment the contact events for each foot.

7 Conclusion

We developed a multi-modal deep learning-based contact estimation method that does not require contact/force sensors and works well with different robot gaits on distinct terrains. We present open-source contact data sets that require minimal supervision for labeling contact events of an MIT Mini Cheetah robot. We show that the trained contact network can achieve 94% of accuracy across different terrains. The estimated contacts are employed in a contact-aided invariant state observer for quadruped robots, and the resulting trajectory is comparable to a modern visual SLAM system.

References

[1] M. H. Raibert, H. B. Brown Jr, M. Chepponis, J. Koechling, and J. K. Hodgins. Dynamically stable legged locomotion. Technical report, Massachusetts Inst of Tech Cambridge Artificial Intelligence Lab, 1989.

[2] M. Bijelic, T. Gruber, and W. Ritter. A benchmark for LIDAR sensors in fog: Is detection breaking down? In *IEEE Intelligent Vehicles Symposium*, pages 760–767. IEEE, 2018.

[3] D. Belter and M. R. Nowicki. Optimization-based legged odometry and sensor fusion for legged robot continuous localization. *Robotics and Autonomous Systems*, 111:110–124, 2019.

[4] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari. Bayesian spatial kernel smoothing for scalable dense semantic mapping. *IEEE Robotics and Automation Letters*, 5(2):790–797, April 2020. ISSN 2377-3774.

[5] D. Wisth, M. Camurri, and M. Fallon. Preintegrated velocity bias estimation to overcome contact nonlinearities in legged robot odometry. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 392–398. IEEE, 2020.

[6] Y. Gong and J. Grizzle. Angular momentum about the contact point for control of bipedal locomotion: Validation in a lip-based controller. *arXiv preprint arXiv:2008.10763*, 2020.

[7] G. Bledt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, and S. Kim. MIT Cheetah 3: Design and control of a robust, dynamic quadruped robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2245–2252. IEEE, 2018.

[8] B. Katz, J. Di Carlo, and S. Kim. Mini cheetah: A platform for pushing the limits of dynamic quadruped control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 6295–6301. IEEE, 2019.

[9] R. Hartley, M. Ghaffari, R. M. Eustice, and J. W. Grizzle. Contact-aided invariant extended kalman filtering for robot state estimation. *International Journal of Robotics Research*, 39(4):402–430, 2020.

[10] A. S. Huang, E. Olson, and D. C. Moore. Lcm: Lightweight communications and marshalling. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4057–4062. IEEE, 2010.

[11] D. Moore, E. Olson, and A. Huang. Lightweight communications and marshalling for low-latency interprocess communication. 2009.

[12] M. Focchi, V. Barasuol, I. Havoutis, J. Buchli, C. Semini, and D. G. Caldwell. Local reflex generation for obstacle negotiation in quadrupedal locomotion. In *Nature-Inspired Mobile Robotics*, pages 443–450. World Scientific, 2013.

[13] S. A. Fakoorian, D. Simon, H. Richter, and V. Azimi. Ground reaction force estimation in prosthetic legs with an extended kalman filter. In *2016 Annual IEEE Systems Conference (SysCon)*, pages 1–6. IEEE, 2016.

[14] G. Fink and C. Semini. Proprioceptive sensor fusion for quadruped robot state estimation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 10914–10920. IEEE, 2020.

[15] A. De Luca, A. Albu-Schaffer, S. Haddadin, and G. Hirzinger. Collision detection and safe reaction with the dlr-iii lightweight manipulator arm. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1623–1630. IEEE, 2006.

[16] S. Haddadin, A. Albu-Schaffer, A. De Luca, and G. Hirzinger. Collision detection and reaction: A contribution to safe physical human-robot interaction. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3356–3363. IEEE, 2008.

[17] J. Hwangbo, C. D. Bellicoso, P. Fankhauser, and M. Hutter. Probabilistic foot contact estimation by fusing information from dynamics and differential/forward kinematics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3872–3878. IEEE, 2016.

[18] F. Jenelten, J. Hwangbo, F. Tresoldi, C. D. Bellicoso, and M. Hutter. Dynamic locomotion on slippery ground. *IEEE Robotics and Automation Letters*, 4(4):4170–4176, 2019.

[19] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, et al. Anymal-a highly mobile and dynamic quadrupedal robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 38–44. IEEE, 2016.

[20] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim. Contact model fusion for event-based locomotion in unstructured terrains. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4399–4406. IEEE, 2018.

[21] 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. *IEEE Robotics and Automation Letters*, 2(2):1023–1030, 2017.

[22] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[23] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

[24] N. Rotella, S. Schaal, and L. Righetti. Unsupervised contact learning for humanoid estimation and control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 411–417. IEEE, 2018.

[25] S. Piperakis, S. Timotheatos, and P. Trahanias. Unsupervised gait phase estimation for humanoid robot walking. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 270–276. IEEE, 2019.

[26] A. Barrau. *Non-linear state error based extended Kalman filters with applications to navigation*. PhD thesis, Mines Paristech, 2015.

[27] M. Bloesch, M. Hutter, M. A. Hoepflinger, S. Leutenegger, C. Gehring, C. Remy, and R. Siegwart. State estimation for legged robots: Consistent fusion of leg kinematics and IMU. In *Proceedings of the Robotics: Science and Systems Conference*, 2012.

[28] A. Hereid and A. D. Ames. FROST: Fast robot optimization and simulation toolkit. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vancouver, BC, Canada, 2017. IEEE.

[29] D. Kim, J. D. Carlo, B. Katz, G. Bledt, and S. Kim. Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control, 2019.

[30] R. Mur-Artal and J. D. Tardós. ORB-SLAM2: An open-source slam system for monocular, stereo, and RGB-D cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017.

[31] Y. Gong and J. Grizzle. Zero dynamics, pendulum models, and angular momentum in feedback control of bipedal locomotion, 2021.

[32] Y. Zhang and Q. Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1, 2021. doi:10.1109/TKDE.2021.3070203.

[33] T. Evgeniou and M. Pontil. Regularized multi–task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117, 2004.