



HAL
open science

A Data-driven Contact Estimation Method for Wheeled-Biped Robots

Ü. Bora Gökbakan, Frederike Dümbgen, Stéphane Caron

► **To cite this version:**

Ü. Bora Gökbakan, Frederike Dümbgen, Stéphane Caron. A Data-driven Contact Estimation Method for Wheeled-Biped Robots. 2024. hal-04726386v2

HAL Id: hal-04726386

<https://hal.science/hal-04726386v2>

Preprint submitted on 16 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Data-driven Contact Estimation Method for Wheeled-Biped Robots

U. Bora Gökbakan¹, Frederike Dümbgen¹ and Stéphane Caron¹

Abstract—Contact estimation is a key ability for limbed robots, where making and breaking contacts has a direct impact on state estimation and balance control. Existing approaches typically rely on gate-cycle priors or designated contact sensors. We design a contact estimator that is suitable for the emerging wheeled-biped robot types that do not have these features. To this end, we propose a Bayes filter in which update steps are learned from real-robot torque measurements while prediction steps rely on inertial measurements. We evaluate this approach in extensive real-robot and simulation experiments. Our method achieves better performance while being considerably more sample efficient than a comparable deep-learning baseline.

I. INTRODUCTION

Legged robots with feet can navigate rough terrains or confined spaces, yet at a relatively higher cost of transport and mechanical complexity than their wheeled counterparts. At the intersection between legs and wheels, wheeled-legged robots [1], [2] combine legs for active suspension and locomotion over varied terrains. They can drive, walk, jump [1] or skate [3], [4].

Wheeled-legged robots make and break contact with their environment, for instance, when jumping or stepping down-stairs. In the resulting hybrid dynamics, the presence or absence of contacts changes the equations of motion, thus impacting state estimation and motion control. Hence, the question of *contact estimation* or contact detection, where the system determines its contacts from sensory measurements, has been studied in itself as a key component of legged robotic systems.

One straightforward approach to contact estimation is to add dedicated contact sensors to the system. Direct measurements of contact forces by a collocated force/torque (F/T) sensor were spearheaded by the Honda P2 [5] and were kept in following iterations of the humanoid robotics project [6]. This design choice is still actively studied today, with recent works including an unsupervised approach by fuzzy c -means clustering [7] and a supervised approach training a contact-detection multi-layer perceptron (MLP) [8].

Alternative direct contact sensors include load cells and additional inertial measurement units (IMUs). The Solo open-source quadruped robot, in its first iteration [9], included a load cell in each foot, while measurements from collocated IMUs were considered in [10]. Overall, the upsides of dedicated contact sensors include the reliability and simplicity of estimating contact from direct measurements. The downsides include increased leg inertia, design complexity, and hardware costs.

¹ The authors are with Inria and the Computer Science Department of ENS (DI ENS), PSL Research University, Paris, France. Corresponding author: umit-bora.gokbakan@inria.fr




Fig. 1. Robustly detecting the moments when a wheeled-biped robot makes and breaks contact is crucial for successful estimation and control. This paper proposes a contact estimator based only on inertial and torque measurements. The measurements are fed into a novel Bayesian filter formulation to robustly estimate the binary contact state. We validate our results extensively both in simulation and real-world experiments, as depicted in the bottom figure.

More recent humanoid [11] and quadruped [12], [13], [14] robots explored a different design path, using integrated joint torque sensors rather than dedicated contact sensors. This *indirect* approach made the question of *contact estimation* more pressing, as fault-prone contact detection was identified as a major cause of drift by ensuing state estimators [15]. Following up on this observation, [16] proposed a probabilistic state machine to estimate each leg’s contact state separately by fusing the system’s kinematic and dynamic models. The approach was later extended in [17] to distinguish between sticking and sliding contact modes. Nevertheless, these methods rely on gait priors and kinematic models, which are not informative for stiff-legged wheeled-bipeds.

In this work, we estimate the ground contact state of a wheeled biped robot using only IMU and joint torque sensors that are readily available on current robots. Our main contributions are 1) the design of a Bayes filter with parameters learned from real-robot data, 2) its noise analysis in simulation and real-robot performance evaluation compared with a deep-learned baseline, and 3) the real-time, open-source implementation of the estimator, packaged as a new agent for the Upkie wheeled-biped platform.




Fig. 2. Overview of the estimation pipeline. We use Bayesian filtering to estimate the posterior probability $P(S_t | \mathbf{m}_{0:t})$ of being in contact state S_t given measurements $\mathbf{m}_{0:t}$. Internally, measurement probabilities $P(\mathbf{m}_t | S_t)$ are estimated by kernel density estimation (KDE) from knee and wheel torque sensors, while transition probabilities are estimated from IMU measurements through frequency analysis.

II. RELATED WORKS

Our approach borrows the probabilistic state machine of [16], [17], which we cast into the framework of Bayesian filtering. Rather than using a kinodynamic model of the system, however, we rely directly on real-robot data to learn the underlying measurement model distributions and to derive transition probabilities.

In a more data-driven approach, [18] used logistic regression models to estimate the probability of ground contact, with parameters learned from real-robot data. This pipeline also relies on a kinodynamic model to compute ground reaction forces, while our proposal is to work directly with collected sensory data. In a more recent work, [10] instrumented the feet of their robots with additional IMU sensors, fitting a kernel density estimator (KDE) on each of the six IMU measurements independently. Both [18] and [10] are purely measurement-based approaches, whereas in the probabilistic framework of Bayesian filtering we will consider two separate measurement and transition processes. We evaluate the effect of this separation as an ablation in our experiments and show that considering transition probabilities is important for high performance.

Training contact estimation from extensive simulation data has also been explored by [19], who suggested early stopping as an alternative to domain randomization for closing the sim-to-real gap more effectively. We reproduce the contact estimation network from their work and discuss it further in our experiments. We show that we achieve higher performance with significantly lower sample complexity by learning from a few real-world experiments.

III. METHOD

We consider the question of contact estimation from a probabilistic standpoint. Our proposal is to model our system as a *Markov process* and use Bayesian filtering to recursively update the contact estimate based on incoming measurements. Figure 2 gives an overview of the pipeline.

A. Probabilistic Framework

We represent the contact state at time t as a binary random variable S_t that can take values $\{C, \neg C\}$, where C indicates contact with the ground and $\neg C$ indicates no contact. We aim to calculate the posterior probability distribution of S_t , which we will refer to as the *belief* following the convention

of [20]. It is defined as

$$\text{bel}(S_t) := P(S_t | \mathbf{m}_{0:t}), \quad (1)$$

where $\mathbf{m}_{0:t} = \{\mathbf{m}_0, \dots, \mathbf{m}_t\}$ contains all measurements made up to time t . The belief at time $t > 0$ can be recursively computed by alternating between a prediction and an update step, defined respectively by

$$\overline{\text{bel}}(S_t) = \sum_{S \in C, \neg C} P(S_t | S_{t-1} = S) \text{bel}(S_{t-1}), \quad (2a)$$

$$\text{bel}(S_t) = \eta_t^{-1} P(\mathbf{m}_t | S_t) \overline{\text{bel}}(S_t), \quad (2b)$$

where we have introduced the intermediate prior belief $\overline{\text{bel}}(S_t) := P(S_t | \mathbf{m}_{0:t-1})$, and the normalization constant η_t . The prediction step updates the current belief based on the *transition model* $P(S_t | S_{t-1})$ (also called *process* or *motion* model), while the update step incorporates new observations through the *measurement model* $P(\mathbf{m}_t | S_t)$. We discuss the form of these two models in the next sections. The normalization factor η_t in (2b) is estimated by marginalizing over both possible contact states:

$$\eta_t = \sum_{S_t \in \{C, \neg C\}} p(\mathbf{m}_t | S_t = S) \overline{\text{bel}}(S_t = S). \quad (3)$$

In the absence of other information, we define a flat (non-informative) prior $P(S_0 = C) = P(S_0 = \neg C) = 0.5$ for the initial contact state.

B. Measurement Model

The measurement model $P(\mathbf{m}_t | S_t)$ is given by the likelihood of observing the measurements \mathbf{m}_t given the contact state S_t . Rather than using a parametric model, such as a kinodynamic robot model, we follow a non-parametric approach by fitting a KDE to torque measurements from the knee and wheel joints of each leg:

$$\mathbf{m}_t = (\tau_{\text{knee}}, \tau_{\text{wheel}}) \in \mathbb{R}^2 \quad (4)$$

KDEs estimate the probability density function (PDF) of a random variable by placing a kernel function K over all data points and summing them up. The kernel function is a non-negative function with a bandwidth parameter h that trades off smoothness and accuracy of the estimated PDF, whose formula is given by:

$$f_S(\mathbf{x}) = \frac{1}{N_S} \sum_{i=1}^{N_S} K(\mathbf{x} - \hat{\mathbf{x}}_i^S, h) \quad (5)$$




Fig. 3. Results of modeling torque measurements \mathbf{m}_t in contact (C) and no-contact ($-C$) states. The left two plots show the Gaussian KDEs fit to vectors of absolute knee and wheel torques $\mathbf{m}_t \in \mathbb{R}^2 : (\tau_{\text{knee}}, \tau_{\text{wheel}})$. The contact likelihood model is normalized by the marginal likelihood by assuming a flat $\text{bel}(S_t)$ at each timestep. This gives a simple non-recursive “Measurement Only” model to estimate the contact probability directly, without any dependence on the history or prior beliefs (right-most plot).

where K is the kernel function, $\hat{\mathbf{x}}_i^S$ are the data points collected for a given contact state S , and N_S is the number of data points in the measurement set.

This PDF provides our measurement model:

$$P(\mathbf{m}_t | S_t = S) = f_S(\mathbf{m}_t), \quad (6)$$

where f_S is the KDE fitted on contact state S . As we model contact as a binary variable, we only have two KDEs f_C and f_{-C} to fit. We use Gaussian kernels, as they are smooth functions with infinite support that allow us to cover the entire space of torque measurements. Figure 3 illustrates the KDEs estimated from real-robot data.

C. Transition Model

Given a pair $A, B \in \{C, -C\}$, we estimate the transition probability $P(S_t = A | S_{t-1} = B)$ empirically from accelerometer measurements. It is possible to detect when a *transition event* occurs thanks to nonstationarities in accelerometer measurements. It is nevertheless not straightforward to predict the *direction* (landing or takeoff) of the transition, as empirically, both takeoff and landing yield high upward and downward accelerations indifferently. We approach this by decomposing transition probabilities into two factors:

- 1) *Switch* (non-directional) probabilities $P(S_t \neq S_{t-1})$ that treat takeoff and landing indifferently, and
- 2) *Directional* probabilities $P(S_t = C | S_t \neq S_{t-1})$ and $P(S_t = -C | S_t \neq S_{t-1})$ that estimate the direction of a given switch.

Transition probabilities are then a product of these two. For instance, in case of a landing, we have:

$$\begin{aligned} & P(S_t = C, S_{t-1} = -C) \\ &= P(S_t = C, S_t \neq S_{t-1}) \\ &= P(S_t = C | S_t \neq S_{t-1}) P(S_t \neq S_{t-1}). \end{aligned} \quad (7)$$

1) *Switch Probabilities*: We use the vertical accelerometer measurements for calculating switch probabilities as follows. Let us denote by $\gamma_t \in \mathbb{R}^3$ the accelerometer measurement from the IMU at a given time step t . We consider a slice of




Fig. 4. The Bullet environment used to collect contact data. The robot was driven down a flight of steps, each with a height of 0.25. The captured IMU and proprioceptive readings were used to evaluate the contact estimator against noise.

the N most recent accelerometer measurements projected on the vertical axis in the world frame:

$$\mathbf{a}_t := [\gamma_{t-N+1}^z, \gamma_{t-N+2}^z, \dots, \gamma_t^z] \in \mathbb{R}^N \quad (9)$$

and compute its power spectrum $A_t(\omega)$:

$$A_t(\omega) = \frac{\|\hat{\mathbf{a}}_t(\omega)\|_2^2}{N} \quad (10)$$

where $\hat{\mathbf{a}}_t(\omega)$ is obtained by performing the short-time Fourier transform (STFT) of the windowed signal, and ω is the frequency bin.

Both take-off and landing events are impulsive, meaning their respective windows will have higher power densities than when neither occurs. We thus formulate the transition probability with a sigmoid function σ_1 of the power density in the window:

$$P(S_t \neq S_{t-1}) = \sigma_1(A_t(\omega)) \quad (11)$$

where σ_1 is a sigmoid with offset and slope set to $(8, 1)$ in all experiments.

2) *Directional Transition Probabilities*: On real-robot data, we identified that vertical accelerations were higher during landing, by a significant margin, compared with take-off. These Dirac-like accelerations are reflected in a more uniform energy distribution across the frequency spectrum. Takeoffs, on the other hand, are smoother and display relatively more energy at lower frequencies. We then distinguish the two with a heuristic based on the median frequency ω_{median} splitting the energy mass of the sliding window into two halves¹:

$$\omega_{\text{median}}(A_t(\omega)) = \arg \min_{\omega^*} \int_0^{\omega^*} A_t(\omega) d\omega - \int_{\omega^*}^{+\infty} A_t(\omega) d\omega. \quad (12)$$

We convert ω_{median} to a directional probability using a logistic function σ_2 with offset and slope given by $(7.5, 2.5)$:

$$P(S_t = C | S_t \neq S_{t-1}) = \sigma_2(\omega_{\text{median}}(A_t(\omega))). \quad (13)$$

¹As we deal with discrete time and frequencies, we interpolate linearly between the two frequencies closest to the true median.




Fig. 5. Robustness study with respect to measurement noise in simulation. The heatmap shows the overall success rate of the estimator under different noise levels (standard deviation of zero-mean Gaussian measurement noise). On its right, we provide a detailed study for two representative noise settings, highlighted with black rectangles. We zoom in on the robot rolling down one step (z-coordinate in world frame shown in green) and show, from left to right, the ground truth (top) and noisy (bottom) torque measurements, the resulting contact probabilities, the STFT of the ground truth (top) and noisy (bottom) accelerometer data (including power and median frequency), and the resulting transition probabilities.

IV. SIMULATION STUDY

We study the sensitivity of the designed filter to sensory noise in simulation. The simulation environment, shown in Figure 4, is built on *Bullet* and publicly available. In the environment, the open source Upkie wheeled biped robot [21] is driven down a flight of steps of height 0.25 m each. We collect two datasets: one to fit the measurement model (one run with no additional noise) and one with many runs to test the estimator’s performance by adding zero-mean Gaussian noise of varying magnitudes to knee torques, wheel torques, and accelerometer measurements.

Using the ground-truth contact state, we measure the success rate as the number of correctly classified contact states, where we use a threshold of 0.8 on the probabilities to define contact. This resulting success rate is visualized as a function of the noise levels in Figure 5. Across all acceleration noise levels, we observe a near-perfect success rate if the noise on torques is below 1 Nm. This is far above the noise levels we observed in real-world experiments (up to 0.12 Nm for wheel torques and 0.25 Nm for knee torques). The performance with respect to the acceleration noise is also near-perfect up to noise levels of about 7 m/s^2 , and real-world data was always below 0.23 m/s^2 .

We finally provide qualitative plots of the measurement and transition model performance to investigate how the different elements affect the drop in performance as we increase the noise. At the considered torque noise level, we observe many false non-contact detections, however this affects performance only little, possibly because the transition probability smoothes out such false detection. For the transition model, the added noise particularly affects the median frequency, which detects contact switch directions. This sensitivity leads to a drastic performance drop as we

increase the noise (see, for example, the acceleration noise 6 in the the left heatmap in Figure 5).

V. EXPERIMENTS

We implement and test the performance of the Bayesian filter to detect contacts on the wheeled-biped Upkie. Our study focuses on drops from unforeseen steps, where the robot alternates between contact and free-fall phases. We set up a test course where the robot travels over a horizontal floor, climbs a meter-long ramp with an 8° incline, drops 3 cm from the ramp onto a platform, rolls forward and drops again from a 12-cm step, which we repeated for a total number of 10 drops or 5 laps.²

A. Data collection

We fit measurement models and transition functions using real-robot data collected during a 20-minute session. For measurement likelihoods, we collected data for 7 min where the robot either (i) balanced and moved around the room, or (ii) was hung on a tether without making contact with the ground. Figure 3 illustrates the resulting distributions of torque measurements for both contact states.

In the second 13-minute long part of the data collection, we put the robot on an elevated platform (Fig. 1, top) and drove it over the edge to induce a drop. After each drop, we wait for the robot to stabilize (helping it with a tether if needed), and reposition it on the platform. We use robot measurements and a relatively high-speed camera to segment each trial into three consecutive phases: *rolling*, *falling*, and *stabilization*. We define *takeoff* and *landing* as the start and end times of each *falling* phase.

²Falls took slightly longer than a true free-fall because of the rolling down phase.

B. Training

a) Measurement Model: As we collected *contact* and *no-contact* joint measurements separately, we can use them directly without annotation to fit our two KDEs. We estimate the bandwidth parameter h (≈ 12.5) of our kernel functions following Scott’s rule of thumb, based on the number of samples and dimensions [22].

As the KDE method has a time and memory complexity of $\mathcal{O}(nd)$ in the number of samples n and dimensions d , it is common to use efficient search algorithms like balltrees or k-d trees and sample only near the input point³. In our case, we build 200×200 regular lookup grids and interpolate them linearly to achieve even faster prediction times, which is possible as our data has low-enough dimensionality.

b) Ablation: We integrate KDEs to predict the measurement likelihoods $P(\mathbf{m}_t | S_t)$ in the overall filter, but it would also be possible to normalize them to estimate a contact probability directly as:

$$P_\tau(S_t | \mathbf{m}_t) = \frac{P(\mathbf{m}_t | S_t = C)}{\sum_{S \in \{C, -C\}} P(\mathbf{m}_t | S_t = S)}, \quad (14)$$

We will call this indicator “measurements only”, as it is an ablation of the Bayes filter without transitions.

C. Neural Measurement Network

We reproduce the proposal from [19] of a *neural measurement network* (NMN) trained from simulated data produced by running an independent locomotion policy. Practical benefits from training in simulation include the ability to work with a large dataset and compare against perfect ground truth, the downside being a corresponding sim-to-real gap. While the adaption from point-foot quadruped to wheeled biped requires some changes, we strive to follow the original proposal whenever possible.

a) NMN inputs: In the formulation of [19] for quadrupeds, the NMN receives as input a measurement vector $\mathbf{m}_t = [\gamma_t, \omega_t, \mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{u}_{quad,t-1}]$, where γ_t denotes the linear acceleration measured by the accelerometer of the IMU, ω_t is the body angular velocity from the IMU frame, \mathbf{q}_t is the vector of measured joint angles with $\dot{\mathbf{q}}_t$ its time-derivative, and $\mathbf{u}_{quad,t-1} = \mathbf{q}_{des,t-1}$ is the vector of joint position targets sent at the previous time step to actuators. Rather than a position-controlled quadruped robot, in these experiments, our system is a wheeled biped with locked hip and knee positions and velocity-controlled wheels. We thus adapted the input measurement vector to:

$$\mathbf{m}_t = [\gamma_t, \omega_t, \mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{u}_{wheel,t-1}] \quad (15)$$

where IMU and joint measurements have the same definition as above, and $\mathbf{u}_{wheel,t-1}$ is the commanded ground velocity at the previous time step. In both cases, $\mathbf{u}_{quad,t-1}$ and $\mathbf{u}_{wheel,t-1}$ represent the action sent by the agent to its environment.

³This is an approximation when the kernel has infinite support.

b) NMN outputs: The NMN outputs (1) a probability $P_{\text{NMN}}(S_t = C | \mathbf{m}_t)$ that the robot’s feet are in contact with the ground, and (2) an estimate of the body linear velocity ${}_b\mathbf{v}_t(\mathbf{m}_t)$ from the IMU frame to the world frame. Although we do not use the latter thereafter, we included it in our reproduction as it can be beneficial for the estimator to be trained to predict both outputs rather than training to detect contacts only.

c) Simulation: We train the neural measurement network (NMN) in the open-source Bullet physics simulator [23]. The terrain consists of a flat horizontal floor, with a lateral friction coefficient of $\mu = 0.1$ and a rolling friction coefficient of $\mu_R = 0.01$. As recommended for wheels in the simulator documentation [23], we set the “stiffness” and “damping” parameters (of the underlying frictional-contact linear complementarity problem) to $3 \cdot 10^4$ and 10^3 respectively. During each episode, we randomly apply vertical external forces to lift the robot up or let it fall down, generating a variety of *contact* and *no contact* states.

d) Training: The network architecture from [19] consists of a one-layer gated recurrent unit (GRU) with 128 hidden features, followed by a 256×128 MLP. We generate a dataset of 4,400,000 simulation steps (≈ 5.5 hours of simulation data), each step lasting 5 ms, then split it into a training set of 4,000,000 steps and a validation set of 400,000 steps. As in [19], at each iteration of the NMN training loop, we use a batch size of 400 parallel environments rolling out episodes of 400 simulation steps using an independent locomotion policy. We select a sequence length of 10 to train the recurrent layer, trading off between the value of 1 used at inference (faster training, higher validation loss) and larger values such as 100 (slower training, lower validation loss). We use the Adam optimizer [24] with a learning rate of $5 \cdot 10^{-4}$ and 32 epochs.

As advocated in [19], we rely on early stopping rather than domain randomization over geometric and inertial parameters. We tracked training and validation losses to determine the number of iterations to stop at, which we identified as 20 in our use case, as opposed to 200 for a quadruped walking over varied terrains. This difference in scale can be readily explained by the lower number of degrees of freedom and terrain complexity of our application.

D. Results

We evaluate the complete Bayes filter, the simulation-trained NMN and the measurements-only ablation on the final task of detecting contacts on the real robot.

a) Contact estimation: We validate contact estimation as a binary classification task over a test set of real-robot data manually annotated by a human operator. We threshold the scalar output of each method to obtain a corresponding binary classification variable $\hat{y}_t \in \{C, -C\}$. We then evaluate “point-wise” *precision* and *recall* for each, as reported in Table I.

All methods achieve (near-)perfect results when the robot is in contact, as the robot stays in contact for longer periods,

where delays in transitions and label noise have very little effect. Nevertheless, predictions made during the free-flight phase are much more sensitive to errors and delays, as they are much shorter in duration. Despite this, we observe that both the Bayes filter and the measurement-only method maintain their ability to detect flight phases while predicting contact.

Method	$S_t = C$		$S_t = -C$	
	Precision	Recall	Precision	Recall
Bayes filter	1.0	1.0	0.55	0.37
Meas. only	1.0	1.0	0.53	0.40
NMN [19]	0.99	1.0	0.28	0.15

TABLE I

PRECISION AND RECALL OF POINT-WISE CONTACT-STATE DETECTION.

b) Transition Detection: Evaluating contact detection as a classification problem does not take into account the time-series nature of the collected data. In particular, delays in takeoff detection have a considerable impact on no-contact *recall* and *precision*, as delays account for a significant portion of the short flight duration. Consequently, we also evaluate the (mis)-detection rate of each method for landings and takeoffs. Results are reported in Table II. After each transition, we employ a period of 250 ms to exclude irrelevant (mis)detections.

Method	Takeoff		Landing	
	Precision	Recall	Precision	Recall
Bayes filter	0.91	1.0	0.71	1.0
Meas. only	0.59	1.0	0.55	1.0
NMN [19]	0.53	0.90	0.56	0.90

TABLE II

DETECTION OF TRANSITION EVENTS OVER 10 DROPS DURING A SESSION OF 4 MINUTES.

Both the measurement ablation and the Bayes filter achieve perfect recall on both transition types. In other words, they never fail to detect a transition within the admitted period, while the NMN baseline misses one. It is also important to consider the precision of the transitions, as misdetections can potentially lead to inappropriate actions by the robot. The Bayes method achieves a considerably higher precision rate, meaning we can ascribe more confidence to its predictions.

We witness that using measurements only can lead to more frequent misdetections. A transition can be triggered due to a spontaneous change in leg torques, which is frequent for a wheeled bicycle during balancing or crouching.

In Table III, the average time each method takes to detect a contact transition event is given. It is shown that, despite being a more conservative estimator that is not triggered as often as the measurement-only baseline, the Bayes method achieves favorable latencies.

Method	Takeoff	Landing
Bayes filter	77.1 ± 31.8	17.9 ± 9.1
Meas. only	83.1 ± 31.9	10.1 ± 7.37
NMN [19]	125.0 ± 59.6	22.4 ± 7.4

TABLE III

TRANSITION LATENCIES OF TESTED METHODS (MS). ONLY CORRECTLY IDENTIFIED TRANSITIONS WERE CONSIDERED.

VI. CONCLUSION AND DISCUSSION

We have presented a Bayesian filter for contact estimation on wheeled-biped robots, where all measurements come from IMU and joint torque sensors that are commonly available on experimental platforms. We learn measurement likelihoods from a small real-robot dataset of labeled torque measurements, and estimate transition probabilities online using accelerometer data. We have evaluated this contact estimator in real-robot experiments, where we observed better accuracy while requiring considerably fewer samples than an existing alternative trained on extensive simulation data.

The proposed method can be improved in several ways. First, because of the Markov assumption, the filter ignores longer-term effects, such as the low-frequency acceleration oscillations that can follow a harsh landing. Augmenting the state or using transition models with memory capabilities, such as long-short-term memory, could be a way to capture longer contexts. On another front, our method requires retraining for different types of robots. A future study could investigate the variability of learned distributions across different robots.

ACKNOWLEDGEMENTS

This work was supported by the European Union through the AGIMUS project (GA No. 101070165).

REFERENCES

- [1] V. Klemm, A. Morra, C. Salzmann, F. Tschopp, K. Bodie, L. Gulich, N. Kung, D. Mannhart, C. Pfister, M. Vierneisel, F. Weber, R. Deuber, and R. Siegwart, "Ascento: A Two-Wheeled Jumping Robot," in *2019 International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE, May 2019, pp. 7515–7521. [Online]. Available: <https://ieeexplore.ieee.org/document/8793792/>
- [2] M. Bjelonic, V. Klemm, J. Lee, and M. Hutter, "A survey of wheeled-legged robots," in *Robotics in Natural Settings*, J. M. Cascalho, M. O. Tokhi, M. F. Silva, A. Mendes, K. Goher, and M. Funk, Eds. Cham: Springer International Publishing, 2023, pp. 83–94.
- [3] M. Geilinger, R. Poranne, R. Desai, B. Thomaszewski, and S. Coros, "Skaterbots: optimization-based design and motion synthesis for robotic creatures with legs and wheels," *ACM Transactions on Graphics*, vol. 37, no. 4, pp. 1–12, Aug. 2018. [Online]. Available: <https://dl.acm.org/doi/10.1145/3197517.3201368>
- [4] M. Bjelonic, C. Dario Bellicoso, M. Efe Tiryaki, and M. Hutter, "Skating with a Force Controlled Quadrupedal Robot," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid: IEEE, Oct. 2018, pp. 7555–7561. [Online]. Available: <https://ieeexplore.ieee.org/document/8594504/>
- [5] K. Hirai, M. Hirose, Y. Haikawa, and T. Takenaka, "The development of honda humanoid robot," in *Proceedings. 1998 IEEE international conference on robotics and automation (Cat. No. 98CH36146)*, vol. 2. IEEE, 1998, pp. 1321–1326.

- [6] S. Kajita, "HRP-4 and other HRP robots," in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Springer Netherlands, 2019, pp. 101–116. [Online]. Available: http://link.springer.com/10.1007/978-94-007-6046-2_11
- [7] N. Rotella, S. Schaal, and L. Righetti, "Unsupervised contact learning for humanoid estimation and control," *CoRR*, vol. abs/1709.07472, 2017. [Online]. Available: <http://arxiv.org/abs/1709.07472>
- [8] S. Piperakis, M. Maravagakis, D. Kanoulas, and P. Trahanias, "Robust contact state estimation in humanoid walking gaits," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 6732–6738.
- [9] F. Grimmering, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols, *et al.*, "An open torque-controlled modular robot architecture for legged locomotion research," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3650–3657, 2020.
- [10] M. Maravagakis, D.-E. Argiropoulos, S. Piperakis, and P. Trahanias, "Probabilistic Contact State Estimation for Legged Robots using Inertial Information," Mar. 2023, arXiv:2303.00538 [cs]. [Online]. Available: <http://arxiv.org/abs/2303.00538>
- [11] J. Engelsberger, A. Werner, C. Ott, B. Henze, M. A. Roa, G. Garofalo, R. Burger, A. Beyer, O. Eiberger, K. Schmid, *et al.*, "Overview of the torque-controlled humanoid robot toro," in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 916–923.
- [12] M. Hutter, C. Gehring, M. Bloesch, M. A. Hoepflinger, C. D. Remy, and R. Siegwart, "Starleth: A compliant quadrupedal robot for fast, efficient, and versatile locomotion," in *Adaptive mobile robotics*. World Scientific, 2012, pp. 483–490.
- [13] S. Seok, A. Wang, M. Y. Chuah, D. Otten, J. Lang, and S. Kim, "Design principles for highly efficient quadrupeds and implementation on the mit cheetah robot," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 3307–3312.
- [14] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, and M. Hoepflinger, "ANYmal - a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Daejeon, South Korea: IEEE, Oct. 2016, pp. 38–44. [Online]. Available: <http://ieeexplore.ieee.org/document/7758092/>
- [15] M. Bloesch, M. Hutter, M. Hoepflinger, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, "State estimation for legged robots - consistent fusion of leg kinematics and IMU," in *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012.
- [16] J. Hwangbo, C. D. Bellicoso, P. Fankhauser, and M. Hutter, "Probabilistic foot contact estimation by fusing information from dynamics and differential/forward kinematics," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Daejeon: IEEE, Oct. 2016, pp. 3872–3878. [Online]. Available: <https://ieeexplore.ieee.org/document/7759570/>
- [17] F. Jenelten, J. Hwangbo, F. Tresoldi, C. D. Bellicoso, and M. Hutter, "Dynamic Locomotion on Slippery Ground," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4170–4176, Oct. 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8772165/>
- [18] 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*, vol. 2, no. 2, pp. 1023–1030, Apr. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7815333/>
- [19] D. Youm, H. Oh, S. Choi, H. Kim, and J. Hwangbo, "Legged robot state estimation with invariant extended kalman filter using neural measurement network," 2024, arXiv preprint arXiv:2402.00366.
- [20] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT press, 2005.
- [21] S. Caron, N. Perrin-Gilbert, V. Ledoux, U. B. Gökbakan, P.-G. Raverdy, A. Raffin, and V. Tordjman-Levasseur, "Upkie open source wheeled biped robot," 2024. [Online]. Available: <https://github.com/upkie/upkie>
- [22] D. W. Scott, "On optimal and data-based histograms," *Biometrika*, vol. 66, no. 3, pp. 605–610, 12 1979. [Online]. Available: <https://doi.org/10.1093/biomet/66.3.605>
- [23] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.