

Simultaneous Calibration of Noise Covariance and Kinematics for State Estimation of Legged Robots via Bi-level Optimization

Denglin Cheng^{1,*}, Jiarong Kang^{2,*}, and Xiaobin Xiong³

Abstract—Accurate state estimation is critical for legged and aerial robots operating in dynamic, uncertain environments. A key challenge lies in specifying process and measurement noise covariances, which are typically unknown or manually tuned. In this work, we introduce a bi-level optimization framework that jointly calibrates covariance matrices and kinematic parameters in an estimator-in-the-loop manner. The upper level treats noise covariances and model parameters as optimization variables, while the lower level executes a full-information estimator. Differentiating through the estimator allows direct optimization of trajectory-level objectives, resulting in accurate and consistent state estimates. We validate our approach on quadrupedal and humanoid robots, demonstrating significantly improved estimation accuracy and uncertainty compared to hand-tuned baselines. Our method unifies state estimation, sensor, and kinematics calibration into a principled, data-driven framework applicable across diverse robotic platforms. Video available at: <https://youtu.be/1zFORUMdLbg>.

I. INTRODUCTION

State estimation is fundamental to the autonomy of robotic systems, providing the basis for planning and control. Robots must infer their states—such as position, velocity, and orientation—from noisy sensor data, often under challenging real-world conditions. Classical approaches such as the Kalman filter and its extensions [1], [2] remain central, and have been successfully deployed in applications ranging from ground vehicles [3] and aerial robots [4] to quadrupeds and humanoids [5]–[8]. More recently, factor graph-based methods and smoothing techniques [9], [10] have enabled large-scale estimation for simultaneous localization and mapping (SLAM) [11] and visual-inertial odometry [12], [13].

A persistent challenge in these methods is the specification of noise covariance matrices. The process noise covariance captures uncertainty in the dynamics and actuation, while the measurement noise covariance characterizes sensor error models. In practice, both are difficult to obtain: manufacturer datasheets provide only partial information, while individual sensor calibration and physical system identification are both expensive and task-dependent. As a result, practitioners often rely on manual tuning [7], [14], [15] - a heuristic, time-consuming process that can produce estimators that are either sub-optimal or inconsistent [3] over out-of-distribution tasks. This problem is particularly acute for highly dynamic robots

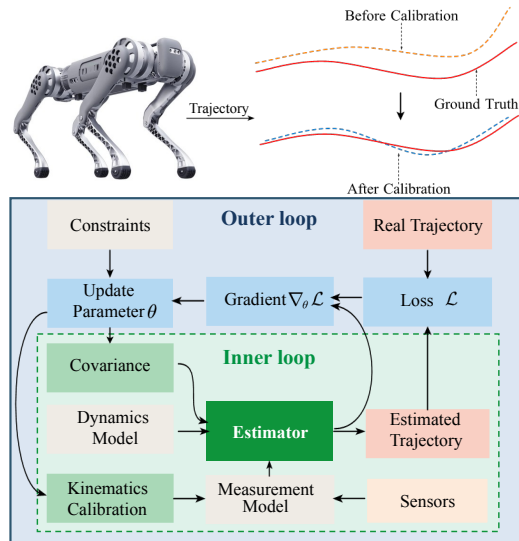


Fig. 1. Overview of the work with its application to a quadrupedal robot. such as quadrupeds and humanoids, whose complex, high-dimensional hybrid dynamics and kinematics, and multimodal sensor nonlinearities amplify the challenges to obtain precise state estimates [6]–[8], [14], [15], especially under kinematic uncertainties.

Several lines of work have been proposed to address this issue. Adaptive filtering methods [16] adjust covariances online based on innovation statistics, while expectation-maximization (EM) frameworks [17]–[19] iteratively refine noise parameters to maximize likelihood [20], [21]. Learning sensor parameters or noises parameters [19], [22], [23] is another common approach. However, these methods often either assume no availability of ground truth data or only focus on decoupled calibrations. More importantly, their treatments have not been examined on modern legged robots that have multimodal sensor suites and complex dynamics.

In this work, we propose a bi-level optimization framework for simultaneous covariance and kinematics calibration for robotic state estimation. At the upper level, we treat covariance matrices and uncertain kinematics as optimization variables. At the lower level, we solve a full-information estimator using the Maximum a Posteriori (MAP) formulation given the current covariance estimates and kinematics information. By differentiating through the estimator [24], [25], we directly optimize parameters to be calibrated with respect to trajectory-level objectives of minimizing the errors between the estimates and the measurable ground truth. This approach eliminates manual tuning, enforces physical consistency through structured parameterization and constraints,

Code is available at <https://github.com/DLinC3/LegBiCal>

*Equal contribution.

All authors are with the Legged AI Lab.

¹D. Cheng is now with the LCSR, Johns Hopkins University, MD, USA.

²J. Kang is now with the University of Wisconsin–Madison, WI, USA.

³X. Xiong is now with the Shanghai Innovation Institute (SII), Shanghai, China, and was with the University of Wisconsin–Madison during this work. Corresponding to X. Xiong (xiaobin.xiong@sii.edu.cn)

and generalizes across sensor modalities and robots.

We validate our approach on quadrupedal and bipedal robots, comparing against hand-tuned baselines. Across both platforms, our method calibrates the parameters to physically plausible values and significantly reduces estimation error and improves uncertainty consistency. These results highlight bi-level optimization for calibration as a principled and practical tool for bridging the gap between theoretical estimation frameworks and the demands of real-world robotics.

II. RELATED WORK

Legged Robotic State Estimation: State estimation for legged robots has been extensively studied due to the challenges posed by intermittent ground contacts, actuator uncertainties, and highly dynamic motions. Early approaches utilize Extended Kalman Filters (EKFs) to fuse proprioceptive sensing - Inertial Measurement Unit (IMU), contact sensors, and joint encoders - with kinematic constraints to achieve drift-reduced estimates on quadrupedal [5], bipedal [26] and humanoid robots [6]. Recent focuses of using EKFs on legged robots employ the invariant properties of the SE(3) manifold, and thus yield more accurate estimates [7], [14]. Another line of work leveraged factor graphs [10], [27] and smoothing or windowed formulations such as Moving Horizon Estimation (MHE) [15] to integrate proprioceptive and exteroceptive sensing data at scale. Despite these advances, most approaches rely on fixed or manually tuned noise covariances, limiting the accuracy of estimates, especially in the presence of kinematic uncertainty. Our work addresses this gap by introducing a bi-level optimization framework that calibrates covariance matrices and uncertain kinematics simultaneously in a principled manner.

Sensor Calibration: Sensor calibration and noise characterization have also been studied extensively in robotics, since the accuracy of estimation strongly depends on the quality of sensor models. Traditional approaches focus on extrinsic and intrinsic calibration of individual sensors, such as cameras [28], LiDARs [29], and IMUs [30]. At the theoretical level, several works address the identification of noise statistics. For instance, expectation-maximization (EM) algorithms iteratively refine noise models to maximize likelihood [17]–[19]. Recent efforts explored data-driven estimation of uncertainty, including machine learning to optimize covariances from raw sensory inputs [23], [31]. Despite these advances, most existing techniques commonly are tailored to specific sensors rather than complex legged robots; they often assume restrictive statistical models, or require careful initialization. In contrast, our bi-level optimization jointly calibrates both process and measurement noise with uncertain kinematics across diverse and complex legged robotic platforms.

System Identification: System identification (SysID) has long been a cornerstone of robotics, enabling accurate modeling of dynamics, kinematics, and actuation. Classical techniques for parameter estimation of rigid-body kinematics or dynamics exploit regressions of the geometric or Lagrangian models [32]–[34] subject to physical constraints. Gaussian processes or neural networks approximate resid-

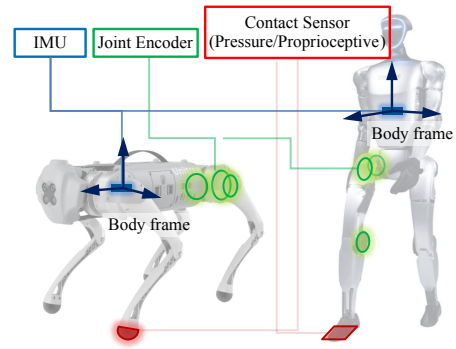


Fig. 2. Illustration of the common sensing capabilities on legged robots: Quadrupedal Go1 and Humanoid G1 from Unitree Robotics.

uals not captured by physics-based models [35] are commonly used in calibration. Recent works [36], [37] focus on online kinematics calibration via state estimators. Despite significant progress, SysID typically treats noise statistics, model parameters, and estimator design separately. Our bi-level optimization approach bridges this gap by jointly calibrating positive definite covariance matrices and kinematic parameters, highlighting the natural synergy between system identification and state estimation.

III. PRELIMINARIES

We begin by introducing the common modeling methods and techniques to address the state estimation of legged robots. The critical states to be estimated rather than measured are the linear velocity and orientation of their torso. We assume the robot is equipped with common proprioceptive sensors: joint encoders, IMU, and contact sensing/detection on the feet. Without loss of generality, we do not include additional exteroceptive sensors such as LiDAR or Cameras; yet, our method does not necessarily exclude fusing the exteroceptive sensing with the proprioceptive sensors.

We first introduce the models of legged state estimation that are widely used in the literature [5]–[7], [14], [15]. Then, we will elaborate the preliminaries of the mathematics on formulating state estimation as an optimization problem that we will employ in our calibration work.

A. Modeling of Legged Robot Estimation

Legged robot state estimation is generally formulated by a nonlinear process and measurement model that jointly describe the system dynamics and sensor observations.

1) *Process Model:* We define the state as the composition of the floating-base kinematics, the absolute positions of the footholds, and the IMU biases in a discrete-time model:

$$\mathbf{x}_k := (\mathbf{p}_k, \mathbf{v}_k, \mathbf{q}_k, \mathbf{p}_{\text{foot},k}, \mathbf{b}_{a,k}, \mathbf{b}_{\omega,k}). \quad (1)$$

Here \mathbf{p}_k and \mathbf{v}_k denote the position and velocity of the main body expressed in inertial frame W . The orientation is represented by a unit quaternion \mathbf{q}_k that parameterizes the rotation from W to the body frame B . To incorporate leg kinematics, the state is augmented with the absolute positions of the footholds $\mathbf{p}_{\text{foot},k}$ in W , the accelerometer and gyroscope biases $\mathbf{b}_{a,k}, \mathbf{b}_{\omega,k}$ expressed in B . Formally, the state \mathbf{x}_k resides on the product manifold $\mathcal{G} := \text{SE}(3) \times \mathbb{R}^{n_{\text{aug}}}$,

where the pose $(\mathbf{p}_k, \mathbf{q}_k)$ constitutes the SE(3) component and the remaining Euclidean sub-states are stacked in $\mathbb{R}^{n_{\text{aug}}}$.

The IMU provides measurements of the proper acceleration $\tilde{\mathbf{a}}_k$ and the angular rate $\tilde{\boldsymbol{\omega}}_k$ of the robot's main body. These quantities are modeled as the true values corrupted by slowly varying biases and additive Gaussian white noise [5], [14], [15]:

$$\tilde{\mathbf{a}}_k = \mathbf{a}_k + \mathbf{b}_{a,k} + \mathbf{w}_{a,k}, \quad \tilde{\boldsymbol{\omega}}_k = \boldsymbol{\omega}_k + \mathbf{b}_{\omega,k} + \mathbf{w}_{\omega,k},$$

where \mathbf{a}_k is the absolute acceleration rotated into frame B minus the gravitational vector, and $\mathbf{w}_{(\cdot)}$ represents the discrete Gaussian noise terms. Assuming a zero-order hold (ZOH) for the bias-corrected IMU measurements [5]. The discrete-time prediction equations are formulated as follows:

$$\begin{aligned} \mathbf{p}_{k+1} &= \mathbf{p}_k + \mathbf{v}_k \Delta t + \frac{1}{2} (\mathbf{C}(\mathbf{q}_k)(\tilde{\mathbf{a}}_k - \mathbf{b}_{a,k}) + \mathbf{g}) \Delta t^2 + \mathbf{w}_p, \\ \mathbf{v}_{k+1} &= \mathbf{v}_k + (\mathbf{C}(\mathbf{q}_k)(\tilde{\mathbf{a}}_k - \mathbf{b}_{a,k}) + \mathbf{g}) \Delta t + \mathbf{w}_v, \\ \mathbf{q}_{k+1} &= \zeta((\tilde{\boldsymbol{\omega}}_k - \mathbf{b}_{\omega,k}) \Delta t) \otimes \mathbf{q}_k + \mathbf{w}_q, \end{aligned}$$

$$\mathbf{P}_{\text{foot},k+1} = \mathbf{P}_{\text{foot},k} + \mathbf{w}_{\text{foot}}, \quad \mathbf{b}_{(\cdot),k+1} = \mathbf{b}_{(\cdot),k} + \mathbf{w}_{b_{(\cdot)}},$$

where $\mathbf{C}(\mathbf{q}_{(\cdot)})$ extracts the rotation matrix from quaternion, and \mathbf{g} is the gravity vector in W, \otimes denotes the quaternion multiplication operator and $\zeta(\cdot)$ represents the map from a rotation vector to its corresponding unit quaternion [5], [15]. Although footholds are assumed stationary, a random walk is introduced to accommodate potential slippage, with IMU biases likewise modeled as Brownian motions. The aggregated process noise $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$ follows the uncertainty modeling architecture detailed in Section IV.B of [15].

2) *Measurement Model*: We formulate the measurement model for a generic leg by relating the derived kinematics $\tilde{\mathbf{y}}_k$ to the state-predicted observations $\mathbf{h}(\mathbf{x}_k)$. The full multi-leg model follows by stacking individual residuals, implying a block-diagonal covariance structure under the assumption of independent sensor noise across legs. This relationship accounts for the mapping of raw sensor uncertainty into the task space via an additive effective noise term \mathbf{n}_k [5]:

$$\tilde{\mathbf{y}}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{n}_k, \quad \mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k).$$

Kinematic Measurements: Joint encoders and the IMU gyroscope provide raw measurements $\tilde{\mathbf{z}} := [\tilde{\boldsymbol{\alpha}}^\top, \tilde{\boldsymbol{\alpha}}^\top, \tilde{\boldsymbol{\omega}}^\top]^\top$ consisting of joint angles/velocities $\tilde{\boldsymbol{\alpha}}/\tilde{\boldsymbol{\alpha}}$ and body angular velocities $\tilde{\boldsymbol{\omega}}$. We model these raw inputs as the true values corrupted by additive Gaussian noise $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\eta)$, where $\boldsymbol{\Sigma}_\eta = \text{blkdiag}(\mathbf{R}_\alpha, \mathbf{R}_{\dot{\alpha}}, \mathbf{Q}_\omega)$. Leveraging the forward kinematics $\text{fk}(\cdot)$ and the analytical leg Jacobian $\mathbf{J}(\cdot)$, $\mathbf{h}(\mathbf{x}_k)$ and $\tilde{\mathbf{y}}_k$ derived from raw sensor inputs are defined as follows:

$$\begin{aligned} \mathbf{h}_p(\mathbf{x}_k) &= \mathbf{C}(\mathbf{q}_k)^\top (\mathbf{p}_{\text{foot},k} - \mathbf{p}_k), \\ \tilde{\mathbf{y}}_{p,k} &= \text{fk}(\tilde{\boldsymbol{\alpha}}_k), \\ \mathbf{h}_v(\mathbf{x}_k) &= -\mathbf{C}(\mathbf{q}_k)^\top \mathbf{v}_k, \\ \tilde{\mathbf{y}}_{v,k} &= \mathbf{J}(\tilde{\boldsymbol{\alpha}}_k) \tilde{\boldsymbol{\alpha}}_k + [\tilde{\boldsymbol{\omega}}_k - \mathbf{b}_{\omega,k}]_\times \text{fk}(\tilde{\boldsymbol{\alpha}}_k), \end{aligned} \quad (2)$$

where $[\cdot]_\times$ denotes the skew-symmetric operator, \mathbf{h}_v models negative body velocity under the static contact assumption, while $\tilde{\mathbf{y}}_v$ captures the actual proprioceptive measurement.

Uncertainty Modeling: A critical aspect of our formulation is the analytic propagation of sensor noise into the task space.

Following the linearization approach in, the covariance matrices $\mathbf{R}_{p,k}$ and $\mathbf{R}_{v,k}$ for the position and velocity residuals are derived via a first-order expansion of the kinematic functions w.r.t. the noise vector $\boldsymbol{\eta}$: $\mathbf{R}_k = \mathbf{G}_k \boldsymbol{\Sigma}_\eta \mathbf{G}_k^\top + \mathbf{R}_{\text{add}}$, where \mathbf{R}_{add} accounts for unmodeled kinematic calibration errors, and the mapping matrix $\mathbf{G}_k := \partial \tilde{\mathbf{y}}_k / \partial \tilde{\mathbf{z}}_k$ is the Jacobian of the kinematic measurements $\tilde{\mathbf{y}}_k := [\tilde{\mathbf{y}}_{p,k}^\top, \tilde{\mathbf{y}}_{v,k}^\top]^\top$ w.r.t $\tilde{\mathbf{z}}_k$:

$$\mathbf{G}_k = \begin{bmatrix} \mathbf{J}(\tilde{\boldsymbol{\alpha}}_k) & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \text{DJ}(\tilde{\boldsymbol{\alpha}}_k)[\dot{\tilde{\boldsymbol{\alpha}}}_k] + [\tilde{\boldsymbol{\omega}}_k^c]_\times \mathbf{J}(\tilde{\boldsymbol{\alpha}}_k) & \mathbf{J}(\tilde{\boldsymbol{\alpha}}_k) & -[\text{fk}(\tilde{\boldsymbol{\alpha}}_k)]_\times \end{bmatrix}, \quad (3)$$

where $\tilde{\boldsymbol{\omega}}_k^c := \tilde{\boldsymbol{\omega}}_k - \mathbf{b}_{\omega,k}$ and $\text{DJ}(\boldsymbol{\alpha})[\mathbf{u}]$ denotes directional derivative of the matrix-valued map \mathbf{J} at $\boldsymbol{\alpha}$ along $\mathbf{u} \in \mathbb{R}^{n_\alpha}$. This formulation explicitly capture coupling between IMU gyro noise and leg odometry uncertainty. Finally, a binary contact flag $c_k \in \{0, 1\}$ gates the velocity residual, such that the effective information matrix is $\boldsymbol{\Omega}_v = c_k \mathbf{R}_{v,k}^{-1}$, ensuring velocity updates occur only during reliable stance phases.

B. State Estimation as a Maximum A Posteriori Problem

We formulate the optimal state estimation problem within a Maximum A Posteriori (MAP) framework over a finite horizon T , the optimal state trajectory $\mathbf{x}_{0:T} := \{\mathbf{x}_0, \dots, \mathbf{x}_T\}$ is inferred by maximizing the posterior density given the sequence of inputs $\mathbf{u}_{0:T-1}$ and measurements $\mathbf{y}_{0:T}$:

$$\begin{aligned} \hat{\mathbf{x}}_{0:T}^{\text{MAP}} &= \arg \max_{\mathbf{x}_{0:T}} p(\mathbf{x}_{0:T} \mid \mathbf{y}_{0:T}, \mathbf{u}_{0:T-1}) \\ &= \arg \min_{\mathbf{x}_{0:T}} (-\ln p(\mathbf{y}_{0:T} \mid \mathbf{x}_{0:T}) - \ln p(\mathbf{x}_{0:T} \mid \mathbf{u}_{0:T-1})). \end{aligned}$$

Under the Markov assumption and additive Gaussian noise models, it reduces to a deterministic Full Information Estimation (FIE). We formulate this as a constrained nonlinear optimization over the decision variables $\mathcal{Z}_{0:T}$:

$$\begin{aligned} \min_{\mathcal{Z}_{0:T}} \quad & \|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|_{\mathbf{P}_0^{-1}}^2 + \sum_{k=0}^{T-1} \|\mathbf{w}_k\|_{\mathbf{Q}_k^{-1}}^2 + \sum_{k=0}^T \|\mathbf{n}_k\|_{\mathbf{R}_k^{-1}}^2, \\ \text{s.t.} \quad & \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k, \quad \forall k \in \{0, \dots, T-1\}, \\ & \mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{n}_k, \quad \forall k \in \{0, \dots, T\}, \\ & \mathbf{g}(\mathbf{x}_k) \leq \mathbf{0}, \quad \forall k \in \{0, \dots, T\}. \end{aligned} \quad (4)$$

Here, the notation $\|\mathbf{v}\|_{\boldsymbol{\Sigma}^{-1}}^2 := \mathbf{v}^\top \boldsymbol{\Sigma}^{-1} \mathbf{v}$ denotes the squared Mahalanobis distance. The prior state distribution is given by $\mathbf{x}_0 \sim \mathcal{N}(\tilde{\mathbf{x}}_0, \mathbf{P}_0)$. The decision variable vector $\mathcal{Z}_{0:T}$ explicitly includes the process noise \mathbf{w}_k and measurement noise \mathbf{n}_k as slack variables to satisfy the equality constraints:

$$\mathcal{Z}_{0:T} := [\mathbf{x}_{0:T}^\top, \mathbf{w}_{0:T-1}^\top, \mathbf{n}_{0:T}^\top]^\top.$$

$\mathbf{g}(\mathbf{x}_k) \leq \mathbf{0}$ represents general state constraints (physical limits or obstacle avoidance). The major drawback of FIE is that its computational cost grows with the horizon T , making online use practically intractable. MHE [15] instead uses a fixed past window of measurements and repeatedly solves the problem on a shifting horizon (as in MPC), yielding bounded and predictable cost. Because legged estimation is real-time, EKF variants with analytical updates are commonly used instead of windowed optimization. For offline calibration, we exploit all available data and therefore adopt FIE.

IV. PROBLEM STATEMENT

With properly selections of the noise covariance matrices $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$, the commonly applied estimators such as Extended Kalman Filters (EKFs) [5], [6], invariance EKFs [7], [14], and MHE [15] can yield satisfying estimates. However, it is widely known that extensive expert tuning of $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$ is inevitable on complex robots, and it can easily lead to sub-optimal, inconsistent or even diverging estimates on real robots. Additionally, inaccurate kinematics can further degrade estimate stability and precision. Therefore, we aim to provide an automatic and tuning-free method to optimally calibrate $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$ along with uncertain kinematics so that any reasonable choices of the estimators, regardless of EKFs or MHEs depending on the computational budget, can yield precise and consistent estimation on diverse set of tasks.

A. Assumptions and Objective

Sensors and Ground Truth Measurements: We assume the robot is equipped with common proprioceptive sensors that are specified in previous section. The robot has a functional controller that allows it to locomote in an environment that has motion capture devices to provide accurate measurements of the torso pose and velocities, which will be used as ground truth, denoted at \mathbf{x}_{GT} . Note that \mathbf{x}_{GT} does not cover all the states such as the IMU biases. We assume that the robot is controlled to realize certain tasks with which all the sensor information $\mathbf{Y} := \{\mathbf{y}_k\}_{k=0}^T$ along with the ground truth measurements \mathbf{x}_{GT} are provided to us.

Objectives: Our objective is thus to calibrate the optimal covariance matrices $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$ along with uncertain kinematic parameters so that the estimated results will be as close as possible to the ground truth using our process and measurement models. Without losing generality, for legged robots, we assume the last link on each leg has uncertain kinematic information; other linkage can be assumed to have uncertainties as well. This kinematic uncertainty on the last link (shin/calf) results in an uncertain foot position. Additionally, the motion capture device will capture markers that are placed on the torso. Therefore, besides uncertain shin kinematics, \mathbf{x}_{GT} will not directly measure the robot state but with a constant kinematic offset, which will need to be calibrated along with the calibration of $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$.

B. Bi-level Optimization

We propose a bi-level optimization to solve the calibration problem. The optimization is given by:

$$\begin{aligned} \min_{\boldsymbol{\theta} \in \mathbb{R}^{n_\theta}} \quad & \mathcal{L}(\hat{\mathbf{x}}(\boldsymbol{\theta}), \mathbf{x}_{\text{GT}}, \boldsymbol{\theta}), \\ \text{s.t.} \quad & \boldsymbol{\theta} \in \mathcal{C}, \\ & \hat{\mathbf{x}}(\boldsymbol{\theta}) = \arg \min_{\mathbf{x} \in \mathbb{R}^{n_x \times (T+1)}} \mathcal{J}(\mathbf{x}, \boldsymbol{\theta}). \end{aligned} \quad (5)$$

At the upper level, we minimize a differentiable loss \mathcal{L} that measures the deviation between the estimated trajectory $\hat{\mathbf{x}}$ and the ground truth trajectory \mathbf{x}_{GT} . We decompose the decision vector into covariance parts and geometric calibration parts $\boldsymbol{\theta} = (\boldsymbol{\theta}_{\text{cov}}, \boldsymbol{\theta}_{\text{foot}}, \boldsymbol{\theta}_{\text{base}})$, where $\boldsymbol{\theta}_{\text{cov}}$ parameterizes the covariance matrices, $\boldsymbol{\theta}_{\text{foot}}$ stacks the offsets on the foot kinematics, and $\boldsymbol{\theta}_{\text{base}} \in \mathbb{R}^3$ is the constant offset from the

base frame of the torso B to the motion-captured body frame rigidly attached to the torso M, expressed in B. In particular, each $\boldsymbol{\theta}_{\text{foot},j} \in \mathbb{R}^3$ is a constant vector expressed in the calf frame of leg j C_j that points from the foot to the contact point on the real robot. Thus, the actual foot position is:

$$p_{\text{real},j}^{\text{W}} = p_{\text{foot},j}^{\text{W}} + \mathbf{C}_{\text{WC}_j} \boldsymbol{\theta}_{\text{foot},j}, \quad (6)$$

in which the offset $\boldsymbol{\theta}_{\text{foot},j}$ is rotated from C_j to the world frame by \mathbf{C}_{WC_j} and added to the nominal foot position $p_{\text{foot},j}^{\text{W}}$ to yield the actual contact location $p_{\text{real},j}^{\text{W}}$. This time-invariant offset represents kinematic errors that come from mounting tolerances or manufacture inaccuracy. Note that, the observation vector $\mathbf{y}(\boldsymbol{\theta})$ (2) and the kinematic mapping matrices $\mathbf{G}(\boldsymbol{\theta})$ (3) depend explicitly on $\boldsymbol{\theta}_{\text{foot}}$.

At the lower level, we fix the vector $\boldsymbol{\theta}$ and solve the MAP problem, where \mathcal{J} is the MAP objective in (4). Given a covariance specification $\Sigma(\boldsymbol{\theta})$, where $\Sigma(\boldsymbol{\theta})$ is the collection of $\mathbf{Q}_{(\cdot)}$ and $\mathbf{R}_{(\cdot)}$, the solution of the MAP yields an estimated trajectory $\hat{\mathbf{x}}(\boldsymbol{\theta}) \in \mathcal{X}$. Following (1), we represent the estimated state as $\hat{\mathbf{x}}_k := (\mathbf{T}_k, \mathbf{e}_k) \in \mathcal{G}$, $\mathbf{X} := (\hat{\mathbf{x}}_0, \dots, \hat{\mathbf{x}}_T) \in \mathcal{G}^{T+1}$, where $\mathbf{T}_k \in \text{SE}(3)$ is the base pose and $\mathbf{e}_k \in \mathbb{R}^{n_{\text{aug}}}$ stacks the Euclidean components in (1).

C. Constraint Set \mathcal{C}

Given a positive-definite matrix $\Sigma_i \in \mathbb{R}^{d_i \times d_i}$, we define its vectorization $\boldsymbol{\theta}_{\text{cov}} = \text{vec}(\Sigma_i) \in \mathbb{R}^{d_i(d_i+1)/2}$ by stacking the independent entries (upper triangular part including the diagonal) in a fixed order. During optimization, the parameter vector must remain in the constraint set $\mathcal{C} = \left\{ \boldsymbol{\theta} \in \mathbb{R}^{n_\theta} \mid \underline{\boldsymbol{\theta}} \leq \boldsymbol{\theta} \leq \bar{\boldsymbol{\theta}}, \Sigma_i(\boldsymbol{\theta}_{\text{cov}}) \in \mathbb{S}_{++}^{d_i}, \forall i \right\}$, where $\underline{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}} \in \mathbb{R}^{n_\theta}$ denote element-wise lower and upper bounds on each parameter, including both variances, correlations, and kinematic parameters. The set $\mathbb{S}_{++}^{d_i}$ denotes the cone of symmetric positive definite matrices of size $d_i \times d_i$.

To ensure statistical and physical feasibility, the constraint set \mathcal{C} imposes two requirements. First, each covariance block Σ_i is constrained to lie in $\mathbb{S}_{++}^{d_i}$ by the property of covariance matrices. Second, all parameter entries are bounded within predefined intervals $[\underline{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}]$, reflecting prior sensor knowledge in the form of individual bounds on variances and correlations. These bounds also serve to improve numerical stability by avoiding ill-conditioned solutions during optimization.

The convexity and compactness of the set \mathcal{C} : The cone $\mathbb{S}_{++}^{d_i}$ is convex but open cone. By imposing element-wise box constraints $[\underline{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}]$, we obtain a closed and bounded subset \mathcal{C} , which remains convex and is therefore compact.

V. METHOD

We now present our method to solve the bi-level optimization problem in (5). We adopt an iterative optimization approach: the lower-level problem (4) is solved using a standard interior-point method [38], while the upper-level problem leverages its solution and Karush–Kuhn–Tucker (KKT) condition to update the parameters $\boldsymbol{\theta}$ via an adaptive Frank–Wolfe algorithm [39] that iteratively refines the estimate until convergence. The index $t \in \mathbb{N}$ denotes the Frank–Wolfe iteration counter, and superscript $(\cdot)^t$ denotes the value of variable at iteration t .

Algorithm 1 Frank–Wolfe Algorithm

- 1: **Input:** \mathcal{L} , θ^t , \mathcal{C}^t , t
 - 2: **Gradient compute:** $\nabla_{\theta}\mathcal{L}(\theta^t)$ \triangleright as in Sec. (V-B.2)
 - 3: **Direction finding (LMO):**

$$\begin{aligned} \mathbf{s}^* &= \operatorname{argmin} \mathbf{s}^{\top} \nabla_{\theta}\mathcal{L}(\theta^t) \\ &\text{subject to } \mathbf{s} \in \mathcal{C}^t \subset \mathcal{C} \end{aligned}$$
 - 4: **Step-size:** $\gamma^t \in (0, 1] \leftarrow \text{LINESEARCH}(\mathcal{L}, \nabla\mathcal{L}, \theta^t, \mathbf{s}^*)$
 - 5: **Update:** $\theta^{t+1} \leftarrow (1 - \gamma^t)\theta^t + \gamma^t\mathbf{s}^t$
-

A. Frank–Wolfe Algorithm for the Upper Level Optimization

Given that the feasible set \mathcal{C} (IV-C) is convex and compact, we minimize the upper-level objective over \mathcal{C} using the Frank–Wolfe algorithm, which proceeds by optimizing a first-order (linear) surrogate of \mathcal{L} at the current iterate.

With the upper-level gradient $\nabla_{\theta}\mathcal{L}(\theta^t)$ computed as in Sec. (V-B.2), each iteration solves the Linear Minimization Oracle (LMO) and updates the candidate solution via a convex combination of the current point and the LMO solution, with the step size chosen by line search. Equivalently, upper level optimization can be viewed as a sequential convex programming scheme: at iteration t , we minimize the linearization model of \mathcal{L} over the convex set \mathcal{C}^t , yielding a convex LMO subproblem.

To enhance numerical stability, each Frank–Wolfe subproblem is restricted to a local trust region around the current iterate θ^t : $\mathcal{C}^t = \mathcal{C} \cap \left\{ \theta \in \mathbb{R}^{n_{\theta}} \mid \|\theta - \theta^t\|_{\infty} \leq \Delta^t \right\}$, where $\Delta^t > 0$ denotes the trust-region radius. This intersection preserves convexity and compactness, ensuring that the linear minimization oracle remains well-posed.

Linear Minimization Oracle: The convex LMO can be solved using two alternative approaches:

(i) *Constraint-based approach:* As \mathcal{L} is nonlinear in θ , in each Frank–Wolfe iteration we linearize the objective at θ^t , which turns the direction-finding step into the optimization of a linear function over the intersection of the positive semidefinite cone with box constraint. This can be efficiently solved by standard SDP solvers. Concretely, we embed \mathbf{s} into symmetric block matrices via fixed symmetric basis matrices $\{F_i\}$ and enforce strict positive definiteness:

$$\begin{aligned} \min_{\mathbf{s}} \quad & \mathbf{s}^{\top} \nabla_{\theta}\mathcal{L}(\theta^t), \\ \text{s.t.} \quad & \underline{\mathbf{s}} \leq \mathbf{s} \leq \bar{\mathbf{s}}, \quad \sum_i s_i F_i \succeq \epsilon I. \end{aligned}$$

(ii) *Parameterization-based approach:* Alternatively, we parameterize each covariance block as $\Sigma_i = L_i L_i^{\top}$, where L_i stands for the lower triangular matrix with positive entries. This enforces Σ_i being positive definite. In this approach, \mathcal{C} reduces to box bounds on the free entries of $\{L_i\}$, and the LMO becomes a lightweight box-constrained linear program in those entries.

Armijo Backtracking Step Size: Fixed, pre-defined decreasing step-size schedules can cause oscillations in \mathcal{L} and slow the decay of $\|\nabla_{\theta}\mathcal{L}(\theta)\|_2$. We therefore use an Armijo backtracking line search to choose an adaptive step size $\gamma^t \in (0, 1]$ along the feasible Frank–Wolfe direction $\mathbf{d}^t \triangleq \mathbf{s}^* - \theta^t$, returned by the LMO. Specifically, for given

parameters $\rho \in (0, 1)$ and $\beta \in (0, 1)$, we initialize $\gamma \leftarrow 1$ and repeatedly shrink $\gamma \leftarrow \beta\gamma$ until the following sufficient-decrease condition is satisfied at γ^t

$$\mathcal{L}(\theta^t + \gamma^t \mathbf{d}^t) \leq \mathcal{L}(\theta^t) + \rho \gamma^t (\nabla_{\theta}\mathcal{L}(\theta^t))^{\top} \mathbf{d}^t.$$

The iterate is then updated as $\theta^{t+1} = \theta^t + \gamma^t(\mathbf{s}^* - \theta^t)$.

Regarding convergence, under good initialization, the iterates typically remain in a locally smooth, nearly convex region, and we often observe an $O(1/t)$ -like decrease in objective value, consistent with classical behavior for smooth convex objectives [39]. Under poorer initialization, the problem is effectively nonconvex, so one generally expects at best a stationarity-type rate of order $O(1/\sqrt{t})$.

B. Cost Definition and Gradient Computation

Given that the full trajectory lies on the Lie-group manifold $\mathcal{M} = \text{SE}(3)^{T+1} \times \mathbb{R}^{n_{\text{aug}}(T+1)}$ with $\mathbf{X} \in \mathcal{M}$, we denote the per-time-step state manifold by $\mathcal{G} = \text{SE}(3) \times \mathbb{R}^{n_{\text{aug}}}$ with $\mathbf{x} \in \mathcal{G}$. The smooth operators on the manifold \mathcal{G} are defined as $\boxplus : \mathcal{G} \times \mathbb{R}^{6+n_{\text{aug}}} \rightarrow \mathcal{G}$, $\boxminus : \mathcal{G} \times \mathcal{G} \rightarrow \mathbb{R}^{6+n_{\text{aug}}}$. We write $\log : \text{SE}(3) \rightarrow \mathfrak{se}(3)$ and $\exp : \mathfrak{se}(3) \rightarrow \text{SE}(3)$ for the Lie logarithm and exponential maps.

1) *Upper level loss function \mathcal{L} :* For each time step in the state trajectory, the upper-level loss function $\mathcal{L} : \mathcal{G} \rightarrow \mathbb{R}$ is defined by measuring trajectory deviation through residual operator that respects the product-manifold geometry. For the $\text{SE}(3)$ component, we use the geodesic distance induced by the left-invariant logarithmic map, while the Euclidean component relies on the Euclidean norm. The cost for the entire trajectory from time index 0 to T is defined as:

$$\begin{aligned} \mathcal{L}(\hat{\mathbf{x}}, \mathbf{x}_{\text{GT}}) &= \frac{1}{2} \sum_{k=0}^T \|\hat{\mathbf{x}}_k \boxminus \mathbf{x}_{\text{GT},k}\|_2^2, \\ &= \frac{1}{2} \sum_{k=0}^T \left(\left\| \left(\log(\mathbf{T}_k \mathbf{T}_{\text{GT},k}^{-1}) \right)^{\vee} \right\|_2^2 + \|\mathbf{e}_k - \mathbf{e}_{\text{GT},k}\|_2^2 \right), \end{aligned} \quad (7)$$

where $(\cdot)^{\vee} : \mathfrak{se}(3) \rightarrow \mathbb{R}^6$ maps a matrix in $\mathfrak{se}(3)$ to the twist.

2) *Gradient:* Each Frank–Wolfe iteration in Sec. (V-A) requires the gradient w.r.t. the calibration parameters, i.e., $\nabla_{\theta}\mathcal{L}(\theta)$, so we begin by formalizing gradients on the product manifold via left-invariant Lie-group Jacobians. The left Jacobian of \mathcal{L} at $\hat{\mathbf{x}} \in \mathcal{G}$ is the linear map, which is:

$$\frac{{}^e D\mathcal{L}(\hat{\mathbf{x}})}{D\hat{\mathbf{x}}} = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}(\hat{\mathbf{x}} \boxplus (\epsilon\tau)) - \mathcal{L}(\hat{\mathbf{x}})}{\epsilon},$$

for any $\tau \in T_e(\mathcal{G}) \cong \mathbb{R}^{6+n_{\text{aug}}}$. For the $\text{SE}(3)$ component, with canonical basis $\{\xi_i\} \subset \mathfrak{se}(3)$,

$$\left[\frac{{}^e D\mathcal{L}(\mathbf{T})}{D\mathbf{T}} \right]_i = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{L}(\exp(\epsilon\xi_i) \cdot \mathbf{T}) - \mathcal{L}(\mathbf{T})}{\epsilon}, \quad \forall i = 1 \dots 6.$$

Similarly, the left Jacobian of the estimated solution $\hat{\mathbf{x}}(\theta)$ w.r.t. θ is ${}^e D\hat{\mathbf{x}}(\theta)/D\theta$ where is defined as

$$\frac{{}^e D\hat{\mathbf{x}}(\theta)}{D\theta} = \lim_{\epsilon \rightarrow 0} \frac{\hat{\mathbf{x}}(\theta + \epsilon\nu) \boxminus \hat{\mathbf{x}}(\theta)}{\epsilon},$$

for any $\nu \in \mathbb{R}^m$. For the $\text{SE}(3)$ component,

$$\left[\frac{{}^e D\mathbf{T}}{D\theta} \right]_j = \lim_{\epsilon \rightarrow 0} \frac{\left(\log(\mathbf{T} \cdot \tilde{\mathbf{T}}_j^{-1}) \right)^{\vee}}{\epsilon}, \quad \forall j = 1 \dots n_{\theta},$$

where \tilde{T}_j is the perturbed pose obtained from $\theta \mapsto \theta + \epsilon e_j$. **Analytical Gradient Computation:** The limit definitions above enable multi-threaded finite-difference approximations on modern CPUs. However, for long horizons and high-dimensional upper-level parameters, numerical differentiation is still prohibitively expensive, even with CPU parallelization. We therefore compute *analytical* gradients.

The upper-level objective (7) aggregates residuals on the product manifold \mathcal{G}^{T+1} . Applying the chain rule yields:

$$\nabla_{\theta} \mathcal{L}(\theta) = \left(\frac{D\mathcal{L}}{D\hat{\mathbf{x}}} \right) \left(\frac{D\hat{\mathbf{x}}}{D\theta} \right) + \frac{\partial \mathcal{L}}{\partial \theta}, \quad (8)$$

where the Jacobians $\frac{D\mathcal{L}}{D\hat{\mathbf{x}}} \in \mathbb{R}^{1 \times r}$ and $\frac{D\hat{\mathbf{x}}}{D\theta} \in \mathbb{R}^{r \times n_{\theta}}$ are stacked over the horizon with $r = (6 + n_{\text{aug}})(T + 1)$.

To differentiate through the lower-level estimator, we analyze the first-order optimality conditions of (4). Let $\mathbf{z} := (\hat{\mathbf{x}}, \mathbf{w}, \mathbf{n}, \boldsymbol{\lambda})$ collect the primal variables and Lagrange multipliers. We define the KKT residual mapping as

$$\mathcal{F}(\mathbf{z}^*(\theta), \theta, \mathbf{y}(\theta), \mathbf{G}(\theta)) = 0, \quad (9)$$

where \mathbf{y} and $\mathbf{G} = \{G_i\}$ are the observation term (2) and the kinematics-to-measurement mapping (3), respectively. Due to the kinematic offset parameters $\theta_{\text{foot},j}$ in (6), both \mathbf{y} and \mathbf{G} depend explicitly on θ . Assuming \mathcal{F} is continuously differentiable in a neighborhood of the solution and the Jacobian $\partial \mathcal{F} / \partial \mathbf{z}$ is nonsingular at $\mathbf{z}^*(\theta)$, the *implicit function theorem* guarantees the existence of a locally unique differentiable solution map $\mathbf{z}^*(\theta)$. Differentiating (9) with respect to θ yields the linear system for trajectory sensitivity:

$$\left(\frac{\partial \mathcal{F}}{\partial \mathbf{z}} \right) \frac{\partial \mathbf{z}^*}{\partial \theta} = - \left(\frac{\partial \mathcal{F}}{\partial \theta} + \frac{\partial \mathcal{F}}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \theta} + \frac{\partial \mathcal{F}}{\partial \mathbf{G}} \frac{\partial \mathbf{G}}{\partial \theta} \right), \quad (10)$$

where $\partial \mathcal{F} / \partial \theta$ denotes the partial derivative with (\mathbf{y}, \mathbf{G}) held fixed. The trajectory sensitivity $\frac{D\hat{\mathbf{x}}}{D\theta}$ is obtained by extracting the block of $\partial \mathbf{z}^* / \partial \theta$ corresponding to $\hat{\mathbf{x}}^*$. In practice, (10) is solved by performing a sparse factorization of the KKT matrix $\partial \mathcal{F} / \partial \mathbf{z}$ with PyPardiso, yielding a negligible per-iteration computational cost at the upper level.

VI. RESULTS

We evaluate our calibration approach on three robots: a bipedal robot STRIDE [40] and two quadrupedal robot Go1 and B1 from Unitree. The data used from STRIDE and Go1 are obtained from simulation in Matlab and MuJoCo environment in their open source repository [40]. We obtained the data of B1 on hardware in the motion capture room with 12 Opti-track Cameras, with a combination of Prime 13 and 22. The bi-level optimization is implemented in Python and C++. The outer loop iteration as LMO is solved by MOSEK; the inner loop FIE is solved with IPOPT [38]. The analytical gradients are generated via CasADi [41] and Pinocchio [42].

We concisely focus on presenting three core results as follows. All the results are consistent across these robots, regardless of their morphology and dynamics, showing successful calibrations of noise covariances and kinematic parameters with significantly improved robot estimates.

Joint Sensor Noise & Kinematics Calibration: For the robot STRIDE, we mainly evaluate the joint calibration on




Fig. 3. The robot STRIDE (a), quadrupedal robot Go1 (b), and B1 (c) are used in the evaluation (Pictures are used with permission).

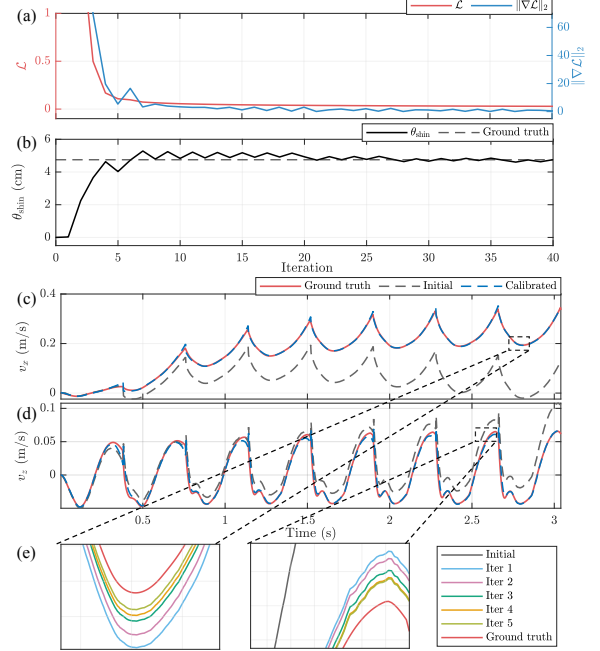


Fig. 4. Calibration results on STRIDE: (a) the convergence of the upper loss function and norm of the gradients, (b) the convergence of kinematics bias w.r.t. ground truth, (c) and (d) the linear velocities of the ground truth, initial estimates, and calibrated estimates.

the sensor noises and robot kinematics since their dynamics processes are deterministic in Matlab. We injected additive white Gaussian noise into the joint encoder angle/velocity measurements, and corrupt the IMU with non-diagonal correlated Gaussian noise, along with an articulated shin length offset to the estimator. Fig. 4 shows the calibration results. The loss successfully reduced through the iterations of the bi-level optimization, and estimated velocity converges to the ground truth; the kinematic parameters converge as well.

Joint Process & Sensor Noise Calibration: For the robot Go1, we evaluate the calibration using data generated by MuJoCo’s time-stepping, multi-contact simulation. Sensor noise is defined in the MJCF/XML and evaluated at every step as a function of the simulated state. The loss and gradient norm successfully converge through the iterations of the bi-level optimization. Additionally, we inject an offset between the torso and ground truth to model motion-capture offset. As optimization proceeds, the estimated offset converges to the pre-defined ground-truth value.

Joint Noise & Kinematics Calibration: Last and most importantly, we calibrate the robot B1 using hardware data. The robot is controlled to walk in the motion-capture room. The torso motion is measured by the motion capture system, taken as ground truth with a constant unknown offset to the base frame position of the torso. The orientation mea-




Fig. 5. Calibration results on Go1: convergences of cost and gradient (top), and kinematic offset (bottom).




Fig. 6. Convergence of the kinematics calibration on B1 robot hardware: (top) the foot offset of the rear right leg and (bottom) the base offset.

surement of the torso is assumed to be unbiased, because the torso and the marked rigid body are well-aligned in the beginning of the experiment. Additionally, we assume all the shins of B1 have kinematic errors from the factory URDF. After we obtained the sensory data along with the ground truth measurements, we applied our calibration to B1. Fig. 7 shows the improved estimates of the torso’s orientation and linear velocity, and Fig. 6 shows the kinematics converge to reasonable values: in our experimental setup, the marked torso location is indeed about centimeters away from the base frame location. Table I shows the quantitative results in terms of RMSE; Fig. 8 shows the new evaluation on a different segment of data using calibrated covariances and kinematics, which still yields highly accurate estimates.

Computational Performance: On an Intel Core Ultra 9 275HX CPU (24C/24T) with a 15s window at 200Hz ($N = 3000$ states, $n_\theta = 69$ parameters), the proposed bi-level calibration typically converges within ~ 15 – 25 outer iterations and takes about 0.5–1.5 hour wall-clock time, excluding the one-time CasADi function generation for analytical derivatives. We believe this runtime can be further reduced by exploiting the log-det barrier structure of the covariance variables and implementing a tailored interior-point solver, which may yield locally quadratic convergence. To sum up, our joint calibration significantly improves the estimates by very large margins.

VII. CONCLUSION AND FUTURE WORK

We presented a bi-level optimization framework for jointly calibrating noise covariances and kinematics in legged robot

TABLE I
RMSEs OF THE ESTIMATES ON B1 ROBOT HARDWARE.

Estimation Metric	Before Calib.	After Calib.	New Eval.
RMSE v [m/s]	0.2658	0.0610	0.0706
RMSE $Euler$ [rad]	0.3457	0.0151	0.0514

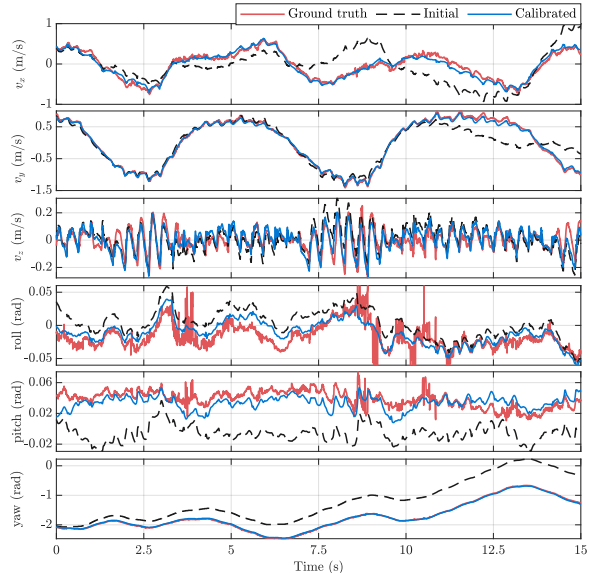


Fig. 7. Calibration results on B1 hardware in terms of linear velocity and torso orientation.

state estimation. By differentiating through a full-information estimator, our method removes manual tuning and improves accuracy and consistency across simulated and real robots. Future work will focus on reducing computational cost, analyzing observability, and extending the approach to dynamics and inertia identification for fully self-calibrating systems.

REFERENCES

- [1] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of basic Engineering*, 1960.
- [2] Y. Bai, B. Yan, C. Zhou, T. Su, and X. Jin, “State of art on state estimation: Kalman filter driven by machine learning,” *Annual Reviews in Control*, vol. 56, p. 100909, 2023.
- [3] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.
- [4] H. A. O. Mohamed, G. Nava, G. L’Erario, S. Traversaro, F. Bergonti, L. Fiorio, P. R. Vanteddu, F. Braghin, and D. Pucci, “Momentum-based extended kalman filter for thrust estimation on flying multibody robots,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 526–533, 2021.
- [5] M. Bloesch, M. Hutter, M. A. Hoepflinger, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, “State estimation for legged robots-consistent fusion of leg kinematics and imu,” *Robotics*, vol. 17, pp. 17–24, 2013.
- [6] N. Rotella, M. Bloesch, L. Righetti, and S. Schaal, “State estimation for a humanoid robot,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 952–958.
- [7] Z. He, S. Teng, T.-Y. Lin, M. Ghaffari, and Y. Gu, “Invariant filtering for full-state estimation of ground robots in non-inertial environments,” *IEEE/ASME Transactions on Mechatronics*, 2025.
- [8] X. Xiong and A. Ames, “3-d underactuated bipedal walking via h-lip based gait synthesis and stepping stabilization,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2405–2425, 2022.
- [9] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, “isam2: Incremental smoothing and mapping using the bayes tree,” *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [10] F. Dellaert, M. Kaess *et al.*, “Factor graphs for robot perception,” *Foundations and Trends® in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.

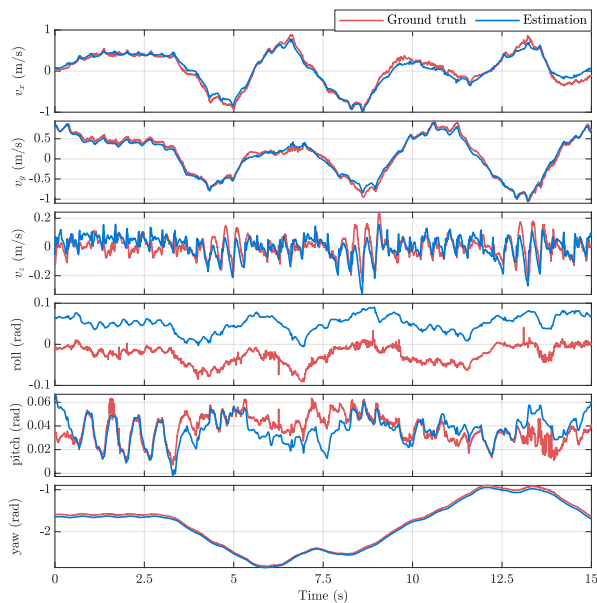


Fig. 8. Evaluation results on B1 hardware in terms of linear velocity and torso orientation.

- [11] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on robotics*, vol. 32, no. 6, pp. 1309–1332, 2017.
- [12] 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*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [13] V. Usenko, N. Demmel, D. Schubert, J. Stückler, and D. Cremers, "Visual-inertial mapping with non-linear factor recovery," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 422–429, 2019.
- [14] R. Hartley, M. Ghaffari, R. M. Eustice, and J. W. Grizzle, "Contact-aided invariant extended kalman filtering for robot state estimation," *The International Journal of Robotics Research*, vol. 39, no. 4, pp. 402–430, 2020.
- [15] J. Kang, Y. Wang, and X. Xiong, "Fast decentralized state estimation for legged robot locomotion via ekf and mhe," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 10914–10921, 2024.
- [16] L. Zhang, D. Sidoti, A. Bienkowski, K. R. Pattipati, Y. Bar-Shalom, and D. L. Kleinman, "On the identification of noise covariances and adaptive kalman filtering: A new look at a 50 year-old problem," *IEEE Access*, vol. 8, pp. 59362–59388, 2020.
- [17] M. Ge and E. C. Kerrigan, "Noise covariance identification for nonlinear systems using expectation maximization and moving horizon estimation," *Automatica*, vol. 77, no. C, p. 336–343, Mar. 2017.
- [18] X. Guo, W. Li, Y. Cui, C. Wang, and Z. Ding, "Expectation-maximization based disturbance identification and velocity tracking for gimbal servo systems with dynamic imbalance," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 71, no. 7, pp. 3357–3367, 2024.
- [19] T. N. Yap and C. R. Shelton, "Simultaneous learning of motion and sensor model parameters for mobile robots," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 2091–2097.
- [20] M. Bloesch, M. Hutter, C. Gehring, M. A. Hoepflinger, and R. Siegwart, "Kinematic batch calibration for legged robots," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 2542–2547.
- [21] A. Censi, A. Franchi, L. Marchionni, and G. Oriolo, "Simultaneous calibration of odometry and sensor parameters for mobile robots," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 475–492, 2013.
- [22] J. N. Wong, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Variational inference with parameter learning applied to vehicle trajectory estimation," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5291–5298, 2020.
- [23] M. Qadri, Z. Manchester, and M. Kaess, "Learning covariances for estimation with constrained bilevel optimization," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 15951–15957.
- [24] B. Amos and J. Z. Kolter, "Optnet: Differentiable optimization as a layer in neural networks," *International Conference on Machine Learning (ICML)*, 2017.
- [25] S. Gould *et al.*, "Implicit differentiation of discrete-time fixed-point iterations," in *International Conference on Machine Learning (ICML)*, 2022.
- [26] S. Teng, M. W. Mueller, and K. Sreenath, "Legged robot state estimation in slippery environments using invariant extended kalman filter with velocity update," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3104–3110.
- [27] S. Yang, "Real-time optimization for robust state estimation and control of legged robots," Ph.D. dissertation, Carnegie Mellon University, 2024.
- [28] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 4304–4311.
- [29] X. Li, Y. Xiao, B. Wang, H. Ren, Y. Zhang, and J. Ji, "Automatic targetless lidar-camera calibration: a survey," *Artificial Intelligence Review*, vol. 56, no. 9, pp. 9949–9987, 2023.
- [30] B. Fu, F. Han, Y. Wang, Y. Jiao, X. Ding, Q. Tan, L. Chen, M. Wang, and R. Xiong, "High-precision multicamera-assisted camera-imu calibration: Theory and method," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–17, 2021.
- [31] B. Wang, Z. Ma, S. Lai, and L. Zhao, "Neural moving horizon estimation for robust flight control," *IEEE Transactions on Robotics*, vol. 40, pp. 639–659, 2023.
- [32] Q. Leboutet, J. Roux, A. Janot, J. R. Guadarrama-Olvera, and G. Cheng, "Inertial parameter identification in robotics: A survey," *Applied Sciences*, vol. 11, no. 9, p. 4303, 2021.
- [33] P. M. Wensing, S. Kim, and J.-J. E. Slotine, "Linear matrix inequalities for physically consistent inertial parameter identification: A statistical perspective on the mass distribution," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 60–67, 2018.
- [34] S. Khorshidi, M. Dawood, B. Nederkorn, M. Bennewitz, and M. Khadiv, "Physically-consistent parameter identification of robots in contact," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 677–683.
- [35] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [36] S. Yang, H. Choset, and Z. Manchester, "Online kinematic calibration for legged robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8178–8185, 2022.
- [37] C. Burgul, W. Lee, P. Geneva, and G. Huang, "Online determination of legged kinematics," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 9043–9049.
- [38] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [39] M. Jaggi, "Revisiting Frank-Wolfe: Projection-free sparse convex optimization," in *Proceedings of the 30th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, S. Dasgupta and D. McAllester, Eds., vol. 28, no. 1. Atlanta, Georgia, USA: PMLR, 17–19 Jun 2013, pp. 427–435.
- [40] Y. Huang, Y. Zeng, and X. Xiong, "Stride: An open-source, low-cost, and versatile bipedal robot platform for research and education," in *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*. IEEE, 2024, pp. 402–409.
- [41] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [42] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, "The Pinocchio C++ library – A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *SII 2019 - International Symposium on System Integrations*, Paris, France, Jan. 2019.