

PRIME: Physically-conconsistent Robotic Inertial and Motion Estimation for Legged and Humanoid Robots

Jiarong Kang¹, Kunzhao Ren¹, Tao Pang, and Xiaobin Xiong^{1,2,†}

¹University of Wisconsin-Madison, ²Current affiliation: Shanghai Innovation Institute

[†]Corresponding author: xiaobin.xiong@sii.edu.cn

Project website: <https://jkangkjr.github.io/PRIME-project>

Abstract—Humanoid and legged robots interact with the environment through intermittent contacts, making accurate motion estimation fundamentally dependent on reasoning about contact dynamics. However, standard sensing pipelines—whether based on onboard proprioception with Extended Kalman Filters (EKFs) or external motion capture systems—recover only kinematics, while contact forces, contact timing, and inertial parameters remain unobserved. As a result, purely kinematic reconstructions often violate rigid-body dynamics, particularly during contact-rich motions. To enable accurate motion estimation from onboard kinematics in real-world deployment, we propose PRIME (Physically-consistent Robotic Inertial and Motion Estimation), a Maximum A Posteriori (MAP) formulation that refines measured kinematics and actuator commands into a dynamically consistent trajectory while jointly estimating frictional contact forces and physically consistent inertial parameters. Our approach incorporates differentiable contact dynamics with smoothed complementarity constraints and an Anitescu-style friction model, yielding a smooth optimization problem that remains tractable across versatile contact transitions. We evaluate PRIME on contact-rich locomotion with quadrupedal robots and the Unitree G1 humanoid, demonstrating improved trajectory consistency and accurate inertial parameter identification. Beyond improving state estimation and feedback control with calibrated inertial parameters, PRIME produces force- and contact-annotated motion reconstructions from real robots in deployment, which can be used to provide high-quality data for downstream learning applications, including large-scale behavior modeling and robot foundation models.

I. INTRODUCTION

Humanoid and legged robots interact with their environment through intermittent contacts, where accurate motion estimation is fundamentally coupled with contact dynamics [1], [2]. During locomotion and manipulation, impacts, stance transitions, and frictional interactions induce forces that directly shape the motion of the robot and the environment therein. Reliable reasoning about such behaviors therefore requires estimates that are not only kinematically accurate but also dynamically consistent. This challenge becomes especially pronounced in contact-rich settings, which are central to real-world deployment of legged and humanoid systems.

Despite significant advances in sensing and estimation, most existing pipelines remain limited to kinematic reconstruction. Standard onboard proprioceptive estimators, typically based on Extended Kalman Filters (EKFs) [3] and various smoothers [4], [5], as well as high-precision external motion capture systems [6], recover only the robot’s configuration and velocity, while contact forces, contact timing, and inertial

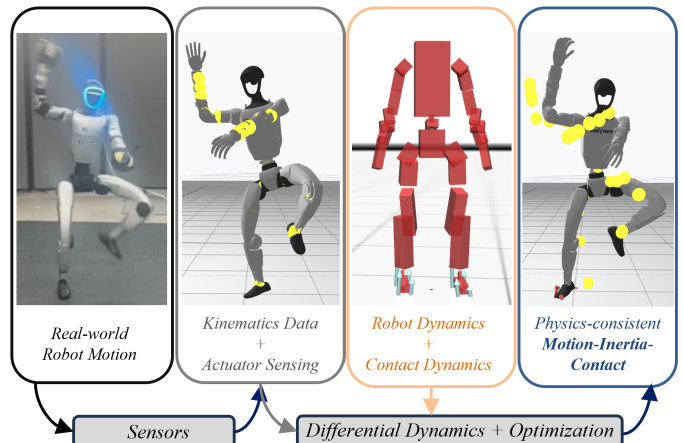


Fig. 1. Overview of PRIME: PRIME reconstructs *physically consistent robot trajectories* and *hardware-matched inertial parameters* from kinematic measurements and onboard actuator sensing using a parameter full-information estimation framework with differentiable contact dynamics.

parameters remain latent. As a result, reconstructed motions often violate rigid-body dynamics, particularly during phases involving intermittent contact events. These inconsistencies in the recovered robot state/data not only hinder downstream tasks such as planning and control [7]–[9], but also prevents data-driven methods such as imitation learning [10]–[12], vision-language-action models [13], [14] and robot foundation models [15], [16] to function in real-world.

To address this problem, we propose PRIME (Physically-consistent Robotic Inertial and Motion Estimation) as illustrated in Fig. 1, a framework for refining measured kinematics into dynamically consistent motion trajectories using IMU, joint sensors and motion capture measurements. PRIME formulates motion and parameter estimation as a Maximum A Posteriori (MAP) problem that jointly estimates state trajectories, frictional contact forces, and physically consistent inertial parameters. Rather than treating contacts as exogenous or implicitly absorbed into noise, PRIME explicitly reasons about contact dynamics, enabling physically grounded reconstruction of contact-rich motions. While our work here focuses on legged and humanoid locomotion, the formulation naturally extends to other contact-rich behaviors, such as robot manipulation, where PRIME could in principle estimate dynamically consistent robot-object motion, contact forces, and inertial parameters under intermittent and unobserved interactions.

PRIME incorporates differentiable contact dynamics with smoothed complementarity constraints [17] and an Anitescu-

style friction model [18], yielding a smooth optimization problem that remains tractable across contact transitions. This formulation enables PRIME to handle switching contact modes robustly without requiring explicit contact sensing or discontinuous dynamics. The resulting problem is solved efficiently as a Parameter Full-Information Estimation [19] task using Feasibility-Driven Differential Dynamic Programming (FDDP) [20], refining noisy or dynamically inconsistent kinematics into trajectories that satisfy rigid-body dynamics and contact constraints, while simultaneously identifying inertial parameters consistent with the observed motion.

We evaluate PRIME on contact-rich locomotion across diverse legged platforms, including the Unitree G1 humanoid. The results show that PRIME produces dynamically consistent motion trajectories and accurate inertial-parameter estimates from kinematic sensing, while reconstructing high-quality contact interactions over long horizons (10 s to 20 s). These improvements enable more reliable state estimation and control, and provide physically grounded motion reconstructions for downstream analysis and learning methods that depend on high-fidelity data. Overall, the results suggest that explicitly reasoning about contact dynamics is a key enabler for accurate motion estimation in real-world robotic systems.

II. RELATED WORK

State estimation and system identification for legged robots have been studied extensively, motivated by the challenges posed by intermittent ground contact, actuator uncertainties, and highly dynamic motions. For state estimation, classical approaches commonly rely on Kalman filters and their variants [3] to fuse onboard sensors, including IMU measurements, joint encoders, and contact sensing, to achieve drift-reduced estimation. More recent optimization-based methods leverage factor-graph formulations or windowed optimization, such as fixed-lag smoothing [4] and moving-horizon estimation [5], and improve accuracy by exploiting a finite history of heterogeneous measurements from both proprioceptive sensors and exteroceptive modalities such as vision and Lidar [21].

However, much of the existing work remains largely kinematics-centric. Contact is typically incorporated either through direct contact sensing [3] or through decoupled inference from actuator torque measurements. Such sensor-level cues usually provide only binary contact indicators, rather than contact-force estimates. Recent efforts have begun to incorporate dynamics-related factors into the estimation, for example by combining momentum terms with contact constraints to recover ground-reaction forces from onboard measurements [22], or by integrating factor-graph estimation with a decoupled momentum observer to estimate torso disturbances [23]. Nevertheless, contact is still often treated in a decoupled manner without enforcing full physical consistency, which can yield kinematic reconstructions that violate rigid-body dynamics and contact physics, thereby limiting the quality of data used in downstream data-driven pipelines.

Building on state estimation, system identification is typically formulated as a least-squares problem, with physical con-

sistency enforced through constrained inertial parameterizations [24]. However, when state estimation is largely restricted to kinematics-based reconstruction, contact interactions are not explicitly modeled and must therefore be addressed separately in the identification process. Existing approaches [25] either rely on high-resolution force/torque sensing at contact interfaces or perform identification within subspaces that reduce sensitivity to unknown external forces [26]. These approaches often assume relatively high-quality data and address noise primarily through pre-processing such as low-pass filtering, which can be insufficient under frequent contact transitions and unmodeled interaction forces.

Recent works have begun to combine state estimation and system identification in a unified framework, allowing both tasks to benefit from physical consistency. For example, [27] formulates a parameter moving-horizon estimation problem for legged robots and jointly estimates trajectories and inertial parameters that are consistent with a measured contact sequence, but the approach remains limited by decoupled contact measurements. To enforce dynamics-consistent contact, [19] proposes a joint parameter and contact estimation method via batch optimization with explicit linear complementarity constraints, yielding physically consistent contact forces and inertial parameters. However, due to the numerical difficulties induced by complementarity constraints, the method is only validated in a single-contact scenario of planar systems. To mitigate these numerical issues, [28] introduces differentiable contact dynamics and demonstrates identification in box sliding experiments. More broadly, contact models that support derivative-based optimization have been widely used for locomotion and manipulation control [17], [29], [30], but remain less explored for joint estimation and identification, particularly in complex, distributed-contact settings such as humanoid locomotion. Another line of work improves contact fidelity by using sampling [31], which performs identification by querying a simulator. However, due to the dimensionality limits of sampling, these methods typically sample only inertial parameters without refining the motion trajectory.

III. MODELING

These gaps in related works motivate our work PRIME. By leveraging smoothed contact dynamics with informative analytic gradients, PRIME robustly reconstructs physically-consistent trajectories, contact interactions, and inertial parameters over long horizons for high-dimensional legged systems. In this section, we present the modeling of dynamics and measurements that form the basis of PRIME.

A. Robot Dynamics Modeling

The legged robot system is described by $\mathbf{x} = [\mathbf{q}^\top \ \mathbf{v}^\top]^\top$, where $\mathbf{q} \in SE(3) \times \mathbb{R}^n$ denotes the generalized coordinates, including the floating-base pose \mathbf{p} and \mathbf{R} , and the joint positions $\boldsymbol{\alpha}$, with n representing the number of articulated joints. The generalized velocity is given by $\mathbf{v} \in \mathfrak{se}(3) \times \mathbb{R}^n$, comprising both the base twist and joint velocities. The

continuous-time dynamics is approximated by a discrete-time system using a semi-implicit Euler integration scheme:

$$\begin{aligned} \mathbf{q}^+ - \mathbf{q} &= \Delta t \mathbf{v}^+, \\ \mathbf{M}(\mathbf{q})(\mathbf{v}^+ - \mathbf{v}) &= \Delta t(\mathbf{B}\mathbf{u} - \mathbf{h}(\mathbf{q}, \mathbf{v}) + \sum_i \mathbf{J}_i(\mathbf{q})^\top \mathbf{f}_i), \end{aligned} \quad (1)$$

where the term $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{(6+n) \times (6+n)}$ represent the inertia matrix, and $\mathbf{h}(\mathbf{q}, \mathbf{v})$ is the bias vector including Coriolis, centrifugal and gravitational effects. Δt is the time discretization, and the superscript \square^+ denotes the states at the next time step. The input matrix $\mathbf{B} \in \mathbb{R}^{(6+n) \times n}$ maps the joint torque vector \mathbf{u} to the generalized force space. $\mathbf{f}_i \in \mathbb{R}^3$ represents the contact force at the i -th contact point, and $\mathbf{J}_i(\mathbf{q}) \in \mathbb{R}^{3 \times (6+n)}$ is the associated contact Jacobian; here, we model collisions between rigid bodies on corresponding points and planes. The contact impulse $\boldsymbol{\lambda}_i = \Delta t \mathbf{f}_i$ is computed implicitly using time-stepping formulation based on a rigid contact model with anisotropic friction [32]. The contact constraints are expressed with a nonlinear friction cone and enforced via a smoothed complementarity condition [17] using a log-barrier relaxation, akin to the treatment in interior-point methods.

B. Inertia Parameterization

The robot dynamics are encoded through the inertial parameters of each link. The Lagrangian dynamics (1) is linear in a suitable inertial parameterization [24]: $\mathbf{Y}(\mathbf{q}, \mathbf{v}, \mathbf{v}^+) \boldsymbol{\pi} = \boldsymbol{\tau}$, where \mathbf{Y} is the classical regressor, and $\boldsymbol{\pi}$ collects the inertial parameters of each link and can be written as:

$$\boldsymbol{\pi} = [m \ h_x \ h_y \ h_z \ I_{xx} \ I_{yy} \ I_{zz} \ I_{xy} \ I_{yz} \ I_{xz}]^\top \in \mathbb{R}^{10}. \quad (2)$$

Here, m is the link mass, $\mathbf{h} = m\mathbf{c}$ is the first mass moment with \mathbf{c} being the center-of-mass (COM) position expressed in the link frame. The remaining terms I_{\square} compose the rotational inertia matrix $\mathbf{I} \in \mathbb{R}^{3 \times 3}$, which is defined at the link-fixed reference point expressed in the body frame. To enforce physical consistency of the inertia parameters, we formulate the pseudo-inertia matrix [24] as:

$$\mathcal{I} = \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{h} \\ \mathbf{h}^\top & m \end{bmatrix}, \quad (3)$$

and require it to be positive definite, where $\boldsymbol{\Sigma} = \frac{1}{2} \text{Tr}(\mathbf{I})\mathbf{1}_3 - \mathbf{I}$. In nonlinear optimization, we employ the unconstrained Log-Cholesky inertial parameterization [33], parameterized by $\boldsymbol{\theta}$:

$$\mathcal{I} = \mathbf{U}\mathbf{U}^\top, \quad \mathbf{U} = e^\alpha \begin{bmatrix} e^{d_1} & s_{12} & s_{13} & t_1 \\ 0 & e^{d_2} & s_{23} & t_2 \\ 0 & 0 & e^{d_3} & t_3 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (4)$$

$$\boldsymbol{\theta} = [\alpha \ d_1 \ d_2 \ d_3 \ s_{12} \ s_{23} \ s_{13} \ t_1 \ t_2 \ t_3]^\top \in \mathbb{R}^{10}. \quad (5)$$

This defines a smooth mapping from $\boldsymbol{\theta}$ to the standard inertial parameters $\boldsymbol{\pi}$, with a closed-form Jacobian $\frac{\partial \boldsymbol{\pi}}{\partial \boldsymbol{\theta}}$ [33].

C. Process Modeling with Contact Dynamics

In time-stepping contact dynamics, physical consistency imposes additional constraints that couple the next-step configuration \mathbf{q}^+ , velocity \mathbf{v}^+ , and contact impulse $\boldsymbol{\lambda}$.

Complementarity Condition: At the next time step $(\cdot)^+$, the signed distance $\phi(\mathbf{q}^+)$ between two bodies and the normal contact impulse λ^n over the time interval must be nonnegative. Moreover, the normal contact force can only be nonzero when the bodies are in contact. These conditions yield the classical complementarity condition:

$$\phi(\mathbf{q}^+) \geq 0 \perp \lambda^n \geq 0. \quad (6)$$

Maximum dissipation: The maximum-dissipation principle selects tangential friction forces that maximize kinetic-energy dissipation. For a single contact i , this can be written as

$$\min_{\boldsymbol{\lambda}_i} (\boldsymbol{\lambda}_i)^\top \mathbf{v}_i^+ \quad (7)$$

$$\text{s.t. } \|\boldsymbol{\lambda}_i^t\|_2 \leq \mu \lambda_i^n, \quad (8)$$

where μ is the coefficient of friction, $\boldsymbol{\lambda}_i^t$ is the tangential component of the contact impulse, and \mathbf{v}_i^+ denotes the next-step relative velocity at contact i , e.g., $\mathbf{v}_i^+ = \mathbf{J}_i(\mathbf{q}^+) \mathbf{v}^+$.

Mathematical Program with Complementarity Constraints: The maximum dissipation principle (7), together with the discrete-time dynamics (1) and non-penetration constraint (6), defines a complementarity-based contact model formulated as a Mathematical Program with Complementarity Constraints (MPCC) [32]. Solving the MPCC yields the next-step robot state, compactly written as:

$$\mathbf{x}^+ \leftarrow \mathbf{Dyn}(\mathbf{x}, \mathbf{u}). \quad (9)$$

Contact Dynamics as Convex Optimization: The complementarity constrained time-stepping dynamics can be reformulated as a convex problem using Anitescu's relaxation [18]. This convexification, however, may admit nonphysical behavior, where tangential sliding can induce a nonzero normal separation (i.e., spurious lift-off) [34]. In locomotion, this artifact is often acceptable because contacts are typically firm and short-lived, and most interactions occur near the sticking regime. In particular, friction can be expressed either as a second-order impulse cone \mathcal{K} :

$$\mathcal{K} = \{(\lambda_i^n, \boldsymbol{\lambda}_i^t) \in \mathbb{R} \times \mathbb{R}^2 : \|\boldsymbol{\lambda}_i^t\|_2 \leq \mu \lambda_i^n\}, \quad (10)$$

or, equivalently, by its dual cone in contact velocity \mathcal{K}^* :

$$\mathcal{K}^* = \left\{ (v_i^n, \mathbf{v}_i^t) \in \mathbb{R} \times \mathbb{R}^2 : \|\mathbf{v}_i^t\|_2 \leq \frac{1}{\mu} v_i^n \right\}. \quad (11)$$

Here $v_i^n = \mathbf{J}_i^n \mathbf{v}^+$ and $\mathbf{v}_i^t = \mathbf{J}_i^t \mathbf{v}^+$ denote the post-impact contact velocity in the normal and tangential directions, respectively. The impulse cone encodes the Coulomb friction bound, while the dual velocity cone encodes the admissible post-impact velocity with unilateral separation. The Anitescu's cone complementarity conditions are then expressed as:

$$(\lambda_i^n, \boldsymbol{\lambda}_i^t) \in \mathcal{K}, \quad (12)$$

$$(v_i^n, \mathbf{v}_i^t) \in \mathcal{K}^*, \quad (13)$$

$$(v_i^n, \mathbf{v}_i^t)^\top (\lambda_i^n, \boldsymbol{\lambda}_i^t) = 0. \quad (14)$$

These conditions are equivalent to imposing the non-penetration condition together with maximum dissipation, and

correspond to the KKT conditions of the following second-order cone program (SOCP):

$$\min_{\mathbf{v}^+} \frac{1}{2} \|\mathbf{v}^+ - \mathbf{v}_{\text{free}}^+\|_{\mathbf{M}(\mathbf{q})}^2, \quad (15)$$

$$\text{s.t.} \quad \left(\frac{\phi_i}{\Delta t} + \mathbf{J}_i^n \mathbf{v}^+ \right)^2 \geq \mu^2 \|\mathbf{J}_i^t \mathbf{v}^+\|_2^2, \quad (16)$$

where $\mathbf{v}_{\text{free}}^+$ denotes the contact-free integrated velocity. This SOCP can be efficiently solved by interior-point methods [35]. **Smoothing of contact dynamics:** Although Anitescu's relaxation yields a convex formulation, the resulting contact dynamics remains only piecewise smooth. This induces discontinuous (or ill-defined) gradients, which hinder the direct use of gradient-based optimization methods. Prior work has addressed this issue in different ways: [8] performs local linearization of the contact dynamics around a pre-planned trajectory, while [36] modifies the discontinuous analytic gradients by introducing relaxed conditions. More recently, [17] proposed a log-barrier relaxation based smoothing approach of the second-order cone constraints in (15), which converts the SOCP into an unconstrained convex optimization problem:

$$\min_{\mathbf{v}^+} \frac{1}{2} \|\mathbf{v}^+ - \mathbf{v}_{\text{free}}^+\|_{\mathbf{M}(\mathbf{q})}^2 - \sum_i \frac{1}{\kappa} \log \left(\frac{(\phi_i/\Delta t + \mathbf{J}_i^n \mathbf{v}^+)^2}{\mu^2} - \|\mathbf{J}_i^t \mathbf{v}^+\|_2^2 \right), \quad (17)$$

with κ being the weight of the log-barrier relaxation. This smoothing yields well-behaved and informative gradients around transitions of contact modes, enabling the optimizer to explore motions that involve changes in the contact modes. In this paper, the smoothed contact dynamics are solved using a customized Newton solver for fast computation, warm-started with a feasibility initializer that projects the previous time-step solution onto the normal feasibility constraints.

Approximation of Latent Impulses: Although the smoothed contact dynamics are cast as an unconstrained optimization problem, the contact impulses—interpretable as the Lagrange multipliers of the original friction-cone (SOC) constraints (16)—can still be recovered from the first-order optimality conditions. In particular, for each contact indexed by i , the normal and tangential impulses can be approximated as:

$$\lambda_i^n = \frac{2(\phi_i/\Delta t + \mathbf{J}_i^n \mathbf{v}^+)}{\mu^2 \kappa s_i}, \quad (18)$$

$$\lambda_i^t = -\frac{2\mathbf{J}_i^t \mathbf{v}^+}{\kappa s_i}, \quad (19)$$

where:

$$s_i = \frac{(\phi_i/\Delta t + \mathbf{J}_i^n \mathbf{v}^+)^2}{\mu^2} - \|\mathbf{J}_i^t \mathbf{v}^+\|_2^2. \quad (20)$$

As a result, the impulse as a latent variable is computed from the estimated trajectory that serves as the impulse estimation.

We will denote the smoothed contact dynamics by $\mathcal{F}(\cdot)$ in what follows. For the estimation task, the control input \mathbf{u} is assumed to be known from joint torque sensors. We introduce

an additive process disturbance $\delta^{\mathbf{x}} \in \mathbb{R}^{n+6}$, which includes a floating-base component that is penalized more heavily, and model it as zero-mean Gaussian noise, $\delta^{\mathbf{x}} \sim \mathcal{N}(\mathbf{0}, \mathbf{Cov}_{\mathbf{x}})$, where $\mathbf{Cov}_{\mathbf{x}}$ denotes the covariance matrix. With this, the process dynamics of the robot are then written as:

$$\mathbf{x}^+ \leftarrow \mathcal{F}(\mathbf{x}, \mathbf{u}) \oplus \delta^{\mathbf{x}}, \quad (21)$$

where \oplus represents additive stochastic disturbances in the state propagation. (21) models contact dynamics while incorporating process uncertainty, such as actuation noise and unmodeled effects, which are assumed to follow the Gaussian distribution.

D. Measurement Modeling

To complement the increased fidelity of the contact dynamics, we incorporate all available measurements at the kinematic and actuation levels that are either onboard or readily accessible. The multi-sensor measurement vector \mathbf{y} is modeled as a function of the state \mathbf{x} and the control input \mathbf{u} :

$$\mathbf{y} = \mathbf{Meas}(\mathbf{x}, \mathbf{u}) \oplus \delta^{\mathbf{y}}, \quad (22)$$

where $\delta^{\mathbf{y}} \sim \mathcal{N}(\mathbf{0}, \mathbf{Cov}_{\mathbf{y}})$ represents measurement uncertainty, modeled as zero-mean Gaussian noises. The measurement model fuses heterogeneous sensing modalities, including encoder-measured joint positions and velocities, floating-base measurements from the motion-capture system, and actuator torque measurements that are commonly available on modern legged robots, e.g., via Quasi-Direct-Drive actuation. The measurement model does not require direct contact force or torque measurements; these quantities are inferred through the contact dynamics.

IV. PRIME VIA OPTIMIZATION-BASED ESTIMATION

Now we present the problem formulation of PRIME via optimization, and then we present our solution method that solves the optimization problem robustly and efficiently.

A. Optimization-based Estimation

Given a history of sensor measurements, optimal state estimation can be cast as a Maximum A Posteriori (MAP) problem. Full Information Estimation (FIE) solves this MAP problem over the entire time window by minimizing the corresponding negative log-likelihood [37], thereby producing optimal estimates that leverage all available measurements. The formulation incorporates the system dynamics (Section II.C), measurement model (Section II.D), and physical constraints:

$$\begin{aligned} & \min_{\mathbf{x}_{[0,T]}, \delta_{[0,T]}} \Gamma(\mathbf{x}_0) \\ & + \sum_{k=0}^{T-1} \|\delta_k^{\mathbf{x}}\|_{\mathbf{Cov}_{\mathbf{x}}^{-1}}^2 + \sum_{k=0}^{T-1} \|\delta_k^{\mathbf{y}}\|_{\mathbf{Cov}_{\mathbf{y}}^{-1}}^2 \end{aligned} \quad (\text{FIE})$$

$$\text{s.t.} \quad \mathbf{x}_{k+1} \leftarrow \mathcal{F}(\mathbf{x}_k, \mathbf{u}_k) \oplus \delta_k^{\mathbf{x}}, \quad \forall k \in \{0, \dots, T-1\}$$

$$\mathbf{y}_k = \mathbf{Meas}(\mathbf{x}_k) \oplus \delta_k^{\mathbf{y}}, \quad \forall k \in \{0, \dots, T\},$$

$$\mathbf{0} \leq \mathbf{Constr}(\mathbf{x}_k), \quad \forall k \in \{0, \dots, T\},$$

in which dynamics and measurement uncertainties are minimized. Here, the subscript \square_k denotes the time index, and

$\square_{[0,T]}$ denotes the trajectories from the start of time 0 to the end of time T . $\Gamma(\cdot)$ denotes the prior term that anchors the initial state of the trajectory, and δ_k^x and δ_k^y represent process and measurement noise, respectively, modeled as zero-mean Gaussian random variables with covariances \mathbf{Cov}_x and \mathbf{Cov}_y .

B. Parameter and State Estimation via FIE

During the estimation process, system parameters such as linkage inertia can be estimated, in other words, identified jointly with the robot state. Accordingly, the parameter-dependent dynamics in (21) are written as

$$\mathbf{x}^+ \leftarrow \mathcal{F}(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}) \oplus \boldsymbol{\delta}^x, \quad (23)$$

where $\boldsymbol{\theta}$ denotes the inertial-parameter representation introduced in Section II.B, including a subset of link inertial parameters targeted for identification. Assuming a Gaussian prior on the initial parameter estimate, $\boldsymbol{\theta} \sim \mathcal{N}(\hat{\boldsymbol{\theta}}, \mathbf{Cov}_\theta)$, we reformulate the original estimation problem in (FIE) to jointly estimate states and parameters, yielding the Parameter Full-Information Estimation (PFIE) formulation of PRIME:

$$\begin{aligned} \min_{\mathbf{x}_{[0,T]}, \boldsymbol{\delta}_{[0,T]}, \boldsymbol{\theta}} \quad & \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_{\mathbf{Cov}_\theta^{-1}}^2 + \Gamma(\mathbf{x}_0) \\ & + \sum_{k=0}^{T-1} \|\boldsymbol{\delta}_k^x\|_{\mathbf{Cov}_x^{-1}}^2 + \sum_{k=0}^T \|\boldsymbol{\delta}_k^y\|_{\mathbf{Cov}_y^{-1}}^2 \end{aligned} \quad (\text{PFIE})$$

$$\begin{aligned} \text{s.t.} \quad & \mathbf{x}_{k+1} \leftarrow \mathcal{F}(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\theta}) \oplus \boldsymbol{\delta}_k^x, \quad \forall k \in \{0, \dots, T-1\} \\ & \mathbf{y}_k = \mathbf{Meas}(\mathbf{x}_k) \oplus \boldsymbol{\delta}_k^y, \quad \forall k \in \{0, \dots, T\}. \end{aligned}$$

By adopting the Log-Cholesky inertia parameterization, the inertial parameters remain feasible without explicitly imposing constraints, i.e., $\mathbf{0} \leq \mathbf{Constr}(\boldsymbol{\theta})$.

C. Optimal Estimation via DDP

Both the (PFIE) and (FIE) can be solved using Differential Dynamic Programming (DDP), by leveraging the differential properties of the contact dynamics (17), even in the presence of complementarity constraints introduced by contact dynamics. We illustrate the derivations for the (FIE) problem first. By analyzing the Bellman equation of (FIE), the value function can be recursively expressed as:

$$V(\mathbf{x}) = \min_{\boldsymbol{\delta}} Q(\mathbf{x}, \boldsymbol{\delta}), \quad (24)$$

$$Q(\mathbf{x}, \boldsymbol{\delta}) = V^+(\mathbf{x}^+) + L(\mathbf{x}, \boldsymbol{\delta}), \quad (25)$$

where $V(\cdot)$, $Q(\cdot)$, and $L(\cdot)$ are the value function, the local Q-function, and the stage cost, respectively. In DDP, the model uncertainty $\boldsymbol{\delta}$ is computed to minimize the local second-order approximation of $Q(\cdot)$ at the current estimated trajectory ($\hat{\cdot}$):

$$Q(\mathbf{x}, \boldsymbol{\delta}) \approx Q(\hat{\mathbf{x}}, \hat{\boldsymbol{\delta}}) + \Delta Q, \quad (26)$$

$$\Delta Q = \frac{1}{2} \begin{bmatrix} 1 \\ \delta \mathbf{x} \\ \delta \boldsymbol{\delta} \end{bmatrix}^\top \begin{bmatrix} 0 & \nabla_{\mathbf{x}} Q^\top & \nabla_{\boldsymbol{\delta}} Q^\top \\ \nabla_{\mathbf{x}} Q & \nabla_{\mathbf{x}}^2 Q & \nabla_{\mathbf{x}\boldsymbol{\delta}} Q \\ \nabla_{\boldsymbol{\delta}} Q & \nabla_{\boldsymbol{\delta}\mathbf{x}} Q & \nabla_{\boldsymbol{\delta}}^2 Q \end{bmatrix} \begin{bmatrix} 1 \\ \delta \mathbf{x} \\ \delta \boldsymbol{\delta} \end{bmatrix}. \quad (27)$$

Minimizing ΔQ w.r.t. $\delta \boldsymbol{\delta}$ yields the optimal uncertainty:

$$\delta \boldsymbol{\delta}^* = \mathbf{K} \cdot \delta \mathbf{x} + \alpha \mathbf{k}, \quad (28)$$

Algorithm 1 PRIME via DDP

Require: Initial trajectory $\tilde{\mathbf{x}}_t$, $\tilde{\boldsymbol{\delta}}_t$, and parameters $\tilde{\boldsymbol{\theta}}$
while not converged **do**

Backward pass:

$$\nabla_{\tilde{\mathbf{x}}} V := \nabla_{\tilde{\mathbf{x}}} L_N, \quad \nabla_{\tilde{\mathbf{x}}}^2 V := \nabla_{\tilde{\mathbf{x}}}^2 L_N$$

for $i = N - 1 \rightarrow 0$ **do**

$$A = \partial_{\tilde{\mathbf{x}}} \mathcal{F}(\tilde{\mathbf{x}}, \tilde{\boldsymbol{\delta}}, \tilde{\boldsymbol{\theta}}),$$

$$B = \partial_{\tilde{\boldsymbol{\delta}}} \mathcal{F}(\tilde{\mathbf{x}}, \tilde{\boldsymbol{\delta}}, \tilde{\boldsymbol{\theta}}),$$

$$\nabla_{\tilde{\mathbf{x}}} Q = \nabla_{\tilde{\mathbf{x}}} L + A^\top \nabla_{\tilde{\mathbf{x}}} V,$$

$$\nabla_{\tilde{\boldsymbol{\delta}}} Q = \nabla_{\tilde{\boldsymbol{\delta}}} L + B^\top \nabla_{\tilde{\mathbf{x}}} V,$$

$$\nabla_{\tilde{\mathbf{x}}}^2 Q = \nabla_{\tilde{\mathbf{x}}}^2 L + A^\top (\nabla_{\tilde{\mathbf{x}}}^2 V) A,$$

$$\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q = \nabla_{\tilde{\boldsymbol{\delta}}}^2 L + B^\top (\nabla_{\tilde{\mathbf{x}}}^2 V) B,$$

$$\nabla_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{x}}} Q = \nabla_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{x}}} L + B^\top (\nabla_{\tilde{\mathbf{x}}}^2 V) A.$$

Choose $\eta_i > 0$ s.t. $\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q := \nabla_{\tilde{\boldsymbol{\delta}}}^2 Q + \eta_i I > 0$.
 Compute:

$$k_i = -\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q^{-1} \nabla_{\tilde{\boldsymbol{\delta}}} Q$$

$$K_i = -\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q^{-1} \nabla_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{x}}} Q,$$

$$\nabla_{\tilde{\mathbf{x}}} V = \nabla_{\tilde{\mathbf{x}}} Q + K_i^\top \nabla_{\tilde{\boldsymbol{\delta}}} Q,$$

$$\nabla_{\tilde{\mathbf{x}}}^2 V = \nabla_{\tilde{\mathbf{x}}}^2 Q + K_i^\top \nabla_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{x}}} Q.$$

end for

Forward pass:

Choose $\eta > 0$, s.t. $\nabla_{\tilde{\mathbf{x}}}^2 V := \nabla_{\tilde{\mathbf{x}}}^2 V + \eta I > 0$.

Compute:

$$\delta \tilde{\mathbf{x}}_0^* = -[\nabla_{\tilde{\mathbf{x}}}^2 V]^{-1} \nabla_{\tilde{\mathbf{x}}} V,$$

$$\tilde{\mathbf{x}}_0' = \tilde{\mathbf{x}}_0 + \alpha \delta \tilde{\mathbf{x}}_0^*, \quad V' := 0.$$

for $i = 0 \rightarrow N - 1$ **do**

$$\boldsymbol{\delta}'_i = \boldsymbol{\delta}_i + \alpha k_i + K_i \delta \tilde{\mathbf{x}}_i,$$

$$\tilde{\mathbf{x}}'_{i+1} = \begin{bmatrix} \mathcal{F}(\tilde{\mathbf{x}}'_i, \boldsymbol{\delta}'_i, \boldsymbol{\theta}') \\ \boldsymbol{\theta}' \end{bmatrix},$$

$$\delta \tilde{\mathbf{x}}'_{i+1} = \tilde{\mathbf{x}}'_{i+1} - \begin{bmatrix} \tilde{\mathbf{x}}_i \\ \boldsymbol{\theta} \end{bmatrix},$$

$$V' = V' + L_i(\tilde{\mathbf{x}}'_i, \boldsymbol{\delta}'_i).$$

end for

$$V' = V' + L_N(\tilde{\mathbf{x}}'_N).$$

end while

return optimal estimates of the physically-consistent trajectories \mathbf{x}_t and parameters $\boldsymbol{\theta}$

$$\mathbf{K} = -\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q^{-1} \nabla_{\tilde{\boldsymbol{\delta}}\tilde{\mathbf{x}}} Q, \quad \mathbf{k} = -\nabla_{\tilde{\boldsymbol{\delta}}}^2 Q^{-1} \nabla_{\tilde{\boldsymbol{\delta}}} Q,$$

where α is a chosen step size via Armijo backtracking line search. When applying DDP to state estimation, the initial state is assumed to follow a prior distribution $\Gamma(\cdot)$, as indicated in (FIE). During DDP, the optimal perturbation to the initial state

$\delta \mathbf{x}_0$ is generally non-zero and is computed as:

$$\delta \mathbf{x}_0^* = - [\nabla_{\mathbf{x}_0}^2 V]^{-1} \nabla_{\mathbf{x}_0} V. \quad (29)$$

Remark: For parameter and state estimation in (PFIE), the robot dynamics is augmented with the static parameter $\boldsymbol{\theta}$ dynamics within the horizon, and is similarly updated based on the optimal condition of the initial node:

$$\bar{\mathbf{x}}^+ = \begin{bmatrix} \mathbf{x}^+ \\ \boldsymbol{\theta}^+ \end{bmatrix} \leftarrow \begin{bmatrix} \mathcal{F}(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}) \oplus \delta \mathbf{x} \\ \boldsymbol{\theta} \end{bmatrix}. \quad (30)$$

The implementation of our DDP is detailed in Algorithm 1.

D. Multiple Shooting Differential Dynamic Programming

To robustly solve the (PFIE) problem, we employ a multiple shooting variant of DDP, Feasibility-driven differential dynamic programming (FDDP). Unlike classical DDP, which typically assumes a dynamically feasible rollout, FDDP allows initialization and iteration with state-control trajectories that violate the dynamics. It achieves this by modifying the backward pass to account for defect variables and by permitting controlled dynamics infeasibility during early forward passes, progressively driving these defects toward zero across iterations. This ability to handle dynamics infeasibility is particularly valuable for contact estimation, where high-fidelity contact models can be numerically fragile—even with smoothing—and may otherwise hinder convergence. A detailed algorithmic description can be found in [20].

E. Analytic Gradient of Contact Dynamics

As (F)DDPs rely on local approximations of the dynamics, the relationship among the perturbations $\delta \bar{\mathbf{x}}^+$, $\delta \bar{\mathbf{x}}$, and $\delta \boldsymbol{\delta}$ is governed by the infinitesimal dynamics induced by (21), i.e.,

$$\delta \bar{\mathbf{x}}^+ = \mathbf{A} \delta \bar{\mathbf{x}} + \mathbf{B} \delta \boldsymbol{\delta}, \quad (31)$$

where $\mathbf{A} = \partial_{\bar{\mathbf{x}}} \mathcal{F}(\bar{\mathbf{x}}, \boldsymbol{\delta})$ and $\mathbf{B} = \partial_{\boldsymbol{\delta}} \mathcal{F}(\bar{\mathbf{x}}, \boldsymbol{\delta})$ are the linearized dynamics matrices used in Algorithm 1 iterations. Since the smoothed contact dynamics are posed as an unconstrained convex optimization, these derivatives can be obtained analytically via sensitivity analysis and the implicit function theorem.

As an example, consider differentiation with respect to the parameter vector $\boldsymbol{\theta}$. The first-order optimality condition of (17) can be written as:

$$\nabla_{\mathbf{v}^+} c(\mathbf{v}^+, \boldsymbol{\theta}) \triangleq \mathbf{g}(\mathbf{v}^+, \boldsymbol{\theta}) = \mathbf{0}, \quad (32)$$

where $c(\cdot)$ denotes the objective and $\mathbf{g}(\cdot)$ denotes its first-order derivative with respect to \mathbf{v}^+ . Linearizing the optimality condition to first order yields:

$$\frac{\partial \mathbf{g}(\mathbf{v}^+, \boldsymbol{\theta})}{\partial \mathbf{v}^+} \delta \mathbf{v}^+ + \frac{\partial \mathbf{g}(\mathbf{v}^+, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \delta \boldsymbol{\theta} = \mathbf{0}, \quad (33)$$

which gives the sensitivity of the post-impact generalized velocity with respect to the parameters:

$$\frac{\partial \mathbf{v}^+}{\partial \boldsymbol{\theta}} = - \left(\frac{\partial \mathbf{g}(\mathbf{v}^+, \boldsymbol{\theta})}{\partial \mathbf{v}^+} \right)^{-1} \frac{\partial \mathbf{g}(\mathbf{v}^+, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}. \quad (34)$$

For further discussion of the properties of this gradient, see [17] and the Supplementary Material. These analytic gradients are the key for efficient computation to solve the (PFIE).

V. RESULTS

We evaluate PRIME in both simulation and hardware of several legged robots to simultaneously estimate their motion and uncertain inertia properties. First, we develop a custom MATLAB simulator to validate PRIME on a hopping robot. In this simulator, planar contact dynamics are formulated and solved as a linear complementarity problem (LCP) without smoothing. The required derivatives are obtained using the MATLAB Symbolic Math Toolbox, and the resulting (PFIE) is solved using an open-source DDP implementation [38].

We then utilize off-the-shelf high-fidelity MuJoCo simulation to demonstrate that PRIME generalizes to more complex and high-dimensional systems, including both quadrupedal Go2 and humanoid robot G1. For the hardware experiments, we obtained the data in the motion capture system that has 12 Opti-track cameras with a combination of Prime 13 and Prime 22. In these experiments, the robots are controlled using open-source reinforcement-learning policies [39].

PRIME is fully implemented in C++ for fast computation. Analytic kinematic and dynamic derivatives are computed with Pinocchio [40]; notably, Pinocchio also provides the kinematic Hessians required for our analytic differentiation. The smoothed contact dynamics are solved using a customized Newton solver, and the resulting (PFIE) problem is solved using FDDP as implemented in Crocoddyl [41].

A. Hopper Experiments

In the hopper experiments, we configure PFIE with a 2.5 s horizon and a time discretization of 0.025 s. The dataset is generated using a Contact-Implicit MPC controller [36]. Results on the PRIME hopping robot, as in Fig. 2, show that the proposed method is able to filter the noise-corrupted kinematic and actuator measurements to reconstruct the contact sequence and identify inertial parameters, starting from a biased initial inertia model. The iteration-wise convergence of the inertia estimates in Fig. 2 further indicates consistent convergence toward the ground-truth values.

B. PRIME for Quadrupedal Robots

In the experiments on quadrupedal robots Go2, we jointly estimate robot motion and identify uncertain robot inertia in both simulation and hardware. Here, we assume the torso inertia is uncertain due to added payload onto the torso. In MuJoCo, we conduct two trials: (a) adding a 3 kg payload to




Fig. 2. (a) Trajectories of the Hopper. (b) Convergence of the inertial parameters over iterations. The estimated inertial parameters converge to the ground-truth value, and the state trajectories align with the true trajectories with corrected contact modes.




Fig. 3. Quadrupedal robot Go2 motion estimation and inertia identification experiments with a 4.6 kg payload attached beneath the torso. (a) A 10 s segment comprising multiple locomotion behaviors, including torso rotations and forward hopping. (b) PRIME reconstructs the motion trajectory with dynamically consistent contact generated without using any contact measurements. The yellow spheres indicate the link poses from the raw kinematics. (c) The raw kinematics are physically inconsistent—showing non-physical contact penetrations—due to motion-capture orientation bias and jitters during locomotion.




Fig. 4. Humanoid G1 motion estimation experiment. (a) Simulation results showing the reconstructed trajectory closely aligned with the ground truth. (b,c) Real-robot results comparing raw kinematics (left) and PRIME reconstruction (right). Yellow spheres overlaid on the reconstructed motion indicate link positions from the raw kinematics, and red arrows denote the reconstructed contact forces at the foot corners. Despite noisy and contact-inconsistent kinematic measurements during the dancing sequence, PRIME reconstructs dynamically consistent motion with spatially distributed planar foot contact.

the torso, and (b) shifting the torso center of mass downward by 0.1 m. In hardware, we attach two 2.3 kg barbell plates to the belly of the robot and then execute dynamic motions on the robot, including pronounced torso rotations and consecutive hopping with lateral velocity. For both simulation and hardware, PRIME is configured with a 10 s horizon and a 0.01 s time discretization (1000 samples). As summarized in Table I, the results show that the imposed inertial changes are accurately reflected in the estimated parameters; additionally, the reconstructed trajectories and contact forces closely match the ground truth. In the real-world experiments, PRIME similarly recovers the payload effect, estimating an added mass of approximately 4.8 kg and a downward shift of the torso COM consistent with the placement of the plates. The reconstructed state trajectory also remains consistent with the real experi-

ments, even though the raw motion-capture kinematics exhibit orientation bias and jitter, which are clearly not contact- or dynamics-feasible—especially during the consecutive hopping behaviors. Despite these nonphysical measurement artifacts, PRIME still provides high-quality contact-consistent reconstructions as shown in Fig. 3. Moreover, PRIME is quite efficient, converging within 200 s for a horizon of 1000 steps on a PC equipped with an Intel Core i9-13900 CPU.

C. PRIME for Humanoids

In the humanoid experiments, we collect diverse Unitree G1 locomotion behaviors, including walking, running, and dancing. Ground-truth contact forces for a dancing sequence are collected using a Bertec-4060 six-axis force plate (1000 Hz, ± 0.4 N resolution). For all humanoid trials, PRIME uses a 15 s

TABLE I
IDENTIFICATION OF TORSO INERTIA AND COM ON Go2.

Parameter	Original Values	Simulation m (+3 kg)	Simulation c_z (-0.1 m)	Real-world m (+4.6 kg)
m [kg]	6.927	9.975	6.865	11.74
c_x [m]	0.022	0.034	0.242	0.037
c_y [m]	0.000	0.138	0.089	0.004
c_z [m]	-0.005	-0.012	-0.105	-0.013
I_{xx}	0.025	0.033	0.024	0.014
I_{yy}	0.098	0.138	0.089	0.274
I_{zz}	0.108	0.150	0.097	0.274

TABLE II
IDENTIFICATION OF TORSO INERTIA AND COM ON G1.

Parameter	Original Values	Real-world Total weight (+2.914kg)
m [kg]	9.60	13.02
c_x [m]	0.00332	0.03186
c_y [m]	0.00026	-0.01649
c_z [m]	0.1798	0.1775
I_{xx}	0.1241	0.2669
I_{yy}	0.1121	0.3075
I_{zz}	0.0327	0.1097

horizon and a time discretization of 0.01 s (1500 samples), sufficient to capture diverse motions and multiple contact transitions. PRIME solves each problem in approximately 400 s on the same PC. Planar foot contact is approximated by a four-point contact model at the corners of each foot, consistent with the rigid-foot design of the G1.

Although the ground-truth inertial parameters of the G1 are unavailable, our experiments reveal a measurable discrepancy between the nominal model and the physical robot. The scale-measured total robot mass is 38.029 kg, whereas the nominal model mass is 35.115 kg, yielding a difference of 2.914 kg. This discrepancy may be attributed to the battery, which weighs 2.496 kg. As shown in Table II, PRIME estimates an additional mass of approximately 3.4 kg, while preserving a similar COM location and increasing the rotational inertia accordingly. The corresponding equivalent inertia boxes further suggest a reasonable localized load distribution, as illustrated in Fig. 5. This identification result is also reflected in the contact-force comparison. With system identification enabled, PRIME achieves a lower cost in (FIE) and produces force estimates that more closely match the force-plate ground truth, with smaller measurement residuals and lower RMSEs, as shown in Fig. 6 and Table III. These results indicate that the improved inertial model leads to more accurate force reconstruction. PRIME further reconstructs smooth planar contact transitions, including heel-to-toe rolling and diverse contact patterns during dancing on a complex humanoid platform as illustrated in Fig. 4. To the best of our knowledge, this level of contact-force reconstruction has not been previously demonstrated for such humanoid motions. Additional analysis and comparisons with alternative contact-modeling choices are provided in the Supplementary Material.

VI. CONCLUSION AND FUTURE WORK

This paper presented PRIME, an optimization-based estimation framework that refines measured kinematics into dynam-

TABLE III
RMSEs AND FIE COSTS ON G1 HARDWARE EXPERIMENTS.

Estimation Metric	With ID	W/O ID
RMSE $_F$ [N]	24.486	26.141
Cost [$\times 10^3$]	1.016	1.880




Fig. 5. Equivalent inertia-box of hardware experiments (a) Go2; (b) G1.




Fig. 6. Force comparison for the G1 humanoid (left foot) over a dancing motion. The identification module estimates inertial discrepancies and corrects the reconstructed contact forces, improving accuracy against the force-plate ground truth.

ically consistent motion trajectories by explicitly reasoning about contact dynamics and inertial properties. By formulating motion and parameter estimation as a Maximum A Posteriori problem with differentiable contact dynamics, PRIME jointly reconstructs state trajectories, frictional contact forces, and physically consistent inertial parameters. Experimental results on diverse legged robot hardware demonstrate that PRIME produces contact-consistent motion reconstructions and accurate inertial identification over long horizons, even when raw kinematic measurements are noisy or physically inconsistent. These results highlight the importance of enforcing physical consistency for reliable motion estimation in contact-rich robotic systems.

While this work focuses on legged and humanoid locomotion, the underlying formulation of PRIME is not specific to locomotion and naturally extends to other contact-rich behaviors, including robotic manipulation and loco-manipulation, where intermittent contacts and unobserved interaction forces pose similar challenges. An important direction for future work is to study the observability and identifiability of motion, contact forces, and inertial parameters under limited sensory information, and to understand how task structure and contact excitation influence estimation performance. In addition, by transforming raw robot logs into force- and contact-consistent motion reconstructions, PRIME suggests a pathway for enriching real-robot datasets with physically grounded supervision, which could support downstream data-driven approaches, including recent vision-language-action and large-scale behavior models.

REFERENCES

- [1] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. Del Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2023.
- [2] J. Carpentier and N. Mansard, "Multicontact locomotion of legged robots," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1441–1460, 2018.
- [3] R. Hartley, M. G. Jadidi, J. W. Grizzle, and R. M. Eustice, "Contact-aided invariant extended kalman filtering for legged robot state estimation," *arXiv preprint arXiv:1805.10410*, 2018.
- [4] Z. Yoon, J.-H. Kim, and H.-W. Park, "Invariant smoother for legged robot state estimation with dynamic contact event information," *IEEE Transactions on Robotics*, vol. 40, pp. 193–212, 2023.
- [5] 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.
- [6] Y. Ze, S. Zhao, W. Wang, A. Kanazawa, R. Duan, P. Abbeel, G. Shi, J. Wu, and C. K. Liu, "Twist2: Scalable, portable, and holistic humanoid data collection system," *arXiv preprint arXiv:2511.02832*, 2025.
- [7] W. Yang and M. Posa, "Impact-invariant control: Maximizing control authority during impacts," *Autonomous Robots*, vol. 49, no. 4, pp. 1–24, 2025.
- [8] S. Le Cleac'h, T. A. Howell, S. Yang, C.-Y. Lee, J. Zhang, A. Bishop, M. Schwager, and Z. Manchester, "Fast contact-implicit model predictive control," *IEEE Transactions on Robotics*, vol. 40, pp. 1617–1629, 2024.
- [9] V. Kurtz, A. Castro, A. Ö. Önel, and H. Lin, "Inverse dynamics trajectory optimization for contact-implicit model predictive control," *The International Journal of Robotics Research*, vol. 45, no. 1, pp. 23–40, 2026.
- [10] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, vol. 44, no. 10-11, pp. 1684–1704, 2025.
- [11] T. Tsuji, Y. Kato, G. Solak, H. Zhang, T. Petrič, F. Nori, and A. Ajoudani, "A survey on imitation learning for contact-rich tasks in robotics," *arXiv preprint arXiv:2506.13498*, 2025.
- [12] J. Sun, Q. Zhang, Y. Duan, X. Jiang, C. Cheng, and R. Xu, "Prompt, plan, perform: Llm-based humanoid control via quantized imitation learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16236–16242.
- [13] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, "Openvla: An open-source vision-language-action model," *arXiv preprint arXiv:2406.09246*, 2024.
- [14] Y. Ma, Z. Song, Y. Zhuang, J. Hao, and I. King, "A survey on vision-language-action models for embodied ai," *arXiv preprint arXiv:2405.14093*, 2024.
- [15] D. Li, Y. Jin, Y. Sun, Y. A. H. Yu, J. Shi, X. Hao, P. Hao, H. Liu, X. Li *et al.*, "What foundation models can bring for robot learning in manipulation: A survey," *The International Journal of Robotics Research*, p. 02783649251390579, 2024.
- [16] R. Firoozi, J. Tucker, S. Tian, A. Majumdar, J. Sun, W. Liu, Y. Zhu, S. Song, A. Kapoor, K. Hausman *et al.*, "Foundation models in robotics: Applications, challenges, and the future," *The International Journal of Robotics Research*, vol. 44, no. 5, pp. 701–739, 2025.
- [17] T. Pang, H. J. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4691–4711, 2023.
- [18] M. Anitescu, "Optimization-based simulation of nonsmooth rigid multi-body dynamics," *Mathematical Programming*, vol. 105, no. 1, pp. 113–143, 2006.
- [19] N. Fazeli, R. Kolbert, R. Tedrake, and A. Rodriguez, "Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1437–1454, 2017.
- [20] C. Mastalli, W. Merkt, J. Marti-Saumell, H. Ferrolho, J. Solà, N. Mansard, and S. Vijayakumar, "A feasibility-driven approach to control-limited ddp," *Autonomous Robots*, vol. 46, no. 8, pp. 985–1005, 2022.
- [21] D. Wisth, M. Camurri, and M. Fallon, "Robust legged robot state estimation using factor graph optimization," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4507–4514, 2019.
- [22] J. Kang and X. Xiong, "Simultaneous ground reaction force and state estimation via constrained moving horizon estimation," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 7080–7086.
- [23] J. Kang, H. Kim, and K.-S. Kim, "View: Visual-inertial external wrench estimator for legged robot," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8366–8373, 2023.
- [24] P. M. Wensing, G. Niemeyer, and J.-J. E. Slotine, "A geometric characterization of observability in inertial parameter identification," *The International Journal of Robotics Research*, vol. 43, no. 14, pp. 2274–2302, 2024.
- [25] Y. Ogawa, G. Venture, and C. Ott, "Dynamic parameters identification of a humanoid robot using joint torque sensors and/or contact forces," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 457–462.
- [26] 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.
- [27] S. Martinez, R. Griffin, and C. Mastalli, "Multi-contact inertial parameters estimation and localization in legged robots," *IEEE Robotics and Automation Letters*, 2025.
- [28] Q. Le Lidec, I. Kalevatykh, I. Laptev, C. Schmid, and J. Carpentier, "Differentiable simulation for physical system identification," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3413–3420, 2021.
- [29] Y. Jiang, M. Yu, X. Zhu, M. Tomizuka, and X. Li, "Contact-implicit model predictive control for dexterous in-hand manipulation: A long-horizon and robust approach," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5260–5266.
- [30] T. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa, "Predictive sampling: Real-time behaviour synthesis with mujoco," *arXiv preprint arXiv:2212.00541*, 2022.
- [31] N. Sobanbabu, G. He, T. He, Y. Yang, and G. Shi, "Sampling-based system identification with active exploration for legged robot sim2real learning," *arXiv preprint arXiv:2505.14266*, 2025.
- [32] D. Stewart and J. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with coulomb friction," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, vol. 1, 2000, pp. 162–169 vol.1.
- [33] C. Rucker and P. M. Wensing, "Smooth parameterization of rigid-body inertia," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2771–2778, 2022.
- [34] T. Pang and R. Tedrake, "A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6614–6620.
- [35] E. D. Andersen and K. D. Andersen, "The mosek interior point optimizer for linear programming: an implementation of the homogeneous algorithm," in *High performance optimization*. Springer, 2000, pp. 197–232.
- [36] G. Kim, D. Kang, J.-H. Kim, S. Hong, and H.-W. Park, "Contact-implicit model predictive control: Controlling diverse quadruped motions without pre-planned contact modes or trajectories," *The International Journal of Robotics Research*, vol. 44, no. 3, pp. 486–510, 2025.
- [37] C. K. Enders, "The performance of the full information maximum likelihood estimator in multiple regression models with missing data," *Educational and Psychological Measurement*, vol. 61, no. 5, pp. 713–740, 2001.
- [38] Y. Tassa, "iLQG/DDP Trajectory Optimization," MATLAB Central File Exchange, 2022, accessed: 2026-01-31.
- [39] M. Mittal, P. Roth, J. Tigue, A. Richard, O. Zhang, P. Du, A. Serrano-Munoz, X. Yao, R. Zurbrugg, N. Rudin *et al.*, "Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning," *arXiv preprint arXiv:2511.04831*, 2025.
- [40] 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 *2019 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2019, pp. 614–619.
- [41] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact

optimal control,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.

- [42] S. Martinez, S. Tonneau, and C. Mastalli, “System identification under constraints and disturbance: A bayesian estimation approach,” *arXiv preprint arXiv:2602.16358*, 2026.
- [43] S. Caron, Q.-C. Pham, and Y. Nakamura, “Stability of surface contacts for humanoid robots: Closed-form formulae of the contact wrench cone for rectangular support areas,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5107–5112.
- [44] F. Liu, Z. Gu, Y. Cai, Z. Zhou, H. Jung, J. Jang, S. Zhao, S. Ha, Y. Chen, D. Xu *et al.*, “Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation,” *IEEE Robotics and Automation Letters*, 2025.

VII. SUPPLEMENTARY MATERIALS

The supplementary materials present additional results and analysis for the proposed PRIME framework, including multimedia demonstrations, supplementary plots, and extended experimental evaluations. All experiments are conducted on the Unitree Go2 and G1 robots, as illustrated in Fig. 7.

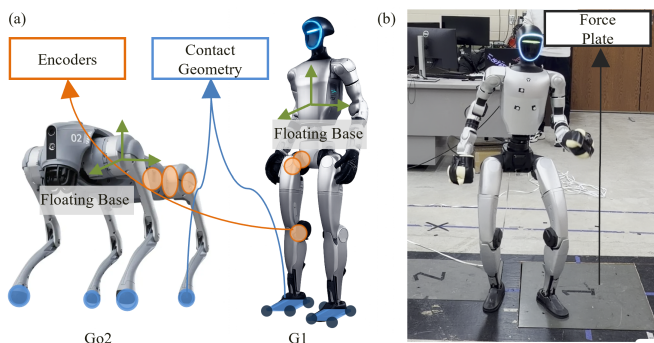


Fig. 7. (a) Robot experimental configurations for Unitree Go2 and G1. Point contact model is used to approximate the underlying contact geometry, which makes the framework invariant under point- or planar-contact for both robots. (b) G1 experimental setup with contact force ground truth from force plate.

A. Multimedia Materials

Table IV summarizes the multimedia materials used to visualize the reconstructed motion and contact force trajectories for experiments on both Unitree Go2 and G1.

TABLE IV
INDEX TO MULTIMEDIA MATERIALS.

Extension	Media Type	Description
1	Video	Hardware results on Unitree Go2 with a 4.6kg payload attached to the belly, demonstrating torso rotations and forward hopping.
2	Video	Hardware results on Unitree G1, demonstrating reconstructed motion for a dancing sequence.
3	Video	Hardware results on Unitree G1 during force-plate data collection, demonstrating reconstructed motion for an additional dancing sequence.
4	Video	Simulation results on Unitree Go2 with a 3kg payload added to the torso, demonstrating torso rotations and forward hopping.
5	Video	Simulation results on Unitree G1, demonstrating reconstructed motion for a similar dancing sequence as 2.

TABLE V
DEFAULT COST WEIGHTS USED BY PRIME IN THE GO2 AND G1 EXPERIMENTS.

Cost Weights	Expression	Coefficient
Floating base cost		
Position noise	$\ p - \tilde{p}\ ^2$	4×10^2
Velocity noise	$\ v - \tilde{v}\ ^2$	1×10^1
Orientation noise	$\ \log(R\tilde{R}^\top)\ ^2$	3×10^1
Angular velocity noise	$\ \omega - \tilde{\omega}\ ^2$	1.5×10^2
Joint cost		
Position noise	$\ \alpha - \tilde{\alpha}\ ^2$	2×10^2
Velocity noise	$\ \dot{\alpha} - \tilde{\dot{\alpha}}\ ^2$	4×10^1
Torque noise	$\ \tau - \tilde{\tau}\ ^2$	2×10^1
Parameter regularization		
Log-Cholesky	$\ \theta - \theta_0\ ^2$	4×10^{-2}

B. Weight Setting

Table V summarizes the default weight settings used in PRIME. These weights are chosen intuitively based on the observed measurement quality. In particular, the floating-base linear velocity is obtained by differentiating the mocap position measurements and is therefore more noise-sensitive, so it is assigned a lower weight. Moreover, due to calibration bias, motion jitter, and marker occlusion, the mocap-based orientation measurements are of degraded quality, whereas the angular-velocity measurements from the IMU are more robust. To mitigate the effect of noisy data, we include a regularization term that penalizes deviations of the inertial parameters from their nominal values provided by the manufacturer.

C. PRIME for Quadrupedal Robots

Additional figures are provided for the hardware experiments on the quadrupedal robot Go2 in Section V.B, where reconstructed trajectories are compared against dynamically infeasible raw kinematic measurements (Figs. 13-16). The roll-direction discrepancy in Fig. 13 and Supplementary Video 1 illustrates a common practical error arising from imperfect mocap calibration. More broadly, even with careful calibration, mocap remains a purely kinematic measurement system and does not enforce dynamic or contact consistency by construction; this can lead to physically inconsistent contact states, such as artificial foot-ground separation or penetration, as shown in Fig. 3(c).

D. PRIME for Humanoid Robots

Additional figures are provided here to complement the G1 humanoid experiments in Section V.C and provide a more complete view of the reconstructed motion and contact behavior. Since ground-truth contact forces at individual foot corners are unavailable in the real-world experiments, we further assess the reconstructed contact forces using simulation ground truth and by checking consistency between hardware and simulation. Fig. 8 compares the reconstructed right-foot forces for the same dancing motion sequence in hardware and simulation. The force profiles and contact sequence are consistent across hardware and simulation, suggesting that the




Fig. 8. Force estimation for the G1 humanoid right foot during a dancing motion. (a) PRIME reconstruction on real hardware. (b) PRIME reconstruction in simulation, compared against MuJoCo ground-truth forces. The simulated force profiles and contact sequence closely match those reconstructed from the real robot under the same motion, suggesting that the chosen contact geometry provides a consistent representation across simulation and hardware.

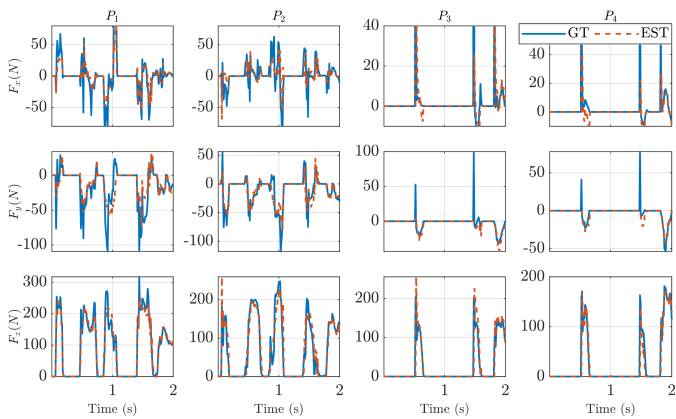


Fig. 9. Humanoid G1 simulation force estimation using a four-point contact model on each foot. This four-point approximation accurately reconstructs planar foot contact, enabling high-fidelity contact reconstruction.

chosen contact geometry, which approximates planar foot contact using corner points, provides a reasonable representation in both simulation and hardware. Fig. 9 shows the per-corner force reconstruction in simulation, where the estimated contact schedule and force profiles closely match the ground truth.

To further assess the reconstructed motion, Figs. 17-20 compare the estimated trajectories with raw kinematic measurements. The reconstructed kinematic trajectories align with the motion-capture measurements and raw kinematics while correcting orientation misalignment and reducing jitter. Following the identification analysis in Section V.C, the torque reconstruction also aligns with the sensor readings after correcting the torso inertia. The remaining torque mismatches in Fig. 20 occur primarily at the upper-limb joints, where the three-finger hands are not modeled.

E. Baseline Comparison with Contact-Constrained Estimation

We compare PRIME with a contact-constrained estimator that optimizes the motion trajectory under a fixed contact sequence, following formulations similar to those in [27], [42]. The baseline uses the same full-information estimation objective, measurement model, and cost terms as in (FIE). The key difference lies in the dynamics model and contact constraints. Instead of using the smoothed contact dynamics in (17), the

Method	RMSE _F [N]	RelativeError _F [%]	Cost [$\times 10^3$]
PRIME	19.833	10.494	1.127
Baseline	73.103	52.693	1.36

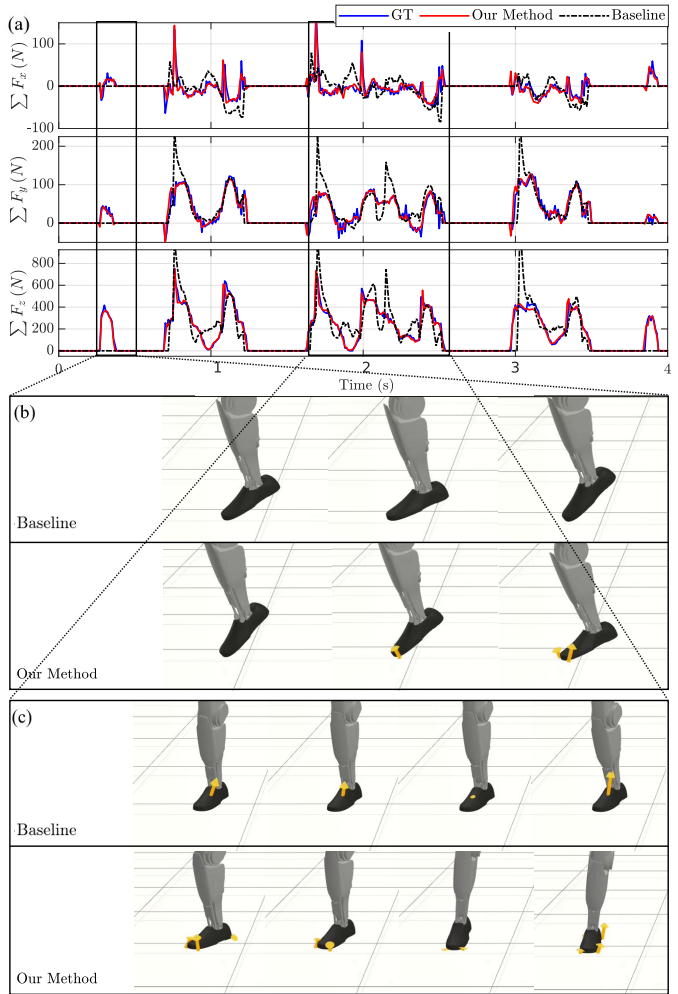


Fig. 10. (a) Force-estimation comparison between our method and the contact-constrained estimation baseline for the G1 humanoid (right foot) over a simulated dancing motion. (b) Loss of contact fidelity during partial support at the foot tip. (c) Inaccurate contact reconstructions during rolling and spatially varying foot contact.

baseline first assigns contact flags by thresholding the end-effector height computed from the measured kinematics. This is consistent with common hardware pipelines that infer end-effector contacts from pressure sensors, estimated external disturbances, or raw kinematic signals, which typically provide only a binary contact indicator. Given the prescribed contact sequence, each active foot contact in the humanoid experiments is modeled by a rigid contact constraint and a spatial friction wrench cone at the mid-foot. The constrained dynamics are written as:

$$\mathbf{M}(\mathbf{q})\dot{\mathbf{v}} + \mathbf{h}(\mathbf{q}, \mathbf{v}) = \mathbf{B}^\top \boldsymbol{\tau} + \sum_i \mathbf{J}_i(\mathbf{q})^\top \mathbf{w}_i, \quad (35)$$

with the rigid contact acceleration constraint for each active contact i :

$$\mathbf{J}_i(\mathbf{q})\dot{\mathbf{v}} + \dot{\mathbf{J}}_i(\mathbf{q}, \mathbf{v})\mathbf{v} = \mathbf{0}. \quad (36)$$




Fig. 11. Effects of the smoothing parameter of contact dynamics on the optimization convergence (blue) and the estimation accuracy of contact forces (red) for G1 humanoid in simulation.




Fig. 12. Effects of the smoothing parameter of contact dynamics on the optimization convergence (blue) and the identification accuracy of torso mass (red) for Go2 quadrupedal in simulation.

Here $\mathbf{J}_i(\mathbf{q}) \in \mathbb{R}^{6 \times (6+n)}$ is the spatial frame Jacobian of contact i , located at the center of the footpad in the humanoid experiments. The vector $\mathbf{w} = [f_x \ f_y \ f_z \ \tau_x \ \tau_y \ \tau_z]^T$ denotes the equivalent spatial wrench applied at the corresponding contact frame. The wrench \mathbf{w} is constrained by the friction wrench cone constraint [43], which can be written as:

$$\begin{aligned} |f_x| &\leq \mu f_z, & |f_y| &\leq \mu f_z, & f_z &> 0, \\ |\tau_x| &\leq Y f_z, & |\tau_y| &\leq X f_z, & \tau_{z,\min} &\leq \tau_z \leq \tau_{z,\max}. \end{aligned} \quad (37)$$

$$\begin{aligned} \tau_{z,\min} &= -\mu(X + Y)f_z + |Yf_x - \mu\tau_x| + |Xf_y - \mu\tau_y|, \\ \tau_{z,\max} &= \mu(X + Y)f_z - |Yf_x + \mu\tau_x| - |Xf_y + \mu\tau_y|, \end{aligned}$$

where X and Y denote the length and width of the supporting rectangle. Compared with PRIME, this baseline decouples contact detection from estimation by assigning contact modes from kinematic thresholding outside the estimation problem. PRIME instead reconstructs contact status jointly with motion and inertial parameters through its contact dynamics model, avoiding externally prescribed contact flags. Under the fixed-mode formulation, distributed foot contact is further approximated by a single equivalent wrench at the midfoot. Although this approximation is common in control-oriented formulations with prescribed contacts [44], it reduces fidelity for estimation, especially during rolling, partial support, and spatially varying foot contact. As shown in Fig. 10 and Table VI, the baseline produces substantially worse contact-force reconstruction despite accurate inertial parameters, highlighting the benefit of explicitly modeling contact dynamics.

F. Smoothed Contact Dynamics Analysis

PRIME enables robust reconstruction of latent contact modes by leveraging the smoothed contact dynamics. In particular, the log-barrier approximation of the second-order contact constraints (SOC) as in Section III.C can be interpreted

as inducing an implicit “force field” [17]: the barrier introduces a repulsive potential whose influence increases as the solution approaches the constraint boundary, thereby smoothly and differentiably shaping contact engagement and disengagement. As $\kappa \rightarrow \infty$, the approximation approaches hard SOC enforcement and recovers rigid, physically consistent contact. For moderate κ , the relaxation can introduce a force-at-a-distance effect: contact forces may appear smoothly before strict contact activation. This behavior should be interpreted differently in the proposed estimation problem than in its control-side counterpart, contact-implicit control. In control, smoothing artifacts are usually undesirable because the primary goal is task execution with physically feasible trajectories, where contact should be rigidly enforced; the relaxation is therefore often reduced or removed over iterations, mainly to aid numerical convergence. In contrast, PRIME aims to reconstruct motion, contact, and inertial parameters that best explain noisy kinematic measurements and onboard actuator sensing. In this setting, moderate smoothing can improve robustness by reducing sensitivity to contact switching, measurement noise, and data down-sampling.

We quantify these effects in PRIME by varying κ under the same cost settings and evaluating the resulting convergence behavior, force-estimation accuracy, and identification accuracy using simulation data from the G1 and Go2 experiments in Section V. As shown in Fig. 11 and Fig. 12, a moderate choice of $\kappa = 500$ promotes convergence by providing informative gradients, while maintaining physical consistency with rigid contact dynamics and tolerating noise and contact misalignment introduced by data downsampling. Unless otherwise stated, we use the same κ across all experiments. When the hardware measurements are particularly noisy, PRIME can be initialized with a smaller κ (stronger smoothing) and then gradually increase κ over iterations to the target value, thereby improving convergence while ultimately enforcing physically consistent contact behavior.

G. Computation Profiling

Although PRIME is currently formulated as an offline long-horizon Full-Information Estimation (FIE) problem for system identification and motion reconstruction, its single-step computation is efficient due to the fully C++ implementation. On a PC equipped with an Intel Core i9-13900 CPU, one quadruped simulation rollout with differentiation takes $80 \mu\text{s}$ on average. This is comparable to the $70 \mu\text{s}$ differentiation time reported in [36] for a similar contact-rich setting with real-time contact-implicit MPC, although the rollout time is not separately profiled in [36]. Given the similar computational structure of MPC and Moving Horizon Estimation (MHE), this timing suggests that PRIME may be amenable to a real-time receding-horizon formulation. Such an extension would require an appropriate arrival cost to propagate posterior information across sliding windows, especially for the coupled uncertainty among motion states, contact forces, and inertial parameters. We leave this MHE implementation and arrival-cost design for future work, following related frameworks [5].

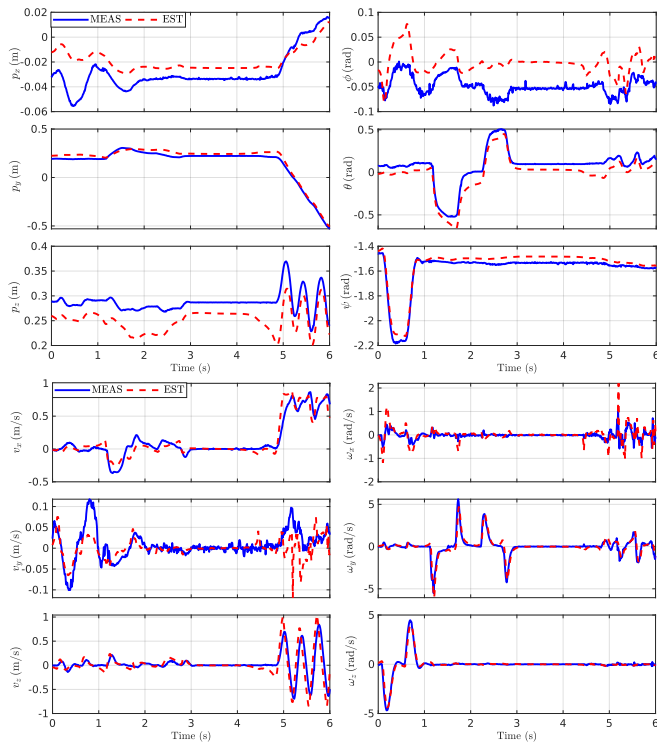


Fig. 13. Floating-base estimation of Go2 over diverse motions. The roll component of the mocap measurement orientation is biased due to imperfect calibration.

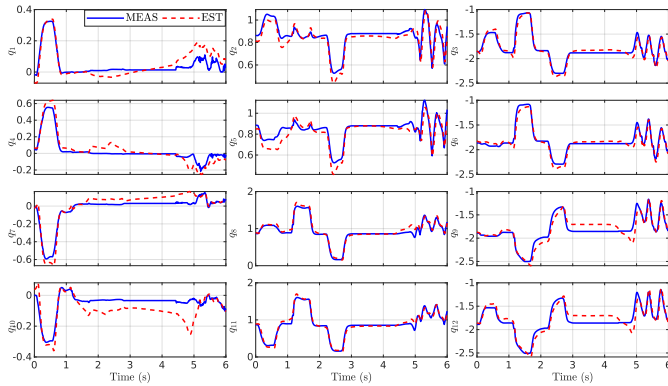


Fig. 14. Joint position estimation of Go2 over diverse motion.




Fig. 15. Joint velocity estimation of Go2 over diverse motion.




Fig. 16. Joint torque estimation of Go2 over diverse motion.

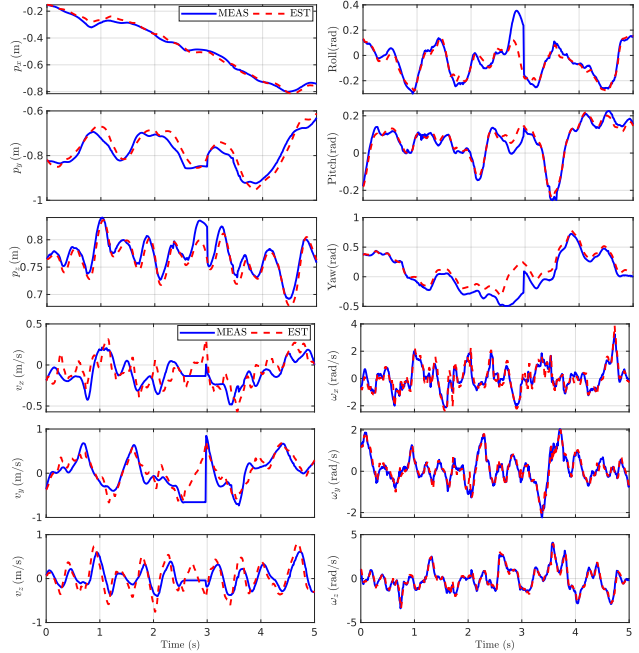


Fig. 17. Floating base estimation of G1 humanoid over a dancing motion.




Fig. 18. Joint position estimation of G1 humanoid over a dancing motion.




Fig. 19. Joint velocity estimation of G1 humanoid over a dancing motion.




Fig. 20. Joint torque estimation of G1 humanoid over a dancing motion.