Title:
Automated calibration of robot digital twins through the discovery of interpretable physical phenomena

Authors:
Xinping GU, Université Paris-Saclay, ENS Paris-Saclay, LURPA, 91190 Gif-sur-Yvette
Olivier Bruneau, Université Paris-Saclay, ENS Paris-Saclay, LURPA, 91190 Gif-sur-Yvette
Charyar Mehdi-Souzani, Université Paris-Saclay, ENS Paris-Saclay, LURPA, Université Sorbonne Paris Nord, 91190 Gif-sur-Yvette

Keywords:
Digital twin, Industrial robot, Dynamics calibration, Physics discovery, Uncertainty quantification

DOI: 10.14733/cadconfP.2026.154-158

Introduction:
Digital twins (DTs) have become an active research topic in recent years, including in industrial robotics [8]. In industrial robotic applications, DTs are useful for offline verification tasks such as robotic machining and quality control [4,10]. However, the practical value of a robot DT strongly depends on the fidelity of its dynamics layer, and the reality gap between nominal models and physical systems remains a central challenge [1]. Torque predictions may deviate from the physical system because of payload variations, parameter drift, or other unmodeled physical effects [7]. These discrepancies create persistent, trajectory-dependent errors that can reduce confidence in DT-enabled offline assessments [6].

Related work has also addressed this issue through uncertainty-aware parameter updating [3] and data-driven inverse-dynamics compensation [5]. More recently, interpretable physics-discovery methods based on sparse Bayesian learning have shown that compact Lagrangian models can be identified together with uncertainty information [2,9]. However, robot DT calibration still lacks a practical method that can improve torque prediction while keeping the correction physically interpretable and providing usable uncertainty information.

In this work, we propose a calibration method for robot DT torque layers based on time-synchronized joint trajectories and actuation data. We model the discrepancy between measured and nominal dynamics as a residual term and, in the present offline setting, assume that its main contribution can be approximated by conservative rigid-body effects. We then use a physics-based candidate library together with sparse Bayesian inference to identify a small set of residual correction terms and estimate their posterior uncertainty. The method also provides two useful uncertainty outputs for the DT: (i) posterior inclusion probabilities (PIPs), which indicate which residual terms are supported by the data, and (ii) predictive intervals for the residual torque. Tests on a simulated two-degree-of-freedom manipulator show that the proposed correction can reduce the reality gap relative to the nominal model, while remaining interpretable and providing uncertainty estimates.

Main Idea:
We formulate the auto-calibration of the DT dynamics layer as a structural discovery problem, where the goal is to model the discrepancy between the deployed robot and its nominal dynamics. Let $\boldsymbol{q}, \dot{\boldsymbol{q}}, \ddot{\boldsymbol{q}} \in \mathbb{R}^d$ denote the joint states derived from time-synchronized trajectory data. We define the residual torque τ_{res} as the difference between the measured actuation τ_{meas} and the nominal rigid-body torque τ_{nom}:

$$\tau_{\text{res}}(t) = \tau_{\text{meas}}(t) - \tau_{\text{nom}}(t, \boldsymbol{q}, \dot{\boldsymbol{q}}, \ddot{\boldsymbol{q}}) \tag{1}$$

In the present offline calibration setting, we assume that the main reality-gap effects, such as inertial parameter drift, payload variations, or unmodeled gravitational biases, can be approximated as conservative. Accordingly, we represent this discrepancy by a scalar residual Lagrangian function, $\Delta\mathcal{L}(\boldsymbol{q},\dot{\boldsymbol{q}})$, which augments the nominal dynamics in an energy-consistent manner [9], while the remaining non-conservative effects (e.g., friction and actuator losses) are treated as stochastic uncertainty. The schematic overview of the framework is shown in Fig. 1.

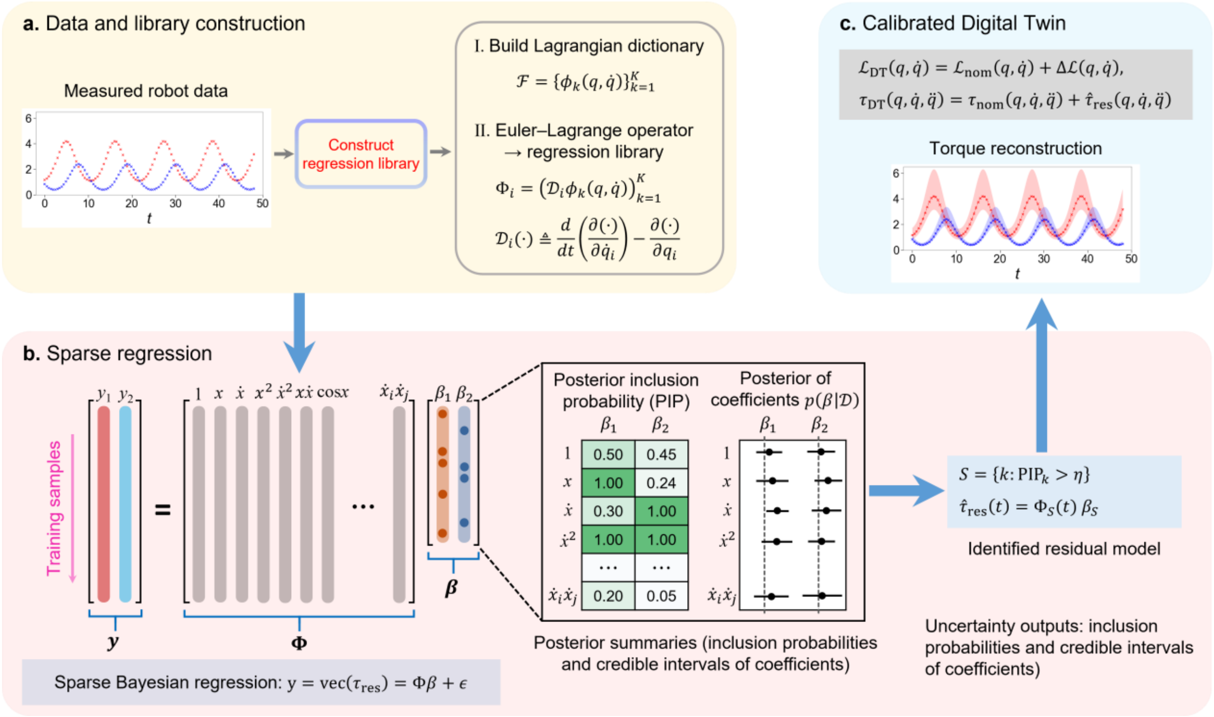

Fig. 1: Schematic overview of a digital-twin calibration framework based on conservative residual Lagrangian discovery.

To preserve interpretability, we express the unknown correction $\Delta\mathcal{L}$ as a sparse linear combination of K candidate basis functions from a physics-informed library $\mathcal{F}=\{\phi_k\}_{k=1}^{K}$:

$$\Delta\mathcal{L}(\boldsymbol{q},\dot{\boldsymbol{q}})=\sum_{k=1}^{K}\beta_k\ \phi_k(\boldsymbol{q},\dot{\boldsymbol{q}}) \tag{2}$$

where $\beta\in\mathbb{R}^K$ is the vector of unknown coefficients. The library $\mathcal{F}$ is designed to capture the dominant rigid-body interactions expected in the target robot, including kinetic-like terms (e.g., $\dot{q}_i^2$ and $\dot{q}_i\dot{q}_j$) and potential-like terms (e.g., $\sin(q_i)$ and $\cos(q_i)$) for gravitational effects. Here, $\Delta\mathcal{L}$ is a scalar residual Lagrangian correction, whereas the observations τ_{res} are vector-valued generalized forces. We connect them through the Euler–Lagrange operator

$$\mathcal{D}_i(\cdot)\triangleq\frac{d}{dt}\left(\frac{\partial(\cdot)}{\partial\dot{q}_i}\right)-\frac{\partial(\cdot)}{\partial q_i}. \tag{3}$$

Applying $\mathcal{D}_i$ to the basis expansion maps the scalar candidates into the torque domain:

$$\tau_{\mathrm{res},i} \approx \mathcal{D}_i(\Delta\mathcal{L}) = \sum_{k=1}^{K} \beta_k \underbrace{\mathcal{D}_i(\phi_k(\boldsymbol{q}, \dot{\boldsymbol{q}}))}_{\text{EL Regressors}} \tag{4}$$

This step converts the nonlinear identification problem into a linear regression form, where the transformed features serve as regressors consistent with rigid-body mechanics.

The calibration task is then written as a probabilistic linear regression. By stacking the evaluations of the Euler–Lagrange regressors across all d joints and N time samples, we obtain the design matrix $\boldsymbol{\Phi} \in \mathbb{R}^{(Nd)\times K}$ and the target vector $\boldsymbol{y} \in \mathbb{R}^{Nd}$, which represents the flattened residual torque τ_{res}:

$$\boldsymbol{y} = \boldsymbol{\Phi}\beta + \epsilon, \qquad \epsilon \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}) \tag{5}$$

Here, σ^2 is the noise variance, $\boldsymbol{I}$ is the identity matrix, and the error term ϵ aggregates measurement noise and unmodeled non-conservative effects, treated as aleatoric uncertainty. To enforce parsimony and retain only the physically necessary corrections, we place a Spike-and-Slab prior on β. This hierarchical prior uses latent binary indicators to separate relevant physical terms (the slab) from irrelevant ones (the spike at zero). Inference is performed by a Gibbs sampler, which alternates between updating the sparse coefficients and estimating the noise variance.

Finally, the framework yields a calibrated dynamics core $\tau_{\mathrm{calib}} = \tau_{\mathrm{nom}} + \boldsymbol{\Phi}\hat{\beta}$, where $\hat{\beta}$ denotes the posterior mean coefficients. The calibrated model is also accompanied by two uncertainty outputs that are useful for DT assessment. Posterior Inclusion Probabilities (PIPs) indicate which residual terms are supported by the data, while posterior predictive intervals provide bounds on the residual torque along a trajectory.

Simulation:

We validate the proposed framework on a simulated two-degree-of-freedom planar arm, which serves here as a simple representative industrial manipulator. Numerical experiments were conducted in a physics-based simulator (MuJoCo), and nominal torques were computed using a rigid-body dynamics model. Bayesian sparse inference was performed using Markov Chain Monte Carlo (MCMC) sampling, and uncertainty summaries were reported. A nominal rigid-body model provides the baseline torque, $\tau_{\mathrm{nom}}(t)$. To emulate practically relevant reality-gap effects, we generate ground-truth data from a parameter-perturbed rigid-body model with inertial and gravity drifts, together with a hidden conservative correction used only for data generation. By applying the Euler–Lagrange equations to this perturbed system and injecting measurement noise, we synthesize the measured torques, $\tau_{\mathrm{meas}}(t)$. The identification task is cast as residual supervision, i.e., we regress the deviation between measured and nominal torques. This formulation lets the learned dynamics augment the digital twin as an additive correction, without changing the structure of the nominal model.

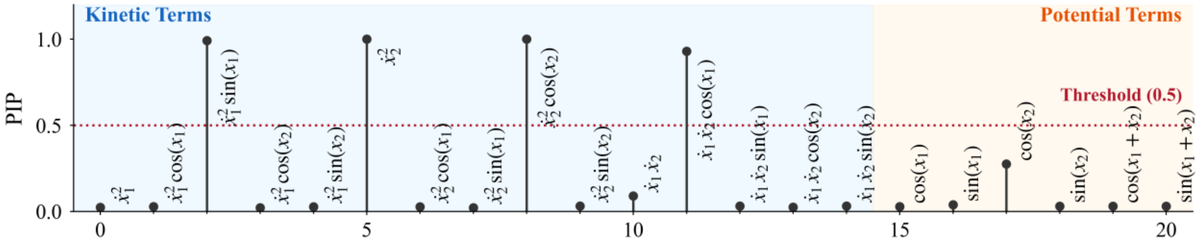

Fig. 2. Structural discovery using posterior inclusion probabilities (PIP) of candidate terms in the residual-Lagrangian dictionary. Terms with PIP ≥ 0.5 (majority posterior support) are selected; the horizontal axis denotes the dictionary term index.

Using the Bayesian residual-Lagrangian formulation, we express $\Delta\mathcal{L}$ as a sparse linear combination of candidate basis functions containing both kinetic- and potential-energy terms. The Euler–Lagrange operator maps this expansion into a linear regression problem on the residual torques with additive Gaussian noise. We then employ a spike-and-slab prior with Gibbs sampling to jointly perform structure learning and uncertainty quantification. As shown in Fig. 2, the PIP identifies a compact subset of active terms, highlighting the dominant conservative corrections. The posterior distributions of the retained coefficients (Fig. 3) further characterize parameter uncertainty and coupling in the recovered model. On held-out data (Fig. 4), the calibrated twin combines the nominal baseline with the posterior predictive residual and closely tracks the measured torques. The associated predictive intervals provide uncertainty bounds along the trajectory, showing that the method can support offline DT validation in this simulated setting.

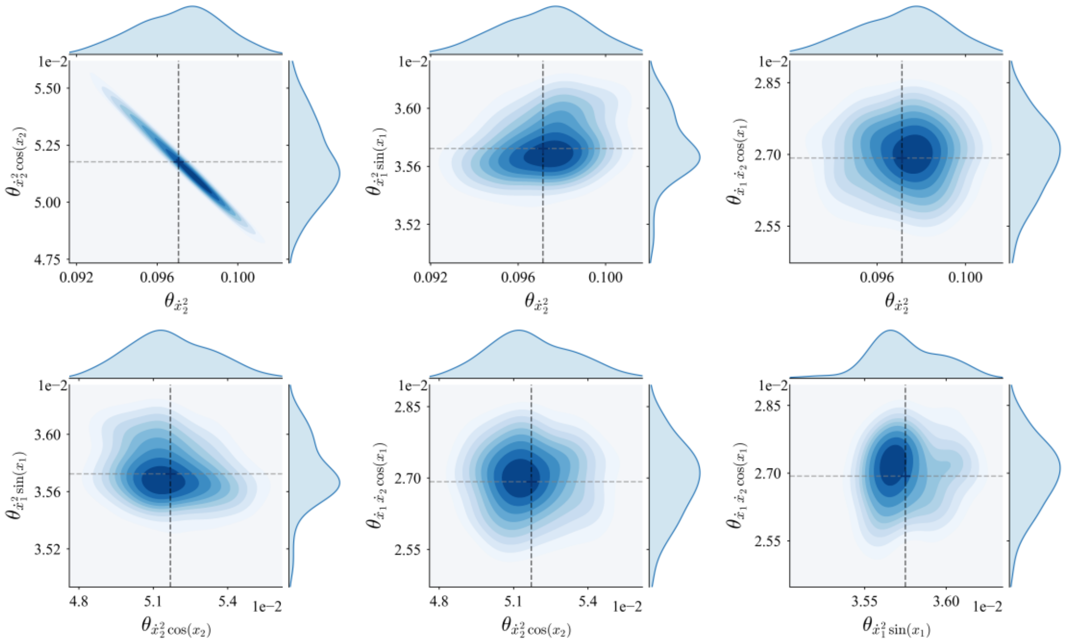

Fig. 3: Coefficient posterior distributions. Joint and marginal plots visualize parameter uncertainty and coupling in the recovered model.

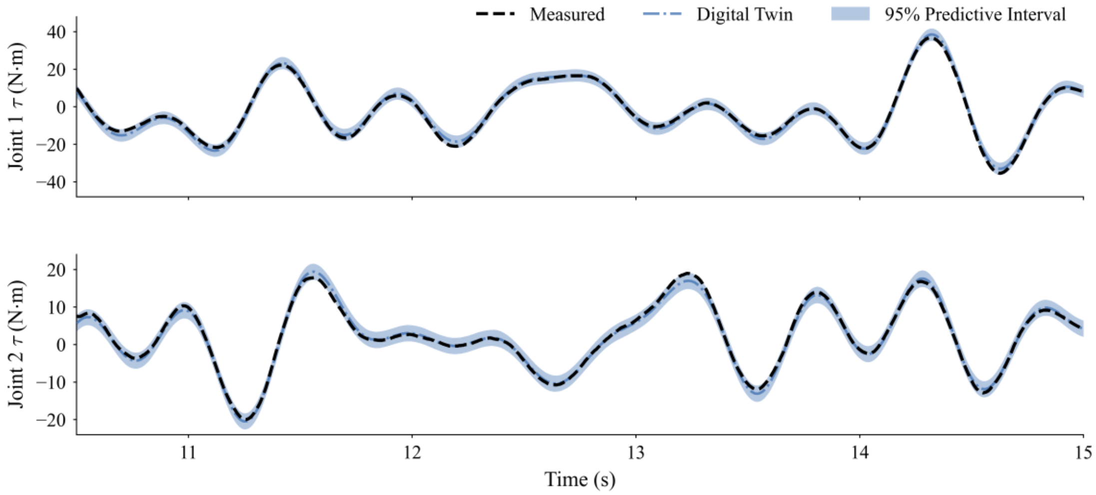

Fig. 4: Torque prediction by the proposed model.

Conclusions:

We have presented an automated calibration framework for the dynamics layer of industrial robot DTs, together with uncertainty quantification. By enforcing Euler–Lagrange structure within a sparse Bayesian learning process, the method identifies interpretable rigid-body residual corrections from torque observations, instead of relying on black-box compensation models. In the simulated two-DOF study, the recovered conservative corrections are consistent with the perturbations introduced during data generation, and the calibrated twin provides both PIP-based support for the identified terms and posterior predictive intervals for the residual torque.

Although the present validation is limited to offline calibration of conservative dynamics, the proposed formulation provides a practical basis for offline simulation and verification with quantified uncertainty. Future work will test the framework in broader validation settings and extend it to non-conservative effects such as friction and damping.

Acknowledgements:

The project (22DIT01 ViDiT) has received funding from the European Partnership on Metrology, co-financed from the European Union's Horizon Europe Research and Innovation Programme and by the Participating States.

References:

[1] Aljalbout, E.; Xing, J.; Romero, A.; Akinola, I.; Garrett, C.R.; Heiden, E.; Gupta, A.; Hermans, T.; Narang, Y.; Fox, D.; Scaramuzza, D.; Ramos, F.: The Reality Gap in Robotics: Challenges, Solutions, and Best Practices, Annual Review of Control, Robotics, and Autonomous Systems, 9, 2025, In press. https://doi.org/10.1146/annurev-control-031924-100130

[2] Hirsh, S.M.; Barajas-Solano, D.A.; Kutz, J.N.: Sparsifying Priors for Bayesian Uncertainty Quantification in Model Discovery, Royal Society Open Science, 9(2), 2022, 211823. https://doi.org/10.1098/rsos.211823

[3] Kessels, B.M.; Fey, R.H.B.; van de Wouw, N.: Uncertainty Quantification in Real-Time Parameter Updating for Digital Twins Using Bayesian Inverse Mapping Models, Nonlinear Dynamics, 113(8), 2025, 7613-7637. https://doi.org/10.1007/s11071-024-10608-9

[4] Ni, H.; Hu, T.; Deng, J.; Chen, B.; Luo, S.; Ji, S.: Digital Twin-Driven Virtual Commissioning for Robotic Machining Enhanced by Machine Learning, Robotics and Computer-Integrated Manufacturing, 93, 2025, 102908. https://doi.org/10.1016/j.rcim.2024.102908

[5] Pikuliński, M.; Malczyk, P.; Aarts, R.: Data-Driven Inverse Dynamics Modeling Using Neural-Networks and Regression-Based Techniques, Multibody System Dynamics, 63(3), 2025, 341-366. https://doi.org/10.1007/s11044-024-10024-2

[6] Shao, G.; Hightower, J.; Schindel, W.: Credibility Consideration for Digital Twins in Manufacturing, Manufacturing Letters, 35, 2023, 24-28. https://doi.org/10.1016/j.mfglet.2022.11.009

[7] Taie, W.; ElGeneidy, K.; AL-Yacoub, A.; Sun, R.: Online Identification of Payload Inertial Parameters Using Ensemble Learning for Collaborative Robots, IEEE Robotics and Automation Letters, 9(2), 2024, 1350-1356. https://doi.org/10.1109/LRA.2023.3346268

[8] Tao, F.; Zhang, M.; Liu, Y.; Nee, A.Y.C.: Digital Twin in Industry: State-of-the-Art, IEEE Transactions on Industrial Informatics, 15(4), 2019, 2405-2415. https://doi.org/10.1109/TII.2018.2873186

[9] Tripura, T.; Chakraborty, S.: A Bayesian Framework for Discovering Interpretable Lagrangian of Dynamical Systems from Data, Mechanical Systems and Signal Processing, 221, 2024, 111737. https://doi.org/10.1016/j.ymssp.2024.111737

[10] Wright, L.; Davidson, S.: Digital Twins for Metrology; Metrology for Digital Twins, Measurement Science and Technology, 35, 2024, 051001. https://doi.org/10.1088/1361-6501/ad2050