

TwinTrack: Bridging Vision and Contact Physics for Real-Time Tracking of Unknown Objects in Contact-Rich Scenes

Wen Yang^{1,2}, Zhixian Xie¹, Yiting Wang¹, Abhijit Tadepalli¹, Heni Ben Amor³, Shan Lin⁴, Wanxin Jin^{1*}




Fig. 1: **TwinTrack** is a physics-aware perception framework for robust, real-time tracking of unseen, dynamic objects in contact-rich scenes with motion blur (Cols. 1–4) and heavy occlusion (Cols. 5–6). The framework integrates two key components: **Real2Sim**, which combines the complementary strengths of vision and contact dynamics to jointly estimate an object’s geometry and physical properties; and **Sim2Real**, which performs robust object pose estimation through adaptive fusion of visual tracking and contact simulation. [video link](#)

Abstract—Real-time tracking of previously unseen, highly dynamic objects in contact-rich scenes, such as during dexterous in-hand manipulation, remains a major challenge. Pure vision-based approaches often fail under heavy occlusions due to frequent contact interactions and motion blur caused by abrupt impacts. We propose *TwinTrack*, a physics-aware perception system that enables robust, real-time 6-DoF pose tracking of unknown dynamic objects in contact-rich scenes by leveraging contact physics cues. At its core, *TwinTrack* integrates *Real2Sim* and *Sim2Real*. *Real2Sim* combines vision and contact physics to jointly estimate object geometry and physical properties: an initial reconstruction is obtained from vision, then refined by learning a geometry residual and simultaneously estimating physical parameters (e.g., mass, inertia, and friction) based on contact dynamics consistency. *Sim2Real* achieves robust pose estimation by adaptively fusing a visual tracker with predictions from the updated contact dynamics. *TwinTrack* is implemented on a GPU-accelerated, customized MJX engine to guarantee real-time performance. We evaluate our method on two contact-rich scenarios: object falling with environmental contacts and multi-fingered in-hand manipulation. Results show that, compared to baselines, *TwinTrack* delivers significantly more robust, accurate, and real-time tracking in these challenging settings, with tracking speeds above 20 Hz. [Project page](#)

I. INTRODUCTION

Real-time 6-DoF pose tracking of previously unseen, highly dynamic objects in contact-rich scenes is fundamental to robotic manipulation [1], [2] and scene understanding [3]. Contact-rich scenes typically involve frequent object-robot-environment interactions and abrupt motion changes, such as jumps in the object velocity profile, causing large occlusions and motion blur. Pure vision-based tracking [4] struggles with the incomplete and/or degraded observations.

Integrating contact physics cues [5] can overcome these challenges. Simulated contact dynamics, when aligned with

reality, anticipate object poses in highly occluded frames and compensate for degraded visual tracking. Furthermore, incorporating contact dynamics also bridges the gap between visual perception and digital physics simulation, enabling downstream applications such as model-based dexterous manipulation control [6]–[8].

Building physics-informed perception system from monocular RGB-D vision for unknown objects is challenging with real-time constraints. The system must jointly estimate object pose, geometry, and physical properties (e.g., mass, inertia, friction coefficients) from RGB-D inputs. The high-dimensional optimization space requires a modular optimization pipeline with GPU-accelerated implementation. Specifically, three major challenges need to be addressed: (i) Efficiently integrating the complementary benefits of vision and contact dynamics cues for geometry estimation. Vision provides dense geometric reconstructions, but these are often corrupted by noise and occlusion. In contrast, contact events imply highly accurate geometric information but are sparse and intermittent [9]. (ii) Contact interactions introduce non-smooth motion, which challenges gradient-based learning and optimization [10]. (iii) Real-time performance requires highly efficient implementation.

To address these challenges, we propose *TwinTrack* (Fig. 1), a physics-aware perception framework that bridges vision and contact dynamics for robust, real-time 6-DoF tracking of dynamic unknown objects in contact-rich environments. *TwinTrack* takes monocular RGB-D images and robot proprioceptive data as input and outputs real-time 6-DoF object poses that are resilient to occlusions and degraded vision observations. The core contributions of *TwinTrack* include:

(I) An architecture that closes the loop between *Real2Sim* and *Sim2Real*. *Real2Sim* reconstructs object geometry and aligns a contact dynamics simulator with temporal RGB-D observations. *Sim2Real* then leverages the aligned simulator to perform physics-aware 6-DoF pose tracking that remains robust under occlusions and degraded visual input.

(II) Geometry reconstructed from RGB-D vision and refined

¹School for Engineering of Matter, Transport, and Energy.

²School of Computing and Augmented Intelligence.

³School of Electrical, Computer and Energy Engineering.

All authors are with Arizona State University, Tempe, AZ 85281, USA. Wen Yang² is also with Department of Robotics, Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), Abu Dhabi, UAE.

*Corresponding author: Wanxin Jin (wanxin.jin@asu.edu).

by *contact dynamics*. In Real2Sim, a dense geometry representation of the object is first extracted from pure vision. Due to noise and occlusion, the geometry is often inaccurate or incomplete. To address this, we propose refining it by learning a geometry residual in the alignment of contact dynamics simulator, together with estimating physical properties of the contact dynamics. This two-phase process enables effective integration of visual and reliable contact-dynamics cues to derive contact-consistent geometry.

(III) *Real-time implementation and optimization*. Real2Sim (slow backend) and Sim2Real (fast frontend) are implemented separately for real-time tracking performance. To handle the non-smooth contact dynamics in Real2Sim optimization, we use GPU-parallelized, sampling-based optimization. We developed a deeply customized MJX-based physics engine [11] for collision detection and contact dynamics alignment.

We evaluate TwinTrack in two challenging contact-rich scenarios: object free-fall and in-hand manipulation, both involving highly dynamic motion, frequent occlusions, and contact impacts. Compared with the state-of-the-art visual-only tracker [12], TwinTrack achieves physically plausible and robust tracking across diverse real-world objects while maintaining real-time performance ($>20\text{Hz}$).

II. RELATED WORK

A. Visual Tracking and Geometry Reconstruction

Vision-based estimation of object pose and geometry has been a long-standing focus in computer vision. Methods with strong priors (CAD models, template matching) [13]–[15] achieve high accuracy but lack generalization to unseen objects. More recent approaches aim to track previously unseen objects without any prior knowledge, either by integrating pre-trained image descriptors [16] or by jointly learning 3D representations during tracking via neural fields [4], [12].

While vision-based approaches perform well in occlusion-free and moderately dynamic conditions [4], [12], [16], they struggle in highly dynamic, contact-rich scenarios involving fast motion, abrupt changes, and frequent occlusion. Visual observations become degraded or incomplete, leading to unreliable tracking. Object geometry from visual data may be uncertain due to noise, causing inconsistencies in motion prediction [9]. This underscores the need to incorporate physics knowledge (contact dynamics) to enhance robustness and ensure physically plausible estimates.

B. Physics-Informed Tracking and Reconstruction

Physics-informed tracking incorporates physics priors to enforce physical plausibility. Human motion capture (e.g., [17]) uses simulation-based constraints and soft physics priors for improved realism, but often assumes known physical properties limiting generalization. More recent approaches (e.g., [18]) focus on jointly learning geometry and physical parameters by minimizing simulation/prediction loss from differentiable simulators. For example, [3] proposes using Real2Sim to extract collision geometry from robot sensors and cameras, while [19] reconstructs interactive digital twins from object interaction videos. However, their scenarios involve

relatively slow and less occluded motion. In contact-rich scenes with high occlusion and motion blur, visual geometry may be inaccurate, degrading dynamics learning. Thus, [20] proposes combining visual and contact cues for geometry inference under occlusion, but they focus on reconstruction rather than real-time tracking.

C. Learning Contact Dynamics

Accurate physics models are essential for physics-informed perception [19]. However, real-world contact interactions are inherently non-smooth with stiff transitions, posing learning challenges. Standard neural networks struggle with discontinuities due to smooth bias [10]. Differentiable simulators (e.g., MJX [11]) provide gradients that are often discontinuous and local, limiting global optimization [10]. Studies [6], [8], [10] relax stiff constraints for smooth gradients. Our method avoids contact gradients by adopting sampling-based optimization. Furthermore, built on the neural signed distance function (SDF) representation of contact geometry in the simulator, we propose learning geometry residuals to refine the functional geometry used in contact dynamics, enhancing simulation-reality alignment.

III. PROBLEM STATEMENT AND TWINTRACK OVERVIEW

We consider RGB-D image sequences of a contact-rich dynamic scene where an unknown object frequently interacts with robots or/and environments. Denote RGB and depth image at frame k as (I_k, D_k) . TwinTrack takes RGB-D images and robot proprioceptive data (when available) as input. The goal is to estimate the 6-DoF object pose \mathbf{q}_k in real-time despite degraded visual observations from occlusion and motion blur due to contact impacts.

As shown in Fig. 2, TwinTrack handles the problem by a backend–frontend loop between Real2Sim and Sim2Real. Sim2Real runs online and produces a per-frame pose estimate for the current observations. Real2Sim runs asynchronously on accumulated observational history to update the object model and keyframes that Sim2Real can reliably use. The two components exchange information continuously: Sim2Real provides keyframes and tracking history, while Real2Sim returns updated keyframe anchors and an improved contact dynamics model. At the beginning, the system relies mainly on vision; as Real2Sim becomes reliable, the learned physical model is progressively incorporated to stabilize tracking under occlusion. The details will be presented in Section IV and V.

IV. REAL2SIM

This section describes the Real2Sim (backend), whose goal is to produce (i) optimized keyframes with consistent poses and (ii) a simulation-ready contact dynamics model. Real2Sim is a two-phase process: the object visual geometry is first estimated using RGB-D keyframes from a keyframe buffer (Sec. IV-A); we then fine the geometry to align better with observed contact behavior and identify the physical parameters (Sec. IV-B) using tracking data provided by Sim2Real (Sec. V). Inputs to Real2Sim: keyframe buffer and historical tracking data from Sim2Real. Outputs: the simulation-ready contact dynamics and pose-optimized keyframes.

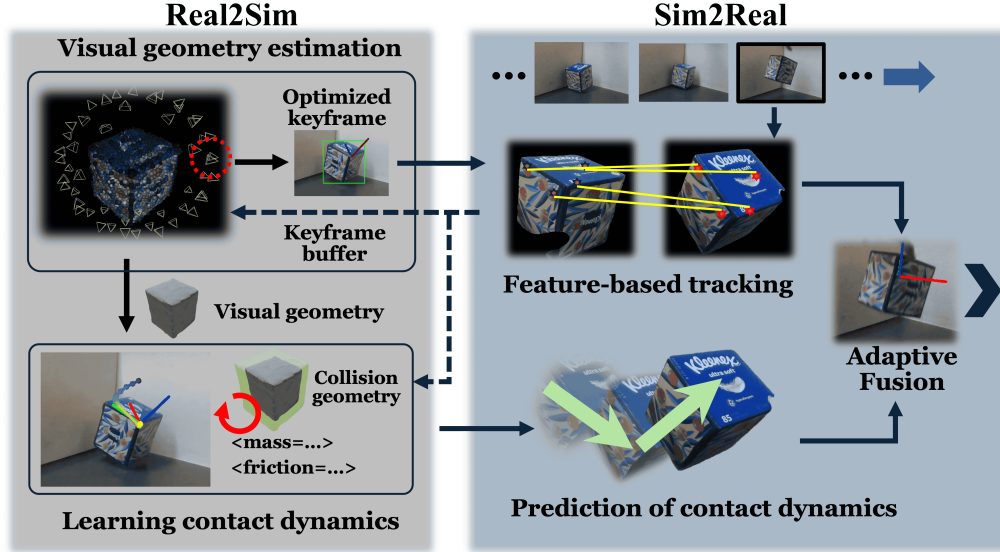


Fig. 2: Overview of **TwinTrack** framework. Our framework includes two main components: **Real2Sim** for learning object geometry and simulation-ready contact physics, and **Sim2Real** for physics-aware real-time pose tracking. In **Real2Sim**, object visual geometry, represented as the Gaussian Splats, is obtained from a selection of RGB-D keyframes (Sec. IV-A). The obtained geometry is updated in the next stage of learning contact dynamics, together with identifying other physical properties (Sec. IV-B). In **Sim2Real**, feature-based tracking is performed for each new frame (Sec. V-A) relative to optimized keyframes from **Real2Sim**; meanwhile the contact physics simulation also predicts the current object pose. The final object pose is an adaptive fusion of both visual tracker and simulation (Sec. V-B).

A. Visual Geometry Reconstruction




Fig. 3: Visual geometry reconstruction. Left: joint optimization of Gaussian Splatting (GS) model \mathcal{G}_{gs} and keyframe poses. Middle: depth rendering from the obtained GS model. Right: learning neural SDF $\mathbf{SDF}_{\text{vision}}$ from the rendered depth images.

This goal is to jointly reconstruct the visual geometry of the unknown object from keyframes and simultaneously optimize the keyframe poses. The reconstructed geometry will be refined in contact dynamics learning (Sec. IV-B), while the optimized keyframes serve as pose anchors for visual tracking in Sim2Real (Sec. V-A).

As in Fig. 3, we use the Gaussian Splatting (GS) primitives \mathcal{G}_{gs} [21] to represent the spatial model of the object, jointly encoding geometry and appearance. Following [22], we jointly optimize the GS model \mathcal{G}_{gs} and the poses of keyframes $\{\mathbf{q}_i\}_{i \in \mathcal{K}}$ from a keyframe buffer \mathcal{K} (detailed in Section V-A) based on the following loss

$$\min_{\{\mathbf{q}_i\}_{i \in \mathcal{K}}, \mathcal{G}_{\text{gs}}} \sum_{i \in \mathcal{K}} \lambda_{\text{pho}} E_i^{\text{pho}} + (1 - \lambda_{\text{pho}}) E_i^{\text{geo}} \quad (1)$$

with E_i^{pho} the photometric residual and E_i^{geo} the geometric loss for the keyframe i , respectively defined as

$$\begin{aligned} E_i^{\text{pho}} &= \|I(\mathcal{G}_{\text{gs}}, \mathbf{T}(\mathbf{q}_i)) - I_i\|_1, \\ E_i^{\text{geo}} &= \|D(\mathcal{G}_{\text{gs}}, \mathbf{T}(\mathbf{q}_i)) - D_i\|_1, \end{aligned} \quad (2)$$

where $I(\mathcal{G}_{\text{gs}}, \mathbf{T}(\mathbf{q}_i))$ renders the Gaussians \mathcal{G}_{gs} for pose \mathbf{q}_i by α -blending; $D(\mathcal{G}_{\text{gs}}, \mathbf{T}(\mathbf{q}_i))$ is the depth rendering; I_i and D_i are the observed image and depth of keyframe i ; and $\mathbf{T}(\cdot)$ converts pose to transformation matrix. For the optimization and implementation details, refer to [22].

The obtained \mathcal{G}_{gs} can be difficult to be directly integrated in physics engine. Thus, as the final step in this phase, we learn a neural signed distance function (SDF) representation, denoted as

$$\phi = \mathbf{SDF}_{\text{vision}}(\mathbf{p}), \quad \mathbf{p} \in \mathbb{R}^3 \quad (3)$$

via the ray-casting sampling on the rendered depth from the GS model \mathcal{G}_{gs} , as in Fig. 3. The obtained $\mathbf{SDF}_{\text{vision}}$ will then be used in the contact dynamics (simulator) in Section IV-B.

B. Contact Dynamics and Geometry Compensating

Our contact dynamics model is built on the MJX physics engine [11], [23], [24] with substantially redesigned collision detection routine to integrate the obtained neural SDF from the previous phase. The JAX implementation of the MJX physics engine enables GPU-parallel simulation and sampling-based optimization in this subsection. In generic notation, the contact dynamics [11] can be modeled as

$$\begin{aligned} \mathbf{M}(\mathbf{q}_k)(\mathbf{v}_{k+1} - \mathbf{v}_k) &= h\mathbf{c}(\mathbf{q}_k, \mathbf{v}_k) + h\mathbf{u}_k + \sum_i \mathbf{J}_i(\mathbf{q}_k)^\top \boldsymbol{\lambda}_i, \\ \mathbf{0} \leq \boldsymbol{\lambda}_i \perp \phi_i(\mathbf{q}_{k+1}) &\geq \mathbf{0}, \quad i \in \mathcal{C}(\mathbf{q}_k, \mathbf{SDF}_{\text{collision}}). \end{aligned} \quad (4)$$

Here h is the time interval, $\mathbf{q}_k \in \mathbb{R}^{n_q}$ and $\mathbf{v}_k \in \mathbb{R}^{n_v}$ are the generalized configuration and velocity of all degrees of freedom (object and robot), the inertia matrix $\mathbf{M}(\mathbf{q}_k) \in \mathbb{R}^{n_v \times n_v}$, the bias force $\mathbf{c}(\mathbf{q}_k, \mathbf{v}_k) \in \mathbb{R}^{n_v}$ includes gravity and inertial effects, applied force is $\mathbf{u}_k \in \mathbb{R}^{n_v}$. $\boldsymbol{\lambda}_i \in \mathbb{R}^{m_i}$ is the i -th frictional contact impulse with corresponding contact Jacobian $\mathbf{J}_i(\mathbf{q}_k) \in \mathbb{R}^{m_i \times n_v}$. $\phi_i(\mathbf{q}_{k+1}) \in \mathbb{R}^{m_i}$ defines the

general contact distance (i.e., gaps in normal direction and for dual cone faces of friction cone), with m_i being the contact dimension. $\mathcal{C}(\mathbf{q}_k, \mathbf{SDF}_{\text{collision}})$ is the collision detection module which returns the active contact set based on the collision SDF $\mathbf{SDF}_{\text{collision}}$ queries. The second inequality denotes complementarity constraint enforcing $\lambda_i \geq \mathbf{0}$, $\phi_i(\mathbf{q}_{k+1}) \geq \mathbf{0}$, and $\lambda_i^\top \phi_i(\mathbf{q}_{k+1}) = \mathbf{0}$. The configuration update follows the semi-implicit Euler step $\mathbf{q}_{k+1} = \mathbf{q}_k \oplus h\mathbf{v}_{k+1}$, with \oplus including quaternion integration.




Fig. 4: Collision geometry is a sum of visual geometry representation and a learnable geometry residual.

In our method, we modify the following aspects of the MJX physics engine to better integrate with neural geometry and contact dynamics learning.

a) *Geometry Residual*: In each physics simulation step (4), collision detection routine requires querying *collision geometry* $\mathbf{SDF}_{\text{collision}}$ to obtain necessary collision information. Visual geometry $\mathbf{SDF}_{\text{vision}}$ reconstructed in (3) can be imperfect to be directly used as collision SDF $\mathbf{SDF}_{\text{collision}}$, due to noisy and/or partial vision observation. Thus, we propose to obtaining the collision SDF $\mathbf{SDF}_{\text{collision}}$ by refining the visual geometry with a residual term:

$$\mathbf{SDF}_{\text{collision}}(\mathbf{p}) = \mathbf{SDF}_{\text{vision}}(\mathbf{p}) - \delta(\mathbf{p}). \quad (5)$$

The term $\delta(\mathbf{p})$ is called *geometry residual*, which compensates for visual geometry inaccuracies. The geometry residual $\delta(\mathbf{p})$ is learned along with other physics parameters (mass, inertia, friction coefficient) in (4) in Sec. IV-C.

b) *Collision Detection with Neural SDFs*: MJX physics engine mainly supports convex meshes and primitive shapes for collision detection. To handle arbitrary non-convex neural geometry represented by $\mathbf{SDF}_{\text{collision}}$, we customize MJX's collision detection routine following [6]. Specifically, collision detection with SDF geometry involves querying $\mathbf{SDF}_{\text{collision}}(\mathbf{p})$ with surface points \mathbf{p} sampled from surrounding entities (ground, walls, robot fingers) as shown in Fig. 4. JAX-GPU acceleration enables parallel distance queries with reasonably low computational overhead.

C. Learning Contact Dynamics

Learning contact dynamics involves estimating physical parameters in (4),

$$\theta = \{m_{\text{obj}}, \text{diag}[I_x, I_y, I_z]_{\text{obj}}, \mu, \delta\}, \quad (6)$$

including object mass, diagonalized inertia matrix, friction coefficient, and collision geometry residual (5). We update

all those parameters by minimizing the one-step prediction (simulation) loss of the contact dynamics:

$$\mathcal{L}(\theta) = \sum_{(\mathbf{q}_k, \mathbf{v}_k, \mathbf{q}_{k+1}) \in \mathcal{D}_{\text{tracking}}} \|\text{sim}_{\theta}(\mathbf{q}_k, \mathbf{v}_k) - \mathbf{q}_{k+1}\|_W^2, \quad (7)$$

where $\mathcal{D}_{\text{tracking}}$ is history tracking data from Sim2Real (Section V); sim_{θ} denotes one-step simulation based on the contact dynamics model (4) with integration $\mathbf{q}_{k+1} = \mathbf{q}_k \oplus h\mathbf{v}_{k+1}$; and the weight matrix W balances translation and rotation losses.

Note that for geometry residual learning in (7), we adopt a constant residual δ rather than a more expressive residual function $\delta(\mathbf{p})$. This is because the contact events are inherently sparse and intermittent, fine-tuning an expressive residual function (e.g., neural networks) can be unstable and prone to overfitting, although we acknowledge that this could be addressed by using regularization and careful parameterization. We will explore them in our future work. Here, a constant residual δ means correcting the visual geometry by uniformly expanding or shrinking by a margin. This provides a lightweight yet practically effective mechanism (in our cases) to refine contact behavior.

Due to the non-smoothness of contact dynamics (4), the gradient-based optimization is hard to apply to (7). While the MJX engine supports differentiable simulation but the gradients are limited in local contact modes, not useful to provide global optimization guidance [10]. On the other hand, MJX enables highly parallelized simulations across randomized environments, making it suitable for sample-based optimization. Thus, we use a variant of cross-entropy method (CEM) [25] to solve (7), where unlike standard CEM methods that update based on elite samples, we use softmax-weighted average over all samples (similar to [26]). Specifically, let j denote learning iteration, and N the sample size at each iteration. The particle $\theta_i^{(j+1)}$, $i = 1, 2, \dots, N$, at the $(j+1)$ -th iteration is sampled from $\theta_i^{(j+1)} \sim \mathcal{N}(\boldsymbol{\mu}^{(j+1)}, \boldsymbol{\Sigma})$, with

$$\boldsymbol{\mu}^{(j+1)} = \frac{\sum_{i=1}^N \exp(\gamma \mathcal{L}(\theta_i^{(j)})) \theta_i^{(j)}}{\sum_{i=1}^N \exp(\gamma \mathcal{L}(\theta_i^{(j)}))}. \quad (8)$$

where γ is the temperature (hyperparameter), and $\boldsymbol{\Sigma}$ is set as a diagonal constant matrix with entries depending on physical unit scales. This update is similar to the MPPI method [27].

V. SIM2REAL

This section presents Sim2Real (frontend), an online module that estimates object pose at each frame by combining two complementary cues: a visual pose estimate obtained by matching the current frame to nearby optimized keyframes, and a simulation state from the learned contact dynamics. We introduce a feature-based visual tracker in Sec. V-A, and then describe the pose optimization and adaptive fusion in Sec. V-B. Besides outputting the real-time pose, Sim2Real maintains a keyframe buffer and logs tracking history, which are used by Real2Sim for continuous updates.

A. Feature-Based Visual Tracker

Feature-based visual tracker computes relative pose between the current frame and an optimized keyframe (pose anchor) by detecting and matching visual features. The pose anchor is the nearest optimized keyframe from Real2Sim (Sec. IV-A). The process follows [4], summarized below.

a) *Feature Matching*: We segment object region using SAM2 [28] with prompt given only at the first frame. SuperPoint [29] extracts 3D object features from the current frame and reference keyframe. LightGlue [30] performs feature matching, followed by RANSAC-based pose estimator to filter wrong correspondences.

b) *Keyframe Buffer*: We use keyframes as pose anchors instead of the previous frames to avoid long-term drift. The keyframe is selected following the setup of [4], [16], and a keyframe buffer \mathcal{K} is built to contain multi-view but sparse keyframes, whose poses are continuously optimized in Real2Sim (Sec. IV-A). The first frame is always a keyframe. New frames are added if rotational geodesic distances to existing keyframes exceed threshold [4], [16].

c) *Feature-Based Pose Estimation*: For each new frame, the closest keyframe is selected from the keyframe buffer. Object pose is estimated from matched (and filtered) features between the current frame and selected keyframe. We use a single frame pair instead of multi-keyframe pose graph optimization [4], [16] for real-time performance, with accuracy refined by fusion with contact physics prediction.

B. Pose Optimization by Adaptive Fusion

While object pose estimation from feature-based tracking and contact dynamics can be formulated as a Kalman filter [31], [32], we propose a simple yet effective strategy that avoids the complexity of optimization and update. At frame k , we estimate the object pose \mathbf{q}_k by jointly minimizing visual feature correspondence loss and one-step prediction loss of the contact dynamics (simulation):

$$\min_{\mathbf{q}_k} w_k \mathcal{L}_{\text{vision}}(\mathbf{q}_k) + (1 - w_k) \mathcal{L}_{\text{dynamics}}(\mathbf{q}_k), \quad (9)$$

where $\mathcal{L}_{\text{vision}}(\mathbf{q}_k)$ is the visual feature correspondence loss; $\mathcal{L}_{\text{dynamics}}(\mathbf{q}_k)$ is the contact dynamics prediction loss; the weight $w_k \in [0, 1]$ balances the contributions of the two and will be *adaptively set* based on the quality of feature correspondences. We explain the three terms below.

Visual feature correspondence loss $\mathcal{L}_{\text{vision}}(\mathbf{q}_k)$ quantifies sum of squared distances of matched 3D feature points between the current frame and selected keyframes:

$$\mathcal{L}_{\text{vision}}(\mathbf{q}_k) = \sum_j \left\| \mathbf{T}(\mathbf{q}_k) \mathbf{p}_k^j - \mathbf{T}(\mathbf{q}_n) \mathbf{p}_n^j \right\|_2^2, \quad (10)$$

where $(\mathbf{p}_k^j, \mathbf{p}_n^j)$ is the j -th matched feature pair from current frame pose \mathbf{q}_k and selected keyframe pose \mathbf{q}_n .

One-step contact dynamics prediction loss is defined as

$$\mathcal{L}_{\text{dynamics}}(\mathbf{q}_k) = \left\| \mathbf{q}_k - \text{sim}(\mathbf{q}_{k-1}, \mathbf{v}_{k-1}) \right\|_2^2, \quad (11)$$

where sim is (4) learned in Real2Sim (Section IV-C), \mathbf{q}_{k-1} is previous frame’s robust pose estimate, and \mathbf{v}_{k-1} is velocity




Fig. 5: Adaptive weighting parameter

estimated via finite difference on visual pose (if reliable) or from simulator dynamics evolution.

Weight $w_k \in [0, 1]$ controls balance between physics-based prediction and feature-based estimation. We set w_k adaptively based on feature matching quality, quantified by number of matched feature pairs after filtering (Section V-A). Higher reliable feature matches indicate favorable visual conditions (less occlusion and motion blur), so visual estimate gets more weight. We define w_k as sigmoid-like function of matched feature pairs (Fig. 5).

VI. EXPERIMENTS

We evaluate TwinTrack across diverse contact-rich scenarios, focusing on its two key components: (1) Real2Sim: how effectively it learns contact dynamics to align with and predict real-world contact-rich motion, and (2) Sim2Real: how much it improves tracking robustness in highly dynamic, contact-rich environments. These evaluations highlight TwinTrack’s ability to bridge sim-to-real gap and enable robust perception.

A. Tasks and Dataset



Fig. 6: Two contact-rich scenarios with different objects.

Our evaluation centers on two contact-rich scenarios across various objects (Fig. 6):

Scenario I: Object falling with rich impacts: The scene consists of two walls and ground forming a corner with three orthogonal planar surfaces. Objects are released from human hand with initial velocity, undergo free fall, and collide with walls and ground. The scene is impact-rich, and the instantaneous impulses cause abrupt velocity changes of object motion, leading to frequent motion blur. Objects used include snack box, tissue box, milk bottle, Spam can, 3D-printed duck and mustard bottle.

Scenario II: In-hand manipulation with occlusion: LEAP four-fingered robotic hand [33] manipulates objects with random policies. The motion involves rich contact between the object and hand fingers and palm. The primary challenge arises from partial occlusions of the object. The object also undergoes abrupt impacts from randomly moving fingers. We use the Spam can and 3D-printed duck, as their sizes make them well-suited for handheld interaction.




Fig. 7: First row: Open-loop simulation using random physical parameters (prior to Real2Sim). Second row: Open-loop simulation of the contact dynamics learned but without geometry refinement, i.e., directly using visual geometry for simulation. Third row: Open-loop simulation of the learned contact dynamics, including geometry refinement. Fourth row: object’s actual trajectory in real world.




Fig. 8: Convergence of all physics parameters in contact dynamics learning for all scenario-object pairs. Each learning runs for 5 trials with different random seeds. The solid lines represent the mean values, and the shaded areas indicate the standard deviation across the five trials.

We collect RGB-D videos using an Intel RealSense D435i camera for both scenarios. In Scenario II, we additionally record the joint states of the robotic hand. For each scenario-object pair, we capture four trajectories under varying motion conditions: three with mild contact interactions (such as planar sliding and lateral shaking) to facilitate the learning of contact dynamics, and one with highly dynamic interactions involving abrupt and rich impacts for robustness evaluation. In Scenario I, data collection begins from the free-motion phase following object release. The object’s initial velocity is estimated with a visual tracker applied to the first few frames.

For ground-truth tracking, we offline reconstruct high-quality textured mesh models using Scaniverse app [34], then use FoundationPose [12] to generate pseudo ground-truth object poses for all video frames.

B. Evaluation of Real2Sim

For each scenario-object pair, Real2Sim learns contact dynamics using multiple trajectories with diverse contact interactions: bouncing against walls, sliding along ground, and rich contact in robot hand.

TABLE I: One-step prediction error (7) ($100\times$ position error + quaternion error) of contact dynamics model before/after learning.

	snack box	tissue box	bottle	spam	duck	mustard	LEAP spam	LEAP duck
Before	18.78	9.57	9.82	7.24	5.13	10.18	4.18	23.03
After w/o δ	0.29	0.47	0.87	1.16	0.38	0.56	0.18	0.38
After with δ	0.20	0.22	0.70	0.43	0.21	0.15	0.13	0.28

Table I shows the prediction error (7) of the contact dynamics model before and after learning. "Before": initial contact dynamics model (4) with random physical parameters shows substantial gap from reality, with simulated objects exhibiting exaggerated rebounds/unrealistic sliding due to inaccurate mass/inertia parameters. "After w/o δ ": learning all physical parameters (6) except geometry residual δ (i.e., only using visual geometry) reduces gap but suffers from inaccurate contact behaviors. "After with δ ": learning all parameters including δ further lowers prediction error, indicating better alignment between the contact dynamics model and reality. Learned δ may be imprecise due to camera calibration errors but remains physically plausible for downstream tracking.

TABLE II: The learned mass [g] of each object versus ground truth.

	snack box	tissue box	bottle	spam	duck	mustard	LEAP spam	LEAP duck
Mass	40	50	24	28	44	65	28	44
Learned	38 ± 3	55 ± 5	32 ± 10	15 ± 5	38 ± 7	46 ± 5	31 ± 5	47 ± 10

Fig. 8 shows the learning convergence of all physical parameters for each scenario-object pair (5 trials with different random seeds). Table II reports learned mass versus ground-truth. Despite convergence trends, noticeable gaps remain between learned and ground-truth parameters due to: (1) weak observability of parameters like friction/inertia when motion lacks sufficient richness/excitation in corresponding degrees of freedom; (2) MJX’s simplified contact handling (point/-

TABLE III: Evaluation for Sim2Real robust tracking [cm].

Metrics	Method	snack box	tissue box	bottle	spam	duck	mustard	Leap Spam	Leap Duck	Avg.±Std.
ADD↓	Baseline	2.20	2.14	3.58	3.54	2.66	7.13	1.88	1.01	3.01±1.86
	w/o dynamics	1.54	9.50	20.29	5.83	3.54	4.81	6.75	5.45	7.21±5.76
	random	>100.0	>100.0	>100.0	>100.0	1.83	2.74	1.70	3.22	>50.0
	w/o δ	1.48	1.05	0.82	0.70	0.89	2.73	1.72	3.11	1.56±0.91
	Proposed	1.37	1.04	0.78	0.67	0.83	2.62	1.67	0.87	1.23±0.65
ADD-S↓	Baseline	1.12	1.11	1.45	1.73	1.36	1.26	0.81	0.41	1.15±0.40
	w/o dynamics	0.78	2.50	12.24	1.26	1.28	2.06	2.12	1.61	2.98±3.78
	random	>100.0	>100.0	>100.0	>100.0	0.89	1.32	0.77	1.19	>50.0
	w/o δ	0.79	0.55	0.40	0.34	0.49	1.31	0.77	1.31	0.74±0.38
	Proposed	0.74	0.54	0.39	0.33	0.39	1.25	0.75	0.39	0.59±0.31

rigid contact) oversimplifying real-world contact interactions, creating inherent modeling bias: the learning converges to parameters fitting observed motion well despite deviating from true ones. We also note that the learned parameters for Spam/Duck differ between falling and manipulation scenarios. This is likely due to LEAP hand model-hardware mismatch (the hand control parameters not updated in simulation). Additionally, friction and mass do not converge to single consistent values across runs. This may be due to (i) trajectory-dependent excitation of contact and motions, yielding different parameters, and (ii) coupling in optimization, where friction and mass can partially trade off to absorb modeling residuals.

Fig. 7 shows multi-step open-loop simulation of the contact dynamics model for tissue box falling alongside real-world video (last row "reality"). Initial position/velocity are aligned. First row ("before learning"): randomly initialized contact dynamics model shows highly unrealistic motion (exaggerated rebounds, missed contact). Second row ("w/o δ "): simulation after learning physical parameters (6) except geometry residual δ . It shows incorrect contact locations due to inaccurate visual geometry, which causes object to "float" without contact. Third row ("with δ "): simulation after learning all parameters including δ . The result closely resembles real-world motion, capturing contact onset, rebound, and settling behavior, demonstrating Real2Sim's effectiveness in learning collision-accurate geometry and contact dynamics.

Despite visually appealing sim-real alignment in Fig. 7, accurate long-horizon contact-rich scene prediction is still challenging due to chaotic/sensitive nature of contact physics. Minor variations in surface geometry can lead to large motion discrepancies. Also, small initial condition perturbations yield dramatically different outcomes. Thus, while we claim Real2Sim enables the learning of an effective simulation model from visual inputs, it is best suited for one-step or few-step prediction. For longer horizons, it is preferable to integrate with real-time visual feedback to continually correct the accumulation of prediction errors, as in our Sim2Real robust tracking below.

C. Evaluation of Sim2Real

We evaluate Sim2Real for robust tracking in both scenarios and use Average Distance of Model Points (ADD) and its symmetric variant (ADD-S) metrics [cm] for tracking accuracy (Table III). As mentioned before, the pseudo Ground-truth poses are obtained from FoundationPose [12] with

high-quality textured mesh models obtained from Scaniverse software. For comparison, we include a baseline, which is model-free tracking also using FoundationPose method [12] but fed with the same set of keyframes used in our Real2Sim.

We perform three ablation studies to assess the design contributions: "w/o dynamics" replaces contact dynamics with constant-velocity prediction (vision-only baseline); "random" uses contact dynamics with random physical parameters; "w/o δ " learns all physical parameters except δ (visual geometry is used in contact dynamics model); "proposed" learns all parameters including δ .

Table III shows the proposed method achieves best overall performance across all scenarios and objects. Learning contact physics and geometry residual δ significantly outperforms all ablated variants. Baseline and "w/o dynamics" rely purely on visual tracking without contact physics, limiting performance under motion occlusions and blur. "w/o dynamics" introduces constant-velocity prior but lacks contact interaction prediction capability, causing drift/tracking errors during environment interaction. "w/o δ " and "proposed" variants highlight the importance of accurate collision geometry for physics-aware tracking. The learned collision geometry residual δ effectively corrects visual geometry errors, leading to more accurate contact dynamics and higher Sim2Real tracking performance.

The tracking time cost for each module is summarized as follows. Experiments are conducted on a PC with a Ryzen 5955WX CPU and an RTX 4090 GPU. For each frame, SAM2 takes around 12ms, and SuperPoint 8ms. The dynamics one-step prediction including neural SDF collision completes in 5ms. Keyframe selection and LightGlue matching consume approximately 15ms, while the adaptive fusion optimization step adds an additional 5ms. Other operations take about 3ms. In total, the system can achieve a processing rate of over 20Hz, enabling real-time performance for 6-DoF tracking.

D. Failure Modes and Limitations

Our system can fail when visual observations are degraded or unavailable for extended periods. In particular, under prolonged vision missing, pose estimation becomes primarily driven by dynamics prediction; despite real2sim alignment, modeling errors can accumulate over long horizons and cause drift. The method also needs a reasonably accurate initial state to start: poor vision-based initialization may lead to incorrect collision geometry, which in turn degrades subsequent dynamics learning and prediction.

VII. CONCLUSIONS AND FUTURE WORK

We present TwinTrack, a physics-aware tracking framework bridging visual observations and contact physics for robust tracking of unknown dynamic objects in contact-rich environments. TwinTrack includes Real2Sim (builds contact dynamics model from visual observations through geometry estimation/refinement and physical property identification) and Sim2Real (uses learned contact dynamics to enhance 6-DoF pose visual tracking robustness/accuracy). Evaluation demonstrates robust real-time performance against occlusion and motion blur in highly dynamic contact-rich environments.

In Real2Sim, we currently use a uniform residual, which cannot handle spatially varying errors in object geometry. We also assume a single rigid object and monocular RGB-D with a first-frame prompt, it cannot track multiple objects or deformable ones. In the future work, we will consider extending to multi-objects tracking scenarios and learning a spatial compensation that varies across the object's surfaces to enhance reconstruction accuracy. Also, our aligned Real2Sim environment may provide a more suitable training substrate for policies learning, potentially improving robustness and real-world transfer; we plan to investigate it in future work.

REFERENCES

- [1] A. M. Okamura, N. Smaby, and M. R. Cutkosky, "An overview of dexterous manipulation," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 255–262.
- [2] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.
- [3] N. Pfaff, E. Fu, J. Binagia, P. Isola, and R. Tedrake, "Scalable real2sim: Physics-aware asset generation via robotic pick-and-place setups," *arXiv preprint arXiv:2503.00370*, 2025.
- [4] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Müller, A. Evans, D. Fox, J. Kautz, and S. Birchfield, "Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 606–617.
- [5] F. Pfeiffer and C. Glocker, *Multibody dynamics with unilateral contacts*. Springer science & business media, 2000, vol. 421.
- [6] W. Yang and W. Jin, "Contactsdf: Signed distance functions as multi-contact models for dexterous manipulation," *IEEE Robotics and Automation Letters*, vol. 10, no. 5, pp. 4212–4219, 2025.
- [7] W. Jin and M. Posa, "Task-driven hybrid model reduction for dexterous manipulation," *IEEE Transactions on Robotics*, vol. 40, pp. 1774–1794, 2024.
- [8] W. Jin, "Complementarity-free multi-contact modeling and optimization for dexterous manipulation," *Robotics: Science and Systems*, 2025.
- [9] S. Pfrommer, M. Halm, and M. Posa, "Contactnets: Learning discontinuous contact dynamics with smooth, implicit representations," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 2279–2291.
- [10] H. J. Suh, M. Simchowitz, K. Zhang, and R. Tedrake, "Do differentiable simulators give better policy gradients?" in *International Conference on Machine Learning*. PMLR, 2022, pp. 20 668–20 696.
- [11] D. Google, "Mujoco-xla (mjax)," <https://github.com/google-deeppmind/mujoco/tree/main/mjx>, 2025.
- [12] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "Foundationpose: Unified 6d pose estimation and tracking of novel objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17 868–17 879.
- [13] Y. Labbé, J. Carpentier, M. Aubry, and J. Sivic, "Cosypose: Consistent multi-view multi-object 6d pose estimation," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*. Springer, 2020, pp. 574–591.
- [14] Y. Labbé, L. Manuelli, A. Mousavian, S. Tyree, S. Birchfield, J. Tremblay, J. Carpentier, M. Aubry, D. Fox, and J. Sivic, "Megapose: 6d pose estimation of novel objects via render & compare," *arXiv preprint arXiv:2212.06870*, 2022.
- [15] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, "Normalized object coordinate space for category-level 6d object pose and size estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 2642–2651.
- [16] B. Wen and K. Bekris, "Bundletrack: 6d pose tracking for novel objects without instance or category-level 3d models," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8067–8074.
- [17] E. Gärtner, M. Andriluka, E. Coumans, and C. Sminchisescu, "Differentiable dynamics for articulated 3d human motion reconstruction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 13 190–13 200.
- [18] J. Z. Zhang, S. Yang, G. Yang, A. L. Bishop, S. Gurumurthy, D. Ramanan, and Z. Manchester, "Slomo: A general system for legged robot motion imitation from casual videos," *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 7154–7161, 2023.
- [19] H. Jiang, H.-Y. Hsu, K. Zhang, H.-N. Yu, S. Wang, and Y. Li, "Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos," *arXiv preprint arXiv:2503.17973*, 2025.
- [20] B. Bianchini, M. Zhu, M. Sun, B. Jiang, C. J. Taylor, and M. Posa, "Vysics: Object reconstruction under occlusion by fusing vision and contact-rich physics," *arXiv preprint arXiv:2504.18719*, 2025.
- [21] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [22] H. Matsuki, R. Murai, P. H. Kelly, and A. J. Davison, "Gaussian splatting slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 18 039–18 048.
- [23] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.
- [24] C. D. Freeman, E. Frey, A. Raichuk, S. Girgin, I. Mordatch, and O. Bachem, "Brax - a differentiable physics engine for large scale rigid body simulation," 2021.
- [25] R. Y. Rubinfeld and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2004.
- [26] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [27] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 1714–1721.
- [28] N. Ravi, V. Gabeur, et al., "Sam 2: Segment anything in images and videos," *arXiv preprint arXiv:2408.00714*, 2024.
- [29] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 337–337 12.
- [30] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, "LightGlue: Local Feature Matching at Light Speed," in *ICCV*, 2023.
- [31] S. J. Julier and J. K. Uhlmann, "New extension of the kalman filter to nonlinear systems," in *Signal processing, sensor fusion, and target recognition VI*, vol. 3068. Spie, 1997, pp. 182–193.
- [32] E. A. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 adaptive systems for signal processing, communications, and control symposium (Cat. No. 00EX373)*. Ieee, 2000, pp. 153–158.
- [33] K. Shaw, A. Agarwal, and D. Pathak, "Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning," *Robotics: Science and Systems (RSS)*, 2023.
- [34] I. Niantic, "Scaniverse: 3d scanner app," 2025. [Online]. Available: <https://scaniverse.com/>