

Few-Shot Neural Differentiable Simulator: Real-to-Sim Rigid-Contact Modeling

Zhenhao Huang, Siyuan Luo, Bingyang Zhou, Ziqiu Zeng and Fan Shi*

Abstract—Accurate physics simulation is essential for robotic learning and control, yet analytical simulators often fail to capture complex contact dynamics, while learning-based simulators typically require large amounts of costly real-world data. To bridge this gap, we propose a few-shot real-to-sim approach that combines the physical consistency of analytical formulations with the representational capacity of graph neural network (GNN)-based models. Using only a small amount of real-world data, our method calibrates analytical simulators to generate large-scale synthetic datasets that capture diverse contact interactions. On this foundation, we introduce a mesh-based GNN that implicitly models rigid-body forward dynamics and derive surrogate gradients for collision detection, achieving full differentiability. Experimental results demonstrate that our approach enables learning-based simulators to outperform differentiable baselines in replicating real-world trajectories. In addition, the differentiable design supports gradient-based optimization, which we validate through simulation-based policy learning in multi-object interaction scenarios. Extensive experiments show that our framework not only improves simulation fidelity with minimal supervision but also increases the efficiency of policy learning. Taken together, these findings suggest that differentiable simulation with few-shot real-world grounding provides a powerful direction for advancing future robotic manipulation and control.

I. INTRODUCTION

From basic grasping and placing to complex assembly and tool use, most manipulation tasks depend on intricate contact interactions that determine whether the operation succeeds or fails. Simulation provides a critical tool for allowing robots to learn and plan these tasks in a safe and controlled environment. Yet, despite decades of research in physics-based simulation, existing methods still struggle to accurately capture the contact dynamics observed in the real world.

In addition, these simulation approaches for robotic manipulation face a fundamental trade-off between physical accuracy and computational efficiency, particularly when modeling complex contact dynamics at scale. Traditional analytical simulators in robotics like IsaacLab [18] and MuJoCo [29] focus on providing physically accurate and stable simulations but suffer from computational bottlenecks when scaling to complex scenes with numerous interacting objects. These simulators rely on constraint-based solvers that, while theoretically sound, become prohibitively expensive for contact-rich scenarios. Furthermore, such solvers are vulnerable and sensitive to parameters such as contact stiffness, damping, and coefficients of friction that are often difficult to measure in real-world

The authors are with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077. huangzhenhao@u.nus.edu

* Corresponding author.

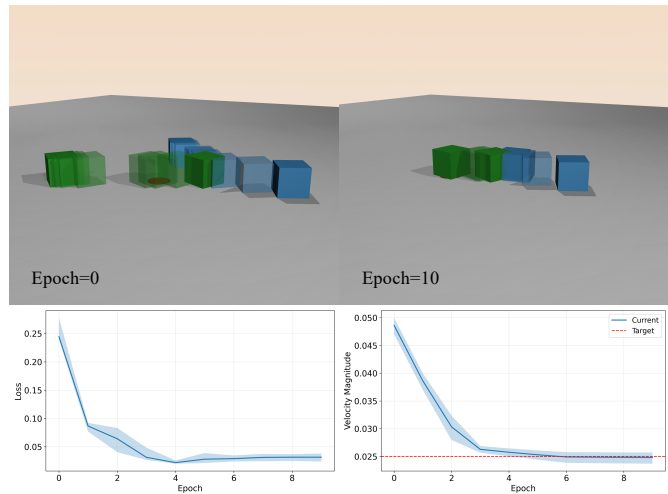


Fig. 1: *Top*: A blue cube is pushed to collide with a green cube, with the objective of stopping it at the red target area. The initial pushing velocity of the blue cube is optimized using stochastic gradient descent in our differentiable simulator, successfully achieving the goal. *Bottom*: Convergence of optimization across multiple runs, showing the loss curve (left) and the magnitude of the initial pushing velocity (right).

scenarios. Recently differentiable analytical simulators have emerged [15, 8, 14], enabling gradient-based optimization for algorithms like motion planning and reinforcement learning [28, 17], but they still rely on the underlying contact modeling. Learning a portion of the simulation such as the signed distance and contact Jacobian [26, 7], alleviate some of these issues. Recent advances in learnable simulators using GNNs have demonstrated promising results in achieving accuracy and computational efficiency, yet they require massive training data obtained via either cheap but inaccurate simulation or accurate but time-consuming real-world collection [5, 27, 16, 24, 3, 4].

In this paper, we propose a novel framework that bridges the gap between real-world contact dynamics and learnable simulation. Our main contributions are as follows:

- We propose a rigid-contact differentiable simulator based on GNN with surrogate gradients of nearest points with respect to object states for collision detection.
- We develop a few-shot real-to-sim data scaling pipeline with contact parameter identification to scale limited real-world data into large, diverse datasets with contact-rich interactions while maintaining physical realism.
- We show that our simulator outperforms the differentiable

simulator Brax and achieves performance comparable to MuJoCo on real-world data, and highlight its potential for generalization to contact-rich scenes and gradient-based optimization in complex rigid-contact scenarios.

II. RELATED WORKS

A. Differentiable Physical Simulation

Differentiable simulators have great potential in robotics, as they can be integrated with gradient-based optimization and learning algorithms for various tasks, such as system identification, trajectory optimization, and reinforcement learning [28]. Recent works have proposed differentiable simulation for rigid-body dynamics [11, 14, 8, 15, 10, 13, 22]. However, these methods often rely on analytical models that may not sufficiently capture real-world contact interactions. In contrast, our approach proposes a learnable GNN-based differentiable simulator that can adapt to real-world contact dynamics.

B. Learnable Simulation

Learnable simulators leverage data-driven methods to model the complex dynamics of physical systems, which are flexible in not only object types but also learnable parameters. The simulator can learn a residual physics neural network such as signed distance functions to model accurate rigid-contact simulation [26, 7, 10], or learn the entire dynamics using graph neural networks (GNNs) [27, 3, 4]. But large amounts of training data are required to train these models effectively. Our method addresses this by developing a few-shot data scaling pipeline which requires minimal real-world data.

C. Real-to-Sim

Some previous works have explored the real-to-sim approach to bridge the gap between simulation and reality. ContactNets [26] utilize AprilTags [12] to obtain accurate 6D poses of objects from real-world images. GradSim [22] employs differentiable simulation with differentiable rendering based on meshes to optimize physical parameters using simulated images, while DANO [14] uses neural density fields learned directly from real-world images for realistic simulation. We follow ContactNets, but with data scaling to reduce the amount of real-world data required for real-to-sim transfer.

D. Collision Detection

Collision detection is a crucial component of contact simulation. While the collision detection algorithms are not the focus of this work, the differentiability of collision detection is particularly relevant. There are several works on differentiable collision detection, either providing approximated gradients [19] or limited to convex primitives [31, 30]. In our work, we derive the gradients of contact points or nearest points regardless of the differentiability of collision detection based on a reasonable assumption.

TABLE I: Comparison of our simulator against existing analytical and learned simulators.

Simulator	Contact Model	Contact Accuracy	Scalability	Differentiable
IsaacLab	PBD [21]	Medium	High	✗
MuJoCo	CCP	High	High	✗
Brax	LCP	Low	Medium	✓
ContactNets	DNN+LCP	Medium	Low	✗
FIGNet	GNN	Medium	Medium	✗
Ours	GNN	High	High	✓

III. METHOD

The overall framework of our approach is illustrated in Figure 2. We begin by identifying key contact parameters that are critical for contact modeling but hard to measure directly (Section III-A). Upon fixing identified contact parameters in the simulation, we proceed to scale the real-world data concerning the initial states and number of objects by generating a large dataset of diverse contact interactions (Section III-B). Finally, we propose a fully differentiable GNN-based simulator that is well-trained by the scaled dataset (Section III-C).

A. Contact Parameter Identification

Given a set of collected real-world trajectories of dynamic rigid bodies, we formulate the contact parameter identification as an optimization problem. Let $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^K$ denote the collected dataset, where $\mathcal{D}_k = \{(\mathbf{x}_t^i, \mathbf{R}_t^i)\}_{i=1}^N\}_{t=0}^{T_k}$ is the set of positions and orientations of N rigid bodies for the k -th demonstration that lasts for $T_k + 1$ time steps. To obtain the missing velocity information, we use backward finite difference to compute the linear and angular velocities from the position and orientation trajectories. Similar to [1], the simulator starts from the same initial state as the real-world trajectory to get the corresponding simulated trajectory $\widehat{\mathcal{D}}_k(\boldsymbol{\theta}) = \{(\widehat{\mathbf{x}}_t^i, \widehat{\mathbf{R}}_t^i)\}_{i=1}^N\}_{t=1}^{T_k}$. The objective is to find the optimal contact parameters $\boldsymbol{\theta}^*$ that minimize the discrepancy between simulated and real-world trajectories:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \sum_{k=1}^K \sum_{t=1}^{T_k} \sum_{i=1}^N \left(\frac{1}{w_i} \|\mathbf{x}_t^{i,k} - \widehat{\mathbf{x}}_t^{i,k}\|_2 + |\text{Angle}(\mathbf{R}_t^{i,k}, \widehat{\mathbf{R}}_t^{i,k})| \right) \quad (1)$$

where w_i represents the positional scaling factor for the i -th body, and $\text{Angle}(\cdot)$ is the angular difference between the real and simulated orientations.

In this work, we focus on identifying contact parameters in a simulator, using the high-fidelity physics simulator MuJoCo as our testbed. The contact solver in MuJoCo is primarily controlled by two parameter sets: *solimp* and *solref*. *solimp* = $\{d_0, d_{\text{width}}, \text{width}, \text{midpoint}, \text{power}\}$ defines the shape of the impedance function, mapping constraint violation to the generated force, while *solref* specifies the time constant and damping ratio of the implicit spring-damper model. In addition, the lateral friction coefficient μ is also a critical parameter for frictional contact interactions.

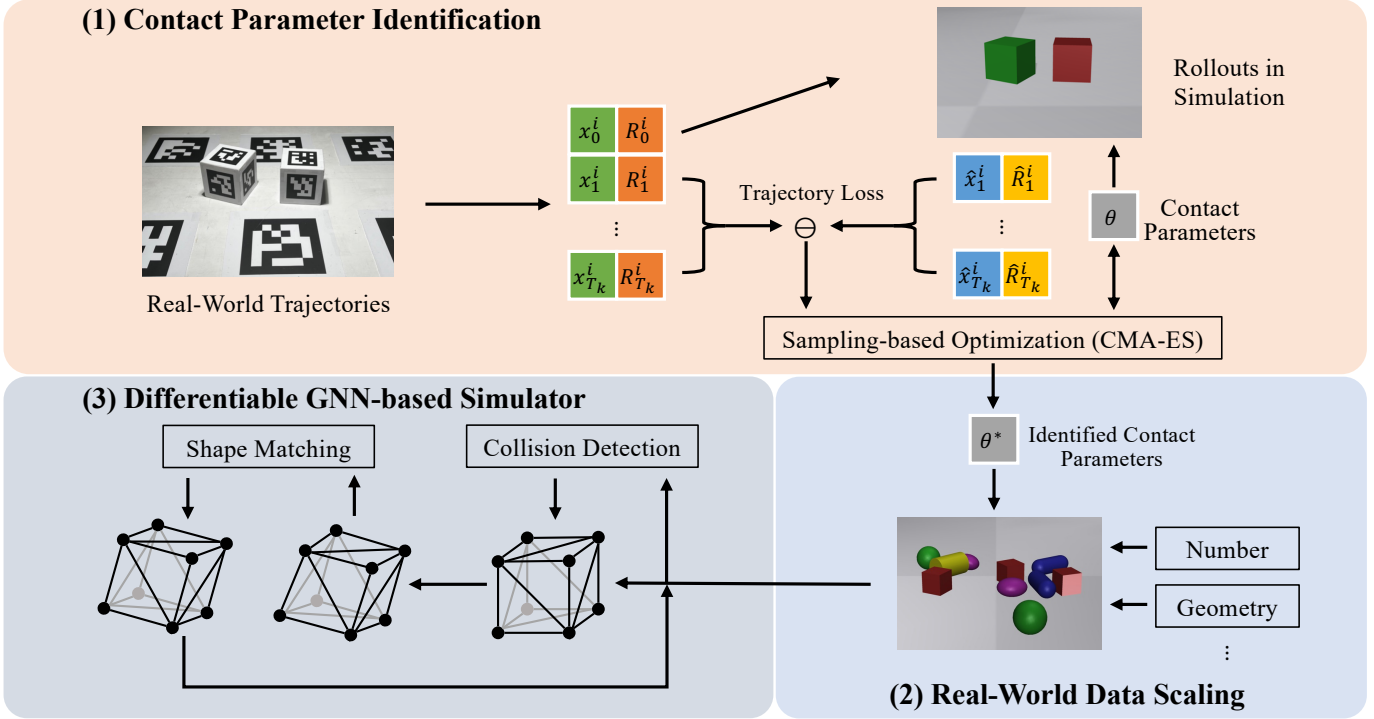


Fig. 2: Illustration of our framework. (1) We utilize sampling-based identification of contact parameters to minimize the trajectory loss between real world and simulation, followed by (2) data scaling in terms of quantity, geometry, *etc.* (3) The GNN is trained on the scaled dataset, equipped with differential collision detection and shape matching to achieve differentiability. The resulting simulator is capable of accurately modeling complex contact interactions.

These parameters $\theta = \{\text{solimp}, \text{solref}, \mu\}$ govern the stiffness, dissipation, and stability of contact interactions, and are thus the primary targets of our contact parameter identification. Furthermore, since MuJoCo is non-differentiable, we employ a gradient-free optimization approach, CMA-ES [2], to solve the optimization problem (1). These identified parameters ensure that the subsequent data scaling produces physically realistic trajectories.

B. Contact-Aware Data Scaling

Training an accurate GNN-based simulator requires a large dataset of diverse contact scenarios [24, 4, 3]. However, constructing such a large real-world dataset is often expensive and time-consuming. Therefore, we propose to scale the scene by generating synthetic data using the contact parameters identified above. Unlike data augmentation techniques that apply transformations to real-world data to prevent overfitting, our data scaling method enriches the distribution of training data, thereby enhancing the GNN’s ability to generalize to unseen scenarios. Since the identified MuJoCo is more physically realistic and stable for contact simulation, the distribution of scaled data is also relatively close to the real world, which can help our simulator better capture real-world contact dynamics.

To scale data, we start with a base scene containing a few objects and systematically change the number and properties (*e.g.*, geometry and mass) of objects with varying initial states in MuJoCo. This allows us to create a rich dataset that captures

a wide range of contact interactions while maintaining realism.

C. Differentiable GNN-based Simulator

Our GNN-based simulator is mainly composed of three parts: a collision detection algorithm, a GNN solver, and a shape matching module. The collision detection algorithm computes the minimum distance between objects for graph construction. The GNN solver predicts the accelerations of nodes and integrates them over time to update the positions, while the shape matching module ensures that the objects maintain their shapes.

1) *Mesh-based GNN solver*: Our GNN solver follows the FIGNet architecture [4]. The core idea is to represent each rigid body as a triangle mesh and model the dynamics through message passing on a graph where nodes correspond to mesh vertices and edges encode spatial relationships.

Graph Construction. For all objects in the scene, a graph is constructed containing two kinds of nodes (mesh and object nodes) and three kinds of bidirectional edges (mesh-mesh, object-mesh, and face-face edges).

Message Passing. The GNN utilizes the standard message-passing framework [24] with L layers. At each layer ℓ , nodes and edges are updated through:

$$\mathbf{e}_{ij}^{\ell+1} = \phi_e^\ell(\mathbf{e}_{ij}^\ell, \mathbf{v}_i^\ell, \mathbf{v}_j^\ell) \quad (2)$$

$$\mathbf{v}_i^{\ell+1} = \phi_v^\ell(\mathbf{v}_i^\ell, \sum_j \mathbf{e}_{ij}^{\ell+1}) \quad (3)$$

where ϕ_e^ℓ and ϕ_v^ℓ are learnable multi-layer perceptrons (MLPs), e_{ij} represents edge features, and v_i represents node features.

Output. The final decoder decodes the node features from the last message-passing layer and outputs accelerations for each mesh node:

$$\mathbf{a}_t^i = \phi_{\text{decoder}}(\mathbf{v}_t^i) \quad (4)$$

These accelerations are then integrated to update positions using a Verlet integrator $\mathbf{p}_{t+1}^i = \mathbf{a}_t^i + 2\mathbf{p}_t^i + \mathbf{p}_{t-1}^i$. During inference, we implement the shape matching algorithm [20] to compute the rigid transformations for each object and project the predicted nodes to the transformed mesh instead of directly using the predicted node positions to keep the mesh topology consistent. Details of the GNN approach can be found in [4].

2) *Surrogate Gradients of Collision Detection:* In our simulator, while the GNN solver is inherently differentiable, achieving end-to-end differentiability during inference requires making both collision detection and shape matching components differentiable. Shape matching involves standard matrix operations such as singular value decomposition and rotation matrix computations, which are readily supported by modern machine learning frameworks like PyTorch with automatic differentiation. Here we focus our attention on deriving gradients through collision detection, which presents a significant challenge as most of the collision detection algorithms either are non-differentiable or provide approximate gradients [19].

We implement the discrete collision detection (DCD) via Coal [23], a state-of-the-art collision detection library that provides efficient and robust collision detection algorithms over various geometric primitives. Since our GNN solver is mesh-based, the Gilbert-Johnson-Keerthi (GJK) algorithm [9] and the Expanding Polytope Algorithm (EPA) [6] are employed to compute the minimum distance or the penetration depth between convex meshes. However, these algorithms are non-differentiable due to their discontinuities. To address this, we derive the surrogate gradients that provides informative gradients and enables backpropagation through DCD.

Generally, given the states of two objects (O_A, O_B), the DCD algorithm outputs a set of contact pairs (o_A^i, o_B^j) together with the corresponding nearest points (p_A^i, p_B^j), where $o_A^i \in O_A$ and $o_B^j \in O_B$ denote the triangles of each object. By its very nature, DCD requires a distance threshold d_ϵ to define potential contact pairs: if the distance between the nearest points of a triangle pair satisfies $|p_A^i - p_B^j|^2 < d_\epsilon$, this pair is regarded as a potential contact. The choice of d_ϵ is crucial. When the threshold is too small, certain true contacts may be missed. Conversely, when the threshold is too large, a substantial number of non-effective contacts are introduced, among which only a fraction will actually contribute in the subsequent contact response stage.

Consequently, by setting the threshold to a slightly larger range, we can make a reasonable assumption: due to the redundancy of detected contact pairs, the original set of contacts is sufficiently rich to support correct contact resolution, even when the motion undergoes small variations within the same time step. In this sense, the redundancy introduced by

a larger threshold serves as a safeguard against local motion perturbations.

Returning to the differentiable simulation framework, and based on the above assumption, we regard the set of contact pairs as fixed within the current time step. Consequently, the requirement to differentiate through the contact detection process itself can be reasonably removed.

To obtain the gradients of the nearest points (p_A^i, p_B^j) with respect to the object states, we assemble the linear velocities of center of masses and angular velocities as $u_A = [v_A^T, \omega_A^T]^T$ and $u_B = [v_B^T, \omega_B^T]^T$, respectively. Similarly, the generalized positions are defined as $q_A = [x_A^T, R_A^T]^T$ and $q_B = [x_B^T, R_B^T]^T$, where x and R denote the position of the center of mass and the orientation expressed by quaternions of each object. Then we define the kinematic mapping between the generalized velocities $u = [u_A^T, u_B^T]^T$ and generalized positions $q = [q_A^T, q_B^T]^T$ such that $u = H\dot{q}$. Let $r_A^i = p_A^i - x_A$ and $r_B^j = p_B^j - x_B$ denote the relative positions from the center of masses to the nearest points, respectively. Now we can get the relative velocities of the nearest points as:

$$v_{\text{rel}}^{ij} = \begin{bmatrix} v_B + \omega_B \times r_B^j - (v_A + \omega_A \times r_A^i) \\ v_A + \omega_A \times r_A^i - (v_B + \omega_B \times r_B^j) \end{bmatrix} \quad (5)$$

We rewrite the above equation in matrix form as $v_{\text{rel}}^{ij} = J^{ij}u = J^{ij}H\dot{q}$, where J^{ij} is the contact Jacobian matrix that maps the generalized velocities to the relative velocities of the nearest points. The contact Jacobian can be computed as:

$$J^{ij} = \begin{bmatrix} -I & [r_A^i]_\times & I & -[r_B^j]_\times \\ I & -[r_A^i]_\times & -I & [r_B^j]_\times \end{bmatrix} \quad (6)$$

where I is the identity matrix and $[\cdot]_\times$ denotes the skew-symmetric matrix. Next, we let $p^{ij} = \begin{bmatrix} p_A^i \\ p_B^j \end{bmatrix}$. Finally, we can compute the surrogate gradients of the nearest points with respect to the generalized positions as:

$$\frac{\partial p^{ij}}{\partial q} = J^{ij}H \quad (7)$$

The surrogate gradients in Eq. (7) can be efficiently computed once the contact pairs and nearest points are obtained from DCD.

Thus, our simulator is fully differentiable, supporting both forward simulation and backward gradient propagation.

IV. EXPERIMENTS

A. Experimental Setup

Our real-world data collection system is shown in Figure 3a. The cubes are 3D-printed and equipped with AprilTag [12] for robust fiducial detection. We utilize four Intel RealSense D435if cameras with 60 Hz to detect the tags and estimate the cubes' poses via TagSLAM [25], shown as Figure 3b. We collect only 3 trajectories, each containing around 20 frames, as the original training set. This small dataset is used for contact parameter identification of all analytical simulators in our comparison. We also augment it to 3000 trajectories

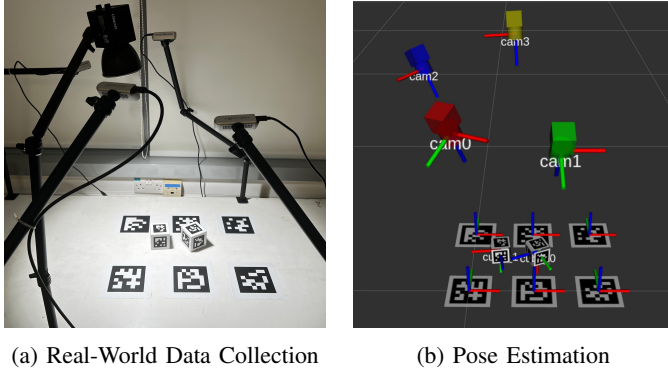


Fig. 3: (a) We collect real-world data on a tabletop setup with two cubes, considering a quasi-planar frictional contact scenario where one cube is pushed towards another cube at rest. (b) The poses of the cubes are estimated using TagSLAM.

TABLE II: Contact Parameters Optimization Range

Simulator	Parameter	Range
All Simulators	μ	[0, 1]
	d_0 (<i>solimp</i>)	[0.9, 0.95]
	d_{width} (<i>solimp</i>)	[0.95, 0.99]
MuJoCo	Width (<i>solimp</i>)	[0.0001, 0.01]
&	Midpoint (<i>solimp</i>)	[0.001, 0.1]
Brax (Generalized)	Power (<i>solimp</i>)	[1, 5]
	Time constant (<i>solref</i>)	[0.001, 0.1]
	Damping ratio (<i>solref</i>)	[0.1, 10]
Brax (Positional & Spring)	Elasticity	[0, 1]

by rotating the trajectories around the z-axis for a learned baseline. Similarly, we use the identified MuJoCo to generate 3000 trajectories as the training data for our GNN-based simulator. In addition, we collect another 14 trajectories as the test set to evaluate the performance of all simulators. The initial positions and orientations of the cubes in the test set are different from those in the training set. The whole framework is performed on a computer with AMD Ryzen 9 7950X3D CPU and single NVIDIA GeForce RTX 4090 GPU.

B. Baselines

We compare our GNN-based simulator with two analytical simulators, MuJoCo [29] and Brax [8]. Brax is a popular differentiable simulator in robotics, which supports three built-in physics pipelines: (i) generalized pipeline resolves dynamics in generalized coordinates, similar to MuJoCo; (ii) positional pipeline uses position based dynamics to stabilize the simulation and resolve collision constraints fast; (iii) spring pipeline leverages an impulse-based method to enable a fast simulation. We also consider training our simulator directly on the real-world data with data augmentation (DA) as mentioned in Section III-B, to evaluate the effectiveness of our data scaling.

C. Contact Parameter Identification Results

We implement the sampling-based method CMA-ES to identify the contact parameters of MuJoCo through Optuna




Fig. 4: Comparison of trajectory errors (see Eq. (1)) in MuJoCo before and after contact parameter identification. The identified parameters significantly enhance the simulation’s accuracy in replicating real-world contact dynamics.




Fig. 5: Illustration of a scenario where the moving cube (green) collides with the static cube (red) before and after contact parameter identification.

Python library [2]. The optimization range of each contact parameter is shown in Table II. The comparison results of MuJoCo are shown in Figure 4, where the average trajectory error on the test set is significantly reduced from 1.14 to 0.73 after contact parameter identification. Since MuJoCo has well-established physical models and formulates the contact dynamics in a convex unconstrained optimization problem, the few-shot identification is effective and efficient for MuJoCo to align with real-world contact dynamics. In Figure 5, we visualize the trajectories simulated in MuJoCo before and after contact parameter identification to the real-world trajectory, with the same initial states. Our identification process significantly improves MuJoCo’s modeling of real-world contact dynamics, ensuring that the simulation environment more accurately reflects actual physical interactions. However, there is still a gap between the identified MuJoCo and the real-world data, which indicates that the well-established analytical models in MuJoCo could not fully reflect the complexities of real-world contact interactions.

D. Simulation Evaluation

To validate the effectiveness of our GNN-based simulator, we evaluate the performance of our GNN-based simulator and all baselines on the test set. The positional and angular errors are shown in Figure 7. Even though our simulator is trained only on the scaled data generated by the MuJoCo baseline, it still achieves similar error levels and even slightly lower positional errors compared to the identified MuJoCo. Thus, data scaling enables our simulator to learn the real-world rigid-contact dynamics well, thanks to the diverse configurations and relatively accurate contact modeling of the identified MuJoCo. Furthermore, both the positional and angular errors




Fig. 6: A rollout example from our GNN-based simulator, illustrating a cube is pushed and collides with a bowling-like array of ten cubes. The ground truth is the scaled data generated by the identified MuJoCo. The output trajectory of our simulator is close to the ground truth, showing that our simulator accurately captures the dynamics and interactions between objects.

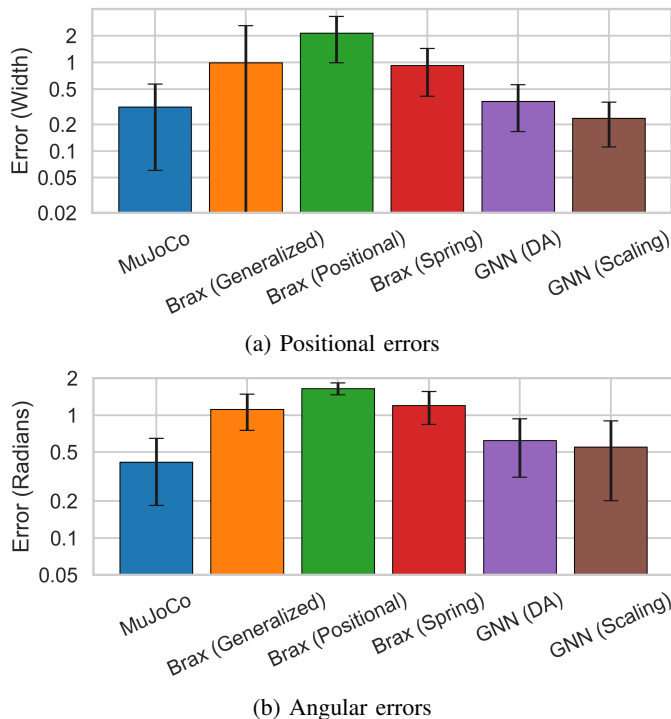


Fig. 7: The positional and angular errors of our GNN-based simulator against analytical and learned baselines. Our simulator achieves comparable performance to identified MuJoCo and outperforms all pipelines of identified Brax, while our data scaling method further enhances the accuracy of our simulator.

of our simulator trained with DA are higher than those of our simulator trained with the scaled data, which also demonstrates the effectiveness of our data scaling method. Compared to all pipelines of Brax, our simulator significantly outperforms them in both positional and angular errors. On the one hand,

the identified Brax pipelines still have a large gap compared to the identified MuJoCo, indicating that Brax’s physical models and contact formulations are less effective in aligning with real-world contact dynamics. On the other hand, our GNN-based simulator, with its data-driven nature and ability to learn complex interactions, can better model the rigid-contact dynamics from the scaled data.

We further conduct an experiment in which a single cube strikes a bowling-like array of ten cubes, shown in Figure 6, to evaluate the complex multi-body interaction capabilities of our simulator. The rollout results demonstrate that our simulator successfully captures near-instantaneous contact behaviors between cubes. This experiment highlights the ability of our GNN-based simulator to accurately simulate contact-rich and computationally heavy scenarios.

We showcase an example of gradient-based optimization in Figure 1, where the initial pushing velocity of the blue cube is optimized to stop the green cube at a target area after colliding with it. The optimization successfully converges within 10 epochs across multiple runs, demonstrating the potential of our differentiable simulator to enable gradient-based optimization for complex rigid-contact scenarios.

V. CONCLUSION

In this work, we propose a novel approach to learning a differentiable GNN-based rigid-body simulator from limited real-world observations. To enable effective training of our GNN simulator, we introduce a few-shot real-to-sim data scaling pipeline that identifies critical contact parameters of a physical simulator without the need of massive real-world data, and further generates a large synthetic dataset for our simulator to learn the diverse contact interactions from the calibrated simulator. Since collision detection is crucial for our simulator to model rigid-contact dynamics but non-differentiable, we derive the surrogate gradients of nearest contact points, allow-

ing our GNN-based simulator to achieve full differentiability. Throughout the experiments, our data scaling alleviates the need for large amounts of real-world data, while it can still facilitate learning-based simulators to capture rigid-contact dynamics. Furthermore, our simulator shows the potential to generalize to novel multi-object interaction scenarios and support gradient-based optimization for contact-rich tasks.

The primary limitation of our approach is its heavy dependence on the accuracy of identification and contact modeling during real-world data scaling. Additionally, our method currently requires real-world object 6D poses, and assumes that the identified contact parameters generalize across different scenarios. In future work, more sophisticated contact representations that describe a broader range of contact dynamics can be investigated, along with the integration of our simulator with vision to enable direct learning from images or videos. Further evaluation with gradient-based control in closed-loop manipulation tasks would strengthen our simulator’s practical utility in robotics.

ACKNOWLEDGMENTS

We acknowledge the support of MOE AcRF Tier 1 24-1234-P0001, Google Gift Fund, and NVIDIA Academic Grant Award. We would also like to express our gratitude to Prana Lab for their valuable support and collaboration on this research.

REFERENCES

[1] Brian Acosta, William Yang, and Michael Posa. Validating robotics simulators on real-world impacts. *IEEE Robotics and Automation Letters*, 7(3):6471–6478, 2022. doi: 10.1109/LRA.2022.3174367.

[2] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631, 2019.

[3] Kelsey R Allen, Tatiana Lopez Guevara, Yulia Rubanova, Kim Stachenfeld, Alvaro Sanchez-Gonzalez, Peter Battaglia, and Tobias Pfaff. Graph network simulators can learn discontinuous, rigid contact dynamics. In Karen Liu, Dana Kulic, and Jeff Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 1157–1167. PMLR, 14–18 Dec 2023. URL <https://proceedings.mlr.press/v205/allen23a.html>.

[4] Kelsey R Allen, Yulia Rubanova, Tatiana Lopez-Guevara, William F Whitney, Alvaro Sanchez-Gonzalez, Peter Battaglia, and Tobias Pfaff. Learning rigid dynamics with face interaction graph networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=J7Uh781A05p>.

[5] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Proceedings of the 30th International*

Conference on Neural Information Processing Systems, NIPS’16, pages 4509–4517, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.

[6] Gino Bergen. Proximity queries and penetration depth computation on 3d game objects. 01 2001.

[7] Bibit Bianchini, Mathew Halm, and Michael Posa. Simultaneous learning of contact and continuous dynamics. In *7th Annual Conference on Robot Learning*, 2023. URL https://openreview.net/forum?id=-3G6_D66Aua.

[8] C. Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax - a differentiable physics engine for large scale rigid body simulation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. URL <https://openreview.net/forum?id=VdvDlnnjzIN>.

[9] E.G. Gilbert, D.W. Johnson, and S.S. Keerthi. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal on Robotics and Automation*, 4(2):193–203, 1988. doi: 10.1109/56.2083.

[10] Eric Heiden, David Millard, Erwin Coumans, Yizhou Sheng, and Gaurav S. Sukhatme. Neursim: Augmenting differentiable simulators with neural networks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9474–9481, 2021. doi: 10.1109/ICRA48506.2021.9560935.

[11] Taylor A. Howell, Simon Le Cleac’h, Jan Brüdigam, Qianzhong Chen, Jiankai Sun, J. Zico Kolter, Mac Schwager, and Zachary Manchester. Dojo: A differentiable physics engine for robotics, 2025. URL <https://arxiv.org/abs/2203.00806>.

[12] Maximilian Krogius, Acshi Haggemiller, and Edwin Olson. Flexible layouts for fiducial tags. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1898–1903, 2019. doi: 10.1109/IROS40897.2019.8967787.

[13] Simon Le Cleac’h, Mac Schwager, Zachary Manchester, Vikas Sindhvani, Pete Florence, and Sumeet Singh. Single-level differentiable contact simulation. *IEEE Robotics and Automation Letters*, 8(7):4012–4019, 2023. doi: 10.1109/LRA.2023.3268824.

[14] Simon Le Cleac’h, Hong-Xing Yu, Michelle Guo, Taylor Howell, Ruohan Gao, Jiajun Wu, Zachary Manchester, and Mac Schwager. Differentiable physics simulation of dynamics-augmented neural objects. *IEEE Robotics and Automation Letters*, 8(5):2780–2787, 2023. doi: 10.1109/LRA.2023.3257707.

[15] Quentin Le Lidec, Igor Kalevatykh, Ivan Laptev, Cordelia Schmid, and Justin Carpentier. Differentiable simulation for physical system identification. *IEEE Robotics and Automation Letters*, 6(2):3413–3420, 2021. doi: 10.1109/LRA.2021.3062323.

[16] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B. Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and

- fluids. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJgbSn09Ym>.
- [17] Jing Yuan Luo, Yunlong Song, Victor Klemm, Fan Shi, Davide Scaramuzza, and Marco Hutter. Residual policy learning for perceptive quadruped control using differentiable simulation, 2024. URL <https://arxiv.org/abs/2410.03076>.
- [18] Mayank Mittal, Calvin Yu, Qinxu Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, Ajay Mandlikar, Buck Babich, Gavriel State, Marco Hutter, and Animesh Garg. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. doi: 10.1109/LRA.2023.3270034.
- [19] Louis Montaut, Quentin Le Lidec, Antoine Bambade, Vladimir Petrik, Josef Sivic, and Justin Carpentier. Differentiable collision detection: a randomized smoothing approach. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3240–3246, 2023. doi: 10.1109/ICRA48891.2023.10160251.
- [20] Matthias Müller, Bruno Heidelberger, Matthias Teschner, and Markus Gross. Meshless deformations based on shape matching. *ACM Trans. Graph.*, 24(3):471–478, July 2005. ISSN 0730-0301. doi: 10.1145/1073204.1073216. URL <https://doi-org.libproxy1.nus.edu.sg/10.1145/1073204.1073216>.
- [21] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *J. Vis. Commun. Image Represent.*, 18(2):109–118, April 2007. ISSN 1047-3203. doi: 10.1016/j.jvcir.2007.01.005. URL <https://doi-org.libproxy1.nus.edu.sg/10.1016/j.jvcir.2007.01.005>.
- [22] J. Krishna Murthy, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Brendan Considine, Jérôme Parent-Lévesque, Kevin Xie, Kenny Erleben, Liam Paull, Florian Shkurti, Derek Nowrouzezahrai, and Sanja Fidler. gradsim: Differentiable simulation for system identification and visuomotor control. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=c_E8kFWfhp0.
- [23] Jia Pan, Sachin Chitta, Dinesh Manocha, Florent Lamiraux, Joseph Mirabel, Justin Carpentier, Louis Montaut, et al. Coal: an extension of the flexible collision library. <https://github.com/coal-library/coal>, 2015–2024.
- [24] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=roNqYLO_XP.
- [25] Bernd Pfrommer and Kostas Daniilidis. Tagslam: Robust slam with fiducial markers, 2019. URL <https://arxiv.org/abs/1910.00679>.
- [26] Samuel Pfrommer, Mathew Halm, and Michael Posa. Contactnets: Learning discontinuous contact dynamics with smooth, implicit representations, 2020. URL <https://arxiv.org/abs/2009.11193>.
- [27] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. Graph networks as learnable physics engines for inference and control. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4470–4479. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/sanchez-gonzalez18a.html>.
- [28] Yunlong Song, Sang bae Kim, and Davide Scaramuzza. Learning quadruped locomotion using differentiable simulation. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=XopATjibyz>.
- [29] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109.
- [30] Kevin Tracy, Taylor A. Howell, and Zachary Manchester. Differentiable collision detection for a set of convex primitives. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3663–3670, 2023. doi: 10.1109/ICRA48891.2023.10160716.
- [31] Simon Zimmermann, Matthias Busenhardt, Simon Huber, Roi Poranne, and Stelian Coros. Differentiable collision avoidance using collision primitives. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8086–8093, 2022. doi: 10.1109/IROS47612.2022.9981093.