Enhancing Physics Simulators for Reliable and Scalable Data Generation

Jeongmin Lee
Seoul National University

I. Introduction

Various rapidly evolving AI systems are primarily designed to infer and explore the solution landscape through data-driven learning. In the context of embodied AI for robotics, this paradigm is also widely investigated. Diverse data sources can be leveraged here, such as pre-trained vision-language models [1, 2], online human video [3, 4], and teleoperation [5, 6]. However, these data sources may present several challenges. From the perspective of the sensorimotor system, humans and various robotic systems possess significantly different modalities, making it challenging to transfer one's policy to another. Moreover, they are more likely to extract relatively coarse-level strategies rather than precise-level control, and the cost of configuring new systems and environments, and collecting data is high. On the other hand, physics simulators [7, 8] can be leveraged to complement this weakness. Various sensor and actuator systems can be implemented with minimal effort, and arbitrary control strategies can be tested within the system. Also, data generation is often much than real-time, especially when utilizing highly parallelized hardware architectures [9].

However, the paradigm of using a physics simulator as a data generator also has fundamental weaknesses. The following problems (**P**) are the ones I aim to address in my research.

P1. Unlike other data sources from the real world, physics simulators are fundamentally *model-based*. These models incorporate various components, including dynamics, contact mechanics, object geometry, and other miscellaneous physical artifacts. While some aspects of these models are well-established (e.g., Newtonian dynamics), many others remain elusive, particularly those related to contact interaction. A physics simulator essentially serves as a solver for the equations governing these models. One crucial aspect here is that, in many cases, modeling and solvers are interdependent. No matter how accurate a model is, if it is numerically intractable, the results produced by the simulator cannot be considered reliable. Therefore, a key problem I aim to address is the development of models for various physical artifacts in simulation, such as intensive contact and local deformation, along with solvers that are effectively tailored to handle these models in an efficient and reliable manner.

P2. Physics simulators themselves only provide *fine-grained* data, as they essentially output state transitions based on action inputs. The characteristics of such fine-grained data can introduce inefficiencies and limitations in its utilization. For example in reinforcement learning, exploration is conducted based on random actions [10], which may struggle to "dig up" useful results. Similarly, in grasping, which is a widely studied problem in manipulation, data are often generated through rejection sampling of the random grasping action in simulation [11]. For more coarse-grained data, a physics simulator should be able to generate data for various inverse problems by leveraging its internal structure, rather than relying on random trials by the user and treating it as a black box. Therefore, a key problem I aim to address is the development of a framework to efficiently search for useful data from a simulator.

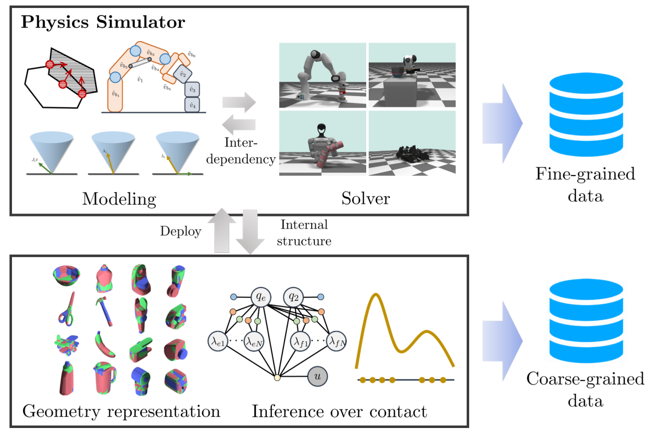

Fig. 1. I am developing physics simulators to generate reliable data by balancing trade-offs while considering the interdependencies between models and solvers. Additionally, I am building a framework to obtain data at a more coarse-grained level through inverse problem-solving, leveraging the internal structure of the simulator.

II. Current Research

A. Reliable Contact Modeling and Solver

Various robotic simulators adopt different models and solvers to handle contact interactions. In most cases, they must balance a subtle trade-off between accuracy, efficiency, and reliability. For example, the soft convex model [12, 13] used in MuJoCo and Drake provides numerical efficiency but may introduce undesirable behaviors such as bouncing and gliding effects. These trade-offs become even more pronounced in scenarios with intensive contact formation and stiff interactions, which frequently occur in dexterous robotic manipulation. In this context, I have developed variations of the augmented Lagrangian (AL) method [14]. This approach modifies the original AL formulation for convex optimization [15] to robotic multi-contact simulation, interpreting the problem as a sequence of feasible subproblems. The underlying philosophy is that *the model and solver operate interdependently*: for each subproblem, the model is iteratively updated to approximate the desired contact condition as closely as possible while remaining within a numerically feasible boundary. At the same time, each subproblem can be efficiently solved using a tailored solver algorithm [14, 16].

This interdependency between contact modeling and the solver can also be leveraged in robotic simulation with deformable objects. One of my representative works in this area is contact nodalization and diagonalization (COND [17]). COND builds on a modeling approach where contact forces act directly on the system's nodal coordinates. The subsequent solver can then precisely enforce contact conditions over nonlinear elastic materials, while computations can be performed highly efficiently without requiring matrix factorization or multiplication. Such a balance of trade-offs can be effectively applied to sim-to-real verification of deformable object manipulation [17, 18].

B. New Geometry Representation

Geometry representation is related to both **P1** and **P2**, as its design impacts both the performance of the simulation itself and the solution process of inverse problems. In standard practice, primitives and meshes are the most widely used models in simulators. However, each approach has its own limitations. Primitives lack the representational power to accurately model a sufficiently wide range of objects in real-world environments. On the other hand, meshes can introduce stability and differentiability issues due to their discrete nature [19]. Meanwhile, tailored functional representations can be developed for specific purposes, such as neural radiance fields [20] for rendering. A key question is: what is a good functional representation for physics simulation and data generation? Motivated by this question, I have developed the concept of a differentiable support function (DSF [21]). DSF defines diverse range of convex geometries by a prescribed functional form of support function [22]. Given the definition, I showed they theoretically guarantee the property of *differentiable contact feature*: contact gap, points, normal between DSFs are always differentiable. Using a set of DSFs, one can represent a diverse range of non-convex objects in the real world [23] while offering several advantages. First, differentiability ensures a smooth transition of contact forces over time steps, enhancing the stability of the simulation. Moreover, gradient-based reasoning for contact interactions can be established. This reasoning includes estimation, planning, and control - which will be more specified in Sec. II-C.

C. Efficient Inference over Contact Interaction

As mentioned in **P2**, extracting useful data from simulation is a challenging problem. Prevalent strategies for this are based on random sampling, nonlinear optimization, or a combination of both [24]. My primary focus is on *generalized and compositional modeling* of reasoning over contact interactions. The design I propose, termed contact factor graphs (CFG [25]), can incorporate diverse factors such as the number of active contacts, their locations, directions, and modes. Based on this, CFG can be utilized to generate solutions (i.e., data) for diverse inverse problems, including planning (grasping, placing, etc.), estimation (assembly, etc.), and control (pushing, pivoting, etc.). Moreover, combined with the geometry module in Sec. II-B, I develop a differentiable probabilistic distribution model on CFG, along with an efficient bi-level inference scheme based on convex optimization and the envelope theorem. Currently, I am working on a method to integrate CFG inference with diverse high-level reasoning (e.g., task-level [26]) and sampling to generate various types of plausible data from a non-convex, multi-modal search space.

III. Future Research

A. Semi-supervised Physics Simulator

As stated above, data generation using a physics simulator is bound to encounter the well-known problem of the sim-to-real gap. Common strategies to addressing this issue involves domain randomization [27] or domain adaptation [28], both of which operate outside the simulator. However, more fundamentally, I believe that solutions within the simulator are necessary. While my current research, as described in Sec. II-A, partially addresses this issue, there remain several aspects of the simulator and solver that are difficult to control and heavily rely on heuristics: such as the amount of damping, clustering rules on the contact manifold, solver hyperparameters, etc.

One promising solution I envision is integrating semi-supervised learning with the simulator. In many cases, physically implausible behaviors can be qualitatively defined: such as excessive penetration, high jerk, or abrupt changes in force. By leveraging such post hoc evaluations, the simulator can be rewarded and gradually learn to adjust its internal parameters and rules accordingly. This semi-supervised approach can be done very efficiently, without any refined data or human intervention. Such learning can be further enhanced by integrating supervision from real-world experimental results [29].

B. Integration with Generative Model

While various emerging studies have been conducted on action/behavior models [6, 30], the integration of diverse data generated from simulations has not yet been extensively explored. One key problem I consider is the integration of data from different sources. Data extracted from simulations (e.g., by inference Sec. II-C) inherently captures pure multimodal aspects of the model. While this facilitates the discovery of various modes, it may also introduce unnecessary complexity by failing to adequately incorporate human biases. In contrast, extracting information as rich as that from simulations is challenging in real-world environments. I believe that the synergy between these data sources will facilitate the high performance intelligence in the future. To this end, I find it promising to study effective theoretical and empirical combinations of simulation and real-world data for improving robot performance.

REFERENCES

[1] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. In *Proceedings of the International Conference on Machine Learning*, 2023.

[2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.

[3] Lin Shao, Toki Migimatsu, Qiang Zhang, Karen Yang, and Jeannette Bohg. Concept2robot: Learning manipulation concepts from instructions and human demonstrations. *The International Journal of Robotics Research*, 40(12-14):1419–1434, 2021.

[4] Kenneth Shaw, Shikhar Bahl, and Deepak Pathak. Videodex: Learning dexterity from internet videos. In *Conference on Robot Learning*, pages 654–665. PMLR, 2023.

[5] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pages 991–1002. PMLR, 2022.

[6] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Robotics: Science and Systems*, 2023.

[7] 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. IEEE, 2012.

[8] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning, 2016.

[9] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.

[10] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[11] Ruicheng Wang, Jialiang Zhang, Jiayi Chen, Yinzhen Xu, Puhao Li, Tengyu Liu, and He Wang. Dexgraspnet: A large-scale robotic dexterous grasp dataset for general objects based on simulation. In *IEEE International Conference on Robotics and Automation*, pages 11359–11366. IEEE, 2023.

[12] Emanuel Todorov. Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in mujoco. In *IEEE International Conference on Robotics and Automation*, pages 6054–6061, 2014.

[13] Alejandro M Castro, Frank N Permenter, and Xuchen Han. An unconstrained convex formulation of compliant contact. *IEEE Transactions on Robotics*, 39(2):1301–1320, 2022.

[14] Jeongmin Lee, Minji Lee, Sunkyung Park, Jinhee Yun, and Dongjun Lee. Variations of augmented lagrangian for robotic multi-contact simulation. *Under Review in IEEE Transactions on Robotics*.

[15] Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999.

[16] Jeongmin Lee, Minji Lee, and Dongjun Lee. Modular and parallelizable multibody physics simulation via subsystem-based admm. In *IEEE International Conference on Robotics and Automation*, pages 10132–10138, 2023.

[17] Jeongmin Lee, Minji Lee, and Dongjun Lee. Large-dimensional multibody dynamics simulation using contact nodalization and diagonalization. *IEEE Transactions on Robotics*, 39(2):1419–1438, 2022.

[18] Jeongmin Lee, Minji Lee, Jaemin Yoon, and Dongjun Lee. A parallelized iterative algorithm for real-time simulation of long flexible cable manipulation. In *IEEE International Conference on Robotics and Automation*, pages 12040–12046, 2021.

[19] Kenny Erleben. Methodology for assessing mesh-based contact point methods. *ACM Transactions on Graphics (TOG)*, 37(3):1–30, 2018.

[20] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

[21] Jeongmin Lee, Minji Lee, and Dongjun Lee. Uncertain pose estimation during contact tasks using differentiable contact features. In *Robotics: Science and Systems*, 2023.

[22] R. Schneider. *Convex bodies: the Brunn–Minkowski theory*. Number 151. Cambridge university press, 2014.

[23] Sunkyung Park, Jeongmin Lee, and Dongjun Lee. Shape abstraction via marching differentiable support functions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

[24] Tao Pang, HJ Terry Suh, Lujie Yang, and Russ Tedrake. Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models. *IEEE Transactions on robotics*, 2023.

[25] Jeongmin Lee, Sunkyung Park, Minji Lee, and Dongjun Lee. Efficient gradient-based inference for manipulation planning in contact factor graphs. In *IEEE International Conference on Robotics and Automation*, 2025.

[26] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual review of control, robotics, and autonomous*

systems, 4(1):265–293, 2021.

[27] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE/RSJ international conference on intelligent robots and systems*, pages 23–30. IEEE, 2017.

[28] Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, et al. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *IEEE international conference on robotics and automation*, pages 4243–4250. IEEE, 2018.

[29] Jaemin Yoon, Minji Lee, Dongwon Son, and Dongjun Lee. Fast and accurate data-driven simulation framework for contact-intensive tight-tolerance robotic assembly tasks. In *arxiv preprint*, 2021.

[30] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023.