Parameterized Quasi-Physical Simulators for Dexterous Manipulations Transfer

Xueyi Liu[1,3], Kangbo Lyu[1], Jieqiong Zhang[1], Tao Du[1,2,3], and Li Yi[1,2,3]

[1] Tsinghua University [2] Shanghai AI Laboratory [3] Shanghai Qi Zhi Institute
https://meowuu7.github.io/QuasiSim

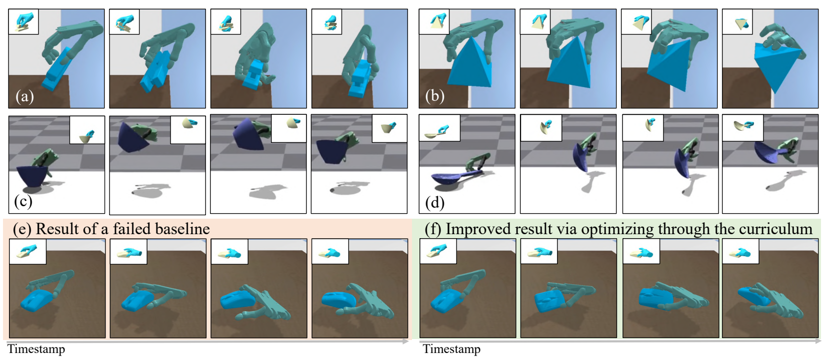

Fig. 1: By optimizing through a **quasi-physical simulator curriculum**, we successfully transfer human demonstrations to dexterous robot hand simulations. We enable accurate tracking of complex manipulations with changing contacts (*Fig. (a)*), non-trivial object motions (*Fig. (b)*) and intricate tool-using (*Fig. (c,d)*). Besides, our physics curriculum can substantially improve a failed baseline (*Fig. (e,f)*).

Abstract. We explore the dexterous manipulation transfer problem by designing simulators. The task wishes to transfer human manipulations to dexterous robot hand simulations and is inherently difficult due to its intricate, highly-constrained, and discontinuous dynamics and the need to control a dexterous hand with a DoF to accurately replicate human manipulations. Previous approaches that optimize in high-fidelity black-box simulators or a modified one with relaxed constraints only demonstrate limited capabilities or are restricted by insufficient simulation fidelity. We introduce **parameterized quasi-physical simulators** and a **physics curriculum** to overcome these limitations. The key ideas are 1) balancing between fidelity and optimizability of the simulation via a curriculum of parameterized simulators, and 2) solving the problem in each of the simulators from the curriculum, with properties ranging from high task optimizability to high fidelity. We successfully enable a dexterous hand to track complex and diverse manipulations in high-fidelity simulated environments, boosting the success rate by 11%+ from the best-performed baseline. The project website is available at QuasiSim.

1 Introduction

Advancing an embodied agent's capacity to interact with the world represents a significant stride toward achieving general artificial intelligence. Due to the substantial costs and potential hazards of setting up real robots to do trial and error, the standard approach for developing embodied algorithms involves learning in physical simulators [9, 11, 17, 19, 23, 34, 35] before transitioning to real-world deployment. In most cases, physical simulators are treated as black boxes, and extensive efforts have been devoted to developing learning and optimization methods for embodied skills within these black boxes. Despite the considerable progress [2,6–8,12,14,15,22,25,27,29,30,36,38,40,41], the question like whether the simulators used are the most suitable ones is rarely discussed. In this work, we investigate this issue and illustrate how optimizing the simulator concurrently with skill acquisition can benefit a popular yet challenging task in robot manipulation – dexterous manipulation transfer.

The task aims at transferring human-object manipulations to a dexterous robot hand, enabling it to physically track the reference motion of both the hand and the object (see Fig. 1). It is challenged by 1) the complex, highly constrained, non-smooth, and discontinuous dynamics with frequent contact establishment and breaking involved in the robot manipulation, 2) the requirement of precisely controlling a dexterous hand with a high DoF to densely track the manipulation at each frame, and 3) the morphology difference. Some existing works rely on high-fidelity black-box simulators, where a small difference in robot control can result in dramatically different manipulation outcomes due to abrupt contact changes, making the tracking objective highly non-smooth and hard to optimize [4, 6, 8, 29, 30]. Other approaches attempt to improve optimization by relaxing physical constraints, with a primary focus on smoothing out contact responses [3, 18, 26, 33, 34]. However, their dynamics models may significantly deviate from real physics [26], hindering skill deployment. Consequently, we ask how to address the optimization challenge while preserving the high fidelity of the simulator.

Our key insight is that a single simulator can hardly provide both high fidelity and excellent optimizability for contact-rich dexterous manipulations. Inspired by the line of homotopy methods [10, 20, 21, 37], we propose a curriculum of simulators to realize this. We start by utilizing a quasi-physical simulator to initially relax physical constraints and warm up the optimization. Subsequently, we transfer the optimization outcomes to simulators with gradually tightened physical constraints. Finally, we transition to a physically realistic simulator for skill deployment in realistic dynamics.

To realize this vision, we propose **a family of parameterized quasi-physical simulators** for contact-rich dexterous manipulation tasks. These simulators can be customized to enhance task optimizability while can also be tailored to approximate realistic physics. The parameterized simulator represents an articulated multi rigid body as a parameterized point set, models contact using an unconstrained parameterized spring-damper, and compensates for unmodeled effects via parameterized residual physics. Specifically, the articulated multi-body dynamics model is relaxed as the point set dynamics model. An articulated object is relaxed into a set of points, sampled from the ambient space surrounding each body's surface mesh. The resulting dynamics model combines the original articulated dynamics with the mass-point dynamics of each individual point. Parameters are introduced to control the point set construction and the dynamics model. The contact model is softened as a parameterized spring-damper model [3, 13, 24, 26, 32] with parameters introduced to control when to calculate contacts and contact spring stiffness. The residual physics network compensate for unmodeled effects from the analytical modeling [16]. The parameterized simulator can be programmed for high optimizability by relaxing constraints in the analytical model and can be tailored to approximate realistic physics by learning excellent residual physics. We demonstrate that the challenging dexterous manipulation transfer task can be effectively addressed through curriculum optimization using a series of parameterized physical simulators.

We demonstrate the superiority of our method and compare it with previous model-free and model-based methods on challenging manipulation sequences from three datasets, describing single-hand or bimanual manipulations with daily objects or using tools. We conduct dexterous manipulation transfer on two widely used simulators, namely Bullet [9] and Isaac Gym [23] to demonstrate the generality and the efficacy of our method and the capability of our quasi-physical simulator to approximate the unknown black-box physics model in the contact-rich manipulation scenario (Fig. 1). We can track complex manipulations involving non-trivial object motions such as large rotations and complicated tool-using such as using a spoon to bring the water back and forth. Our approach successfully surpasses the previous best-performed method both quantitatively and qualitatively, achieving more than 11% success rate than the previous best-performed method. Besides, optimizing through the physics curriculum can significantly enhance the performance of previously under-performed RL-based methods, almost completing the tracking problem from failure, as demonstrated in Fig. 1. This indicates the universality of our approach to embodied AI through optimization via a physics curriculum. Thorough ablations are conducted to validate the efficacy of our designs.

Our contributions are three-fold:

- We introduce a family of parameterized quasi-physical simulators that can be configured to relax various physical constraints, facilitating skill optimization, and can also be tailored to achieve high simulation fidelity.
- We present a quasi-physics curriculum along with a corresponding optimization method to address the challenging dexterous manipulation transfer problem.
- Extensive experiments demonstrate the effectiveness of our method in transferring complex manipulations, including non-trivial object motions and changing contacts, to a dexterous robot hand in simulation.

2 Method

Given a human manipulation demonstration, composed of a human hand mesh trajectory and an object pose trajectory $\{\mathcal{H}, \mathcal{O}\}$, the goal is transferring the

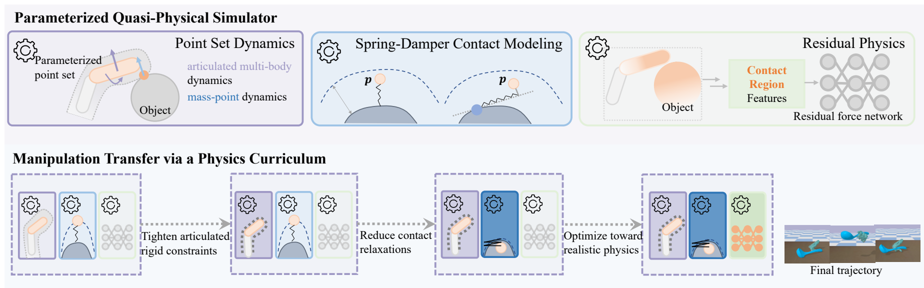

Fig. 2: The **parameterized quasi-physical simulator** relaxes the articulated multi rigid body dynamics as the *parameterized point set dynamics*, controls the contact behavior via an unconstrained *parameterized spring-damper contact model*, and compensates for unmodeled effects via *parameterized residual physics networks*. We tackle the difficult dexterous manipulation transfer problem via **a physics curriculum**.

demonstration to a dexterous robot hand in simulation. Formally, we aim to optimize a control trajectory $\mathcal{A}$ that drives the dexterous hand to manipulate the object in a realistic simulated environment so that the resulting hand trajectory $\hat{\mathcal{H}}$ and the object trajectory $\hat{\mathcal{O}}$ are close to the reference motion $\{\mathcal{H}, \mathcal{O}\}$. The problem is challenged by difficulties from the highly constrained, discontinuous, and non-smooth dynamics, the requirement of controlling a high DoF dexterous hand for tracking, and the morphology difference.

Our method comprises two key designs to tackle the challenges: 1) a family of parameterized quasi-physical simulators, which can be programmed to enhance the optimizability of contact-rich dexterous manipulation tasks and can also be tailored to approximate realistic physics (Section 2.1), and 2) a physics curriculum that carefully adjusts the parameters of a line of quasi-physical simulators and a strategy that solves the difficult dexterous manipulation transfer task by addressing it within each simulator in the curriculum (Section 2.2).

2.1 Parameterized Quasi-Physical Simulators

Our quasi-physical simulator represents an articulated multi-body, *i.e.,* the robotic dexterous hand, as a point set. The object is represented as a signed distance field. The base of the simulator is in an analytical form leveraging an unconstrained spring-damper contact model. Parameters are introduced to control the analytical relaxations on the articulated rigid constraints and the softness of the contact model. Additionally, neural networks are introduced to compensate for unmodeled effects beyond the analytical framework. We will elaborate on each of these design aspects below.

Parameterized point set dynamics. We consider relaxing an articulated multi rigid body into a mass-point set sampled from the ambient space around each body. Each point is regarded as attached to the body where it is sampled from and is both self-actuated and can be actuated via joint motors. A parameter α is introduced to control the point set construction and the dynamics. The point set lets an articulated rigid object behave like a deformable object with a larger action space to adjust the state, thereby easing the optimization problem.

Parameterized spring-damper contact modeling. To ease the optimization challenges posed by contact-rich manipulations, which arise from contact constraints such as the non-penetration requirement and Coulomb friction law [3, 5], as well as discontinuous dynamics involving frequent contact establishment and breaking, we propose a parameterized contact model for relaxing constraints and controlling the contact behavior. Specifically, we leverage a classical unconstrained spring-damper model [13, 24, 32, 35, 39] to model the contacts. This model allows us to flexibly adjust the contact behavior by tuning the contact threshold and the spring stiffness coefficients. By adjusting the contact distance threshold and spring stiffness coefficients, we can modulate the optimizability and fidelity of the contact model.

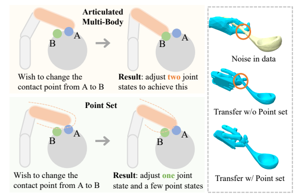

Fig. 3: Point Set can flexibly adjust its states, avoid overfitting to data noise, and ease the difficulty brought by the morphology difference.

Parameterized residual physics. The analytical designs facilitate relaxation but may limit the use of highly sophisticated and realistic dynamics models, deviating from real physics. To address this, the final component of our quasi-physical simulator is a flexible neural residual physics model [1, 16, 28].

Specifically, we propose to employ neural networks to learn and predict residual contact forces and friction forces based on contact-related information. For detailed residual contact force prediction, we introduce a local contact network $f_{\psi_{\text{local}}}$ that utilizes contact information identified in the parameterized contact model and predicts residual forces between each contact pair. To address discrepancies in contact region identification between the parameterized contact model and real contact region, we also incorporate a global residual network $f_{\psi_{\text{global}}}$ that predicts residual forces and torques applied directly to the object's center of mass.

2.2 Dexterous Manipulation Transfer via a Physics Curriculum

Building upon the family of parameterized quasi-physical simulators, we present a solution to the challenging dexterous manipulation transfer problem through a physics curriculum. This curriculum consists of a sequence of parameterized simulators, ranging from those with minimal constraints and the softest contact behavior to increasingly realistic simulators. We address the problem by transferring the manipulation demonstration to the dexterous hand within each simulator across the curriculum progressively.

3 Experiments

We conduct extensive experiments to demonstrate the effectiveness of our method. The evaluation dataset is constructed from three HOI datasets with both single-hand and bimanual manipulations (with rigid objects), with complex manipulations with non-trivial object movements, and rich and changing contacts involved. We use Shadow hand [31] and test in two simulators widely used in the

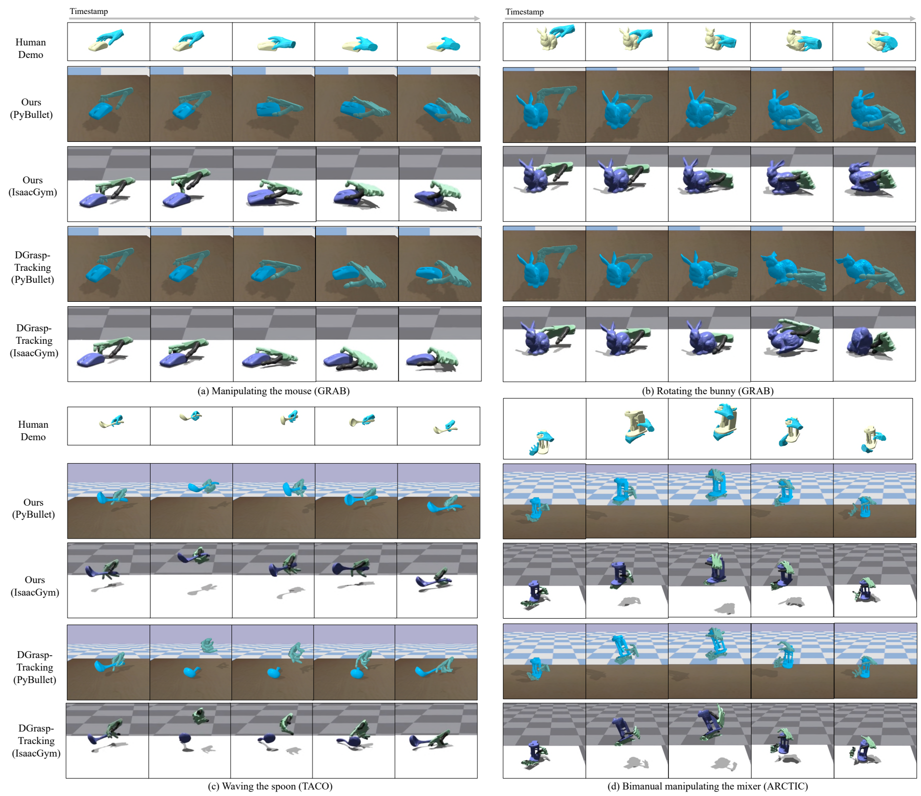

Fig. 4: Qualitative comparisons. Please refer to **our website** and **the accompanying video** for animated results.

embodied AI community: Bullet [9] and Isaac Gym [23]. We compare our method with both model-free approaches and model-based strategies and demonstrate the superiority of our method both quantitatively and qualitatively. We can track complex contact-rich manipulations with large object rotations, back-and-forth object movements, and changing contacts successfully in both of the two simulators, while the best-performed baseline fails (see Section 3.1, Fig. 4). On average, we boost the tracking success rate by 11%+ from the previous best-performed (see Section 3.1).

3.1 Dexterous Manipulating Tracking

We conducted thorough experiments in two widely used simulators [9, 23]. We treat them as realistic simulated physical environments with high fidelity and wish to track the manipulation in them. In summary, we can control a dexterous hand to complete a wide range of the manipulation tracking tasks with non-trivial object movements and changing contacts. As presented in Table 1, we can achieve significantly higher success rates calculated under three thresholds than the best-performed baseline in both tested simulators. Fig. 4 showcases

Table 1: Quantitative evaluations and comparisons to baselines. **Bold red** numbers for best values and *italic blue* values for the second best-performed ones.

Simulator		Method	R_{err} ($°, \downarrow$)	T_{err} ($cm, \downarrow$)	MPJPE ($mm, \downarrow$)	CD ($mm, \downarrow$)	Success Rate ($\%, \uparrow$)
Bullet	Model Free	DGrasp-Base	44.24	5.82	40.55	16.37	0/13.73/15.69
		DGrasp-Tracking	44.45	5.04	37.56	14.72	0/15.69/15.69
		DGrasp-Tracking (w/ curric.)	33.86	4.60	30.47	13.53	7.84/*23.53*/*37.25*
	Model Based	Control-VAE	42.45	*2.73*	25.21	10.94	0/15.68/23.53
		MPC (w/ base sim.)	32.56	3.67	*24.62*	*10.80*	0/15.68/31.37
		MPC (w/ base sim. w/ soften)	*31.89*	3.63	28.26	11.31	0/21.57/*37.25*
		Ours	**24.21**	**1.97**	**24.40**	**9.85**	**27.45**/**37.25**/**58.82**
Isaac Gym	Model Free	DGrasp-Base	36.41	4.56	50.97	18.78	0/7.84 /7.84
		DGrasp-Tracking	44.71	5.57	41.53	16.72	0/0/7.84
		DGrasp-Tracking (w/ curric,)	38.75	5.13	40.09	16.26	0/*23.53*/*31.37*
	Model Based	Control-VAE	*35.40*	4.61	27.63	13.17	0/13.73/29.41
		MPC (w/ base sim.)	37.23	4.73	**23.19**	**9.75**	0/15.69/*31.37*
		MPC (w/ base sim. w/ soften)	36.40	*4.46*	*23.27*	10.34	0/9.80/23.53
		Ours	**25.97**	**2.08**	25.33	*10.31*	**21.57**/**43.14**/**56.86**

qualitative examples and comparisons. Please check out our website and video for animated results.

4 Conclusion

In this work, we investigate creating better simulators for solving complex robotic tasks involving complicated dynamics where the previous best-performed optimization strategy fails. We present a family of parameterized quasi-physical simulators that can be both programmed to relax various constraints for task optimization and can be tailored to approximate realistic physics. We tackle the difficult manipulation transfer task via a physics curriculum.

References

1. Ajay, A., Wu, J., Fazeli, N., Bauza, M., Kaelbling, L.P., Tenenbaum, J.B., Rodriguez, A.: Augmenting physical simulators with stochastic neural networks: Case study of planar pushing and bouncing. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3066–3073. IEEE (2018) 5
2. Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., et al.: Solving rubik's cube with a robot hand. arXiv preprint arXiv:1910.07113 (2019) 2
3. Andrews, S., Erleben, K., Ferguson, Z.: Contact and friction simulation for computer graphics. In: ACM SIGGRAPH 2022 Courses, pp. 1–172 (2022) 2, 3, 5
4. Andrychowicz, O.M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., et al.: Learning dexterous in-hand manipulation. The International Journal of Robotics Research **39**(1), 3–20 (2020) 2
5. Baraff, D.: An introduction to physically based modeling: rigid body simulation ii—nonpenetration constraints. SIGGRAPH course notes pp. D31–D68 (1997) 5
6. Chen, T., Tippur, M., Wu, S., Kumar, V., Adelson, E., Agrawal, P.: Visual dexterity: In-hand reorientation of novel and complex object shapes. Science Robotics **8**(84), eadc9244 (2023). `https://doi.org/10.1126/scirobotics.adc9244`, `https://www.science.org/doi/abs/10.1126/scirobotics.adc9244` 2
7. Chen, T., Xu, J., Agrawal, P.: A system for general in-hand object re-orientation. Conference on Robot Learning (2021) 2

8. Christen, S., Kocabas, M., Aksan, E., Hwangbo, J., Song, J., Hilliges, O.: D-grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20577–20586 (2022) 2
9. Coumans, E., Bai, Y.: Pybullet, a python module for physics simulation for games, robotics and machine learning (2016) 2, 3, 6
10. Dunlavy, D.M., O'Leary, D.P.: Homotopy optimization methods for global optimization. Tech. rep., Sandia National Laboratories (SNL), Albuquerque, NM, and Livermore, CA ... (2005) 2
11. Freeman, C.D., Frey, E., Raichuk, A., Girgin, S., Mordatch, I., Bachem, O.: Brax–a differentiable physics engine for large scale rigid body simulation. arXiv preprint arXiv:2106.13281 (2021) 2
12. Fussell, L., Bergamin, K., Holden, D.: Supertrack: Motion tracking for physically simulated characters using supervised learning. ACM Transactions on Graphics (TOG) **40**(6), 1–13 (2021) 2
13. Geilinger, M., Hahn, D., Zehnder, J., Bächer, M., Thomaszewski, B., Coros, S.: Add: Analytically differentiable dynamics for multi-body systems with frictional contact. ACM Transactions on Graphics (TOG) **39**(6), 1–15 (2020) 3, 5
14. Grandia, R., Farshidian, F., Knoop, E., Schumacher, C., Hutter, M., Bächer, M.: Doc: Differentiable optimal control for retargeting motions onto legged robots. ACM Transactions on Graphics (TOG) **42**(4), 1–14 (2023) 2
15. Gupta, A., Eppner, C., Levine, S., Abbeel, P.: Learning dexterous manipulation for a soft robotic hand from human demonstrations. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3786–3793. IEEE (2016) 2
16. Heiden, E., Millard, D., Coumans, E., Sheng, Y., Sukhatme, G.S.: Neuralsim: Augmenting differentiable simulators with neural networks. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 9474–9481. IEEE (2021) 3, 5
17. Howell, T.A., Le Cleac'h, S., Kolter, J.Z., Schwager, M., Manchester, Z.: Dojo: A differentiable simulator for robotics. arXiv preprint arXiv:2203.00806 **9** (2022) 2
18. Howell, T.A., Le Cleac'h, S., Singh, S., Florence, P., Manchester, Z., Sindhwani, V.: Trajectory optimization with optimization-based dynamics. IEEE Robotics and Automation Letters **7**(3), 6750–6757 (2022) 2
19. Hwangbo, J., Lee, J., Hutter, M.: Per-contact iteration method for solving contact dynamics. IEEE Robotics and Automation Letters **3**(2), 895–902 (2018) 2
20. Liao, S.: On the homotopy analysis method for nonlinear problems. Applied mathematics and computation **147**(2), 499–513 (2004) 2
21. Lin, X., Yang, Z., Zhang, X., Zhang, Q.: Continuation path learning for homotopy optimization (2023) 2
22. Liu, X., Pathak, D., Kitani, K.M.: Herd: Continuous human-to-robot evolution for learning from human demonstration. arXiv preprint arXiv:2212.04359 (2022) 2
23. Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., et al.: Isaac gym: High performance gpu-based physics simulation for robot learning. arXiv preprint arXiv:2108.10470 (2021) 2, 3, 6
24. Marcucci, T., Gabiccini, M., Artoni, A.: A two-stage trajectory optimization strategy for articulated bodies with unscheduled contact sequences. IEEE Robotics and Automation Letters **2**(1), 104–111 (2016) 3, 5

25. Mordatch, I., Popović, Z., Todorov, E.: Contact-invariant optimization for hand manipulation. In: Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation. pp. 137–144 (2012) 2
26. Pang, T., Tedrake, R.: A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 6614–6620. IEEE (2021) 2, 3
27. Peng, X.B., Abbeel, P., Levine, S., Van de Panne, M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. ACM Transactions On Graphics (TOG) **37**(4), 1–14 (2018) 2
28. Pfrommer, S., Halm, M., Posa, M.: Contactnets: Learning discontinuous contact dynamics with smooth, implicit representations. In: Conference on Robot Learning. pp. 2279–2291. PMLR (2021) 5
29. Qin, Y., Wu, Y.H., Liu, S., Jiang, H., Yang, R., Fu, Y., Wang, X.: Dexmv: Imitation learning for dexterous manipulation from human videos. In: European Conference on Computer Vision. pp. 570–587. Springer (2022) 2
30. Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., Levine, S.: Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. arXiv preprint arXiv:1709.10087 (2017) 2
31. ShadowRobot: Shadowrobot dexterous hand (2005), `https://www.shadowrobot.com/dexterous-hand-series/` 5
32. Suh, H., Wang, Y.: Comparing effectiveness of relaxation methods for warm starting trajectory optimization through soft contact (2019) 3, 5
33. Tedrake, R., the Drake Development Team: Drake: Model-based design and verification for robotics (2019), `https://drake.mit.edu` 2
34. Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: 2012 IEEE/RSJ international conference on intelligent robots and systems. pp. 5026–5033. IEEE (2012) 2
35. Wang, Y., Weidner, N.J., Baxter, M.A., Hwang, Y., Kaufman, D.M., Sueda, S.: Redmax: Efficient & flexible approach for articulated dynamics. ACM Transactions on Graphics (TOG) **38**(4), 1–10 (2019) 2, 5
36. Wang, Y., Lin, J., Zeng, A., Luo, Z., Zhang, J., Zhang, L.: Physhoi: Physics-based imitation of dynamic human-object interaction. arXiv preprint arXiv:2312.04393 (2023) 2
37. Watson, L.T., Haftka, R.T.: Modern homotopy methods in optimization. Computer Methods in Applied Mechanics and Engineering **74**(3), 289–305 (1989) 2
38. Wu, Y.H., Wang, J., Wang, X.: Learning generalizable dexterous manipulation from human grasp affordance. In: Conference on Robot Learning. pp. 618–629. PMLR (2023) 2
39. Xu, J., Chen, T., Zlokapa, L., Foshey, M., Matusik, W., Sueda, S., Agrawal, P.: An end-to-end differentiable framework for contact-aware robot design. arXiv preprint arXiv:2107.07501 (2021) 5
40. Yao, H., Song, Z., Chen, B., Liu, L.: Controlvae: Model-based learning of generative controllers for physics-based characters. ACM Transactions on Graphics (TOG) **41**(6), 1–16 (2022) 2
41. Zhang, H., Christen, S., Fan, Z., Zheng, L., Hwangbo, J., Song, J., Hilliges, O.: Artigrasp: Physically plausible synthesis of bi-manual dexterous grasping and articulation. arXiv preprint arXiv:2309.03891 (2023) 2