

ContactExplorer

Contact Coverage-Guided Exploration for General-Purpose Dexterous Manipulation

Ruoyi Qiao^{1,2*}, Zixuan Liu^{1,2*}, Chenrui Tie^{1,2}, Xuanwei Liu¹, Yunfan Lou¹,
Chongkai Gao^{1,2}, Zhixuan Xu^{1,2}, Lin Shao^{1,2†}

¹ School of Computing, National University of Singapore ² RoboScience

* Equal contribution; listed in random order † Corresponding author

<https://contact-explorer.github.io/>

Abstract—Reinforcement learning has achieved remarkable success in domains such as Atari games, navigation, and locomotion, where exploration can often be guided by novelty over states or dynamics. In contrast, dexterous manipulation requires rich physical hand–object interactions, but existing methods often suffer from unstable contact-based novelty signals, inefficient distance novelty signals, or reliance on task-specific priors. We propose ContactExplorer, a general exploration method for dexterous manipulation tasks. ContactExplorer represents contact as the intersection between object surface points and hand keypoints, encouraging dexterous hands to discover diverse and novel contact patterns, namely which fingers contact which object regions. It maintains a contact counter conditioned on discretized object states obtained via learned hash codes, capturing how frequently each finger interacts with different object regions. This counter is leveraged in two complementary ways: (1) to assign a count-based contact coverage reward that promotes exploration of novel contact patterns, and (2) an energy-based reaching reward that guides the agent toward under-explored contact regions. We evaluate ContactExplorer on a diverse set of dexterous manipulation tasks. Experimental results show that ContactExplorer substantially improves sample efficiency and success rates over existing exploration methods, and that the contact patterns learned with ContactExplorer transfer robustly to the real world.

I. INTRODUCTION

Exploration has long been a central challenge in Reinforcement Learning (RL), particularly in tasks where external rewards are sparse or difficult to design. A prominent line of research addresses this challenge through intrinsic motivation, which provides auxiliary rewards for behaviors that are novel, informative, or uncertain [36, 37, 19]. These approaches have led to substantial progress in domains such as Atari games [35, 2, 4, 11, 49], navigation [25, 38, 26], and locomotion [39], where exploration can often be guided by novelty over state or dynamics.

However, dexterous manipulation presents a fundamentally different exploration challenge. Unlike locomotion tasks, which primarily focus on controlling the robot’s own body motion, manipulation requires rich physical interactions with external objects. Prior exploration methods for manipulation have attempted to define novelty using hand–object distance [33], which is an indirect proxy for interaction. More directly, contact captures physical hand–object interaction, and HaC [29]




Fig. 1: **CONTACTEXPLORER** is an exploration method for general-purpose dexterous manipulation, which utilizes contact coverage over the target object to guide dexterous hand fingers towards under-explored object regions. ContactExplorer achieves strong performance with training efficiency in diverse manipulation tasks.

therefore uses contact-force prediction error to compute novelty for manipulation with parallel grippers. Yet in dexterous manipulation, sparse and highly discontinuous contacts make force-based novelty signals noisy and unstable [17]. From the task perspective, many dexterous manipulation methods reduce this exploration challenge by introducing task-specific priors, including carefully designed initializations [45, 10], motion priors [8, 22], or task-specific heuristics [20] that guide the hand toward favorable interaction behaviors. While effective in specific settings, these methods typically rely on substantial human knowledge and lack generality across tasks and objects.

To this end, we propose **ContactExplorer**, a contact-centric exploration framework that explicitly models and incentivizes hand–object interaction on novel contact patterns, namely which fingers contact which object regions. ContactExplorer abstracts objects into surface regions and tracks contact coverage between fingers and object regions. To address the sparsity of contact, ContactExplorer combines two complementary signals: a post-contact count-based reward that encourages novel finger-region contacts and a pre-contact energy-based reaching reward that guides the hand toward regions to yield new interactions. Together, these signals make exploration both contact-focused and consistently guided throughout learning. Moreover, to make the exploration of contact pattern depends

on the task phase and object configuration, ContactExplorer conditions its contact counters on both the current and goal object states. We evaluate ContactExplorer on a diverse suite of dexterous manipulation tasks, including Cluttered Object Singulation, Constrained Object Retrieval, In-Hand Reorientation, Bimanual Object Opening, Bimanual Board Lifting, and Tiled Object Retrieval. Across all tasks, ContactExplorer consistently achieves higher success rates and faster convergence than existing exploration methods, and exhibits robust contact behaviors that transfer effectively to real-world systems.

In summary, this work makes two key contributions. **First**, we introduce ContactExplorer, a contact coverage-guided exploration reward that explicitly models and encourages diverse hand-object contact patterns across task regions. **Second**, through extensive quantitative and qualitative experiments both in simulation and the real world, we demonstrate that ContactExplorer significantly improves training efficiency and final success rates across a wide range of dexterous manipulation tasks. More broadly, ContactExplorer serves as a principled reward exploration for general-purpose dexterous manipulation, providing a general and task-agnostic exploration signal without reliance on handcrafted heuristics or task-specific priors. By guiding robots to systematically discover diverse and meaningful contact patterns, ContactExplorer enables efficient learning of interaction strategies that underpin a wide range of manipulation tasks.

II. THE CONTACT COVERAGE COUNTER DESIGN

A. Problem Formulation

We formulate dexterous manipulation as a Markov Decision Process (MDP), defined by the tuple $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \rho_0, \gamma\}$, where the state space $\mathcal{S}$ includes robot proprioception and object state information, and the action space $\mathcal{A}$ is the continuous low-level control commands for the dexterous hand. The reward function $\mathcal{R}$ consists of a task-specific reward that encourages task completion and an exploration reward designed by this work to promote effective discovery of contact strategies. We use Proximal Policy Optimization (PPO) [32] to solve this MDP for all tasks.

B. Object and Hand Representations

To represent contact interactions in a hand-object manipulation, we first construct a discrete representation of the object surface regions and the hand fingers. For the object, we uniformly sample M surface points $\{(\mathbf{p}_m, \mathbf{n}_m)\}_{m=1}^M$, where each point encodes position and surface normal. These points are grouped into K surface regions based on spatial proximity and normal similarity, denoted by $k = \xi(m)$.

For the hand, we represent the hand as F fingers, where each finger is abstract using a set of predefined surface keypoints and corresponding normal $\{(\mathbf{p}_l, \mathbf{n}_l)\}$ on its finger links. These



Fig. 3: **Hand Keypoints.** Fingers are represented using sparse surface keypoints (red spheres).

keypoints are attached to each finger link and located on their surfaces (Figure 3). We place all the keypoints on the palmar surfaces by default for simplicity, though this is not required; our experiments show robustness when some are perturbed toward the lateral sides of the fingers. This sparse representation follows prior work [42] showing that a limited number of contact points suffices to capture hand-object interaction dynamics across diverse manipulation tasks.

C. Object State Cluster via Learned Hashing

In dexterous manipulation, an identical contact pattern may be useful across different spatial configurations and temporal phases of the object. Maintaining a single global visitation counter causes previously explored patterns to be treated as non-novel regardless of the underlying object state, thereby discouraging the reuse of potentially meaningful interactions under new spatiotemporal contexts.

To enable state-aware exploration, we condition visitation statistics on the object state. We define the object state as $\mathbf{s}^{\text{obj}} = [\mathbf{s}^{\text{cur}}, \mathbf{s}^{\text{goal}}]$, where both terms are point-cloud representations obtained by transforming a fixed set of M canonical surface points according to the object’s current and goal poses. We then partition the continuous state space into discrete clusters $\{s_i\}_{i=1}^S$ and maintain an independent contact counter for each cluster, ensuring decoupled exploration statistics across different object states.

To construct these clusters, we adopt a learned hashing mechanism following Tang et al. [39]. An autoencoder maps $\mathbf{s}^{\text{obj}}$ to a binary latent code $\mathbf{z} \in (0, 1)^D$, trained with a reconstruction loss and a binarization regularizer:

$$\mathcal{L} = \|\hat{\mathbf{s}}^{\text{obj}} - \mathbf{s}^{\text{obj}}\|_2^2 + \frac{\lambda}{D} \sum_{i=1}^D \min\{(1 - z_i)^2, z_i^2\}, \quad (1)$$

where $\hat{\mathbf{s}}^{\text{obj}}$ is the reconstruction of $\mathbf{s}^{\text{obj}}$, z_i is the i -th latent dimension of the latent code. We then binarize $\mathbf{z}$ via thresholding at 0.5 and project it to a compact H -bit hash using SimHash [5] with a fixed random projection matrix. The resulting hash is a discrete state index $s \in \{0, \dots, 2^H - 1\}$. This enables tracking of contact coverage across semantically similar object states. The autoencoder is trained jointly with policy optimization. During the interaction, each contact event updates the counter $\mathbf{C}_{s,f,k}$ corresponding to the inferred state cluster s (Section II-D).

D. Contact Coverage Counter

To encourage exploration of novel contact patterns, we maintain a contact coverage counter $\mathbf{C} \in \mathbb{R}^{S \times F \times K}$, where $\mathbf{C}_{s,f,k}$ records contacts between finger f and surface region k under state cluster s . Contact is detected at the finger level (any keypoint contact triggers the finger), and this merged design empirically outperforms per-keypoint counters. During training, contact detection operates at the level of hand keypoints and object surface points. For each finger f , we identify the




Fig. 2: **Overview of ContactExplorer.** Our contact coverage-guided exploration framework consists of three main components: (1) a learned state hashing module that clusters continuous object states (Sec. II-C); (2) a contact coverage counter that tracks state-conditioned finger-region interactions (Sec. II-D); and (3) a structured exploration reward mechanism, decomposed into contact coverage (Sec. III-A) and energy-based reaching terms (Sec. III-B). The current and goal object states are depicted as distinctly colored point clouds for visualization.

closest interacting pair:

$$(l_f, m_f) = \arg \min_{\substack{l \in \text{finger } f \\ m \in \text{object}}} \|\mathbf{p}_l - \mathbf{p}_m\|_2, \quad (2)$$

where l_f is the finger keypoint nearest to the object and m_f is its closest object surface point. To avoid false positives from transient geometric proximity in simulation, finger-level contact is registered only when *both* geometric proximity and sufficient physical interaction are satisfied:

$$\mathbb{I}^{\text{contact}}(f) = \mathbb{I}[\|\mathbf{p}_{l_f} - \mathbf{p}_{m_f}\|_2 < \delta_{\text{dist}}] \cdot \mathbb{I}[\|\mathbf{F}_{l_f}\|_2 > \delta_{\text{force}}], \quad (3)$$

where $\mathbf{F}_{l_f}$ is the contact force measured at keypoint l_f , and δ_{dist} , δ_{force} are small thresholds. Point m_f is mapped to its surface region $k_f = \xi(m_f)$. We refer to the procedure in Equations 2–3, which computes $\mathbb{I}_t^{\text{contact}}(f)$, k_f , and p_{l_f} as CONTACTQUERY. When $\mathbb{I}^{\text{contact}}(f) = 1$, the counter C_{s,f,k_f} is incremented by one. This discrete approximation avoids expensive fine-grained collision detection while remaining robust to simulation noise. Counters are maintained without reset throughout RL training, enabling a natural transition from exploration to exploitation.

III. CONTACT COVERAGE-GUIDED EXPLORATION

Contact events are inherently sparse during manipulation: rewarding only novel contact patterns provides no guidance for motion in free space, while generic state-novelty objectives often encourage manipulation-irrelevant behaviors in free space. To address this, ContactExplorer decomposes exploration into two complementary signals derived from the same contact coverage statistics. The *post-contact* signal provides a sparse, contact-focused reward that rewards only interaction-relevant behaviors. The *pre-contact* signal provides dense, continuous guidance by shaping motion toward spatial regions likely to yield novel contact patterns. Together, they deliver structured exploration during the training process. Figure 2 illustrates the overall pipeline.

A. Contact Coverage Reward

To ensure exploration remains focused on interaction-relevant events, we provide intrinsic rewards only upon physical contact. At timestep t , for each finger f that contacts the object ($\mathbb{I}_t^{\text{contact}}(f) = 1$), we map the contacted point to its corresponding surface region k under the current object-state cluster s , and compute a contact novelty score:

$$S_{\text{contact}}(t) = \frac{1}{F} \sum_{f=1}^F \mathbb{I}_t^{\text{contact}}(f) \cdot g(C_{s,f,k_f}), \quad (4)$$

where $g(c) = 1/\sqrt{c+1}$ is a monotonically decreasing count-based weighting function.

To avoid repeatedly exploiting previously discovered interaction trajectories, we further adopt a progress-based shaping scheme that rewards only improvements over the best previously achieved score within the current episode:

$$R_{\text{contact}}(t) = \alpha [S_{\text{contact}}(t) - S_{\text{contact}}^{\max}]_+, \quad (5)$$

where $S_{\text{contact}}^{\max}$ denotes the maximum contact novelty score previously achieved in the current episode, α is a scaling coefficient, and $[x]_+ = \max(x, 0)$.

B. Energy-Based Reaching Reward

While the contact coverage reward encourages exploration after contact occurs, it provides no guidance before contact is established. To facilitate pre-contact exploration, we introduce an energy-based reaching objective that guides the policy toward under-explored contact patterns.

For each finger f , we define a contact energy:

$$\Phi_f = \sum_m g(C_{s,f,\xi(m)}) \exp\left(-\frac{\|\mathbf{p}_{l_f} - \mathbf{p}_m\|_2}{\delta}\right), \quad (6)$$

where δ controls the spatial decay, $\mathbf{p}_{l_f}$ denotes the finger keypoint position, and $\mathbf{p}_m$ denotes an object surface point.

TABLE I: Success Rate on All Tasks.

Method	Singulation	Retrieval	InHand	Bimanual Opening	Board Lifting	Grasping	Tiled Retrieval	Avg.
TR	77±33	0±0	78±12	92±5	36±44	68±21	53±12	58
LHCC	77±17	0±0	78±12	90±7	16±31	72±9	59±16	56
HaC	68±28	0±0	77±12	88±11	38±46	61±20	61±15	56
RND-Dist	91±7	0±0	79±11	86±14	18±37	44±31	62±16	54
Ours	96±1	88±6	86±12	96±2	80±9	79±5	77±9	86

The energy-based exploration score is then defined as:

$$S_{\text{energy}}(t) = \frac{1}{F} \sum_{f=1}^F \Phi_f. \quad (7)$$

Similar to the contact reward, we apply episodic progress-based shaping:

$$R_{\text{energy}}(t) = \beta [S_{\text{energy}}(t) - S_{\text{energy}}^{\max}]_+, \quad (8)$$

where $S_{\text{energy}}^{\max}$ denotes the maximum energy score achieved within the current episode and β is a scaling coefficient.

IV. EXPERIMENTS

We evaluate ContactExplorer in simulation and the real world, asking whether it (Q1) improves task performance and sample efficiency over intrinsic-motivation baselines, and (Q2) transfers contact patterns learned in simulation to real-world manipulation. Further studies on task-specific priors and design ablations are reported in Appendix E.

A. Experimental Setup

For Q1, we evaluate seven contact-rich tasks: **Cluttered Object Singulation**, **Constrained Object Retrieval**, **In-Hand Reorientation**, **Bimanual Object Opening**, **Bimanual Board Lifting**, **Grasping**, and **Tiled Object Retrieval**. All simulation experiments run in Isaac Gym [21] with PPO on a UFactory xArm with a 16-DOF LEAP Hand [34], sharing the same observation and action spaces. We report the best success rate (SR) and SR at 40% of training.

B. Comparison with Intrinsic Exploration Baselines

To answer Q1, we compare ContactExplorer with four baselines under identical policy architectures, observations, and hyperparameters: **TR**, using task reward only; **LHCC**, a learned-hash-code count-based method [39, 47]; **HaC**, based on haptic prediction error [29]; and **RND-Dist**, applying random network distillation to hand-object distance [33]. The best SR is reported in Table I, and the learning curves in Figure 9 compare training efficiency across methods.

C. Real-World Experiments

To answer Q2, we validate sim-to-real transfer on a uFactory xArm with a 16-DoF LEAP Hand [34] and two RealSense D435 cameras (Fig. 4), evaluating In-Hand Reorientation and Cluttered Object Singulation. The

TABLE II: Real-World Cluttered Object Singulation Results.

Method	Sing. SR	Grasp SR
TR	36.7	3.3
Ours	76.7	33.3




Fig. 4: Real-World Experiment Setup.

privileged-state teacher is distilled into a point-cloud vision student policy. We conduct 30 real-world trials for each method and report success rates. As shown in Table II, ContactExplorer significantly improves both singulation and task success over TR on Cluttered Object Singulation.

V. LIMITATIONS AND CONCLUSION

ContactExplorer has several limitations. First, tasks with cyclic object states pose a challenge: when identical states recur, conditioning solely on object state may not distinguish contact counts, although the learned hashing can still capture slight state differences. Second, the exploration signal is mainly geometric; incorporating force-torque or tactile sensing could provide richer information about contact quality and interaction dynamics. Third, while sim-to-real transfer is promising, most evaluations remain in simulation, and accelerating RL directly in the real world is an important direction for future work.

In this work, we introduce **ContactExplorer**, a general exploration reward that explicitly incentivizes structured hand-object interactions for dexterous manipulation. By modeling contact coverage between fingers and object regions and combining sparse post-contact rewards with dense pre-contact guidance, ContactExplorer enables efficient, task-agnostic exploration across diverse manipulation tasks without relying on handcrafted shaping. Extensive experimental results demonstrate that ContactExplorer consistently improves learning efficiency and robustness, positioning it as a principled default reward for general-purpose dexterous manipulation.

REFERENCES

- [1] Fengshuo Bai, Yu Li, Jie Chu, Tawei Chou, Runchuan Zhu, Ying Wen, Yaodong Yang, and Yuanpei Chen. Retrieval dexterity: Efficient object retrieval in clutters with dexterous hand. *arXiv preprint arXiv:2502.18423*, 2025.
- [2] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying

- count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [3] Samarth Brahmabhatt, Cusuh Ham, Charles C. Kemp, and James Hays. ContactDB: Analyzing and predicting grasp contact via thermal imaging. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 6 2019.
- [4] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. In *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- [5] Moses S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, STOC '02*, 2002. doi: 10.1145/509907.509965.
- [6] Yuanpei Chen, Chen Wang, Li Fei-Fei, and Karen Liu. Sequential dexterity: Chaining dexterous policies for long-horizon manipulation. In *Proceedings of the 7th Conference on Robot Learning*, pages 3809–3829. PMLR, 2023.
- [7] Yuanpei Chen, Yiran Geng, Fangwei Zhong, Jiaming Ji, Jiechuang Jiang, Zongqing Lu, Hao Dong, and Yaodong Yang. Bi-dexhands: Towards human-level bimanual dexterous manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):2804–2818, 2024. doi: 10.1109/TPAMI.2023.3339515.
- [8] Yuanpei Chen, Chen Wang, Yaodong Yang, and Karen Liu. Object-centric dexterous manipulation from human motion data. In *Proceedings of the 8th Conference on Robot Learning*, pages 3828–3846. PMLR, 2025.
- [9] Zeyuan Chen, Qiyang Yan, Yuanpei Chen, Tianhao Wu, Jiayao Zhang, Zihan Ding, Jinzhou Li, Yaodong Yang, and Hao Dong. Clutterdexgrasp: A sim-to-real system for general dexterous grasping in cluttered scenes. In *Proceedings of the 9th Conference on Robot Learning*, pages 885–905. PMLR, 2025.
- [10] Sudeep Dasari, Abhinav Gupta, and Vikash Kumar. Learning dexterous manipulation from exemplar object trajectories and pre-grasps. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3889–3896, 2023. doi: 10.1109/ICRA48891.2023.10161147.
- [11] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- [12] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J. Black, and Otmar Hilliges. ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [13] Chongkai Gao, Haozhuo Zhang, Zhixuan Xu, Cai Zhehao, and Lin Shao. Flip: Flow-centric generative planning as general-purpose manipulation world model. In *International Conference on Learning Representations*, volume 2025, pages 21927–21948, 2025.
- [14] Irmak Guzey, Yinlong Dai, Georgy Savva, Raunaq Bhirangi, and Lerrel Pinto. Bridging the human to robot dexterity gap through object-oriented rewards. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3344–3351, 2025. doi: 10.1109/ICRA55743.2025.11128690.
- [15] Sandy H Huang, Martina Zambelli, Jackie Kay, Murilo F Martins, Yuval Tassa, Patrick M Pilarski, and Raia Hadsell. Learning gentle object manipulation with curiosity-driven deep reinforcement learning. *arXiv preprint arXiv:1903.08542*, 2019.
- [16] Léonard Hussenot, Robert Dadashi, Matthieu Geist, and Olivier Pietquin. Show me the way: Intrinsic motivation from demonstrations. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, page 620–628. International Foundation for Autonomous Agents and Multiagent Systems, 2021.
- [17] Changwei Jing, Jai Krishna Bandi, Jianglong Ye, Yan Duan, Pieter Abbeel, Xiaolong Wang, and Sha Yi. Contact-aware neural dynamics, 2026. URL <https://arxiv.org/abs/2601.12796>.
- [18] Gagan Khandate, Siqi Shang, Eric T Chang, Tristan L Saidi, Johnson Adams, and Matei Ciocarlie. Sampling-based Exploration for Reinforcement Learning of Dexterous Manipulation. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.020.
- [19] J Zico Kolter and Andrew Y Ng. Near-bayesian exploration in polynomial time. In *Proceedings of the 26th International Conference on Machine Learning*, pages 513–520, 2009.
- [20] Toru Lin, Kartik Sachdev, Linxi Fan, Jitendra Malik, and Yuke Zhu. Sim-to-real reinforcement learning for vision-based dexterous manipulation on humanoids. In *Proceedings of the 9th Conference on Robot Learning*, pages 4926–4940. PMLR, 2025.
- [21] 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.
- [22] Zhao Mandi, Yifan Hou, Dieter Fox, Yashraj Narang, Ajay Mandlekar, and Shuran Song. Dexmachina: Functional retargeting for bimanual dexterous manipulation. *arXiv preprint arXiv:2505.24853*, 2025.
- [23] Mayank Mittal, Pascal Roth, James Tigue, Antoine Richard, Octi Zhang, Peter Du, Antonio Serrano-Muñoz, Xinjie Yao, René Zurbrugg, Nikita Rudin, Lukasz Wawrzyniak, Milad Rakhsha, Alain Denzler, Eric Heiden, Ales Borovicka, Ossama Ahmed, Iretiayo Akinola, Abrar Anwar, Mark T. Carlson, Ji Yuan Feng, Animesh Garg, Renato Gasoto, Lionel Gulich, Yijie Guo, M. Gussert, Alex Hansen, Mihir Kulkarni, Chenran Li,

- Wei Liu, Viktor Makoviychuk, Grzegorz Malczyk, Hammad Mazhar, Masoud Moghani, Adithyavairavan Murali, Michael Noseworthy, Alexander Poddubny, Nathan Ratliff, Welf Rehberg, Clemens Schwarke, Ritvik Singh, James Latham Smith, Bingjie Tang, Ruchik Thaker, Matthew Trepte, Karl Van Wyk, Fangzhou Yu, Alex Milane, Vikram Ramasamy, Remo Steiner, Sangeeta Subramanian, Clemens Volk, CY Chen, Neel Jawale, Ashwin Varghese Kuruttukulam, Michael A. Lin, Ajay Mandlekar, Karsten Patzwaldt, John Welsh, Huihua Zhao, Fatima Anes, Jean-Francois Lafleche, Nicolas Moënnelocoz, Soowan Park, Rob Stepinski, Dirk Van Gelder, Chris Ameyor, Jan Carius, Jumyung Chang, Anka He Chen, Pablo de Heras Ciechowski, Gilles Daviet, Mohammad Mohajerani, Julia von Muralt, Viktor Reutskyy, Michael Sauter, Simon Schirm, Eric L. Shi, Pierre Terdiman, Kenny Vilella, Tobias Widmer, Gordon Yeoman, Tiffany Chen, Sergey Grizan, Cathy Li, Lotus Li, Connor Smith, Rafael Wiltz, Kostas Alexis, Yan Chang, David Chu, Linxi "Jim" Fan, Farbod Farshidian, Ankur Handa, Spencer Huang, Marco Hutter, Yashraj Narang, Soha Pouya, Shiwei Sheng, Yuke Zhu, Miles Macklin, Adam Moravanszky, Philipp Reist, Yunrong Guo, David Hoeller, and Gavriel State. Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025. URL <https://arxiv.org/abs/2511.04831>.
- [24] NVIDIA. Isaac Sim. URL <https://github.com/isaac-sim/IsaacSim>.
- [25] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2778–2787. PMLR, 2017.
- [26] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5062–5071. PMLR, 2019.
- [27] Aleksei Petrenko, Arthur Allshire, Gavriel State, Ankur Handa, and Viktor Makoviychuk. DexPBT: Scaling up Dexterous Manipulation for Hand-Arm Systems with Population Based Training. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.037.
- [28] C. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2016. URL <https://api.semanticscholar.org/CorpusID:5115938>.
- [29] Sai Rajeswar, Cyril Ibrahim, Nitin Surya, Florian Golemo, David Vazquez, Aaron Courville, and Pedro O. Pinheiro. Haptics-based curiosity for sparse-reward tasks. In *Proceedings of the 5th Conference on Robot Learning*, pages 395–405. PMLR, 2022.
- [30] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018. doi: 10.15607/RSS.2018.XIV.049.
- [31] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. URL <https://arxiv.org/abs/2408.00714>.
- [32] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [33] Clemens Schwarke, Victor Klemm, Matthijs van der Boon, Marko Bjelonic, and Marco Hutter. Curiosity-driven learning of joint locomotion and manipulation tasks. In *Proceedings of the 7th Conference on Robot Learning*, pages 2594–2610. PMLR, 2023.
- [34] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *Robotics: Science and Systems (RSS)*, 2023.
- [35] Bradley C Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- [36] Alexander L Strehl and Michael L Littman. A theoretical analysis of model-based interval estimation. In *Proceedings of the 22th International Conference on Machine Learning*, pages 856–863, 2005.
- [37] Alexander L Strehl and Michael L Littman. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74 (8):1309–1331, 2008.
- [38] Adrien Ali Taïga, Aaron Courville, and Marc G Belle-mare. Approximate exploration through state abstraction. *arXiv preprint arXiv:1808.09819*, 2018.
- [39] Haoran Tang, Rein Houthooft, Davis Foote, Adam Stooke, OpenAI Xi Chen, Yan Duan, John Schulman, Filip DeTurck, and Pieter Abbeel. #exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [40] Weikang Wan, Haoran Geng, Yun Liu, Zikang Shan, Yaodong Yang, Li Yi, and He Wang. Unidexgrasp++: Improving dexterous grasping policy learning via geometry-aware curriculum and iterative generalist-specialist learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3891–3902, October 2023.
- [41] Yuhang Wang, Yu Li, Yaodong Yang, and Yuanpei Chen. Dexterous non-prehensile manipulation for un-

graspable object via extrinsic dexterity. *arXiv preprint arXiv:2503.23120*, 2025.

- [42] Zhenyu Wei, Zhixuan Xu, Jingxiang Guo, Yiwen Hou, Chongkai Gao, Zhehao Cai, Jiayu Luo, and Lin Shao. $\mathcal{D}(\mathcal{R}, \mathcal{O})$ grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4982–4988, 2025. doi: 10.1109/ICRA55743.2025.11127754.
- [43] Yueh-Hua Wu, Jiashun Wang, and Xiaolong Wang. Learning generalizable dexterous manipulation from human grasp affordance. In *Proceedings of the 7th Conference on Robot Learning*, pages 3418–3433. PMLR, 2023.
- [44] Lixin Xu, Zixuan Liu, Zhewei Gui, Jingxiang Guo, Zeyu Jiang, Tongzhou Zhang, Zhixuan Xu, Chongkai Gao, and Lin Shao. Dexsingrasp: Learning a unified policy for dexterous object singulation and grasping in densely cluttered environments. *IEEE Robotics and Automation Letters*, 11(2):1346–1353, 2026. doi: 10.1109/LRA.2025.3641152.
- [45] Yinzhen Xu, Weikang Wan, Jialiang Zhang, Haoran Liu, Zikang Shan, Hao Shen, Ruicheng Wang, Haoran Geng, Yijia Weng, Jiayi Chen, Tengyu Liu, Li Yi, and He Wang. Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4737–4746, June 2023.
- [46] Lixin Yang, Kailin Li, Xinyu Zhan, Fei Wu, Anran Xu, Liu Liu, and Cewu Lu. OakInk: A large-scale knowledge repository for understanding hand-object interaction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [47] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. In *Proceedings of the 8th Conference on Robot Learning*, pages 455–472. PMLR, 2025.
- [48] Hui Zhang, Sammy Christen, Zicong Fan, Otmar Hilliges, and Jie Song. Grasppl: Generating grasping motions for diverse objects at scale. In *Proceedings of European Conference on Computer Vision*, pages 386–403. Springer, 2024.
- [49] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E Gonzalez, and Yuandong Tian. Noveld: A simple yet effective exploration criterion. In *Advances in Neural Information Processing Systems*, volume 34. Curran Associates, Inc., 2021.

APPENDIX A RELATED WORK

A. Intrinsic Rewards for RL Exploration

Intrinsic rewards improve exploration by encouraging novelty in the agent’s experience, typically through state novelty or dynamics novelty. State novelty methods reward rarely visited states [36, 37, 19, 4, 2, 39, 38, 11, 49]. In dexterous manipulation, however, their effectiveness depends on the state representation: force-based signals are noisy and unstable due to force spikes and discontinuities [17], while hand-object distance-based novelty [33] may not induce physical interaction. We instead measure state novelty through contact patterns, namely which fingers contact which object regions, yielding a reliable and interaction-centric exploration signal.

Dynamics novelty methods reward transitions that are difficult to predict [35, 25, 26]. For manipulation, HaC [29] uses contact-force prediction error to encourage contact-rich behavior in parallel grippers, but such force prediction is unreliable for dexterous manipulation due to non-smooth contacts [17]. Huang et al. [15] uses a gentleness metric that encourages low-force interactions, but this coarse hand-level signal lacks the spatial resolution required for diverse dexterous contact strategies.

B. Task Priors in RL for Dexterous Manipulation

Beyond exploration strategies, prior work often introduces explicit task priors to improve training efficiency in dexterous manipulation. These approaches can be broadly categorized into reaching guidance and prior knowledge injection. Reaching guidance typically uses dense shaping rewards to encourage the hand to move toward the object [48, 44, 9]. While effective for grasping-oriented tasks, such guidance is often insufficient for more complex dexterous manipulation that requires rich and diverse hand-object interactions. Other approaches incorporate prior knowledge by explicitly encoding assumptions about how the hand should interact with the object. Examples include high-quality initialization [45, 10, 40, 18], manually engineered robot-centric rewards [27, 6, 7, 1, 41, 20, 44], grasp generation modules [43, 45], learning from human videos [14, 13], and expert demonstrations [30, 8, 22, 16]. Although effective in specific settings, these methods rely on strong task- or embodiment-specific priors, which limits their generalizability and provides limited support for structured exploration that can autonomously exploit contact for dexterous manipulation.

APPENDIX B CONTACTEXPLORER IMPLEMENTATION DETAILS

A. Algorithm Overview

Algorithm 1 summarizes the training procedure of ContactExplorer. At each interaction step, the current object state is mapped to a learned hash-based cluster, contact coverage counters are updated from keypoint-level contact queries, and the post-contact coverage reward and pre-contact energy reward are combined with the task reward for PPO training.

The encoder and decoder used for object-state clustering are updated online with Eq. (1).

Algorithm 1 Contact Coverage-Guided Exploration

Require: Surface points $\{\mathbf{p}_m\}_{m=1}^M$, encoder E , decoder f , policy π , value V ; $g(c) = 1/\sqrt{c+1}$, δ , α , β , R_{task} .

- 1: Init counters $\mathbf{C}_s \in \mathbb{R}^{F \times K} \leftarrow 0$, best scores $S_{\text{contact}}^{\text{max}} \leftarrow 0$, $S_{\text{energy}}^{\text{max}} \leftarrow 0$.
- 2: **while not converged do**
- 3: **for step** $t = 1, \dots, T$ **do**
- 4: $s_t \leftarrow \text{SimHash}(\mathbb{I}[E([\mathbf{s}_t^{\text{cur}}, \mathbf{s}_t^{\text{goal}}]) > 0.5])$
- 5: **for finger** $f = 1, \dots, F$ **do**
- 6: $(\mathbb{I}_t^{\text{contact}}(f), \mathbf{p}_{l_f}) \leftarrow \text{CONTACTQUERY}(f)$
- 7: $\mathbf{C}_{s_t, f, k_f} \leftarrow \mathbf{C}_{s_t, f, k_f} + \mathbb{I}_t^{\text{contact}}(f)$
- 8: $\Phi_f \leftarrow \sum_{m=1}^M g(\mathbf{C}_{s_t, f, \xi(m)}) \exp(-\|\mathbf{p}_{l_f} - \mathbf{p}_m\|_2 / \delta)$
- 9: **end for**
- 10: $S_{\text{contact}} \leftarrow \frac{1}{F} \sum_f \mathbb{I}_t^{\text{contact}}(f) g(\mathbf{C}_{s_t, f, k_f})$
- 11: $S_{\text{energy}} \leftarrow \frac{1}{F} \sum_f \Phi_f$
- 12: $R_{\text{contact}} \leftarrow \alpha[S_{\text{contact}} - S_{\text{contact}}^{\text{max}}]_+; S_{\text{contact}}^{\text{max}} \leftarrow \max(S_{\text{contact}}^{\text{max}}, S_{\text{contact}})$
- 13: $R_{\text{energy}} \leftarrow \beta[S_{\text{energy}} - S_{\text{energy}}^{\text{max}}]_+; S_{\text{energy}}^{\text{max}} \leftarrow \max(S_{\text{energy}}^{\text{max}}, S_{\text{energy}})$
- 14: $R(t) \leftarrow R_{\text{task}}(t) + R_{\text{contact}} + R_{\text{energy}}$
- 15: Update π, V with policy gradient; update E, f via Eq. (1)
- 16: **end for**
- 17: **end while**

B. Object Point Cloud and Surface Region Construction

We represent the manipulated object by a canonical surface point cloud $\{\mathbf{p}_m\}_{m=1}^M$ with associated outward normals $\{\mathbf{n}_m\}_{m=1}^M$. To obtain discrete surface regions used by our contact coverage counter, we cluster the canonical points into K regions using farthest-point sampling (FPS) initialization followed by K-means clustering. Point-to-center assignment uses a weighted combination of positional distance and normal disagreement:

$$d(m, k) = (1 - \lambda) \|\mathbf{p}_m - \boldsymbol{\mu}_k\|_2 + \lambda(1 - \mathbf{n}_m^\top \bar{\mathbf{n}}_k),$$

where $\boldsymbol{\mu}_k$ and $\bar{\mathbf{n}}_k$ are the position and mean normal of region k , and λ corresponds to the normal weight. This yields a region label $\xi(m) \in \{1, \dots, K\}$ for each surface point.

C. Contact Coverage Counter

Because Isaac Gym [21] does not provide rigid-body pairwise contact queries at keypoint granularity, we approximate keypoint contact using a distance-force criterion. For keypoint l_f , we compute the distance r_f to the nearest object surface point and the net contact force magnitude $\|\mathbf{F}_f\|_2$. A contact is registered if $r_f < \delta_{\text{dist}}$ and $\|\mathbf{F}_f\|_2 > \delta_{\text{force}}$, where $\delta_{\text{dist}} = 0.5$ cm and $\delta_{\text{force}} = 0.01$ N. This binary signal updates the per-keypoint coverage counters over surface regions $\xi(m)$.

D. Energy-Based Reaching Reward

To encourage physically feasible approach directions, we weight each object surface point m by a directional term computed from the surface normal $\mathbf{n}_m$, the keypoint position $\mathbf{p}_{l_f}$, and the keypoint normal direction $\mathbf{n}_{l_f}$. Let $\mathbf{v}_{l_f,m} = \mathbf{p}_{l_f} - \mathbf{p}_m$ denote the line from surface point m to keypoint l_f . We first suppress back-facing points using

$$w_{l_f,m}^{\text{obj}} = [\cos(\theta_{l_f,m}^{\text{obj}})]_+ = \left[\frac{\mathbf{v}_{l_f,m}^\top \mathbf{n}_m}{\|\mathbf{v}_{l_f,m}\| \|\mathbf{n}_m\|} \right]_+,$$

and further prefer palm-facing configurations via

$$w_{l_f,m}^{\text{keypoint}} = [-\cos(\theta_{l_f,m}^{\text{keypoint}})]_+ = \left[-\frac{\mathbf{d}_{l_f}^\top \mathbf{n}_m}{\|\mathbf{d}_{l_f}\| \|\mathbf{n}_m\|} \right]_+,$$

where $[\cdot]_+ = \min(\max(\cdot, 0), 1)$. The final directional weight is the product $w_{l_f,m}^{\text{dir}} = w_{l_f,m}^{\text{obj}} w_{l_f,m}^{\text{keypoint}}$. Our energy-based reaching reward for keypoint l_f is then computed by summing energy over surface points, modulated by this directional weight and an exponential distance kernel:

$$\Phi_f = \sum_m g(\mathbf{C}_{s,f,\xi(m)}) w_{l_f,m}^{\text{dir}} \exp\left(-\frac{\|\mathbf{p}_{l_f} - \mathbf{p}_m\|_2}{\delta}\right),$$

with δ the kernel scale.

Additionally, we account for line-of-sight occlusions between keypoints l_f and object point m . Let $w_{l_f,m}^{\text{occ}} \in \{0, 1\}$ be a binary visibility mask that equals 1 only if the segment along $\mathbf{v}_{l_f,m} = \mathbf{p}_{l_f} - \mathbf{p}_m$ is not blocked by any obstacle. In the implementation, $w_{l_f,m}^{\text{occ}}$ is computed via a ray-box intersection test against an oriented bounding box: we cast a ray from $\mathbf{p}_{l_f}$ toward $\mathbf{p}_m$ (treating the surface point as the endpoint at $t = 1$) and set $w_{l_f,m}^{\text{occ}} = 0$ if any valid intersection occurs for $t \in (0, 1)$; otherwise $w_{l_f,m}^{\text{occ}} = 1$. We apply this by multiplicatively masking the distance kernel, yielding

$$\Phi_f = \sum_m g(\mathbf{C}_{s,f,\xi(m)}) w_{l_f,m}^{\text{dir}} w_{l_f,m}^{\text{occ}} \exp\left(-\frac{\|\mathbf{p}_{l_f} - \mathbf{p}_m\|_2}{\delta}\right).$$

This occlusion handling is necessary in cluttered object singulation and constrained object retrieval tasks, where nearby geometry should not contribute to the energy-based reaching reward if it is not directly reachable along the approach line.

APPENDIX C

SIMULATION EXPERIMENTS DETAILS

All simulation experiments are conducted with **2048 parallel environments** using Isaac Gym. Policy updates are performed every **16 environment steps**. For each task, results are reported as the **mean and standard deviation over 5 random seeds**.

TABLE III: **Ablation of ContactExplorer Reward Scale** on Constrained Object Retrieval.

Setting	SR (%) $\uparrow$	SR at 40% of training (%) $\uparrow$
$\alpha = 50.0, \beta = 0.32$	17 ± 35	14 ± 28
$\alpha = 100.0, \beta = 0.64$	52 ± 42	30 ± 37
$\alpha = 400.0, \beta = 2.56$	89 ± 4	69 ± 16
$\alpha = 800.0, \beta = 5.12$	56 ± 45	40 ± 35
$\alpha = \mathbf{200.0}, \beta = \mathbf{1.28}$ (Ours)	88 ± 6	67 ± 34

TABLE IV: **Ablation of Learning Parameter for State Clustering** on Constrained Object Retrieval.

Setting	SR (%) $\uparrow$	SR at 40% of training (%) $\uparrow$
$H = 4, \lambda = 0.5$	90 ± 2	68 ± 25
$H = 4, \lambda = 1.0$	85 ± 6	63 ± 32
$H = 4, \lambda = 2.0$	52 ± 43	49 ± 40
$H = 5, \lambda = 0.5$	73 ± 36	64 ± 33
$H = 5, \lambda = 2.0$	72 ± 36	62 ± 32
$H = 6, \lambda = 0.5$	71 ± 36	57 ± 32
$H = 6, \lambda = 1.0$	69 ± 35	46 ± 37
$H = 6, \lambda = 2.0$	91 ± 4	70 ± 10
$H = \mathbf{5}, \lambda = \mathbf{1.0}$ (Ours)	88 ± 6	67 ± 34

A. Task Description

1) *Cluttered Object Singulation*: In this task, the robot is required to extract a single target object from a densely packed shelf. Each scene consists of 5 upright objects of the same size arranged in a single row. At the beginning of simulation creation, the position of the entire book grid is randomized along the edge of the shelf, and the target object is randomly selected among all the books. All non-target books remain fixed, while only the target book is movable. The observation space includes hand root poses and velocities; fingertip poses and velocities; target object positions and velocities; non-target object positions; relative hand-object poses; binary tactile signals at each hand link; and the previous action. The action space consists of delta end-effector poses for the robotic arm and delta joint angles for the dexterous hand, with all joints operating under position control. The task is considered successful when the target book sufficiently reaches the goal position.

2) *Constrained Object Retrieval*: In this task, the robot must retrieve a cube from a top-opening box by sliding it along the interior walls. The cube is lower than the top rim of the box, and the initial gap between the cube and the box is insufficient for inserting the LEAP Hand [34] fingers, making direct grasping infeasible. As a result, successful retrieval requires contact-rich, constrained motions guided by interactions with the box interior. The observation, action spaces and control mode are identical to those used in Cluttered Object Singulation. The task is considered successful when the cube sufficiently reaches the goal position.

3) *In-Hand Reorientation*: In this task, the robotic hand is required to rotate an object to a specified target orientation. The task is evaluated under two settings: up-facing and down-facing. In the up-facing setting, objects include the elephant,

TABLE V: **Qualitative Results of using Allegro Hand.**

Method	SR (%) $\uparrow$		SR at 40% of training (%) $\uparrow$	
	Singulation	Retrieval	Singulation	Retrieval
TR	87 $\pm$ 5	0 $\pm$ 0	48 $\pm$ 24	0 $\pm$ 0
LHCC	72 $\pm$ 18	0 $\pm$ 0	17 $\pm$ 22	0 $\pm$ 0
HaC	8 $\pm$ 16	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0
RND-Dist	55 $\pm$ 45	0 $\pm$ 0	36 $\pm$ 33	0 $\pm$ 0
Ours	93$\pm$4	89$\pm$4	79$\pm$6	59$\pm$22


TABLE VI: **Sensitivity Analysis of Hand Keypoint Selection on Constrained Object Retrieval.**

Setting	SR (%) $\uparrow$	SR at 40% of training (%) $\uparrow$
Low-Level Noise	88 $\pm$ 3	49 $\pm$ 30
High-Level Noise	86 $\pm$ 3	79 $\pm$ 3
Predefined Keypoints (Ours)	88$\pm$6	67$\pm$34

mug, bunny, duck, mouse, and teapot from the ContactDB [3] dataset. In the down-facing setting, we use the elephant, mug, bunny, duck, mouse, and teapot, together with a slender cuboid (16 cm $\times$ 3 cm $\times$ 3 cm). To enable stable learning in the down-facing setting, we rescale the 6 ContactDB objects and assign each object a fixed initial pose and hand joint configuration, such that the object is initially grasped. For all episodes, the object initial pose (except for down-facing) and the goal orientation are randomly sampled. The observation includes hand joint positions, velocities, and forces; object poses and velocities; goal orientation and distance; and the previous action. The action space consists of absolute joint angles for the dexterous hand, with all joints operating under position control. The task is considered successful when the object sufficiently reaches the goal orientation.

4) *Bimanual Object Opening*: In this task, two robotic hands must coordinately manipulate articulated objects, including flipping open the hinged lid of a waffle iron or opening a box from the ARCTIC [12] dataset. Successful execution requires synchronized bimanual control to stabilize the object while actuating its articulated parts. The observation includes hand root poses and velocities for both hands; hand joint positions and velocities; object and articulated-part poses and velocities; relative hand-object and hand-hand poses; fingertip poses and velocities; binary tactile signals; and the previous action. The action space consists of delta end-effector poses for both arms and delta joint angles for both hands, with all joints operating under position control. The task is considered successful when the articulated joint sufficiently reaches the goal position.

5) *Bimanual Board Lifting*: In this task, two robotic hands must jointly lift a large board from a tabletop. Because the board is too large to be stably lifted by one hand, the task requires coordinated bimanual contacts that support the board from both sides. The observation and action spaces follow the bimanual manipulation setup, including both hand root states, hand joint states, object states, relative poses, fingertip states, tactile signals, and previous actions. The task is considered

Fig. 5: **Two Levels of Perturbation on Hand Keypoints.**

successful when the board is lifted to 0.3 m above the table.

6) *Tiled Object Retrieval*: In this task, eight cubes are tightly tiled in a bin whose walls are higher than the cubes, and the gaps between neighboring cubes are smaller than the finger width. A white target cube is randomly selected in each episode, and the robot must retrieve it from the cluttered bin through contact-rich interactions. The observation, action spaces, and control mode are identical to those used in Cluttered Object Singulation. The task is considered successful when the target cube is lifted to 0.3 m above the table.

7) *Grasping*: Grasping is a fundamental dexterous manipulation task in which excessive exploration can be detrimental to task performance. We additionally evaluate ContactExplorer on a grasping task using the rescaled camera, mug, and drill objects from the OakInk [46] dataset. For all episodes, the object initial pose is randomly sampled. The observation space, action space, and control mode follow the same design as the preceding manipulation tasks. The task is considered successful when the object reaches the specified goal position.

B. Task-Specific Rewards

All task rewards share a common additive form

$$r_t = r_t^{\text{reach}} + r_t^{\text{target}} + r_t^{\text{maintain}} + r_t^{\text{success}},$$

where r_t^{reach} is dense hand-to-object shaping, r_t^{target} is contact-gated task progress, and $r_t^{\text{maintain}} + r_t^{\text{success}}$ is a sparse bonus paid out once the goal tolerance is sustained for a fixed number of steps. We omit the per-step accounting of the latter two terms in the equations below for brevity.

1) *In-Hand Reorientation*: Only orientation-tracking rewards are introduced for in-hand reorientation:

$$r_t^{\text{task}} = w_p \|\mathbf{p}_t - \mathbf{p}^*\|_2 + \frac{w_q}{|\theta_t| + \varepsilon} + w_a \|\mathbf{a}_t\|_2^2 + w_s \mathbf{1}[\theta_t \leq \tau], \quad (9)$$

where $\mathbf{p}_t, \mathbf{p}^* \in \mathbb{R}^3$ are the current and target object positions, $\mathbf{q}_t, \mathbf{q}^* \in \mathbb{R}^4$ the corresponding orientation quaternions, $\theta_t = 2 \arcsin(\|[\mathbf{q}_t \otimes \mathbf{q}^{*-1}]_{\text{xyz}}\|)$ the geodesic angle between them, $\mathbf{a}_t$ the action, τ the orientation tolerance, and ε a small constant for numerical stability. We use $w_p = -10$, $w_q = 1$, $w_a = -2 \times 10^{-4}$, $w_s = 250$.

2) *Cluttered Singulation, Constrained Retrieval, Tiled Retrieval, and Grasping*: These four tasks share the same reward template; they differ only in the contact gate $c_t \in \{0, 1\}$ that conditions task progress. Let $\mathcal{P}_t \subset \mathbb{R}^3$ denote the world-frame surface point cloud of the target object, $\{\mathbf{k}_{t,j}\}_{j=1}^K$ the hand keypoints, $\mathbf{g} \in \mathbb{R}^3$ the goal position, and define

TABLE VII: In-Hand Reorientation (Up-Facing) Experiments.

Method	SR (%) $\uparrow$							SR at 40% of training (%) $\uparrow$						
	Elephant	Mug	Bunny	Duck	Mouse	Teapot	Avg.	Elephant	Mug	Bunny	Duck	Mouse	Teapot	Avg.
TR	80 $\pm$ 5	82 $\pm$ 2	80 $\pm$ 4	94 $\pm$ 1	73 $\pm$ 8	68 $\pm$ 9	79 $\pm$ 10	61 $\pm$ 6	62 $\pm$ 9	62 $\pm$ 9	90 $\pm$ 3	42 $\pm$ 16	46 $\pm$ 14	60 $\pm$ 18
LHCC	84 $\pm$ 3	80 $\pm$ 7	78 $\pm$ 2	94 $\pm$ 2	82 $\pm$ 3	72 $\pm$ 4	82 $\pm$ 8	63 $\pm$ 11	61 $\pm$ 11	60 $\pm$ 3	88 $\pm$ 7	56 $\pm$ 4	52 $\pm$ 8	63 $\pm$ 14
HaC	80 $\pm$ 5	81 $\pm$ 3	79 $\pm$ 2	93 $\pm$ 2	77 $\pm$ 11	72 $\pm$ 4	80 $\pm$ 8	64 $\pm$ 6	60 $\pm$ 8	63 $\pm$ 5	84 $\pm$ 10	50 $\pm$ 16	49 $\pm$ 10	62 $\pm$ 15
RND-Dist	86 $\pm$ 2	80 $\pm$ 4	82 $\pm$ 2	94 $\pm$ 1	76 $\pm$ 6	75 $\pm$ 2	82 $\pm$ 7	71 $\pm$ 6	62 $\pm$ 7	67 $\pm$ 4	85 $\pm$ 6	45 $\pm$ 11	57 $\pm$ 5	65 $\pm$ 14
Ours	93$\pm$1	89$\pm$1	88$\pm$3	96$\pm$1	83$\pm$3	81$\pm$2	88$\pm$6	89$\pm$2	82$\pm$2	79$\pm$6	94$\pm$2	62$\pm$4	70$\pm$3	79$\pm$11

TABLE VIII: In-Hand Reorientation (Down-Facing) Experiments.

Method	SR (%) $\uparrow$								SR at 40% of training (%) $\uparrow$							
	Elephant	Mug	Bunny	Duck	Mouse	Teapot	Cube (16x3x3)	Avg.	Elephant	Mug	Bunny	Duck	Mouse	Teapot	Cube (16x3x3)	Avg.
TR	80 $\pm$ 4	90 $\pm$ 8	71 $\pm$ 3	61 $\pm$ 7	80 $\pm$ 5	92 $\pm$ 1	58 $\pm$ 10	76 $\pm$ 14	72 $\pm$ 4	72 $\pm$ 16	65 $\pm$ 2	50 $\pm$ 4	72 $\pm$ 3	85 $\pm$ 2	54 $\pm$ 9	67 $\pm$ 13
LHCC	75 $\pm$ 6	88 $\pm$ 8	71 $\pm$ 3	54 $\pm$ 8	80 $\pm$ 7	92 $\pm$ 2	63 $\pm$ 6	75 $\pm$ 14	67 $\pm$ 6	66 $\pm$ 14	63 $\pm$ 3	46 $\pm$ 5	73 $\pm$ 5	86 $\pm$ 3	56 $\pm$ 5	65 $\pm$ 14
HaC	75 $\pm$ 4	92 $\pm$ 1	70 $\pm$ 6	52 $\pm$ 7	83$\pm$3	89 $\pm$ 2	61 $\pm$ 7	75 $\pm$ 14	65 $\pm$ 4	72 $\pm$ 2	64 $\pm$ 5	47 $\pm$ 4	76$\pm$3	81 $\pm$ 3	54 $\pm$ 6	66 $\pm$ 11
RND-Dist	76 $\pm$ 8	93 $\pm$ 2	67 $\pm$ 4	62 $\pm$ 4	83 $\pm$ 3	91 $\pm$ 2	59 $\pm$ 8	76 $\pm$ 14	69 $\pm$ 6	75 $\pm$ 5	62 $\pm$ 3	48 $\pm$ 3	76 $\pm$ 3	83 $\pm$ 4	52 $\pm$ 6	66 $\pm$ 13
Ours	95$\pm$1	94$\pm$4	79$\pm$9	65$\pm$20	76 $\pm$ 10	98$\pm$0	76$\pm$4	83$\pm$15	83$\pm$9	79$\pm$11	69$\pm$7	51$\pm$10	71 $\pm$ 8	94$\pm$1	68$\pm$4	74$\pm$15

$$d_{t,j} = \min_{\mathbf{p} \in \mathcal{P}_t} \|\mathbf{k}_{t,j} - \mathbf{p}\|_2, \quad d_t^{\text{goal}} = \|\mathbf{x}_t - \mathbf{g}\|_2,$$

where $\mathbf{x}_t \in \mathbb{R}^3$ is the object position. Within the current episode we maintain a running per-keypoint minimum $\bar{d}_{t,j} = \min_{0 \leq t' \leq t} d_{t',j}$ and, gated on satisfied contact, a running object-to-goal minimum $\bar{d}_t^{\text{goal}} = \min\{d_{t'}^{\text{goal}} : 0 \leq t' \leq t, c_{t'} = 1\}$. We then define the per-step reach delta

$$\Delta_t = \frac{1}{K} \sum_{j=1}^K \max(\bar{d}_{t-1,j} - d_{t,j}, 0), \quad (10)$$

which only credits net progress of each keypoint toward the object surface. The task reward is

$$r_t^{\text{task}} = w_r \Delta_t + w_i c_t \max(\bar{d}_{t-1}^{\text{goal}} - d_t^{\text{goal}}, 0) + r_t^{\text{succ}}, \quad (11)$$

where r_t^{succ} aggregates the per-step maintain bonus and the sparse success bonus on the goal indicator $n_t = c_t \mathbf{1}[d_t^{\text{goal}} \leq d_{\text{thres}}]$ with tolerance $d_{\text{thres}} = 0.075$ m. Initial values are $\bar{d}_{0,j} = 0.30$ m and $\bar{d}_0^{\text{goal}} = 0.60$ m. We use $w_r = 20$, $w_i = 60$, and $w_s = 4000$ as the success-bonus magnitude.

The contact gate c_t is task-specific:

$$c_t = \begin{cases} \text{tip}_t \vee \text{table}_t, & \text{Tiled Retrieval, Grasping,} \\ (\text{tip}_t \vee \text{table}_t) \wedge \neg \text{env}_t, & \text{Cluttered Singulation,} \\ \text{tip}_t \vee \text{floor}_t \vee \text{wall}_t^{+y} \wedge \neg \text{wall}_t^{\text{other}}, & \text{Constrained Retrieval.} \end{cases} \quad (12)$$

where tip_t denotes any fingertip-keypoint contact with the target surface, table_t denotes object-table contact, env_t aggregates contact with surrounding clutter, floor_t denotes container-floor contact, wall_t^{+y} the open exit wall of the container, and $\text{wall}_t^{\text{other}}$ the remaining three container walls.

3) *Bimanual Board Lifting and Bimanual Object Opening:* The bimanual reward replaces the single-hand reach delta Δ_t with a per-hand average against the relevant object part, and replaces the running-minimum target of Eq. 11 with a running-maximum progress signal ϕ_t :

$$r_t^{\text{task}} = \frac{w_r^{\text{bi}}}{2} (\Delta_t^r + \Delta_t^l) + w_i^{\text{bi}} c_t \max(\phi_t - \bar{\phi}_{t-1}, 0) + r_t^{\text{succ}}, \quad (13)$$

where Δ_t^r and Δ_t^l are the right- and left-hand instantiations of the reach delta Δ_t defined in 10, computed against the corresponding object part. $\bar{\phi}_t = \max\{\phi_{t'} : 0 \leq t' \leq t, c_{t'} = 1\}$ is the running maximum of progress under satisfied contact, and r_t^{succ} follows the same sustained-success structure as in Eq. 11 with goal indicator n_t defined below. The progress quantity ϕ_t and goal indicator n_t are

$$(\phi_t, n_t) = \begin{cases} (z_t^{\min}, c_t \mathbf{1}[|z_t^{\min} - h^*| < \delta_h]), & \text{Board Lifting,} \\ (\theta_t^{\text{arti}}/\theta^*, c_t \mathbf{1}[|\theta_t^{\text{arti}} - \theta^*| < \delta_\theta]), & \text{Object Opening.} \end{cases} \quad (14)$$

where $z_t^{\min}$ is the lowest world-frame z -coordinate of the board point cloud, θ_t^{arti} is the articulated joint angle, $h^* = 0.50$ m, $\theta^* = 1.57$ rad, $\delta_h = 0.05$ m and $\delta_\theta = 0.10$ rad. The bimanual contact gate c_t requires a keypoint contact from each hand on the relevant object part, and for object opening we instead require table contact, and orientation displacement below 0.20 rad. We use $w_r^{\text{bi}} = 5$, and $(w_i^{\text{bi}}, w_s^{\text{bi}}) = (300, 4000)$ for board lifting and (100, 200) for object opening.

APPENDIX D REAL-WORLD EXPERIMENTS DETAILS

To answer **Q5**, we evaluate our policies on two real-world tasks: Cluttered Object Singulation and In-Hand Reorientation.

A. Teacher-Student Distillation


We adopt a teacher-student distillation pipeline to transfer policies trained in simulation to real-world visuomotor control. First, we train a privileged state-based teacher policy in simulation with domain randomization. We then roll out the teacher policy to collect 1,000 successful trajectories with randomly sampled target objects. These trajectories are used to train a visuomotor policy via behavior cloning.

The student policy receives proprioceptive observations and point cloud observations as input, and predicts the same action space as the teacher policy. Supervision is provided using an L1 loss between the predicted student actions and the teacher actions.

For visual observations, we reconstruct point clouds from two RealSense D435 RGB-D cameras to improve robustness against occlusions. For each camera c , depth pixels are back-projected into a camera-frame point cloud using the camera intrinsics and transformed into the robot base frame using calibrated extrinsics. Point clouds from all camera views are then concatenated, followed by a fixed workspace crop and removal of invalid depth points. Finally, we apply farthest point sampling (FPS) without replacement to downsample the fused point cloud to a fixed size. A comparison between the reconstructed point clouds in simulation and in the real world is shown in Fig. 6.

To identify the target object, we apply SAM2 [31] to the RGB streams and obtain per-view object masks. Each point is augmented with a binary mask indicator and represented as a 4D vector (x, y, z, m) , where $m \in \{0, 1\}$ denotes whether the point belongs to the target object. The masked point cloud is processed by a PointNet encoder [28] to extract a permutation-invariant feature representation, which is concatenated with proprioceptive observations before being fed into the policy network.

The student policy uses a two-step observation history as input. The complete observation space is summarized in Table IX. The action space is identical to that of the teacher policy, consisting of relative commands for a 6-DoF end-effector and 16 hand joints.



(a) Top-View Point Clouds (b) Side-View Point Clouds

Fig. 6: Visualization of point clouds in simulation (blue) and real world (red).

B. Real World Deployment

We deploy the policy at 10 Hz, matching the control frequency used during simulation training. We employ a non-blocking control pipeline in which video streams are processed asynchronously at their native frequencies. At each control step, the policy retrieves the most recent observation from each

TABLE IX: Observation space for student policies.

Name	Dimension
Masked Scene Point Cloud	$2 \times 1024 \times 4$
Proprioceptive	
Arm Joint Position	2×7
Hand Joint Position	2×16

sensor that precedes the current timestep, performs inference, and sends the resulting action command to the robot. This design enables smooth real-time execution while naturally introducing latency from sensing, data processing, and policy inference. Fig. 7 shows the real-world execution of the distilled student policy on the shelf object singulation task.



Fig. 7: **Real-world Policy Execution.** We show a temporal sequence (left to right) of the policy executing the shelf object singulation task.

For the in-hand reorientation task, we evaluate sim-to-real consistency via open-loop trajectory replay. Specifically, we execute in the real world the action sequences generated by the privileged teacher policy in simulation, with the initial object pose and hand configuration aligned to their simulated counterparts. As shown in Fig. 8, real-world rollouts exhibit object pose changes consistent with simulation, suggesting that the manipulation behaviors learned by ContactExplorer in simulation can transfer to the real world.

APPENDIX E ADDITIONAL EXPERIMENT RESULTS

A. Task-Specific Prior and Reward Component Ablations

Table X compares ContactExplorer with task-specific exploration priors (TR and TR-PrePose) on Constrained Object Retrieval, and Table XI ablates the post-contact coverage reward and the pre-contact energy reward on Cluttered Object Singulation and Constrained Object Retrieval. ContactExplorer matches or surpasses the task-specific priors without manual initialization, and the two exploration rewards are complementary, with the full method performing best.



Fig. 8: **Real-World Trajectory Replay.** Simulated action sequences achieve consistent 90° z-axis rotations when replayed in the real world.




Fig. 9: **Learning Curves Across Seven Dexterous Manipulation Tasks.** ContactExplorer achieves higher sample efficiency and success rates, particularly on hard-exploration tasks like Constrained Object Retrieval where baselines fail.

TABLE X: **Comparison with Task-Specific Prior.**

Setting	SR (%) $\uparrow$	SR at 40% of training (%) $\uparrow$
TR	0 $\pm$ 0	0 $\pm$ 0
TR-PrePose	33 $\pm$ 40	25 $\pm$ 32
Ours	88$\pm$6	67$\pm$34

TABLE XI: **Reward Ablation Studies.**

Method	SR (%) $\uparrow$		SR at 40% of training (%) $\uparrow$	
	Sing.	Retr.	Sing.	Retr.
Energy	94 $\pm$ 2	36 $\pm$ 44	88 $\pm$ 4	12 $\pm$ 15
Contact	89 $\pm$ 7	36 $\pm$ 44	73 $\pm$ 23	31 $\pm$ 38
Ours	96$\pm$1	88$\pm$6	94$\pm$3	67$\pm$34

B. Reward Scale Ablation Studies

We study the effect of the ContactExplorer reward scale on the *Constrained Object Retrieval* task by varying the contact coverage reward scale α in Equation 4 and the energy-based reaching reward scale β in Equation 8, as reported in Table III. The results show that an appropriate reward scale is crucial for both final performance and learning efficiency. Small reward scales lead to insufficient exploration and low success rates, while excessively large scales degrade performance and introduce training instability. Our chosen setting ($\alpha = 200.0$, $\beta = 1.28$) achieves a high final success rate while converging fastest to the 70% success threshold, demonstrating a favorable balance between exploration strength and training stability. Under this setting, the per-step exploration reward remains approximately 2 orders of magnitude smaller than the task reward, which we empirically find to be the most effective scale. We observe a consistent trend when tuning other exploration baselines. Therefore, in Table I of the main paper, we apply the same reward scale across all methods to ensure a fair comparison.

C. Object State Clustering Ablation

We test whether state-conditioned counters reduce exploration saturation in a diagnostic Box Push task (Figure 10). The box is initialized on either side of a wall, and solving both initializations requires different pushing directions. We compare ContactExplorer with a **Single-State** variant, which




Fig. 10: **Push Box Task.** (a) The box is initialized at either end of the wall. (b) The actuated ball pushes the box to the goal (right initialization is used as an example). (c) Visualization of two learned object state clusters (state id 0 and 1).

uses one global contact coverage counter without distinguishing object states.

Under the same training setup, ContactExplorer substantially outperforms Single-State, achieving 100% success versus 18%. Single-State often converges to only one initialization because contact counts accumulated in one configuration suppress exploration of the same contact pattern in another configuration where it is still useful. In contrast, ContactExplorer decouples counters across learned state clusters, reducing cross-state interference and preserving state-specific exploration signals, which leads to more robust performance across both initializations.

D. Learning Parameter Ablation Studies for Object State Clustering

We analyze the impact of the learning parameters in the object state clustering module by varying the binary regularization weight λ and the hash length H , which jointly determine

TABLE XII: **Bimanual Object Opening Experiments.**

Method	SR (%) $\uparrow$			SR at 40% of training (%) $\uparrow$		
	Waffle Iron	Box	Avg.	Waffle Iron	Box	Avg.
TR	88 $\pm$ 5	95 $\pm$ 4	92 $\pm$ 5	59 $\pm$ 28	90 $\pm$ 7	74 $\pm$ 25
LHCC	86 $\pm$ 6	95 $\pm$ 3	90 $\pm$ 7	57 $\pm$ 25	77 $\pm$ 20	67 $\pm$ 25
HaC	83 $\pm$ 12	94 $\pm$ 5	88 $\pm$ 11	17 $\pm$ 25	85 $\pm$ 6	51 $\pm$ 38
RND-Dist	76 $\pm$ 13	97$\pm$1	86 $\pm$ 14	47 $\pm$ 17	91$\pm$3	69 $\pm$ 26
Ours	95$\pm$2	97$\pm$1	96$\pm$2	80$\pm$8	90 $\pm$ 6	85$\pm$9

the discretization behavior of the learned state representation. As shown in Equation 1, the regularization term controlled by λ encourages each binary code element b_i to approach either 0 or 1, while the hash length H defines the total number of object state clusters, i.e., $s \in \{0, \dots, 2^H - 1\}$. Specifically, H defines an upper bound on the number of state clusters, while λ implicitly controls the number of clusters that are effectively utilized by encouraging binarization [39].

As shown in Table IV, we observe that the original upper bound ($H = 5$, corresponding to 2^5 clusters) is relatively generous for most of the tasks including Constrained Object Retrieval. Reducing the upper bound to $H = 4$ still yields strong performance, and further decreasing λ (e.g., $H = 4$, $\lambda = 0.5$) improves both success rate and sample efficiency by allowing a larger number of effective state clusters to be used. In contrast, increasing λ overly constrains the representation, causing diverse object states to collapse into fewer clusters and leading to degraded performance. A similar trade-off is observed when increasing the cluster upper bound. For larger H , a stronger regularization is required to prevent over-clustering of the state space. For example, when $H = 6$, increasing λ to 2.0 effectively limits the number of active clusters and restores performance. These results highlight the importance of jointly tuning H and λ to balance state discrimination and sample efficiency.

E. Sensitivity Analysis of Keypoint Selection

Our current hand keypoints are predefined on the *palmar face* of each hand link. To evaluate the robustness of ContactExplorer to keypoint selection, we perturb the predefined keypoints and recompute new keypoints by projecting each perturbed point onto the nearest point on the corresponding link surface. We consider two levels of perturbation shown in Figure 5. For *low-level noise*, noises are sampled from a spherical shell of radius $[0, 1.0]$ cm, resulting into keypoints largely remain on the palmar face. For *high-level noise*, the shell radius is expanded to $[1.0, 2.0]$ cm, which may cause the resulting keypoints to shift to the side face of the link. Performance under these settings is reported in Table VI.

As shown in Table VI, ContactExplorer maintains stable performance under both low- and high-level perturbations, with only minor variations in success rate and learning efficiency. Notably, even when keypoints shift away from the palmar face, the performance degradation remains limited. These results indicate that ContactExplorer does not rely on

precise keypoint placement, but instead benefits from the overall structure of contact coverage, demonstrating robustness to moderate spatial variations in keypoint definition.

F. Cross-Embodiment Experiments

To answer Q5, we conduct cross-embodiment experiments using the Allegro Hand on the Cluttered Object Singulation and Constrained Object Retrieval tasks. All other settings, including observation and action spaces, training procedures, and hyperparameters, remain identical to those used with the LEAP Hand.

As shown in Table V, ContactExplorer consistently improves performance over all baselines on both tasks when transferring to the Allegro Hand. Notably, ContactExplorer achieves substantial gains in success rate for object retrieval, where several baselines fail to solve the task under the same training budget. Meanwhile, ContactExplorer also reduces the number of interaction steps required to reach the success threshold, indicating improved learning efficiency across embodiments. These results suggest that the contact exploration encouraged by ContactExplorer remains effective under changes in different hands, supporting its robustness in cross-embodiment dexterous manipulation.

G. In-Hand Reorientation

Table I shows the average results on the up-facing setting over different objects introduced in Section C. The detailed results of each object are shown in Table VII. Table VIII reports additional results on the down-facing in-hand reorientation setting. Compared with extrinsic exploration baselines, ContactExplorer consistently achieves higher success rates across most object categories, while also reaching the 70% success threshold with fewer interaction steps on average. Notably, ContactExplorer maintains strong performance on geometrically complex objects such as the elephant and teapot, and also shows clear advantages on the slender cube, indicating its effectiveness in handling contact-rich reorientation that requires maintaining stable contacts against gravity. These results further demonstrate the robustness of ContactExplorer in challenging in-hand manipulation scenarios where stable contact exploration is critical.

H. Bimanual Object Opening

Table I shows the average results over 2 different objects introduced in Section C. The detailed results of each object are shown in Table XII.

As shown in Table XII, ContactExplorer achieves the most notable improvement on the geometrically more complex waffle iron, where coordinated bimanual interaction and structured contact exploration are more critical for successful manipulation. In contrast, on the simpler box object, ContactExplorer maintains performance comparable to or better than existing baselines. These results suggest that ContactExplorer is particularly beneficial in contact-rich and geometrically challenging bimanual scenarios.

TABLE XIII: Grasping Experiments.

Method	SR (%) $\uparrow$				SR at 40% of training (%) $\uparrow$			
	Camera	Mug	Drill	Avg.	Camera	Mug	Drill	Avg.
TR	71 $\pm$ 7	76 $\pm$ 2	57 $\pm$ 33	68 $\pm$ 21	43 $\pm$ 26	72 $\pm$ 2	41 $\pm$ 24	52 $\pm$ 25
LHCC	68 $\pm$ 12	75 $\pm$ 2	74 $\pm$ 7	72 $\pm$ 9	31 $\pm$ 23	69 $\pm$ 2	33 $\pm$ 20	44 $\pm$ 25
HaC	70 $\pm$ 2	64 $\pm$ 9	49 $\pm$ 29	61 $\pm$ 20	37 $\pm$ 14	45 $\pm$ 19	5 $\pm$ 8	29 $\pm$ 23
RND-Dist	35 $\pm$ 34	75 $\pm$ 1	22 $\pm$ 12	44 $\pm$ 31	11 $\pm$ 11	69 $\pm$ 2	0 $\pm$ 0	27 $\pm$ 31
Ours	75$\pm$4	77$\pm$2	84$\pm$2	79$\pm$5	66$\pm$10	72$\pm$4	77$\pm$5	72$\pm$8

I. Grasping

As shown in Table XIII, several exploration-based baselines exhibit degraded performance compared to the task-reward-only (TR) baseline, suggesting that overly aggressive exploration may interfere with stable grasp acquisition. In particular, RND-Dist encourages novelty in hand-object distance even in free space, which guides the policy toward exploring non-contact behaviors and conflicts with the prolonged, surface-enveloping finger contact required for stable grasping. In contrast, ContactExplorer consistently improves both success rate and sample efficiency across all objects, demonstrating its ability to encourage structured contact exploration without disrupting task execution.

J. Deformable Object Grasping

We further evaluate ContactExplorer on a deformable object manipulation task with Isaac Lab [23]. As illustrated in Figure 11, we construct a deformable cube in Isaac Sim [24] using a Finite Element Method (FEM) model. The cube is represented by a set of FEM nodes, from which a point cloud is sampled and clustered




Fig. 11: Deformable Object Grasping.

according to node positions. The resulting clusters are used as the object regions. The task requires the robot to grasp the deformable object and transport it to a target position. Table XIV reports the success rates of all methods after 3,000 training iterations. ContactExplorer achieves the highest success rate, demonstrating that contact-based exploration remains effective even when object geometry changes due to deformation.

K. Real-World Experimental Results

We deploy the vision student policy on a real-world shelf object singulation task consisting of two phases: object singulation followed by grasping and transporting the object to a target position. For each policy, we conduct 30 trials. A singulation is considered successful if the object pose becomes pre-grasp feasible and at least half of the object is exposed.

TABLE XIV: Deformable Object Grasping Success Rate.

Method	SR (%) $\uparrow$
TR	66 $\pm$ 29
LHCC	76 $\pm$ 15
HaC	86 $\pm$ 8
RND-Dist	75 $\pm$ 27
Ours	88$\pm$10



(a) Wrong target object

(b) Failed grasp after singulation

Fig. 12: Failure Cases in Real-World Cluttered Object Singulation. The two common failure modes are manipulating the wrong target object and failing to catch or stabilize the target object after singulation.

Final task success is defined as transporting the object to the specified target position. As summarized in Table II, the policy learned with ContactExplorer exhibits more reliable behavior than the student policy trained with the Task Reward (TR) baseline, achieving higher singulation and grasp success rates.

L. Failure Cases

Our failure cases are mainly observed in Cluttered Object Singulation (Figure 12). We observe two representative failure modes. First, the policy manipulates a wrong target object (Figure 12a), indicating that policy distillation in real-world deployment still suffers from a perception gap. Second, even when the target object is correctly singulated, the hand may fail to catch and stabilize the target object at the end of the singulation phase (Figure 12b); this breaks the transition to the grasping phase and causes final failure. We attribute this to a dynamics gap: high-dynamic interactions in singulation introduce accumulated control errors, which then degrade the downstream grasping stage.