

The Curse of Precision: A Data Scaling Law for High-Precision Robotic Manipulation

Cuijie Xu¹, Yuanfan Xu¹, Min Xue¹, Jianjie Lin³, Jian Wang¹, Xudong Zhang¹, Yu Wang¹, Jincheng Yu^{1,2,*}

Abstract—While scaling laws for imitation learning have primarily focused on generalization in open-world settings, the relationship between data and precision in closed-world tasks like robotic assembly remains largely unexplored. This paper systematically investigates this relationship and introduces a novel scaling law. We find that to achieve a fixed success rate, the required number of demonstrations N , grows super-exponentially as the target precision P , approaches a limit c . This relationship is accurately captured by the model $\log(N) \propto 1/(P - c)$. Crucially, we reveal that the limit precision c is not a static physical constant of the task but an emergent property of the entire agent system, including its sensors and expert policy. Through experiments on canonical manipulation tasks, we validate this law and demonstrate that improving system components—such as adding a wrist camera or using a more effective expert—measurably lowers c , thus expanding the system’s achievable precision. Our work provides a new theoretical framework for precision in robotics and a quantitative metric to evaluate system capabilities. Furthermore, these findings provide a practical methodology for guiding the development and debugging of high-precision manipulation systems.

I. INTRODUCTION

The pursuit of scaling has become a central theme driving progress in machine learning, leading to powerful foundation models in natural language processing and computer vision [1–4]. The field of robotics is actively seeking to establish its own scaling laws, with the goal of creating more general and capable agents. Recent landmark studies have begun to uncover how performance scales with data diversity, such as the number of objects and environments, aiming to imbue robots with robust zero-shot generalization capabilities in open-world settings [5–8].

However, a vast and economically significant class of robotic applications, ranging from industrial assembly and electronics manufacturing to surgical automation [9], operates under a different paradigm. In these structured, “closed-world” environments, the primary challenge is not generalization to novelty, but the scaling of precision and reliability.

* Corresponding Author.

¹Department of Electronic Engineering, Tsinghua University, Beijing, China. xcj20@mails.tsinghua.edu.cn, yu-jc@mail.tsinghua.edu.cn

²Institute for Embodied Intelligence and Robotics, Tsinghua University.

³OpenMind (WuHu) Robotics Co., Ltd., Wuhu, China.

This research was supported by National Natural Science Foundation of China (62325405), Tsinghua University Initiative Scientific Research Program, Tsinghua-Efort Joint Research Center for EAI Computation and Perception and SunRisingAI Lab, Beijing National Research Center for Information Science, Technology (BNRist), Beijing Innovation Center for Future Chips, and State Key laboratory of Space Network and Communications.

For these tasks, improving performance from a 99% to a 99.9% success rate, or tightening tolerance from millimeters to sub-millimeters, is of paramount importance. While high precision has traditionally been achieved through analytical control methods [10, 11], the quantitative relationship between the data required for imitation learning and the push for higher precision remains largely uncharacterized. The fundamental data cost of achieving precision is still an open question.

Our own initial experiments highlighted this gap: when applying standard imitation learning techniques to high-precision tasks, performance was often unexpectedly low, yet a slight reduction in the task’s required precision could lead to a dramatic improvement in success rates. This puzzling sensitivity motivated us to ask: What is the data cost required to increase the precision of a robotic skill? Does a predictable scaling law govern this relationship? And perhaps more importantly, how do the intrinsic capabilities of the learning system itself—its sensors, its expert demonstration quality, and its task complexity—modulate this law?

To answer these questions, we conduct a large-scale empirical study in the ManiSkill3 [12] high-fidelity simulation environment, involving over 100 independent model training runs and executing tens of thousands of evaluation rollouts under a rigorous protocol. Our study investigates the three diverse, high-precision manipulation tasks illustrated in Figure 1. Our primary analysis focuses on two image-based tasks, *Peg Insertion* and *Stack Cuboid*. We then present a state-based *Roll Ball* task, as a case study to validate the generality of our findings. Our extensive investigation reveals the following surprising results and key contributions:

- **The “Curse of Precision”:** We uncover a novel scaling law demonstrating that data requirements (N) grow super-exponentially as the target precision (P) increases, accurately modeled by $\log(N) \propto 1/(P - c)$.
- **A New Metric for System Capability:** We identify the law’s key parameter, the limit precision c , not as a fixed physical constant, but as a quantifiable metric for the performance ceiling of the entire learning system.
- **A Roadmap for Surpassing Precision Limits:** Through systematic ablations, we demonstrate that this performance ceiling c can be dramatically improved by enhancing system components like sensors and expert data quality, offering a clear path to overcoming precision bottlenecks.




Fig. 1: **Illustration of the diverse, high-precision manipulation tasks investigated in this study.** *Peg Insertion*: A contact-rich assembly task requiring precise alignment to overcome tight clearances. *Cuboid Stacking*: A precise placement task where stability is challenged by stacking slender cuboids with a small support base. *Roll Ball*: An underactuated dynamics task requiring control over a rolling ball to reach a distant target.

Our work provides a new theoretical framework for understanding the data cost of precision and, based on it, establishes a practical methodology for developing, debugging, and engineering more capable and data-efficient manipulation systems.

II. RELATED WORKS

A. Scaling Laws in Robotics and Beyond

The exploration of scaling laws [4] in robotics is heavily influenced by the revolutionary success of this paradigm in foundational domains like natural language processing (NLP) [1] and computer vision (CV) [3]. Seminal works in these fields established that model performance scales as a predictable power-law with increases in dataset size, model parameters, and compute, leading to the emergent generalization capabilities of large-scale models [1, 2]. Inspired by this, the robotics community has embarked on a large-scale effort to discover analogous laws for embodied intelligence. This effort has predominantly focused on scaling for “Breadth”—improving zero-shot generalization by increasing data diversity across various axes [5–8, 13, 14].

A primary axis of scaling has been data diversity in terms of environments and objects. The work of [5] is the most systematic study in this direction. They empirically demonstrated that a policy’s generalization to novel environments and objects follows a power law with the number of training variations. Their key finding, diversity is all you need, suggests that increasing the variety of settings is far more critical than simply amassing more demonstrations within a fixed setting.

Another significant axis is scaling data heterogeneity across tasks and robot embodiments [8, 15]. The goal here is to train a single, generalist policy on massive, pooled datasets from disparate sources. The Open X-Embodiment project [8] and its derived models like RT-X and Octo [14] epitomize this approach. These works have shown that training on vast,

multi-embodiment, multi-task data can lead to positive cross-skill and cross-embodiment transfer. Similarly, models like Gato [16] aim to create a single agent capable of performing hundreds of diverse tasks.

More recently, scaling has also targeted semantic knowledge by leveraging pre-trained Vision-Language-Models (VLMs). Models like RT-2 [7] co-fine-tune a VLM on both web-scale data and robot trajectories. This allows the model to transfer semantic understanding from the web to the physical world, enabling it to follow novel, abstract instructions.

Taken together, this impressive body of work forms a cohesive research thrust focused on scaling for “Breadth”. These two axes of scaling, Breadth and Depth, are largely orthogonal. A system can excel in generalization, yet its fundamental physical skill precision remains constrained by its manipulation data, as candidly noted in works like RT-2 [7]. While prior work provides invaluable insights into scaling for Breadth, the scaling behavior for achieving “Depth”—i.e., **pushing precision and reliability to their physical limits—remains a critical and open question** that our work directly addresses.

B. High-Precision Robotic Manipulation

Our work also intersects with the extensive literature on high-precision and contact-rich manipulation. Historically, such tasks were the domain of classical control methods [17, 18], which established a paradigm where precision was an outcome of analytical correctness, not data. Foundational concepts such as Operational Space Control [10] and Impedance Control [11] achieve high precision through rigorous analytical models of the robot and its environment, setting a powerful “zero-data” precedent.

When the field transitioned to learning-based approaches like Imitation Learning (IL) and Reinforcement Learning (RL), it inherited this precedent but immediately faced a new, dominant bottleneck: the immense cost, time, and safety risks of real-world data collection [19]. This practical constraint,

combined with the high bar set by classical methods, gave rise to a powerful research imperative: *data efficiency* [20–22].

This pursuit spurred innovations in two main thrusts. The first is algorithmic, aiming to learn more from limited data [23–29] via sample-efficient architectures like Transformers [27], advanced paradigms like adversarial imitation learning [28], or modified objectives like residual RL [29]. The second thrust is system-level, generating richer data by using simulation [30–32], algorithmically expanding datasets with corrective examples [33], or incorporating human feedback [34]. The primary contribution of these works is often the novel, task-specific algorithm or system architecture itself. This frames the research question as “can our novel method solve this hard problem?”, solidifying the treatment of precision as a static goalpost.

The research in this paper adopts a fundamentally different scientific posture. Our objective is not to propose a new data-efficient algorithm to surpass a specific benchmark. Instead, we aim to characterize the universal scaling law that connects data requirements to physical task precision. To achieve this, our methodology is built on a different set of principles. We intentionally fix the learning algorithm to a standard Diffusion Policy implementation. By holding the learning method constant, we can **isolate and measure the intrinsic data cost imposed by “precision” itself** as we treat it as a continuous physical variable. This allows us to uncover the fundamental scaling properties of precision manipulation, which is the precise gap our work aims to fill.

III. METHODOLOGY

This section details the formalisms, core components, and learning paradigm of our study. We first introduce our scaling law formulations and the metric used for evaluation, then describe the robotic learning system that allows us to empirically investigate these laws.

A. Scaling Law Formulation for Precision

Our study investigates the relationship between three core dimensions: *Data Volume* (N), the number of imitation learning trajectories; *Task Precision* (P), the required tolerance for successful task completion; and *System Configuration*, which encompasses the agent’s intrinsic properties like sensors and policies.

Evaluation Metric. The primary metric for our study is the *Success Rate* (SR), a binary measure of task completion. For high-reliability applications such as industrial assembly, where near-perfect performance is often required, SR directly answers the critical question of an agent’s deployment-readiness. This choice is motivated by our use of simulation, which permits a large number of automated rollouts to obtain a statistically robust, binary measure of reliability. This contrasts with real-world experiments, where rollouts are expensive and a more granular, human-assigned score can extract richer information from fewer trials [5]. We also considered the Mean Squared Error (MSE) between the policy’s actions and the expert’s as an alternative metric. However, we

found that MSE does not always reliably correlate with the final task success, a finding consistent with recent large-scale manipulation studies [5]. Therefore, we selected SR as our primary metric due to its direct and unambiguous reflection of a policy’s practical reliability.

Hypothesis 1: The Success Rate Scaling Law. Our first hypothesis posits that for a fixed task precision P , the performance scales with the data volume N . This is inspired by the well-established power-law relationships between performance and data in machine learning. A similar phenomenon was recently observed in robotic manipulation for generalization [5]. We adapt this principle and hypothesize that the failure rate ($1 - SR$) follows a power law with N , which is linear in log-log space:

$$\log(1 - SR) = a \cdot \log(N) + b, \text{ (at target } P\text{)}$$

where a and b are fitting parameters.

Hypothesis 2: The Precision Scaling Law. Our second hypothesis introduces a novel scaling law to describe the data volume N required to achieve a given precision P . For a fixed target success rate, we hypothesize this relationship is governed by the proximity of P to a theoretical limit precision c :

$$\log(N) = m \cdot \frac{1}{P - c} + n, \text{ (at target } SR\text{)}$$

where m and n are fitting parameters specific to this model. The parameter c represents the highest possible precision the system can achieve, which is an intrinsic property of the system configuration, independent of the target success rate. As P approaches this limit, the required data volume is expected to grow super-exponentially.

B. Robotic Learning System

Robot and Environment. All experiments are conducted in the *ManiSkill3* [12] simulation benchmark, which is built upon the *SAPIEN* [35] physics engine. The agent is a *Franka Emika Panda* 7-DoF arm equipped with a parallel-jaw gripper.

Imitation Learning Paradigm. We formulate the task as a behavior cloning problem, where the goal is to learn a policy $\pi(a|o)$ that maps an observation o to an action a .

- **Observation Space (o):** The policy receives multi-modal input consisting of: (1) two 256x256 RGB-D images from a static third-person camera and a wrist-mounted camera, and (2) the robot’s proprioceptive state, including its joint angles and end-effector pose.
- **Action Space (a):** The policy outputs a 7-DoF action $a = (\Delta x, \Delta y, \Delta z, \Delta roll, \Delta pitch, \Delta yaw, gripper)$ representing the desired end-effector delta pose.

Learning Algorithm. To learn the policy $\pi(a|o)$, we use *Diffusion Policy* [36]. We chose this generative approach over simpler regression-based methods for its strong performance and its ability to model complex, multi-modal action distributions, which is crucial for contact-rich precision tasks. Our implementation utilizes a standard architecture with a

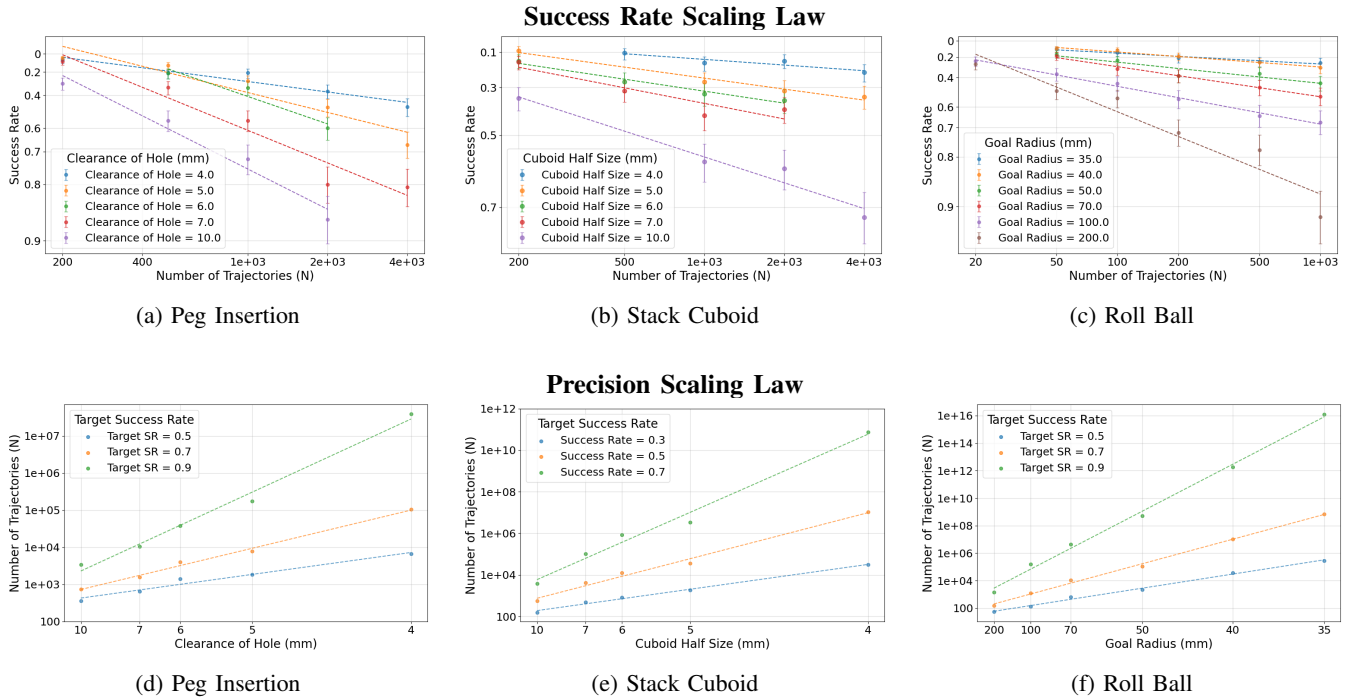


Fig. 2: Empirical Validation of the Proposed Scaling Laws.

ResNet-18 [37] backbone for visual feature extraction and a *1D U-Net* [38] as the denoising network, with action inference performed over 100 *DDPM* [39] steps.

Expert Policy Generation. We employ two types of expert policies depending on the nature of the task. For structured tasks like *Peg Insertion* and *Stack Cuboid*, expert trajectories are generated using privileged, scripted policies that have access to the ground-truth poses of all objects in the simulation. These scripts define a sequence of key waypoints, and a motion planner is used to generate smooth, collision-free trajectories for the robot’s end-effector. Conversely, for a dynamically complex task like *Roll Ball*, where defining an optimal scripted policy is intractable, we instead utilize an expert policy trained via reinforcement learning. This approach also allows us to test the generality of our scaling laws across different expert generation paradigms.

IV. EXPERIMENTS AND RESULTS

This section presents the empirical validation of our proposed scaling laws. We first detail the experimental tasks and evaluation protocol, then present the results that validate our two hypotheses, and finally investigate how the scaling law parameters are influenced by different system configurations.

A. Tasks and Evaluation Protocol

Tasks. We validate our hypotheses on three canonical precision manipulation tasks. To ensure the learned skills are robust, both tasks incorporate extensive domain randomization.

- **Peg Insertion:** The task is to insert a square-profile cuboid peg into a corresponding square hole. The precision P is defined as the clearance between the peg and the hole, which we vary from 4mm to 10mm.

- **Cuboid Stacking:** The task is to stack two slender cuboids. The precision P is defined as the half-side-length of the cuboids’ square base, which we vary from 4mm to 10mm.
- **Roll Ball:** An underactuated dynamics task where the agent must roll a ball such that it passes through a designated circular target region. The precision P for this task is defined by the radius of this target region, which we vary from 35mm to 200mm.

Evaluation Protocol. To empirically measure the scaling laws, we followed a rigorous protocol:

- **Data Generation:** For each precision level P , we generated a data pool consisting only of successful trajectories.
- **Policy Training Sweep:** For a range of data sizes N (e.g., 200, 500, 1000, 2000), we sampled N trajectories to train a Diffusion Policy model. Due to the high computational cost of training (approx. 20 A100-hours per run), each data point (N, P) in our plots corresponds to a single training run.
- **Statistical Evaluation and Checkpoint Selection:** Each policy was trained until convergence (typically 160k to 200k steps), with periodic evaluation. To obtain a stable measure of peak performance while mitigating evaluation noise, we report the mean of the top 3 highest SRs achieved during training. Each evaluation consisted of $M = 100$ episodes. For our direct measurement plots (SR vs. N), we show the 95% confidence interval of the SR estimate as error bars, calculated using the Wilson score interval method over the $3M = 300$ evaluation trials. This provides a robust measure of the evaluation’s statistical reliability.

TABLE I: Parameters for the Success Rate Scaling Law

Task	Precision (mm)	a	b	R^2
Peg Insertion	4.0	-0.19	1.35	0.921
	5.0	-0.35	2.83	0.900
	6.0	-0.49	4.10	0.933
	7.0	-0.58	4.40	0.939
	10.0	-0.72	5.08	0.967
Stack Cuboid	4.0	-0.06	0.35	0.840
	5.0	-0.11	0.70	0.977
	6.0	-0.12	0.67	0.970
	7.0	-0.16	0.92	0.843
	10.0	-0.26	1.41	0.958
Roll Ball	35.0	-0.07	0.19	0.949
	40.0	-0.09	0.38	0.981
	50.0	-0.13	0.42	0.874
	70.0	-0.18	0.69	0.990
	100.0	-0.23	0.61	0.992
	200.0	-0.50	1.88	0.918

TABLE II: Parameters for the Precision Scaling Law

Task	Target SR	m	n	c (mm)	R^2
Peg Insertion	0.5	8.59	7.62	2.35	0.968
	0.7	14.95	7.55		0.993
	0.9	28.62	7.42		0.989
Stack Cuboid	0.3	11.26	5.90	2.75	0.993
	0.5	20.81	6.47		0.988
	0.7	35.30	7.35		0.985
Roll Ball	0.5	198.95	4.77	20.3	0.994
	0.7	346.22	5.72		0.997
	0.9	662.94	7.76		0.996

- **Curve Fitting:** The resulting (N, SR) data points were used to fit our proposed scaling law models. For the Success Rate Scaling Law, we fitted the parameters of the model $\log(1 - SR) = a \cdot \log(N) + b$ via linear regression. For the Precision Scaling Law, we first estimated the data size N required to reach different target success rates (e.g., 0.5, 0.7, 0.9) for each P by interpolating or extrapolating from the fitted linear models of the Success Rate Scaling Law. We then performed a grid search for a single c value that maximized the sum of the coefficients of determination (R^2) across all three target SR curves. This enforces the hypothesis that c is a constant property of the system, independent of the target success rate.

B. Validation of the Proposed Scaling Laws

Our first set of experiments aims to validate the two scaling law hypotheses across all three diverse manipulation tasks. By showing that the laws hold for image-based tasks with scripted experts (Peg Insertion, Stack Cuboid) as well as for a state-based task with an RL-trained expert (Roll Ball), we demonstrate the generality of our findings. The results are presented in Figure 2, with detailed fit parameters in Table I and II.

Success Rate Scaling Law. As predicted by Hypothesis 1, our results confirm that performance scales with data size following a power law. As shown in Figure 2(a-c), all three tasks exhibit a clear power-law relationship between

the failure rate $(1 - SR)$ and the data size N . The detailed fitting parameters in Table I further confirm this with high R^2 values across the board, underscoring the generality of this principle.

Precision Scaling Law. The central result of our paper is the validation of the Precision Scaling Law across all three tasks. As shown in Figure 2(d-f), when we plot $\log(N)$ against $1/(P - c)$ (with fit parameters detailed in Table II), the data points for all tasks and all target success rates align almost perfectly on straight lines. Crucially, for each task, the data points for all three target SRs (e.g., 0.5, 0.7, 0.9) are fitted with a single, shared value of c , yet all yield near-perfect linear fits ($R^2 > 0.97$ as shown in Table II). This strongly **validates our hypothesis that c is an intrinsic property of the system**, independent of the desired performance level.

C. The Limit Precision c as a System Performance Metric

Having validated the basic form of the scaling laws, we now investigate if the *Precision Scaling Law* holds across different system configurations, and whether its key parameter, the limit precision c , acts as a metric for the system’s intrinsic capability. To this end, we use the Peg Insertion task as a case study and conduct a series of ablation studies, summarized in Table III. The results provide strong evidence for both points: the law remains an excellent fit ($R^2 > 0.97$) across all configurations, and the consistent value of c found for each setup validates its role as a system-specific metric, allowing us to analyze how it is influenced by different system components.

Impact of Observation Modality. We first analyze the impact of observation modality by disabling the wrist-mounted camera (Figure 3a, 3b). As shown in Table III (Ablation: Observation), removing this crucial local sensor dramatically degrades system capability, increasing the limit precision c from 2.35 mm to 3.85 mm. This quantifies the critical value of high-quality local observations for precision tasks.

Impact of Expert Policy Strategy. Our baseline system is trained on data from a *conservative expert strategy* (Expert 1). This cautious approach, which performs two additional, corrective alignment maneuvers at the hole’s entrance (Figure 3e), ensures a high raw success rate (about 98% at 5mm clearance) but creates observational ambiguity for the learned policy. While incorporating historical context is a valid and powerful approach to resolve such ambiguities, akin to how large language models leverage longer context windows, our findings reveal a complementary and highly data-efficient principle: that directly reducing the observational ambiguity within the expert demonstrations is a primary lever for improving scaling behavior.

This motivated us to test this principle by designing an *aggressive expert strategy* (Expert 2), which performs a more direct, single-shot insertion (Figure 3e). This design choice represents a trade-off: to obtain observationally unambiguous demonstrations, we forgo the alignment phase, which lowers the expert’s own raw success rate (about 50% at 5mm clear-

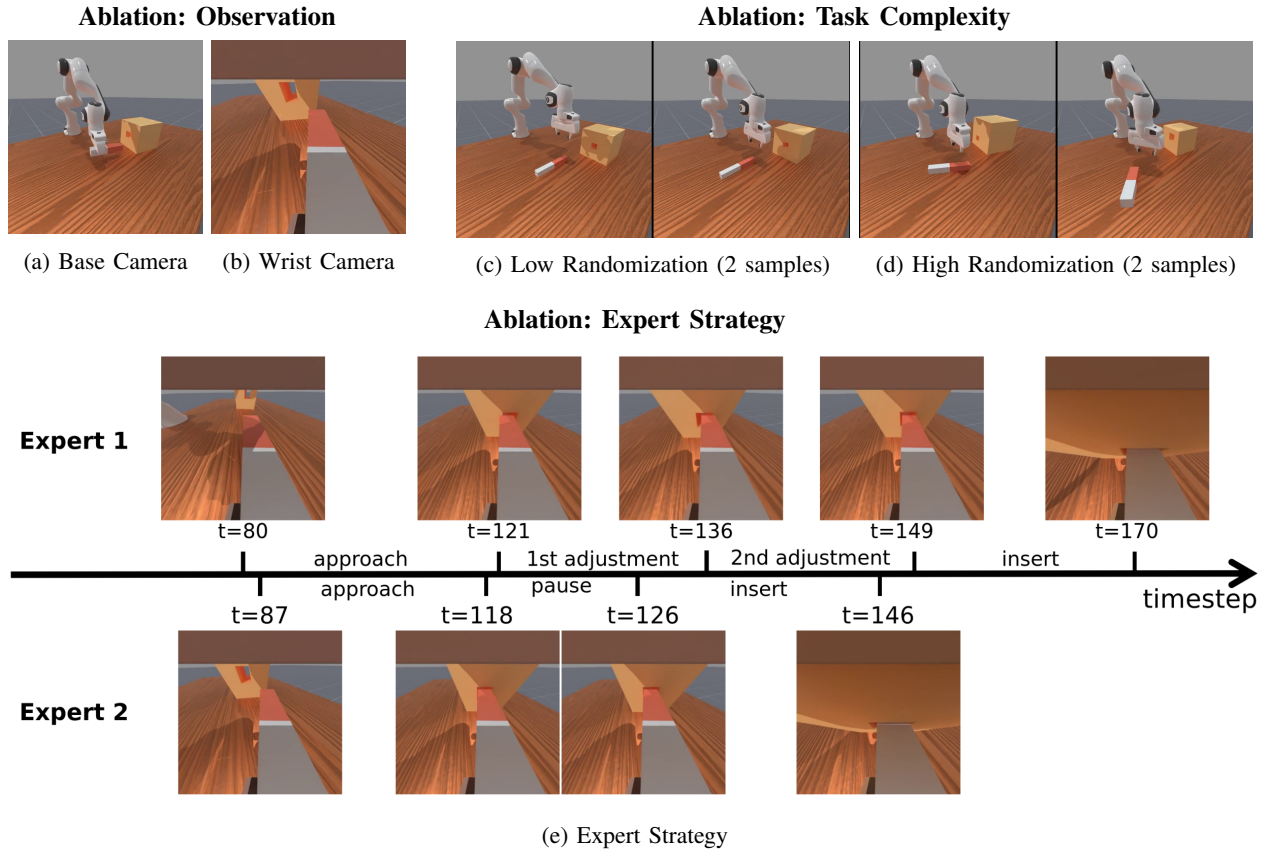


Fig. 3: **Visual comparison of the key system configurations investigated in our ablation studies.** (a-b) **Observation:** The policy’s visual input is composed of a (a) base camera and (b) wrist camera. The baseline system utilizes both, while the ablation relies solely on the base camera. (c-d) **Task Complexity:** We compare two levels of domain randomization. (c) Low Randomization varies only the initial XY positions of the objects. (d) High Randomization, used in our baseline, additionally randomizes object orientations and sizes, hole sizes, and the hole’s position on the box face. (e) **Expert Strategy:** A visual comparison of the two expert strategies. The *Conservative Strategy* (Expert 1, top row) is characterized by several corrective maneuvers prior to insertion. The *Aggressive Strategy* (Expert 2, bottom row) demonstrates a direct strategy, which includes a brief 6-step pause due to the waypoint-based position controller.

TABLE III: The Limit Precision c and Goodness of Fit (R^2) under Different System Configurations.

System Configuration	Limit Precision c (mm)	R^2 ($SR = 0.5$)	R^2 ($SR = 0.7$)	R^2 ($SR = 0.9$)
Baseline	2.35	0.968	0.993	0.989
Ablation: Expert Strategy	1.27	0.986	0.988	0.979
Ablation: Observation	3.85	0.983	0.997	0.987
Ablation: Task Complexity	1.00	0.996	0.992	0.981

* The **Baseline** configuration uses the *Conservative Strategy* (Expert 1) with a wrist camera under high randomization. Each subsequent row represents an ablation study against this baseline. **Expert Strategy** substitutes Expert 1 with the *Aggressive Strategy* (Expert 2); **Observation** removes the wrist camera; **Task Complexity** reduces the level of domain randomization.

ance). However, by filtering for its successful trajectories, we create a dataset of clean demonstrations.

As shown in Table III (Ablation: Expert Strategy), the impact is profound: training on the cleaner data from the aggressive expert results in a dramatically improved limit precision c of 1.27 mm. This key finding demonstrates that for imitation learning, **the clarity and lack of ambiguity in demonstrations can be a more critical factor for achieving high precision** than the expert’s own raw success rate.

Impact of Task Complexity. Finally, we investigate task

complexity by reducing the level of domain randomization (Fig. 3c, 3d). We only randomize object positions, ensuring the agent must still learn a closed-loop policy. As expected, this simplification is reflected in a dramatically improved limit precision c of 1.00 mm, indicating greater data efficiency in a smaller problem space.

These results collectively provide strong evidence that the limit precision c is a sensitive and informative metric for the overall capability of a robotic learning system.

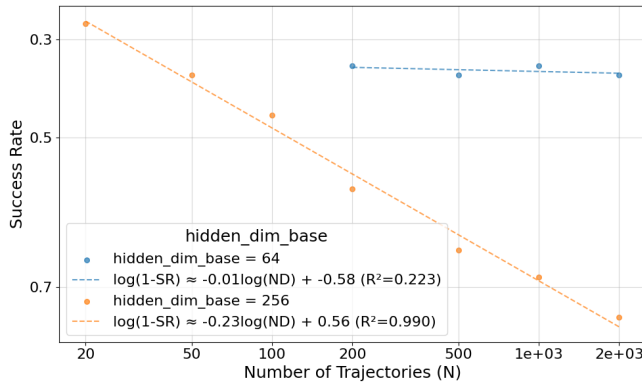


Fig. 4: Impact of model capacity on the Success Rate Scaling Law for the Roll Ball task, comparing a policy trained with a standard-capacity ($\text{hidden_dim_base}=64$) versus a high-capacity ($\text{hidden_dim_base}=256$) denoiser.

D. Further Analysis: On the Precondition of Model Capacity

While validating the scaling laws, the *Roll Ball* task revealed a crucial insight into a precondition for these laws to emerge: the policy network must have sufficient representational capacity.

This phenomenon is illustrated in Figure 4. In our initial experiments using a standard-capacity 1D U-Net denoiser (blue line), a clear scaling law failed to emerge, with a corresponding coefficient of determination of only $R^2 = 0.22$. The policy’s success rate remained low and did not improve predictably with more data. It was only after we substantially increased the capacity of the denoiser network that a distinct scaling law became apparent (orange line), yielding an almost perfect fit with $R^2 = 0.99$.

This finding suggests that for certain dynamically complex tasks, sufficient model capacity is a prerequisite for the scaling law to manifest at all. The system’s representational power must first be adequate to capture the fundamental dynamics of the task before the predictable relationship between data and precision can take hold.

E. Discussion

The empirical findings from our scaling law validation and ablation studies carry significant implications for practical robotic engineering. Our work primarily reveals that for high-precision tasks, “data is not all you need”. When approaching a system’s intrinsic precision limit c , simply scaling up the data volume yields diminishing returns due to the “Curse of Precision”.

More importantly, our scaling law provides a powerful *predictive tool* and a *debugging tool for high-precision tasks*. A key distinction must be made between the law’s one-time, data-intensive *validation* (the focus of this paper) and its far more efficient *application* as a diagnostic. Once the law’s form is assumed, engineers can fit its parameters with relatively few experiments to: 1) Predict the ultimate precision limit (c) of their current system, and 2) Estimate the data cost (N) required to reach a target. This predictive capability is

crucial for debugging. When a task fails at high precision, engineers can evaluate on lower-precision versions to check if performance degrades consistently with the scaling law. A predictable trend suggests the system is hitting its intrinsic limit c , while erratic performance would indicate a bug in the implementation, thus guiding the engineering focus.

V. CONCLUSION

In this work, we shifted the focus of scaling law research in robotics from the open-world paradigm of generalization to the closed-world challenge of precision and reliability. We investigated the fundamental data cost of achieving high precision in manipulation, a critical but underexplored area. Our large-scale empirical study answered the key questions posed in our introduction: we discovered a predictable, super-exponential scaling law governing the relationship between data and precision. We demonstrated that the intrinsic capabilities of the learning system modulate this law through a core parameter, the limit precision c . Finally, we showed that this limit is not a fixed barrier but a quantifiable metric that can be improved through targeted enhancements to system components like sensors and expert data quality.

Our primary contribution is a new theoretical framework and a practical diagnostic tool to guide future efforts away from brute-force data collection and towards more intelligent, system-level solutions. The main takeaway is that for high-precision robotics, the path forward lies not just in scaling data, but in systematically understanding and improving the fundamental limits of our learning systems.

Limitations and Future Work. Our work opens several avenues for future research. First, our ablation study on expert strategies revealed that demonstration clarity, rather than the expert’s raw success rate, is a critical factor for achieving high precision. This suggests a promising research direction in algorithmic data curation. Instead of relying on perfect, unambiguous teleoperation, future work could explore methods to automatically “clean” or “edit” large datasets of imperfect but plentiful demonstrations to extract a smaller set of high-quality, unambiguous trajectories. This could provide a more scalable path to improving a system’s precision limit c . Second, all experiments were conducted in simulation; validating these laws on physical hardware is a crucial next step. Third, our study focuses exclusively on Behavior Cloning. Investigating how these laws change with online methods like Reinforcement Learning or DAgger is another promising direction. Finally, we encourage future work to expand the scope of tasks and system configurations to further validate the universality of the “Curse of Precision”.

REFERENCES

- [1] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, “Scaling laws for neural language models,” *arXiv preprint arXiv:2001.08361*, 2020.
- [2] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford,

- I. Sutskever, and D. Amodei, “Language models are few-shot learners,” *arXiv preprint arXiv:2005.14165*, 2020.
- [3] X. Zhai, A. Kolesnikov, N. Houlsby, and L. Beyer, “Scaling vision transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 12 104–12 113.
- [4] J. Hestness, S. Narang, N. Ardalani, G. Diamos, H. Jun, H. Kianinejad, M. A. Patwary, Y. Yang, and Y. Zhou, “Deep learning scaling is predictable, empirically,” *arXiv preprint arXiv:1712.00409*, 2017.
- [5] F. Lin, Y. Hu, P. Sheng, C. Wen, J. You, and Y. Gao, “Data scaling laws in imitation learning for robotic manipulation,” *arXiv preprint arXiv:2410.18647*, 2024.
- [6] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [7] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choremanski, T. Ding, D. Driess, A. Dubey, C. Finn, P. Florence *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” *arXiv preprint arXiv:2307.15818*, 2023.
- [8] A. Padalkar *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models,” *arXiv preprint arXiv:2310.08864*, 2023.
- [9] B. Siciliano and O. Khatib, Eds., *Springer Handbook of Robotics*, 2nd ed. Springer, 2016.
- [10] O. Khatib, “A unified approach for motion and force control of robot manipulators: The operational space formulation,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [11] N. Hogan, “Impedance control: An approach to manipulation, part i: Theory,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, 1985.
- [12] S. Tao, F. Xiang, A. Shukla, Y. Qin, X. Hinrichsen, X. Yuan, C. Bao, X. Lin, Y. Liu, T. kai Chan, Y. Gao, X. Li, T. Mu, N. Xiao, A. Gurha, V. N. Rajesh, Y. W. Choi, Y.-R. Chen, Z. Huang, R. Calandra, R. Chen, S. Luo, and H. Su, “Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai,” *Robotics: Science and Systems*, 2025.
- [13] A. Zeng, M. Attarian, B. Ichter, K. Choromanski, A. Wong, S. Welker, F. Tombari, A. Purohit, M. S. Ryoo, V. Sindhwani, J. Lee, V. Vanhoucke, and P. Florence, “Socratic models: Composing zero-shot multimodal reasoning with language,” *arXiv preprint arXiv:2204.00598*, 2022.
- [14] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine, “Octo: An open-source generalist robot policy,” *arXiv preprint arXiv:2405.12213*, 2024.
- [15] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn, “Robonet: Large-scale multi-robot learning,” in *Proceedings of the Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, L. P. Kaelbling, D. Kragic, and K. Sugiura, Eds., vol. 100. PMLR, 30 Oct–01 Nov 2020, pp. 885–897.
- [16] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, T. Eccles, J. Bruce, A. Razavi, A. Edwards, N. Heess, Y. Chen, R. Hadsell, O. Vinyals, M. Bordbar, and N. de Freitas, “A generalist agent,” *arXiv preprint arXiv:2205.06175*, 2022.
- [17] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [18] M. H. Raibert and J. J. Craig, “Hybrid position/force control of manipulators,” *Journal of dynamic systems, measurement, and control*, vol. 103, no. 2, pp. 126–133, 1981.
- [19] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [20] S. Schaal, “Learning from demonstration,” in *Advances in Neural Information Processing Systems 9*. Cambridge, MA: MIT Press, 1997, pp. 1040–1046.
- [21] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [22] T. Inoue, G. De Magistris, A. Munawar, T. Yokoya, and R. Tachibana, “Deep reinforcement learning for high precision assembly tasks,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 819–825.
- [23] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet, “Learning latent plans from play,” in *Proceedings of the Conference on Robot Learning (CoRL)*, 2019.
- [24] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 1126–1135.
- [25] L. X. Shi, A. Sharma, T. Z. Zhao, and C. Finn, “Waypoint-based imitation learning for robotic manipulation,” *7th Conference on Robot Learning (CoRL 2023)*, 2023.
- [26] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *CoRL*, 2022.
- [27] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan, “Vima: General robot manipulation with multimodal prompts,” in *International Conference on Machine Learning*, 2023.
- [28] S. Gubbi, S. N. Y. Kolathaya, and B. Amrutur, “Imitation learning for high precision peg-in-hole tasks,” in *2020 6th International Conference on Control, Automation and Robotics (ICCAR)*, 2020, pp. 1–6.
- [29] L. L. Ankile, A. Simeonov, I. Shenfeld, M. T. Villasevil, and P. Agrawal, “From imitation to refinement – residual rl for precise visual assembly,” *arXiv preprint arXiv:2407.16677*, 2024.
- [30] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 23–30.
- [31] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [32] R. Haugaard, J. Langaa, C. Sloth, and A. Buch, “Fast robust peg-in-hole insertion with continuous visual servoing,” in *Proceedings of the 2020 Conference on Robot Learning*, vol. 155. PMLR. PMLR, 2021, pp. 1696–1705.
- [33] L. Ankile, A. Simeonov, I. Shenfeld, and P. Agrawal, “Juicer: Data-efficient imitation learning for robotic assembly,” *arXiv preprint arXiv:2402.14833*, 2024.
- [34] J. Luo, C. Xu, J. Wu, and S. Levine, “Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning,” *arXiv preprint arXiv:2410.21845*, 2024.
- [35] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. Chang, L. J. Guibas, and H. Su, “SAPIEN: A simulated part-based interactive environment,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [36] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion,” in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [38] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, ser. Lecture Notes in Computer Science, vol. 9351. Springer International Publishing, 2015, pp. 234–241.
- [39] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *arXiv preprint arXiv:2006.11239*, 2020.