

From Power to Precision: Learning Fine-grained Dexterity for Multi-fingered Robotic Hands

Jianglong Ye*, Lai Wei*, Guangqi Jiang, Changwei Jing
Xueyan Zou, Xiaolong Wang

UC San Diego

jianglongye.com/power-to-precision




Fig. 1: Our system enables multi-fingered robotic hands to perform a diverse set of manipulation tasks, ranging from precise manipulation of small objects (e.g., pens, nuts, batteries) to power grasps of larger items (e.g., frying pans). This versatility and dexterity are achieved through a co-design framework for both control and fingertip-geometry optimization.

Abstract—Human grasps can be roughly categorized into two types: power grasps and precision grasps. Precision grasping enables tool use and is believed to have influenced human evolution. Today’s multi-fingered robotic hands are effective in power grasps, but for tasks requiring precision, parallel grippers are still more widely adopted. This contrast highlights a key limitation in current robotic hand design: the difficulty of achieving both stable power grasps and precise, fine-grained manipulation within a single, versatile system. In this work, we bridge this gap by jointly optimizing the control and hardware design of a multi-fingered dexterous hand, enabling both power and precision manipulation. Rather than redesigning the entire hand, we introduce a lightweight fingertip geometry modification, represent it as a contact plane, and jointly optimize its parameters along with the corresponding control. Our control strategy dynamically switches between power and precision manipulation and simplifies precision control into parallel thumb–index motions, which proves robust for sim-to-real transfer. On the design side, we leverage large-scale simulation to optimize the fingertip geometry using a differentiable neural-physics surrogate model. We validate our approach through extensive experiments in both sim-to-real and real-to-real settings. Our method achieves an 82.5% zero-shot success rate on unseen objects in sim-to-real precision grasping, and a 93.3% success rate in challenging real-world tasks involving bread pinching. These results demonstrate that our co-design framework can significantly enhance the fine-

grained manipulation ability of multi-fingered hands without reducing their ability for power grasps.

I. INTRODUCTION

A classical taxonomy divides human grasps into two main categories: the *power grasp*, which secures objects against the palm with the fingers, and the *precision grasp*, which places the thumb in opposition to other fingertips [1], [2]. While both are fundamental, the precision grasp is particularly associated with the fine-grained manipulation required for tool use in early humans, which significantly shaped the evolutionary trajectory of our species [3]–[5].

Recent advancements have shown that multi-fingered robotic hands are effective in power grasp, as more contact points provide greater stability [6]–[9]. In contrast, for fine-grained manipulation tasks requiring precision, two-finger parallel grippers are more widely adopted, with impressive applications ranging from folding shirts to gear insertion [10]–[14]. Therefore, replicating the human-level dexterity in multi-fingered hands, especially for precision-oriented manipulation, remains a fundamental challenge.

Our goal is to enable a multi-fingered dexterous hand with reliable precision manipulation while preserving strong power grasp capability. We aim for the following goals: (i) plug-and-play compatibility with existing commercially

* Equal contribution.

available multi-fingered hands by augmenting fingertip geometry rather than designing a new hand from scratch, (ii) achieving both power and precision manipulation in a unified hardware platform and control strategy, and (iii) validation with both sim-to-real transfer using large-scale simulation and real-to-real policies using teleoperation.

To achieve these goals, we adopt a joint control–design optimization framework. On the control side, we switch between power and precise manipulation, where precise manipulation is simplified into coordinated thumb–index actions and parallel finger motions. We also propose a neural switcher that dynamically switch to corresponding grasp mode based on the object. On the design side, we leverage large-scale simulation with a neural physics surrogate model to co-optimize fingertip geometry and control variables, yielding physically effective fingertip designs. Together, these components provide a simple and generalizable solution for both precise and power manipulation.

We validate the proposed system in both sim-to-real and real-to-real settings. In the sim-to-real setting, we focus on grasping tasks, where objects are categorized into power grasps and precise grasps. Demonstrations are optimized with multiple objectives, and a policy is learned from them. Real-world experiment results show that our method outperforms the SOTA method on precise grasps by a large margin. In the real-to-real setting, we tackle more difficult compositional task and precision task. Demonstrations are collected using our switchable teleoperation, where precise manipulation use optimized control. The trained policies are capable of picking up an M4 hex nut with precision and power grasping a pan handle, as shown in Fig. 1.

In summary, we claim the following contributions:

- Control optimization for precise manipulation, including thumb-index motion generation for both grasp synthesis and real-world teleoperation.
- Design optimization with a contact-plane geometry representation and a neural-physics surrogate distilled from large-scale simulations.
- Extensive experiments that demonstrate the effectiveness of our co-design framework in both simulation and the real world, showing that our system improves precise manipulation ability without compromising power grasp.

II. RELATED WORK

Power to Precise Manipulation. *Power manipulation* refers to using whole-hand contacts to apply large, robust forces that move or stabilize objects without requiring high pose accuracy (e.g., holding a hammer, grasping a large bottle). Recent studies [6]–[9], [15]–[19] have shown that multi-fingered robotic hands perform well in power manipulation, mainly because additional contact points increase stability. For example, [6], [7] leverage large-scale simulation to learn dexterous grasping policies and employ point cloud observations for robust sim-to-real transfer. [18], [19] use VR devices to collect high-quality teleoperation data and demonstrate impressive dexterous hand tasks.

In contrast, *precise manipulation* refers to using fingertip-scale, finely controlled motions and contacts to achieve accurate object poses and delicate interactions (e.g., inserting a key, turning a screw). For these precise tasks, parallel grippers are more widely adopted [10]–[14]. How to achieve precise manipulation, especially with multi-fingered hands, is still an open problem for the community. To this end, previous works [20]–[24] either design better hardware or use more accurate sensors such as tactile sensors. [20], [21] optimize hardware design to achieve precise manipulation, while the more general approaches [22]–[24] employ tactile sensing to facilitate precise manipulation. Instead of designing new hardware or using tactile sensing, our approach optimizes both control and finger tip geometry for precise manipulation.

Computational Design and Co-Design. Previous computational methods for mechanical design often use gradient-based optimization and task-aware planning over geometry or topology to create structures for target tasks [25], [26]. However, these pipelines can be computationally expensive and may not generalize well to new objects or tasks. Differentiable simulators address this by enabling the joint optimization of morphology and control within a single framework [27]. Recent work has increasingly utilized co-design as a tool for manipulation, such as in the creation of dexterous hands [28], fingertips [29], and grippers [30], [31]. Rather than designing a completely new dexterous hand, we enhance an existing hand by optimizing its geometry. In our co-design framework, this is achieved by attaching a newly optimized fingertip cover that improves precise manipulation. During optimization, we model the fingertip geometry as a contact plane and leverage large-scale simulation by using a neural dynamics surrogate. **Imitation Learning.** Imitation learning is a common framework for learning robot control policies from demonstrations, and it has been applied in many recent works [18], [19], [32]–[42]. Demonstrations may be obtained through teleoperation [19], [43]–[45], simulation [6], [46]–[48], real-world policy trial-and-error [49], [50], human videos [37], [39], [51], or a mixture of these sources [38], [52], [53]. In this work, we validate our approach on real-to-real tasks using teleoperation data and on sim-to-real tasks using simulation data. We apply existing imitation learning methods [6], [10] to our collected demonstrations. Our co-design framework improves the robustness of both demonstrations and the robotic hand, leading to more effective precise manipulation policies in simulation and the real world.

III. METHOD

A. Overview

This work focuses on learning manipulation skills for multi-fingered hands using imitation learning. In this framework, a control policy π learns from expert demonstrations $D = (\mathcal{O}, \mathcal{A})$, consisting of observations $\mathcal{O}$ and actions $\mathcal{A}$ collected through teleoperation, simulation, or other methods. To achieve precise manipulation, we introduce an approach that jointly optimizes both control (Sec. III-B) and design

(Sec. III-C) of dexterous hands. We cross validate our approach on tasks under both sim-to-real and real-to-real settings, where data is collected through optimization in simulation and teleoperation in the real world.

Control Optimization. Our control method follows two principles. The first is *categorization*, where we switch between power and precision grasps based on the target object. For example, in our *Cooking Setup* task, a power grasp is used for a pot handle, while a precision pinch grasp is used for a thin asparagus. The second principle is *simplification*. We achieve this by reducing the degrees of freedom of the hand and the number of contact points. We also constrain finger motion to be parallel movements during precision grasps. Although it is possible to use more complex motions for precise manipulation, our experiments show that simpler motions are effective for both retargeting and sim-to-real transfer.

Design Optimization. An effective control strategy can still be limited by the hand’s physical design. To address this, we apply a co-design framework to optimize the fingertip geometry for precise manipulation. We model the contact region between the thumb and index finger as a plane and search for the optimal plane. This optimization process uses a learned, differentiable forward dynamics model along with several objectives, which allows us to leverage large-scale simulation to find a better fingertip design.

B. Control Optimization for Precise Manipulation

The robot’s action space $\mathcal{A}$ is defined as the target joint positions $q \in \mathbb{R}^d$, where $d = 19$ (7 for the arm and 12 for the hand). The collected data, in both simulation and the real world, consist of observations $\mathcal{O}$ and trajectories of actions $\{q_t\}_{t=1}^T$, where each $q_t \in \mathcal{A}$.

Sim-to-real Grasping. We proof the effectiveness of the proposed optimization pipeline in Sim-to-real settings. In simulation, a common way to collect demonstrations for dexterous grasping is to first optimize for force closure and then apply motion planning combined with simulation filtering [6], [7], [48]. Although this framework can produce grasps that work even for tiny objects in simulation, the poses are often complex because all fingers are encouraged to apply force to the object. Deploying these grasps to the real world is impractical due to the sim-to-real gap.

We categorize grasping poses based on the type of target object. For simplicity, we focus on two categories: power grasp for large objects, and precision grasp for small or thin objects. For power grasp optimization, we use the original force closure objective from [54]. For precision grasp, we introduce a variant objective. Given object mesh O and joint configuration q , we sample n contacts on the thumb and n on the index fingertip. Each contact has position $x_i \in \mathbb{R}^3$ and normal $c_i \in \mathbb{R}^3$. The grasp map is $G = \begin{bmatrix} I_3 & \cdots & I_3 \\ [x_1]_{\times} & \cdots & [x_{2n}]_{\times} \end{bmatrix}$, where $[x]_{\times}$ is the skew-symmetric matrix of x . We minimize $E_{\text{precise}} = \|Gc\|_2$, with $c = [c_1^\top, \dots, c_{2n}^\top]^\top$, which encourages the net wrench




Fig. 2: **Control Optimization.** For precise grasps control, we optimize for opposite force closure and parallel finger motions, which significantly improve sim-to-real transfer.

to approach zero for the thumb-index grasp. We also fix other fingers during optimization to reduce the degrees of freedom.

After optimizing the grasp pose, a common way for power grasp to obtain pre-grasp and overshoot-grasp poses is based on the object’s signed distance function (SDF), which pushes fingers toward the object surface [6], [48]. We propose to use simple parallel motions for precision grasps. Let d be the normalized direction vector from the thumb contact points to the index contact points. To generate a pre-grasp pose, the thumb and index move apart by a small distance α . The required joint velocity are calculated using the Jacobian pseudoinverse $J^\dagger$ as $\Delta q_{\text{thumb}} = -\alpha J_{\text{thumb}}^\dagger d$, $\Delta q_{\text{index}} = \alpha J_{\text{index}}^\dagger d$. Experimental results show that this simple thumb-index grasp with parallel motion is robust for sim-to-real deployment. The finger motion and its real-world counterpart are visualized in Fig. 2.

All demonstrations are filtered using the ManiSkill simulator [55]. After data collection, we train two DexSimple [6] policies, one on power grasp data and one on precision grasp data. We then add a switcher consisting of PointNet [56] and an MLP to predict whether an object should be grasped with a power grasp or a precision grasp, and apply the corresponding policy actions. The policy is deployed to real world in a zero-shot fashion.

Real-to-real Tasks. We tackle both compositional tasks that combine power and precision grasps, and pure precision tasks in a real-to-real setting. Teleoperation is used to collect demonstrations. The standard position-based retargeting [57] struggles with fine-grained actions such as pinching a nut. To address this, we switch between normal retargeting and a pinch controller. The hand pose is optimized by minimizing $E_{\text{precise}} = \|Gc\|_2$, where c represents contact normals between the thumb and index finger (rather than between hand and object). We apply the same parallel finger motions used in the sim-to-real setting, where the thumb and index move along direction d with joint updates $\Delta q = J^\dagger d$. Finally, an ACT policy [10] is trained on these teleoperated demonstrations for deployment.

C. Design Optimization for Precise Manipulation

We aim to optimize hand geometry to enhance precision manipulation without compromising power manipulation. To this end, we represent geometry using a simple contact plane P , parameterized by a reference point p and a unit normal




Fig. 3: **Design Optimization.** We optimize fingertip geometry (represented as a contact plane) under multiple objectives. The resulting fingertip cover improves precision manipulation in real-world.

vector n : $P = \{x \in \mathbb{R}^3 \mid n^\top(x - p) = 0\}$. This simplifies simulation and improves robustness in sim-to-real transfer. Given P , we project a slightly inflated convex hull of the fingertip onto it and 3D print the resulting union geometry. The fingertip cover is easy to assemble and generalizes well to different multi-fingered hands.

During plane parameter optimization, we jointly optimize P and the joint position q under multiple objectives using gradient descent. (Note that q here is not used for grasp optimization; it assumes that no object is present.) The objectives and the real-world fingertip covers are shown in Fig. 3.

Geometric Objectives. The first two objectives encourage thumb-index contact while preventing penetration:

$$E_{\text{att}} = \sum_{i=1}^N d(x_i, P), \quad E_{\text{rep}} = \sum_{v \in S(F)} [\phi_P(v) < 0] d(v, P). \quad (1)$$

Here, x_i are candidate contact points on the thumb and index, $S(F)$ is the surface point cloud of finger meshes F , $d(\cdot, P)$ is the point-to-plane distance, ϕ_P is the signed distance to P , and $[\cdot]$ denotes the indicator function. The attraction term pulls sampled contacts toward P , while the repulsion term penalizes surface points that cross P .

Manipulability Objective. To further ensure stable motion, we include an objective that encourages the thumb and index to move in parallel along the plane normal direction. We measure this using directional manipulability [58]:

$$E_{\text{mani}} = -(\|J_{\text{thumb}} n\|_2 + \|J_{\text{index}} n\|_2), \quad (2)$$

where J_{thumb} and J_{index} are the Jacobians of the thumb and index, respectively.

Neural Physics Objective. While the previous terms focus on kinematics and geometry, we further introduce a neural physics objective to leverage large-scale simulation. We first optimize 1k plane parameters from different initializations using the above objectives, and then evaluate them in simulation on the grasping task introduced in Sec. III-B. The simulation outcomes are distilled into a neural surrogate model f , implemented as a PointNet encoder followed by an MLP:

$$f : (P, q, o) \mapsto \hat{s}, \quad (3)$$

where (P, q) denotes the plane parameters and joint configuration, o is the object observation (point cloud), and $\hat{s} \in [0, 1]$ is the predicted task success probability. We incorporate this surrogate into the optimization as an energy term:

$$E_{\text{phys}} = -f(P, q, o), \quad (4)$$

which encourages geometries and poses that maximize the predicted success probability. This neural term provides gradient feedback during optimization: we sample a batch of objects and jointly optimize a shared P with diverse q . In this way, P is refined toward geometries that are not only kinematically consistent but also physically effective for manipulation.

IV. EXPERIMENT

A. Experimental Settings

Our method is validated on two platforms: XArm robotic arm + XHand dexterous hand (7+12 DOFs, referred to as xArm below), as well as Unitree G1 Humanoid + Inspire dexterous hand (7+6 DOFs, referred to as G1 below). We cross-validate it on tasks under both sim-to-real and real-to-real settings.

Sim-to-real Tasks. We focus on the *grasping* task in the sim-to-real setting, following the Dex1B [6] benchmark. Our dataset includes 7k Objaverse [59] objects and 1k primitive shapes (spheres, boxes, cylinders) of various sizes. The data are categorized by grasp type. Objaverse objects with successful poses optimized by E_{precise} (Sec. III-B) are assigned to precision grasps, while the rest are used for power grasps. All primitive shapes are used for precision grasps. In total, 6k objects are used for power grasps and 3k for precision grasps. We hold out 30% of the objects for testing. We collect 30k trajectories for power and precision grasps, respectively.

The success criterion in Dex1B for grasping is to lift an object from the table to a certain height while maintaining contact between the fingers and the object. During data collection, Dex1B additionally applies lateral external forces to the objects to ensure that the optimized grasp pose is robust, but these external forces are removed during final evaluation. To ensure better sim-to-real transfer, we keep external forces during all evaluations.

We adopt the DexSimple policy [6] for the grasping policy. The neural switcher (Sec. III-B) consists of a PointNet [56] and an MLP, with hidden dimensions (256, 128). The neural physics model (Sec. III-C) also consists of a PointNet and an MLP with the same hidden dimensions.

Real-to-real Tasks. Compared to sim-to-real experiments, we focus on more difficult compositional task and precision task in the real-to-real setting. The task description are as below:

- **Cooking Setup.** A compositional task in which the robot must sequentially pinch-grasp an asparagus spear from the cutting board, place it into the frying pan, then regrasp the pan handle and lift it off the stove.

Method	Optimization		Power G. SR (%)		Precise G. SR (%)	
	Design	Control	Seen	Unseen	Seen	Unseen
Simulation (xArm)						
Dex1B [6]			59.60	55.88	56.38	53.91
Ours		✓	60.64	54.12	61.54	59.04
Ours	✓	✓	61.52	53.35	64.74	64.17
Simulation (G1)						
Dex1B [6]			60.06	56.46	44.75	44.44
Ours		✓	60.54	55.42	45.62	45.26
Ours	✓	✓	59.93	57.67	49.91	49.32
Zero-shot Sim-to-Real (xArm)						
Dex1B [6]			–	60.00	–	12.50
Ours		✓	–	70.00	–	20.00
Ours	✓	✓	–	80.00	–	82.50

TABLE I: **Main Results for Sim-to-real Grasping Task.**

The top two parts report simulation success rates for G1 and xArm, and the bottom part shows zero-shot sim-to-real results. Our method with joint control and design optimization consistently outperforms Dex1B on precise grasps, especially in real world precision grasps. Power G. and Precise G. denote Power Grasp and Precision Grasp, respectively.

- **Multi-pen Grasp.** The robot is required to grasp two marker pens and place them into a box within a single attempt. Specifically, it first pinch-grasps one pen and dexterously rolls it into the palm, secured by the remaining three fingers. It then pinch-grasp the second pen, before dropping both pens into the box.
- **Nut onto Peg.** In this task, the robot must precisely pinch-grasp an M4 hex nut (inner diameter $\Phi = 4.1$ mm) from the tabletop and accurately insert it onto an upright M3 bolt (outer diameter $\Phi = 2.9$ mm). The small clearance between the nut and bolt requires fine dexterity and precise alignment, making this task particularly challenging for precise manipulation.
- **Bread Pinch.** In this task, the robot is required to pinch-grasp a single slice of toast from the table. Since excessive downward pressing may deform the bread or even trigger the emergency stop, the robot must execute the grasp with precise control.
- **Battery Insert.** This task involves a sequence of precise manipulations in which the robot must grasp a battery from the table, align it with the charging socket, place it in position, and apply a controlled push to ensure it is fully secured. This task is conducted on G1 setup.

We employ the teleoperation framework from [19] to collect demonstration data, where the teleoperator’s wrist pose is mapped to the XArm end-effector, and the teleoperator’s finger motions are retargeted to the corresponding XHand finger positions. For the baseline retargeting method [57], we adopt the dexpilot retargeting scheme. For the scripted baseline, we record joint configurations from manually executed passive grasping motions. In our optimized retargeting approach, the Euclidean distance between the teleoperator’s thumb and index fingertips is mapped to the opening angle of the corresponding XHand fingers. We collected 15 successful



Fig. 4: Real-world test objects used for evaluating zero-shot sim-to-real grasping.

demonstrations for each task.

After data collection, we use [10] for the autonomous policy to verify the policy learning ability in our design. The action trunking transformer (ACT) is implemented by [55], where the input is third-person camera view with XArm-XHand joint position as proprioception information. The output action is joint targets of Robot Arm and Hand. It is worth noting that we do not employ processed or quantified gripper values for hand supervision, but instead rely solely on raw joint command signals. This aims to validate that our control optimization would not affect the performance of the policy. We tested 15 times for each task on the real robot.

B. Sim-to-real Results

Main Results. The main results of the sim-to-real grasping task are reported in Tab. I. We mainly compare against Dex1B [6], a state-of-the-art sim-to-real grasping policy.

The simulation results are presented in the top part of Tab. I, showing policy success rates for power and precision grasps on both training and testing splits. Both Dex1B and our policy are trained with 30K demonstrations per grasp type. The key differences are that our precision grasp demonstrations are collected using the method described in Sec. III-B (Control optimization in the table), and our fingertip geometry is optimized using the method described in Sec. III-C (Design optimization in the table). Since all policies are trained on successful demonstrations, the Dex1B policy still achieves strong performance (53.91%) on tiny objects in the precision grasp category. Our policy outperforms Dex1B on precision grasps by about 10% (64.17% vs 53.91%), primarily because parallel finger motions and flat contact planes provide greater robustness when the policy predicts imperfect grasping poses. At the same time, our policy achieves comparable results on power grasps, indicating that the added design optimization does not compromise the dexterous hand’s capability for power grasps.

The zero-shot sim-to-real results are reported in the bottom part of Tab. I. The real-world testing objects are shown in Fig. 4, and all objects are unseen during training. The policy rollouts are visualized in Fig. 5(a) and (b). We evaluate 5 trials per object, resulting in 40 trials in total for each category. Unlike the simulation results, our method with opti-




Fig. 5: Task policy rollouts across eight tasks. (a)-(b) and (g) illustrate sim-to-real transfer of the precision grasp policy on the xArm and G1 setup. (c)-(f) and (h) show real-to-real executions on five distinct tasks: Cooking Setup, Multi-pen Grasp, Nut onto Peg, Bread Pinch, and Battery Insert on xArm and G1.

mized design outperforms Dex1B by a large margin (82.50% vs. 12.50%) in real-world deployment. The main reason is that all-finger grasping in Dex1B is often too complex for small objects, making it impractical to deploy under sim-to-real gaps in sensing, calibration, and dynamics. We also observe that our policy with control-only optimization does not achieve strong performance (20%), which indicates that design optimization is crucial for robust sim-to-real transfer, as the flat plane provides significantly more contact area.

Ablation Study. We conduct ablation studies in simulation to validate the effectiveness of our proposed techniques. Specifically, we ablate the control optimization, the design optimization, and the E_{phys} objective (see Sec. III). Instead of training an imitation learning policy on successful demonstrations and evaluating performance as in Tab. I, we directly report the optimization success rates for precise grasp data collection, which serve as a more direct metric of effectiveness.

Our control-only optimization (C. only in the table) achieves a lower success rate (0.41% vs. 2.75%), since reducing the DOFs makes the optimization problem more challenging. Nevertheless, this control optimization is still valuable because the resulting parallel motions are more

Method	Optimization		Objective	Opt. SR (%)
	Design	Control	Physics	
Dex1B [6]				2.75
C. only		✓		0.41
C.+D. (w/o E_{phys})	✓	✓		3.77
C.+D.	✓	✓	✓	5.35

TABLE II: **Ablation Study for Precise Grasp Optimization in Simulation.** We report the optimization success rate (Opt. SR; %) for control-only (C.), joint control and design (C.+D.), and the physics objective E_{phys} . Joint C.+D. achieves the highest SR (5.35%). Incorporating the E_{phys} objective improves performance from 3.77% to 5.35%.

robust for learning and sim-to-real transfer. Our joint control and design optimization (C.+D. in the table) achieves the highest success rate (5.35%), highlighting the effectiveness of design optimization. We also evaluate design optimization without the E_{phys} objective (C.+D. w/o E_{phys} in the table). The performance gap (3.77% vs. 5.35%) indicates that the E_{phys} objective distilled from large-scale simulation helps find more physically plausible fingertip geometries.

Method	Setting		Compositional Task		Precision Task		
	Design	Control	Cooking Setup	Multi-pen Grasp	Nut onto Peg	Bread Pinch	Battery Insert
Autonomous Policy							
Baseline	Original	Retargeting [57]	20.0%	53.3%	–	60.0%	13.3%
Ours	Optimized	Optimized	73.3%	66.7%	66.7%	93.3%	66.7%
Teleoperation							
Baseline	Original	Retargeting [57]	41.7%	57.7%	6.6%	57.1%	26.7%
Ours	Optimized	Optimized	88.2%	50.0%	68.2%	93.8%	80.0%

TABLE III: Results under autonomous policy and teleoperation on compositional and precision tasks.

Setting		Success Rate (%)	
Design	Control	Policy	Teleoperation
Original	Retargeting [57]	60.0%	57.1%
Original	Manual Script	73.3%	57.1%
Manual Design	Manual Script	60.0%	50.0%
Optimized (Ours)	Optimized (Ours)	93.3%	93.8%

TABLE IV: Ablation study on the bread pinch task. We evaluate different combinations of fingertip designs and control methods.

C. Real-to-real Results

Main Results. The Table III highlights the effectiveness of the proposed co-optimized control and design framework across both autonomous policy and teleoperation settings, for both compositional and precision manipulation task.

In both teleoperation and autonomous settings, the baseline system shows limited success due to the original fingertip design, which makes it difficult to reliably pinch small or thin objects such as asparagus in the Cooking Setup, pens in the Multi-pen Grasp, and thin-cut toast in the Bread Pinch task. With our optimized control–design, success rates increase significantly: for example, from 20.0% to 73.3% in Cooking Setup and from 0.0% to 66.7% in Nut onto Peg task for the autonomous policies. This demonstrates marked improvements in both compositional and fine-grained precision manipulation skills. Meanwhile, we show that both teleoperation, which maps the operator fingertip distance to the optimized path, and autonomous execution, which learns from joint positions, confirm the effectiveness of our system.

Table IV reports the ablation study results on the Bread Pinch task. Our optimized fingertip design combined with control consistently outperforms all variants, surpassing both the naive implementation and the manually crafted geometry–control design.

Qualitative Results To further illustrate the effectiveness of our proposed method, Fig. 5 presents task rollouts across eight representative scenarios. Subfigures (a), (b) and (g) demonstrate successful sim-to-real transfer of the precision grasp policy, where the robot reliably manipulates small objects like a screwdriver, banana and glue stick using the co-optimized fingertip design and control strategy.

Subfigures (c)-(f) and (h) highlight real-to-real executions of complex and fine-grained tasks: Cooking Setup, Multi-pen

Grasp, Nut onto Peg, Bread Pinch, and Battery Insert. These tasks involve diverse object geometries, contact constraints, dynamic demands and precision demands. In each case, the robot achieves stable and repeatable performance, showcasing the system’s ability to generalize across manipulation contexts. The visual results confirm that the co-optimization approach leads to robust dexterity even in challenging real-world scenarios.

V. CONCLUSION

We introduced a unified framework that enables multi-fingered robotic hands to perform both power and precision grasps through a combination of control-policy learning and fingertip geometry optimization. By simplifying precision control into parallel thumb-index motions and co-designing fingertip covers using a neural-physics surrogate, our method achieves robust, generalizable manipulation without requiring complex hardware modifications.

Extensive evaluations in both sim-to-real and real-world tasks show significant improvements over existing approaches, particularly in fine-grained precision grasping. These results highlight the effectiveness of combining lightweight mechanical design with data-driven control, offering a practical path toward more dexterous and adaptable robotic systems.

REFERENCES

- [1] J. R. Napier, “The prehensile movements of the human hand.” *The Journal of Bone and Joint Surgery. British volume*, 1956.
- [2] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic, “The grasp taxonomy of human grasp types,” *Transactions on Human-Machine Systems*, 2015.
- [3] M. M. Skinner, N. B. Stephens, Z. J. Tsehai, A. C. Foote, N. H. Nguyen, T. Gross, D. H. Pahr, J.-J. Hublin, and T. L. Kivell, “Human-like hand use in australopithecus africanus,” *Science*, 2015.
- [4] T. L. Kivell, “Evidence in hand: recent discoveries and the early evolution of human manual manipulation,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2015.
- [5] F. A. Karakostis, D. Haeufle, I. Anastopoulou, K. Moraitis, G. Hotz, V. Tourloulis, and K. Harvati, “Biomechanics of the human thumb and the evolution of dexterity,” *Current Biology*, 2021.
- [6] J. Ye, K. Wang, C. Yuan, R. Yang, Y. Li, J. Zhu, Y. Qin, X. Zou, and X. Wang, “Dex1b: Learning with 1b demonstrations for dexterous manipulation,” in *Robotics: Science and Systems (RSS)*, 2025.
- [7] J. Chen, Y. Ke, and H. Wang, “Bodex: Scalable and efficient robotic dexterous grasp synthesis using bilevel optimization,” in *International Conference on Robotics and Automation (ICRA)*, 2025.
- [8] R. Singh, A. Allshire, A. Handa, N. Ratliff, and K. Van Wyk, “Dextrah-rgb: Visuomotor policies to grasp anything with dexterous hands,” *arXiv preprint arXiv:2412.01791*, 2024.

- [9] H.-S. Fang, H. Yan, Z. Tang, H. Fang, C. Wang, and C. Lu, "Anydex-grasp: General dexterous grasping for different hands with human-level learning efficiency," *arXiv preprint arXiv:2502.16420*, 2025.
- [10] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *RSS*, 2023.
- [11] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, "Aloha unleashed: A simple recipe for robot dexterity," in *CoRL*, 2024.
- [12] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, *et al.*, " $\pi 0$: A vision-language-action flow model for general robot control," *arXiv*, 2024.
- [13] P. Intelligence, K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, *et al.*, " $\pi 0.5$: a vision-language-action model with open-world generalization," *ArXiv*, 2025.
- [14] G. R. Team, S. Abeyruwan, J. Ainslie, J.-B. Alayrac, M. G. Arenas, T. Armstrong, A. Balakrishna, R. Baruch, M. Bauza, M. Blokzijl, *et al.*, "Gemini robotics: Bringing ai into the physical world," *arXiv preprint arXiv:2503.20020*, 2025.
- [15] H. Zhang, S. Christen, Z. Fan, O. Hilliges, and J. Song, "Graspxl: Generating grasping motions for diverse objects at scale," in *ECCV*, 2024.
- [16] Y. Zhong, X. Huang, R. Li, C. Zhang, Z. Chen, T. Guan, F. Zeng, K. N. Lui, Y. Ye, Y. Liang, *et al.*, "Dexgraspvla: A vision-language-action framework towards general dexterous grasping," *arXiv preprint arXiv:2502.20900*, 2025.
- [17] J. He, D. Li, X. Yu, Z. Qi, W. Zhang, J. Chen, Z. Zhang, Z. Zhang, L. Yi, and H. Wang, "Dexvlg: Dexterous vision-language-grasp model at scale," *arXiv preprint arXiv:2507.02747*, 2025.
- [18] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang, "Open-television: Teleoperation with immersive active visual feedback," *arXiv preprint arXiv:2407.01512*, 2024.
- [19] R. Ding, Y. Qin, J. Zhu, C. Jia, S. Yang, R. Yang, X. Qi, and X. Wang, "Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning," *arXiv preprint arXiv:2407.03162*, 2024.
- [20] J. Zhao and E. H. Adelson, "Gelsight svelte hand: A three-finger, two-dof, tactile-rich, low-cost robot hand for dexterous manipulation," *arXiv preprint arXiv:2309.10886*, 2023.
- [21] Y. Hong, Y. Zhao, J. Berman, Y. Chi, Y. Li, H. Huang, and J. Yin, "Angle-programmed tendril-like trajectories enable a multifunctional gripper with ultradelicacy, ultrastrength, and ultraprecision," *Nature Communications*, vol. 14, no. 1, p. 4625, 2023.
- [22] W. K. Do, A. K. Dhawan, M. Kitzmann, and M. Kennedy, "Densetact-mini: an optical tactile sensor for grasping multi-scale objects from flat surfaces," in *2024 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2024, pp. 6928–6934.
- [23] Z. Yu, W. Xu, S. Yao, J. Ren, T. Tang, Y. Li, G. Gu, and C. Lu, "Precise robotic needle-threading with tactile perception and reinforcement learning," in *Conference on Robot Learning*. PMLR, 2023, pp. 3266–3276.
- [24] A. Bronars, S. Kim, P. Patre, and A. Rodriguez, "Texterity: Tactile extrinsic dexterity," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 7976–7983.
- [25] F. Chen and M. Y. Wang, "Design optimization of soft robots: A review of the state of the art," *IEEE Robotics & Automation Magazine*, vol. 27, no. 4, pp. 27–43, 2020.
- [26] D. Feshbach, W.-H. Chen, L. Xu, E. Schaumburg, I. Huang, and C. Sung, "Algorithmic design of kinematic trees based on csc dubins planning for link shapes," in *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2024.
- [27] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty: An experimental comparison on a dynamic bipedal walker," *Annals of Mathematics and Artificial Intelligence*, vol. 76, pp. 5–23, 2016.
- [28] P. Mannam, X. Liu, D. Zhao, J. Oh, and N. Pollard, "Design and control co-optimization for automated design iteration of dexterous anthropomorphic soft robotic hands," in *2024 IEEE 7th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2024, pp. 332–339.
- [29] K. Ikemura, Y. Dong, D. Blanco-Mulero, A. Longhini, L. Chen, and F. T. Pokorny, "Efficient end-effector co-design by demonstration for deformable fragile object manipulation," in *1st Workshop on Robot Hardware-Aware Intelligence*.
- [30] R. Liu, J. Liang, S. Sudhakar, H. Ha, C. Chi, S. Song, and C. Vondrick, "Paperbot: Learning to design real-world tools using paper," *arXiv preprint arXiv:2403.09566*, 2024.
- [31] S. Yi, X. Bai, A. Singh, J. Ye, M. T. Tolley, and X. Wang, "Co-design of soft gripper with neural physics," *arXiv preprint arXiv:2505.20404*, 2025.
- [32] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," in *Conference on robot learning*. PMLR, 2017, pp. 357–368.
- [33] A. Mandlekar, D. Xu, R. Martín-Martín, S. Savarese, and L. Fei-Fei, "Learning to generalize across long-horizon tasks from human demonstrations," *arXiv preprint arXiv:2003.06085*, 2020.
- [34] Y. Zhu, P. Stone, and Y. Zhu, "Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4126–4133, 2022.
- [35] S. Nasiriany, T. Gao, A. Mandlekar, and Y. Zhu, "Learning and retrieval from prior data for skill-based imitation learning," *arXiv preprint arXiv:2210.11435*, 2022.
- [36] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, "What matters in learning from offline human demonstrations for robot manipulation," *arXiv preprint arXiv:2108.03298*, 2021.
- [37] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang, "Dexmv: Imitation learning for dexterous manipulation from human videos," in *ECCV*, 2022.
- [38] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, "Mimicgen: A data generation system for scalable robot learning using human demonstrations," *arXiv preprint arXiv:2310.17596*, 2023.
- [39] J. Ye, J. Wang, B. Huang, Y. Qin, and X. Wang, "Learning continuous grasping function with a dexterous hand from human demonstrations," *RA-L*, 2023.
- [40] S. Young, D. Gandhi, S. Tulsiani, A. Gupta, P. Abbeel, and L. Pinto, "Visual imitation made easy," in *Conference on Robot Learning*. PMLR, 2021, pp. 1992–2005.
- [41] R.-Z. Qiu, Y. Song, X. Peng, S. A. Suryadevara, G. Yang, M. Liu, M. Ji, C. Jia, R. Yang, X. Zou, and X. Wang, "Wildlma: Long horizon loco-manipulation in the wild," *arXiv preprint arXiv:2411.15131*, 2024.
- [42] L. Wei, J. Ma, Y. Hu, and R. Zhang, "Ensuring force safety in vision-guided robotic manipulation via implicit tactile calibration," *arXiv preprint arXiv:2412.10349*, 2024.
- [43] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, *et al.*, "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Conference on Robot Learning*. PMLR, 2018, pp. 879–893.
- [44] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, *et al.*, "Rt-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.
- [45] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, "Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators," in *IROS*, 2024.
- [46] Y. Wang, Z. Xian, F. Chen, T.-H. Wang, Y. Wang, K. Fragkiadaki, Z. Erickson, D. Held, and C. Gan, "Robogen: Towards unleashing infinite data for automated robot learning via generative simulation," *arXiv preprint arXiv:2311.01455*, 2023.
- [47] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3019–3026, 2020.
- [48] R. Wang, J. Zhang, J. Chen, Y. Xu, P. Li, T. Liu, and H. Wang, "Dexgraspnet: A large-scale robotic dexterous grasp dataset for general objects based on simulation," in *ICRA*, 2023.
- [49] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 3406–3413.
- [50] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, *et al.*, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *Conference on robot learning*. PMLR, 2018, pp. 651–673.
- [51] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, "Affordances from human videos as a versatile representation for robotics," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 778–13 790.
- [52] S. Dass, K. Pertsch, H. Zhang, Y. Lee, J. J. Lim, and S. Nikolaidis,

- “Pato: Policy assisted teleoperation for scalable robot data collection,” *arXiv preprint arXiv:2212.04708*, 2022.
- [53] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6892–6903.
 - [54] T. Liu, Z. Liu, Z. Jiao, Y. Zhu, and S.-C. Zhu, “Synthesizing diverse and physically stable grasps with arbitrary hand structures using differentiable force closure estimator,” *RA-L*, 2021.
 - [55] 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, Z. Huang, R. Calandra, R. Chen, S. Luo, and H. Su, “Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai,” *arXiv*, 2024.
 - [56] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” in *CVPR*, 2017.
 - [57] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox, “Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system,” in *Robotics: Science and Systems (RSS)*, 2023.
 - [58] T. Yoshikawa, “Manipulability of robotic mechanisms,” *The international journal of Robotics Research*, 1985.
 - [59] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. Vander-Bilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi, “Objaverse: A universe of annotated 3d objects,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 13 142–13 153.