Dexterous Hand Co-Design for In-Hand Rotation

K. Fay[1], Sha Yi[1], Michael T. Tolley[1], Xiaolong Wang[1], and Hao Su[1]

[1]University of California San Diego

***Abstract*—In-hand manipulation for robots has recently been possible due to advances in reinforcement learning and ongoing development of new robotic hands. Both advancements have iteratively pushed the frontiers of manipulation with new controllers allow complex manipulators to be effective and hardware advancements allowing true dexterous manipulation. Co-design can be used to create a synergy of these efforts by co-optimizing hardware and control systems for complex tasks. This work aims to explore co-design by optimizing hardware of robotic hands for in-hand rotation. To explore an effective method for codesigned manipulators at scale, we propose a framework for hardware and controller optimization using MPPI and Dial-MPC. The result is a co-design method with fast sampling efficiency and hardware optimized manipulation.**

I. Introduction

Creating dexterous robots capable of complex human tasks remains a long-standing problem in robotics. Recent advances in reinforcement learning have sparked new hardware manipulators to be designed with new dexterous hands and sensors [20][1][29]. These hands in turn have allowed more complex manipulation with exploration into multi-modal manipulation [28], in-hand manipulation with PPO [16][17], and imitation learning for complex tasks. In order to build successful systems capable of sensitive and dexterous tasks, both frontiers of control and hardware design need to converge.

This workshop paper aims to show how co-design of hardware and control policies can be used to develop improved robotic hands for faster in-hand rotation. More specifically, we aim to achieve z-axis rotation with the fingertips with an optimized Allegro hand by altering physical hardware parameters. Co-Design methods are often limited due to the immense search space of both control and hardware parameters. To achieve fast sampling without the need for time intensive training, this work uses Model Predictive Path Integral (MPPI) for testing sampled designs to approximate final Dial-MPC performance. Designs are sampled using Bayesian Optimization across pre-defined search space of hardware parameters.

Specifically, a search over each independent physical parameter is completed with each sample and deployed with MPPI in order to identify which physical parameters have the highest impact. Then a joint search is completed across all identified parameters of impact with the controller deployed for a short time horizon. The best MPPI controller and sampled hardware parameters are saved. In a final stage, Dial-MPC is used from the found best joint parameters to achieve a final controller. Here MPPI is used as a proxy for Dial-MPC [24], a multi step diffusion inspired extension of MPPI.

Primarily, this work aims to contribute:

- A framework for efficient search of hardware parameters and sample efficient training free controllers
- An analysis of what parameters matter most for hardware design of robotic hands
- Showing that co-design can be another tool for achieving dexterous manipulation paired with annealing based

II. Background

Dexterous manipulation remains one of the most challenging areas of robotics. With a wide solution space, complex hardware [26][4] and the need for generalization across human tasks [11], manipulation remains open to exploration for both hardware and software. Over the last decade, reinforcement learning has expanded to manipulation, allowing for drastically finer control than before.

For example, Hora [16] achieved generalized in-hand rotation across objects through rapid motor adaptation and an Allegro hand. By encoding parameters of the object and mimizing the loss between state and privileged information alongside PPO, the authors achieve zero shot sim2real with emergent human-like finger gaits. In a subsequent work, tactile sensors were added, achieving greater levels of success [17]. Likewise, works such as Synesthesia [27] have shown the ability to build successful multi-modal systems for in-hand object rotation. AnyRotate similarly achieved in-hand rotation despite the changing axis of rotation of the hand's wrist orientation [25]. These works' promising results show that higher levels of dexterity can be achieved by both software and hardware improvements.

In addition to PPO, imitation learning and diffusion techniques have helped reduce controller search space by leveraging human data [8][22]. These approaches have shown success, especially in bimanual manipulation where the feasible control space is exceptionally large. Sampling-based methods such as MPC have also been deployed with some success. MPC paired with a VLM has been able to achieve in-hand rotation [10]. This offers potential for robot hands to find sample-efficent methods but rely on additional methods to achieve fine tuning and adaptation or rely on a known model.

Robot hands have existed for well over a century [30] with a vast majority of approaches aiming to mimic human hands [15]. The degree of true biomimicry is varied. Direct drive allows for direct motor feedback of each joint, as implemented with Allegro [1] and LEAP[20]. Alternatively, tendon-driven robot hands aim to more closely mimic human hands and displace motors into the forearm of the robot. Robonaut 2 Hand improved upon the original design by moving motors from the hand to the forearm achieving a wider range of human like grasps and additional sensors to the hand [5]. Other prominent examples of tendon-driven hands include Shadow Hand [2] and tendon driven mimicry by Xu and Todorov [23]. Due to the demand of new reinforcement and imitation learning methods, many new robotic hands have appeared aiming to reduce the gap between teloperation and retargetting such as RUKA [29]. A large gap between true human mimicking and robotic hardware still remains.

Beyond rigid components, soft robots have gained traction in robotics for manipulation. These approaches use soft bodies to leverage compliance often using pneumatics for control. Abondance et. Al [3] created a new soft gripper to replicate primitive rotation showing soft robots are capable of these tasks using pneumatics. Likewise, the Svelete hand [28] uses a mixed rigid and soft components to combine compliance with full tactile sensing. These works show a potential for soft bodied manipulators in the future. However, accurately modeling soft robots is exceptionally challenging often requiring the use of designated soft body simulators, model reduction, or system identification to work reliably.

With the advant of increased hardware design and complex control, co-design has become increasingly explored. In Cheney et. al, a soft body robot using generation of voxels with artificial soft muscle and [6]. The result was a soft body mesh capable of locomotion from expanding and contracting. Similar works have also explored origimi generated structures [14], digital hardware [7], and legged robots [21]. The range of co-design also varies drastically with some works relying on full innovation of designs while others supply starting designs or complete parameter optimization on known starting parameters. For example, DiffAqua [12] generated novel robot fish for fast swimming by finding a new design by sampling features from pre determined expert designs.

While co-design shows promise, few works are able to cross the sim2real gap due to the two fold difficulty of controller sim2real and accurate hardware modeling and manufacturing from simulation. Work deploying sim-to-real include soft robots which deployed sim2real with using model reduction [18]. Work in legged robots has also used optimizing modular link lengths to create complex motion gaits[9]. This transferred successfully on a tested hexapod and quadroped robots, Additionally, the authors were also able to apply this method to gripper arms in simulation. Accurate sim-to-real with co design remains limited with an open space for exploration.

Prior works have explored co-design for locomotion, but few have explored co-design for complex manipulators. Task driven co-design for manipulation has previously explored Bayesian Optimization but for simple manipulators [19]. Initial work in optimizing robot hands was also recently deployed through altering parameters of an existing hand [13]. This work is a very small exploration towards making hardware optimization more tractable for complex tasks.

III. METHOD

Algorithm 1 Hierarchical Design Optimization with MPPI-Guided Hardware Refinement

Require:
- Initial model $\mathcal{M}$, state x_0
- Parameter space $\mathcal{D}$
- Finger and joint sets $\mathcal{F}, \mathcal{J}$
- Horizon H, samples K, iterations N
- Improvement threshold δ
- MPPI parameters: temperature λ, control cost R, noise covariance Σ

1: $\mathbf{d}_0 \leftarrow$ BAYESIANMPPI(OBJECTIVE over $\mathcal{D}$)
2: $\mathcal{I}_\delta \leftarrow$ IDENTIFYACTIVEPARAMETERS($\mathbf{d}_0$, δ)
3: Define reduced subspace $\mathcal{D}_\delta$ with active parameters $\mathcal{I}_\delta$
4: $\mathbf{d}^* \leftarrow$ BAYESIANMPPI(OBJECTIVE over $\mathcal{D}_\delta$)
5: $\pi^* \leftarrow$ DIALMPC($\mathcal{M}, x_0, \mathbf{d}_{\text{refined}}, \mathcal{F}, H$)
6: **return** Refined parameters $\mathbf{d}_{\text{refined}}$ and final control policy π^*

This work aims to identify what hardware features matter for in-hand co-design and identify implications for what kinds of co-design will be most useful in the future. Additionally, this work aims to show how sample-based controllers and co-design may be a good pairing for initial co-design search. For simplicity, the authors chose to keep geometry and body size as a fixed parameter for this workshop, but present preliminary results with linkage lengths and scaling in a later section. The authors limit the search for this workshop paper to friction, inertia, material damping, and similar physical parameters. The authors are currently exploring geometry.

First a Bayesian Optimization sweep over defined parameters was completed in order to identify which physical parameters impacted reward. These parameters were all parameters pre-defined in the Allegro hand model except for contact softness. Physical and control parameters including friction of each body, damping coefficients, and range of motion were explored. For each iteration, MPPI was deployed with a horizon time of 40 over 30 samples. While manipulation is a long horizon task, this simplified run time and was sufficient to determine if reward changed significantly for initial search over each 50 sampled designs per parameter with the exception of control range. Control range was sampled at 50 per joint due to the wide range of solutions. All other parameters were originally sampled individually with the exception of inertia and mass which increased by a constant correlating the two.

As friction relies on both the friction of the hand's surface for each linkage alongside the object's friction a search over 10 objects with 3 different shapes (cube, cylinder, square) and randomized friction between 0.15 and 0.8 for each object. While the total reward per object shape and friction fluctated, higher friction of the fingertips resulted in higher dexterity across all objects.

From the rate of reward change, the parameters most significant in increasing reward were selected. For this work, a minimum reward improvement of 10% was required. The resulting parameters which showed improvement were impedance, linkage lengths, and friction. An additional sweep over these parameters using Bayesian Optimization with 100 samples across the design parameters, 10 random starts, and used expected improvement as the Bayesian Optimization acquisition function. The controller in this loop was set to a horizon of 60 and 20 samples for training efficiency. A final controller, Dial-MPC, is trained for a longer time horizon with the new found optimized parameters as a proxy for longer term efficiency.

Finally, the resulting solved parameters can be used as a new hand model with a more time expensive controller for finer manipulation. In this case, we deploy Dial-MPC with 200 steps, 2048 sampled trajectories for each control step, and MPPI as the update method. Dial-MPC is chosen here as it is training free, extends sample efficient MPPI, and showed initial promise with test-time generalization. In future work, the final controller and hardware pair will be developed and tested across objects.

Algorithm 2 Model Predictive Path Integral Control (MPPI)

```
function MPPI(𝓜, x_0, d, 𝓕, H, K, N, λ, R, Σ)
    Initialize nominal control sequence u ← 0_{H×m}
    for t = 1 to N do
        for k = 1 to K do
            Sample noise sequence ε_k ~ 𝓝(0, Σ)
            u_k ← u + ε_k
            S_k ← TrajectoryCost(𝓜, x_0, u_k, 𝓕, R)
        end for
        Compute weights: w_k ← exp(−S_k/λ) / Σ_{j=1}^{K} exp(−S_j/λ)
        Update nominal control: u ← u + Σ_{k=1}^{K} w_k · ε_k
    end for
    return u
end function
```

TABLE I
REWARD SCALING COEFFICIENTS

Component	Coefficient
Palm Centering	$w_{\text{center}} = 100.0$
Multi-Contact Bonus	$w_{\text{contact}} = 50.0$
Angular Velocity Tracking	$w_{\text{ang}} = 100.0$
Twist Reward	$w_{\text{twist}} = 50.0$
External Force Penalty	$w_{\text{force}} = 50.0$
Control Smoothness	$w_{\text{smooth}} = 10.0$
Finger Distance Penalty	$w_{\text{finger}} = 20.0$

For MPPI's cost function, several in-hand rotation rewards were tried including those from Hora [16]. However, the rewards from reinforcement learning to MPPI did not transfer well and resulted in unstable solution. A temporary and alternative reward was proposed below using palm center, finger contact, smoothness, and multi contact as key parameters. Here the "reward" is negated and minimized to meet MPPI's cost function requirement. MPPI had the tendance to find single contact solutions such as finger juggling even without co-design. This is better stabilized by Dial-MPC which helps solve this contact rich convergence issue with annealing. In future work, the authors will explore this approximation fully and may replace MPPI with a short horizon Dial-MPC depending on hardware parameter convergence and task generalization.

$$
\begin{aligned}
r = \; & - w_{\text{center}} \cdot \|\mathbf{x}_{\text{cube}} - \mathbf{x}_{\text{palm}}\|^2 \\
& - w_{\text{ang}} \cdot \|\boldsymbol{\omega}_{\text{cube}} - \boldsymbol{\omega}_{\text{target}}\| \\
& - w_{\text{force}} \cdot \|\mathbf{f}_{\text{ext}}\| \\
& - w_{\text{smooth}} \cdot \|\mathbf{u}_t - \mathbf{u}_{t-1}\| \\
& - w_{\text{finger}} \cdot \sum_i \max(0, d_i - \delta) \\
& + w_{\text{contact}} \cdot n_{\text{contact}} \\
& + w_{\text{twist}} \cdot (\boldsymbol{\omega}_{\text{cube}} \cdot \hat{\mathbf{z}})
\end{aligned}
\tag{1}
$$

Bayesian Optimization was primarily chosen due to the high cost of evaluation with most controllers. Additionally, manipulation is often noisy and can be treated as a black box. As initial parameters of impact were found individually, the search space was reduced helping improve run time. Additionally, using MPPI as an initial search allows us to determine what truncated range the search should be conducted over. A search range over all parameters was pre-defined based on realistic parameters for the real world.

Parameter	Search Range	Units	Definition
Damping	0 to 10.0	N·m·s/rad	Coefficient for damping DOF
Kp	0 to 200	N·m/rad	Joint pos. control gain
Kv	0 to 300	N·m·s/rad	Joint velocity (damping) control gain
Control Range	$[-\pi, \pi]$	rad	Joint range of motion
FrictionLoss	0 to 20	N·m	Joint torque threshold
Solref	0.1 to 2.0	—	Damping contact coefficient
Solimp	0.5 to 1.0	—	Impedance in contact modeling
Friction	0.1 to 0.85	—	Contact friction coefficient

TABLE II
PARAMETER RANGES AND DEFINITIONS OF CONTROL AND CONTACT PARAMETERS USED IN MUJOCO.

IV. ANALYSIS

From the initial parameter sweep, damping and friction had highest impact on the reward. Position control, frictionloss, and inertia were found be either negligible or already optimized for values of best fit. Across the ten objects tested, higher friction consistnetly improved rate of rotation regardless of object shape of the object's friction.

The resulting highest optimized hand had the parameters of 0.70 for friction, finger tip contact damping of .075, and fingertip linkages of 2.1 scale of the original length were found. Geometry will be further explored in the future, but initial results show long and skinny fingertips were more dexterous. MPPI was very noisy and prone to jitter, leading to a suboptimal final rotation controller. Simulating soft bodies with mujoco can lead to over simplified models due to the difficulty of accuaretly modeling soft body response to soft forces. At the lowest end of soft body parameters, it may be to difficult to deploy this without a specialized soft body simulator. Preliminary results showed impedance shifting towards both extremes depending on the sweep size and number of samples. This may mean both extremely rigid robots which are easier to learn and extremely soft robots which are more complaint are both good options for exploration.

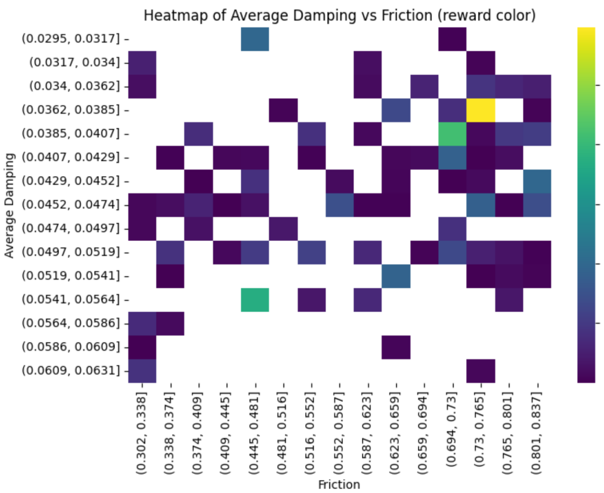

Fig. 1. Results for MPPI joint sweep for Damping and Friction parameters.

For joint control, the Allegro hand already has a wide solution space of possible grasps. Adding additional possible control space did not improve the hand. This is likely due to Allegro already having an optimized control space for in-hand manipulation. Additionally, friction can be easily optimized and deployed in the real world through different textures on robot hands such as finger tips of different material textures.

With a gpu enabled, co-design took over 45 minutes per policy and feature sample using a GTX 4070 and PPO. By contrast, MPPI averaged approximately 10 minutes using 15 samples and a horizon time of 60 without using a GPU. This allowed for a much more rigorous search over design parameters within a reasonable time frame.

One caviate with co-design is solutions may be solved in extremely unintuitive ways. By default, rotating a cube can be achieved by grasping the outsides of the cube and twisting each side with a finger gait. However, under co-design solutions can be found with high stiffness, friction, or softness to try and juggle the cube between rebound of each finger. Similar to reinforcement learning, co design seems to be highly reward sensitive. Further work in exploring stability remains open.

V. Conclusion

Similar parameters for the optimized Allegro hand showing the potential for co-design methods to find designs that work across similar tasks. Additionally, the resulting work is able to find a method which outperforms the baseline. However, Cross Entropy MPC was very unstable. For a final iteration some additional method such as VLMs or another policy should be deployed on the optimized hand. Additionally, friction was

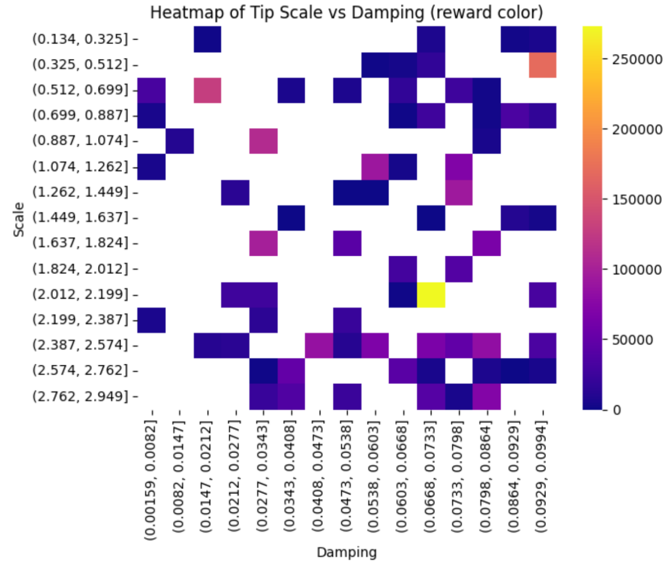

Fig. 2. Finger tip scaling compared to contact damping in joint MPPI Baysian search.

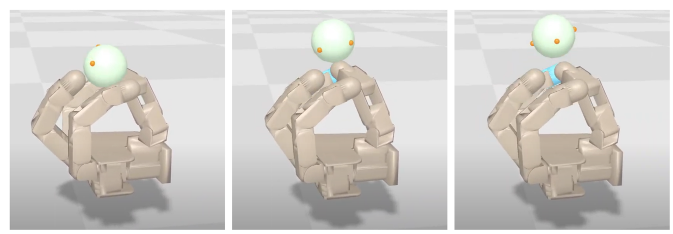

Fig. 3. Example sequence of non ideal in-hand juggling behavior often found by MPPI.

found to be the most important parameter for hand design alongside stiffness.

With a method for cheap sampling with with parameter optimization, more time intensive controllers can be applied after optimization. In future work, we plan to replace the final deployent by training the found hand design with PPO for robustness. This will allow reducing the search time while achieving fine control of reinforcement methods. Additionally, co-design is able to over optimize for specific tasks or controllers. Future work in weighting designs across tasks and exploring how to design for multi-objective functions will be tested.

We plan to deploy sim-to-real with best found manipulators. Optimized physical bodies for found parameters including stiffness, compliance, and texture will be used. For manufacturing, corresponding 3d printed materials and friction textures of fit will be used. Additionally, exploring design morphology through geometry will be explored across generalized objects.

References

[1] Allegro Hand | robot hand. URL https://www.allegrohand.com.

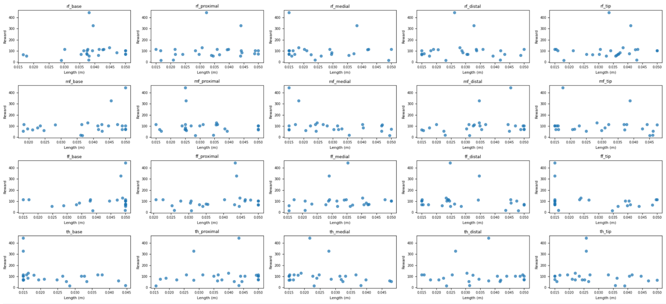

Fig. 4. Individual parameter search for lengths of each fingertip. The authors found the scaling did not matter but individual lengths did.

[2] Shadow Dexterous Hand Series - Research and Development Tool, September 2023. URL https://shadowrobot.com/dexterous-hand-series/. Running Time: 66.

[3] Sylvain Abondance, Clark B. Teeple, and Robert J. Wood. A Dexterous Soft Robotic Hand for Delicate In-Hand Manipulation. *IEEE Robotics and Automation Letters*, 5(4):5502–5509, October 2020. ISSN 2377-3766. doi: 10.1109/LRA.2020.3007411. URL https://ieeexplore.ieee.org/document/9134855.

[4] Aude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364(6446):eaat8414, June 2019. doi: 10.1126/science.aat8414. URL https://www.science.org/doi/full/10.1126/science.aat8414. Publisher: American Association for the Advancement of Science.

[5] L. B. Bridgwater, C. A. Ihrke, M. A. Diftler, M. E. Abdallah, N. A. Radford, J. M. Rogers, S. Yayathi, R. S. Askew, and D. M. Linn. The Robonaut 2 hand - designed to do work with tools. In *2012 IEEE International Conference on Robotics and Automation*, pages 3425–3430, May 2012. doi: 10.1109/ICRA.2012.6224772. URL https://ieeexplore.ieee.org/document/6224772. ISSN: 1050-4729.

[6] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. *SIGEVOlution*, 7(1):11–23, August 2014. doi: 10.1145/2661735.2661737. URL https://dl.acm.org/doi/10.1145/2661735.2661737.

[7] G. De Michell and R.K. Gupta. Hardware/software co-design. *Proceedings of the IEEE*, 85(3):349–365, March 1997. ISSN 1558-2256. doi: 10.1109/5.558708. URL https://ieeexplore.ieee.org/abstract/document/558708.

[8] Bin Fang, Shidong Jia, Di Guo, Muhua Xu, Shuhuan Wen, and Fuchun Sun. Survey of imitation learning for robotic manipulation. *International Journal of Intelligent Robotics and Applications*, 3(4):362–369, December 2019. ISSN 2366-598X. doi: 10.1007/s41315-019-00103-5. URL https://doi.org/10.1007/s41315-019-00103-5.

[9] Sehoon Ha, Stelian Coros, Alexander Alspach, Joohyung Kim, and Katsu Yamane. Joint optimization of robot design and motion parameters using the implicit function theorem. 07 2017. doi: 10.15607/RSS.2017.XIII.003.

[10] Adrian Hess, Alexander M. Kübler, Benedek Forrai, Mehmet Dogar, and Robert K. Katzschmann. Sampling-Based Model Predictive Control for Dexterous Manipulation on a Biomimetic Tendon-Driven Hand, November 2024. URL https://arxiv.org/abs/2411.06183v2.

[11] Oliver Kroemer, Scott Niekum, and George Konidaris. A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms. *Journal of Machine Learning Research*, 22(30):1–82, 2021. ISSN 1533-7928. URL http://jmlr.org/papers/v22/19-804.html.

[12] Pingchuan Ma, Tao Du, John Z Zhang, Kui Wu, Andrew Spielberg, Robert K Katzschmann, and Wojciech Matusik. Diffaqua: A differentiable computational design pipeline for soft underwater swimmers with shape interpolation. *ACM Transactions on Graphics (TOG)*, 40 (4):132, 2021.

[13] Pragna Mannam, Xingyu Liu, Ding Zhao, Jean Oh, and Nancy 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)*, pages 332–339, April 2024. doi: 10.1109/RoboSoft60065.2024.10521927. URL https://ieeexplore.ieee.org/document/10521927. ISSN: 2769-4534.

[14] Ankur Mehta, Joseph DelPreto, and Daniela Rus. Integrated Codesign of Printable Robots. *Journal of Mechanisms and Robotics*, 7(021015), May 2015. ISSN 1942-4302. doi: 10.1115/1.4029496. URL https://doi.org/10.1115/1.4029496.

[15] C. Piazza, G. Grioli, M. G. Catalano, and A. Bicchi. A Century of Robotic Hands. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(Volume 2, 2019):1–32, May 2019. ISSN 2573-5144. doi: 10.1146/annurev-control-060117-105003. URL https://www.annualreviews.org/content/journals/10.1146/annurev-control-060117-105003. Publisher: Annual Reviews.

[16] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-Hand Object Rotation via Rapid Motor Adaptation. In *Proceedings of The 6th Conference on Robot Learning*, pages 1722–1732. PMLR, March 2023. URL https://proceedings.mlr.press/v205/qi23a.html. ISSN: 2640-3498.

[17] Haozhi Qi, Brent Yi, Sudharshan Suresh, Mike Lambeta, Yi Ma, Roberto Calandra, and Jitendra Malik. General In-hand Object Rotation with Vision and Touch. In *Proceedings of The 7th Conference on Robot Learning*, pages 2549–2564. PMLR, December 2023. URL https://proceedings.mlr.press/v229/qi23a.html. ISSN: 2640-3498.

[18] Charles Schaff, Audrey Sedal, and Matthew R. Walter. Soft Robots Learn to Crawl: Jointly Optimizing Design and Control with Sim-to-Real Transfer. In *Robotics: Science and Systems XVIII*. Robotics: Science

and Systems Foundation, June 2022. doi: 10.15607/rss.2022.xviii.062. URL http://www.roboticsproceedings.org/rss18/p062.pdf.

[19] Raphael Schneider, Daniel Honerkamp, Tim Welschehold, and Abhinav Valada. Task-Driven Co-Design of Mobile Manipulators, December 2024. URL https://arxiv.org/abs/2412.16635v2.

[20] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, July 2023. doi: 10.15607/rss.2023.xix.089. URL http://www.roboticsproceedings.org/rss19/p089.pdf.

[21] Tsun-Hsuan Wang, Pingchuan Ma, Andrew Everett Spielberg, Zhou Xian, Hao Zhang, Joshua B. Tenenbaum, Daniela Rus, and Chuang Gan. SoftZoo: A Soft Robot Co-design Benchmark For Locomotion In Diverse Environments, March 2023. URL https://arxiv.org/abs/2303.09555v1.

[22] Fan Xie, Alexander Chowdhury, M. Clara De Paolis Kaluza, Linfeng Zhao, Lawson Wong, and Rose Yu. Deep Imitation Learning for Bimanual Robotic Manipulation. In *Advances in Neural Information Processing Systems*, volume 33, pages 2327–2337. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/18a010d2a9813e91907ce88cd9143fdf-Abstract.html.

[23] Zhe Xu and Emanuel Todorov. Design of a highly biomimetic anthropomorphic robotic hand towards artificial limb regeneration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3485–3492, May 2016. doi: 10.1109/ICRA.2016.7487528. URL https://ieeexplore.ieee.org/document/7487528.

[24] Haoru Xue, Chaoyi Pan, Zeji Yi, Guannan Qu, and Guanya Shi. Full-order sampling-based mpc for torque-level locomotion control via diffusion-style annealing, 2024. URL https://arxiv.org/abs/2409.15610.

[25] Max Yang, Chenghua Lu, Alex Church, Yijiong Lin, Christopher J. Ford, Haoran Li, Efi Psomopoulou, David A. W. Barton, and Nathan F. Lepora. AnyRotate: Gravity-Invariant In-Hand Object Rotation with Sim-to-Real Touch. In *Proceedings of The 8th Conference on Robot Learning*, pages 4727–4747. PMLR, January 2025. URL https://proceedings.mlr.press/v270/yang25c.html. ISSN: 2640-3498.

[26] Chunmiao Yu and Peng Wang. Dexterous Manipulation for Multi-Fingered Robotic Hands With Reinforcement Learning: A Review. *Frontiers in Neurorobotics*, 16, April 2022. ISSN 1662-5218. doi: 10.3389/fnbot.2022.861825. URL https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/fnbot.2022.861825/full. Publisher: Frontiers.

[27] Ying Yuan, Haichuan Che, Yuzhe Qin, Binghao Huang, Zhao-Heng Yin, Kang-Won Lee, Yi Wu, Soo-Chul Lim, and Xiaolong Wang. Robot Synesthesia: In-Hand Manipulation with Visuotactile Sensing. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6558–6565, May 2024. doi: 10.1109/ICRA57147.2024.10610532. URL https://ieeexplore.ieee.org/document/10610532.

[28] Jialiang Zhao and Edward H. Adelson. GelSight Svelte Hand: A Three-finger, Two-DoF, Tactile-rich, Low-cost Robot Hand for Dexterous Manipulation, September 2023. URL http://arxiv.org/abs/2309.10886. arXiv:2309.10886 [cs].

[29] Anya Zorin, Irmak Guzey, Billy Yan, Aadhithya Iyer, Lisa Kondrich, Nikhil X. Bhattasali, and Lerrel Pinto. RUKA: Rethinking the Design of Humanoid Hands with Learning, April 2025. URL https://arxiv.org/abs/2504.13165v1.

[30] Kevin J Zuo and Jaret L Olson. The evolution of functional hand replacement: From iron prostheses to hand transplantation. *Plastic Surgery*, 22(1):44–51, 2014. ISSN 2292-5503. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4128433/.