

EyeSight Hand: Design of a Fully-Actuated Dexterous Robot Hand with Integrated Vision-Based Tactile Sensors and Compliant Actuation

Branden Romero*, Hao-Shu Fang*, Pulkit Agrawal and Edward Adelson
Massachusetts Institute of Technology

brromero@mit.edu, fhs@mit.edu, pulkitag@mit.edu, adelson@csail.mit.edu

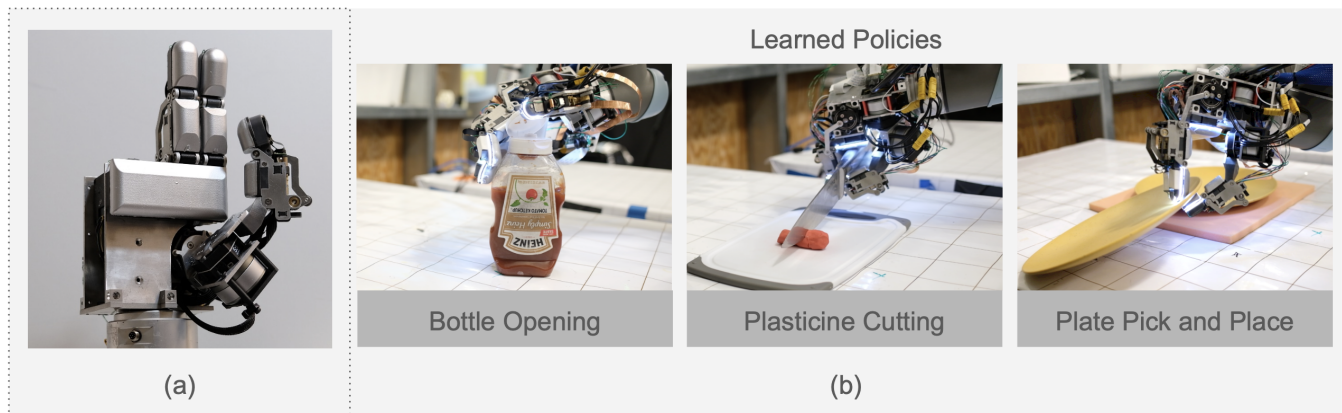


Fig. 1: (a) Picture of the EyeSight Hand. (b) EyeSight hand is used to perform three tasks autonomously using imitation learning: bottle opening, plasticine cutting, and plate pick and place.

Abstract—In this work, we introduce the EyeSight Hand, a 7 degrees of freedom (DoF) humanoid hand featuring integrated vision-based tactile sensors tailored for enhanced whole-hand manipulation. Additionally, we introduce an actuation scheme centered around quasi-direct drive actuation to achieve human-like strength and speed while ensuring robustness for large-scale data collection. We evaluate the EyeSight Hand on three challenging tasks: bottle opening, plasticine cutting, and plate pick and place, which require a blend of complex manipulation, tool use, and precise force application. Imitation learning models trained on these tasks, with a vision dropout strategy, showcase the benefits of tactile feedback in enhancing task success rates. Our results reveal that the integration of tactile sensing dramatically improves task performance, underscoring the critical role of tactile information in dexterous manipulation.

I. INTRODUCTION

The versatility of the human hand is a result of a unique blend of anatomical complexity that combines a complex kinematic structure, an array of muscles and tendons that provide compliance to manipulate objects with varying precision and delicacy, and a high density of tactile sensors in the skin to provide detailed feedback on object geometry and contact forces [18].

In various efforts to replicate the dexterity of the human hand with robotic hands, previous methods have captured only a subset of human hand functionality. For example, current fully actuated hands [30] are usually not compliant,

* Equal contribution. Toyota Research Institute provided funds to support this work.

characterized by low backdrivability and/or high reflected inertia due to high gear reductions. Soft hands [25, 49] that are compliant are often not fully actuated and difficult to control precisely. They also face challenges in incorporating dense tactile sensing. Existing hands equipped with dense tactile sensors [47, 39] typically have a limited number of degrees of freedom (DoF) and do not closely mimic human hand morphology. However, all these properties are crucial for robot learning: compliance makes the hardware robust to large-scale data collection, dense tactile sensing is important for precise manipulation, and a morphology similar to the human hand facilitates easy teleoperation and data collection.

In this work we introduce EyeSight Hand, a low-cost (less than \$2500), 7-DoF humanoid hand that is co-designed with actuation, tactile, and kinematics in mind. The key contributions of the EyeSight Hand are:

- We exploit the benefits of quasi-direct drive actuation and how to integrate it into a humanoid robotic hand, to provide compliant actuation while having human level finger strength and speed. All while being robust enough for large-scale data collection.
- With vision-based tactile sensors emerging as a promising approach to robotic touch, we introduce a new vision-based tactile sensor GelSim(ple) that is easily simulatable and can be adapted to different morphologies while maintaining the sensitivity of previous GelSight sensors. We then integrate 8 of these sensors into the hand to enable sensitive whole-hand manipulation.




Fig. 2: (a-b) Comparison of the dimension of a human hand with the dimensions of the EyeSight Hand. (c-d) Diagram illustrating the kinematics of the (c) index/middle finger and (d) thumb as well as the placements of the motors and transmissions. Axis of rotation is shown with green arrows.

- Given the above two design choices, we manage to make the hand close enough to the morphology of human hand, in terms of kinematics and dimensions, which enables simple and intuitive teleoperation with minimal kinematic retargeting effort.

To evaluate our hand, we propose three challenging tasks that require complex non-prehensile manipulation, forceful tool-use, extrinsic force sensing, and precise dexterous manipulation, including bottle opening, plasticine cutting, and plate pick and place. Our hand can be teleoperated to accomplish these tasks without much effort, and works robustly during hundreds of trials. Finally, we train imitation learning models on these tasks, where a vision dropout strategy is proposed for the model to better utilize the tactile sensing, and quantitatively show the benefits of our dense tactile sensing.

II. BACKGROUND AND RELATED WORK

A. Humanoid Robotic Hands

There is a plethora of research trying to replicate the versatility of human hands to robotic hands. One of the most notable attempts is the commercially available Shadow Hand, a highly dexterous but costly platform. Whose tendon-driven approach reduces the reliability of platform significantly. Hands whose motors directly actuate the joint have been developed like the commercially available Allegro Hand as well as research efforts like [30, 6]. However, they rely on small highly geared servos, that have a tendency to overheat under continuous load, and the high reflected inertia of the gearbox leads to high resistance to external forces which increases risk of breakage. Recently, [19] has shown potential of linkage-based designs particularly in terms of robustness to large forces, however the actuation is slow making it unsuitable for dynamic manipulation and potentially difficult to teleoperate. Although not specifically a humanoid hand, [41] illustrates the potential of quasi-direct drive actuators for robust and dynamic manipulation. We explore such potential for humanoid hand in this paper.

B. Hands with Vision Based Tactile Sensors

While hands like [19, 6] have tactile sensing, they are low resolution compared to vision-based tactile sensors. To enable dense tactile sensing, researchers have proposed incorporating vision-based tactile sensors, such as GelSight [44], into robotic hands. Due to hardware constraints, early attempts [38, 11, 29, 24] were typically flat or round and only mounted on the fingertips of hands. Recently, some works [32, 40, 22, 46] have explored integrating the GelSight sensor across the entire finger. Existing research has focused on parallel grippers with only 2 or 3 DoF. In this work, we propose a design of the GelSight sensor and integrate it with a self-designed 7 DoF hand.

C. Learning for Dexterous Hand

Learning manipulation with a dexterous hand is challenging due to the complexity of managing its many degrees of freedom and the need for precise sensing under severe occlusion. Early algorithms mainly conducted reinforcement learning with ground-truth [50, 28, 21, 23] or estimated state information [3, 1]. To deal with high-dimensional observation, teacher-student learning [9] or representation learning [16] are proposed, showcasing results of transferring to real-world scenarios in hand rotation [10]. Beyond visual observation, some recent work [26, 43, 45, 35] also explored incorporating tactile observation for in-hand rotation tasks. These methods are trained with reinforcement learning in simulation due to the demand for large-scale data. Few work have explored direct imitation learning in the real world since it is unclear how much data is needed to train a good policy, and the hardware issues like non-compliant also make it difficult to teleoperate and collect data. Recent attempts in this direction need additional data source including internet-scale video and augmented simulation demo to pretrain the policy [31, 37], or play data to learn a representation model [13]. In this work, we demonstrate that by incorporating rich tactile perception, we can achieve efficient imitation learning for dexterous manipulation on several challenging contact-rich tasks in the real world.

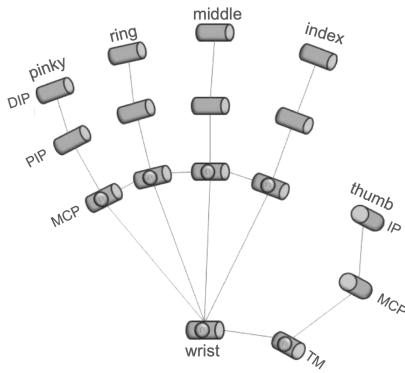


Fig. 3: Simplified kinematics of the human hand [8].

III. MECHANICAL DESIGN

Given the recent achievements in machine learning, the goal is to design a hand that is robust enough for large scale data collection, intuitive enough to perform teleoperation for imitation learning, and a platform that can potentially take advantage of data of humans performing task. In this section we explore the design decisions to achieve that goal.

A. Kinematics

Naturally, to achieve intuitive teleoperation and to best take advantage of human data, we would like to close the embodiment gap and try to closely mimic the kinematics of human hand, as illustrated in Fig. 2, as much as possible. However, in the design of the EyeSight Hand we make several compromises to fit in other features. Mainly, the hand differs from human kinematics (as shown in Fig. 3) in the metacarpophalangeal (MCP) joints and distal interphalangeal (DIP) joints of the index and middle finger, and the MCP joint of the thumb. Where, the DIP joints and the lateral axis of the MCP of the index and middle finger are missing, and the MCP joint of the thumb is missing.

B. Actuation

The actuation scheme is motivated by how compliance of the human hand. In particular, the observation that given an external force to the fingertip, the DIP and proximal interphalangeal (PIP) joints of the finger remain relatively static, while the flexor axis of the MCP joint back drives. Therefore, the hand takes use of two types of actuators, a 30:1 N20 Micro DC Motor driving a 20:1 worm gear reduction, and a TMotor GL40 KV210 brushless DC (BLDC) motor with a 5.23:1 planetary reduction. These actuators are placed in series for design simplicity, where for the middle and index finger the worm drive actuates the proximal interphalangeal PIP joint, while the BLDC motor actuates the flexor axis of the MCP joint. Meanwhile for the thumb the worm drive actuates the interphalangeal (IP) joint, while the BLDC motor actuates both axis of the trapeziometacarpal (TM) joint. These actuator either actuate the joints by directly driving them or via a four bar linkage, making our hand extremely robust compared to tendon driven ones.

The selection of these specific actuators was based on their durability. We choose the worm gear, due to its high reduction in a small factor, and its self-locking/ non-backdriving capability. The high reduction is particularly advantageous as it allows us to use a relatively large tooth size robust to high forces. Then due to self-locking, all of the forces get transmitted to the BLDC motor. Unlike the multistage spur gear servos used in many of dexterous hands like [30], the low reduction of the BLDC motor means less friction in the transmission and lower reflected inertia. This leads to less resistance in the actuator against external forces, resulting in a highly reactive backdrivable actuator with a reduced risk of breakage.

This actuation scheme allows for extremely nimble and strong fingers, with an estimated 19N continuous fingertip force, max fingertip force 57N, and a maximum speed of 420 RPM. This meets the force requirements for activities of daily living, which range between 45 and 68 N, as suggested by [5]. Beyond exerting high forces, high torque actuators have the benefit of being less likely to overheat from sustained use, as the continuous operating torque of these actuators greatly exceed the stall torque of the motors used in [30] and the commercially available Allegro Hand.

C. Electronics

The BLDC actuators are driven by a remotely place MJBOTS Moteus N1 Brushless Controller, and are equipped with AMS AS5047P 14-bit magnetic encoder to provide rotor position. Communication with the BLDC actuators happen over 5MBPS CAN-FD at 1kHz. The BLDC controller is equipped with current sensing, allowing for torque estimation from current [42]. The three micro DC motors are driven by TI DRV8876 controlled by a single Teensy 4.1 and communicate over serial.

IV. TACTILE SENSOR DESIGN

One of the main challenges in designing the EyeSight Hand was in fully integrating vision-based tactile sensing into the hand. The EyeSight Hand consists of seven tactile surfaces, each with varying morphologies, streaming a total of eight images, one for each finger segment, and two for the palm. So the question becomes, how do we come up with a design framework adaptable to different sensor morphologies, that maximizes data throughput, and also enables us to simulate the sensor for sim2real approaches. Therefore, in this work we introduce GelSim(ple) a variant of the GelSight Sensor.

A. Illumination

The initial goal of GelSight was to perform 3D reconstruction of the sensing surface via photometric stereo. This was achieved by illuminating the sensing surface with 3 different colored directional light sources spaced as evenly apart as possible around the sensor while minimizing their angle of elevation to the sensing surface [44]. While this was easy to implement for flat sensors, adapting this strategy to other geometries is difficult and results in various




Fig. 4: (a) Dimension and exploded view of the GelSim(ple) sensor at the tip of the thumb. (b) Example raw outputs of the GelSim(ple) and the approximate simulation of the corresponding deformation.

compromises [36, 47, 29]. However, findings from [38] suggested that a learning-based approaches can infer the 3D geometry of the sensing surface with single colored directional illumination. Furthermore, when it came to simulating GelSight Sensors, directional lighting meant that not only did you have to find a mapping from surface normal to RGB, you also had to accurately simulate shadows [33], which is costly. Given these findings, GelSim(ple) uses a new illumination strategy, which side illuminates the perimeter of a diffuser to provided non-directional overhead lighting to a semi-specular sensing surface as shown in Fig. 4(a), which still clearly gives geometric information about the object in contact, but also eliminates casted shadows (Fig. 4(b)). Furthermore, this allows us to increase data throughput by only transmitting a single channel image. To simulate the sensors, pixel-wise inference is performed via a multilayer perceptron (MLP) that takes as input surface normal and viewing direction augmented with positional encoding of a simulated deformation, and outputs the estimated RGB value [34].

B. Implementation Details

GelSim(ple) (Fig. 4(a)) uses an Arducam B0286 IMX219 Camera Module with a 220 degree fisheye lens. The camera

is press fit into a 3D printed finger segment body, or palm body, that fastens to the sensing module via M2 screws. The sensing module consist of a cut to length SuperLightingLed Narrowest 3mm COB LED White Light Strip formed into the desired geometry and adhered to a 3D printed blinder. The lighting module is then pressed onto a clear rigid 2mm acrylic backing with a spherical view port for the fisheye camera to observe the sensing surface with minimal optical distortion. The acrylic backing has a 3M 3635-70 diffuser on the surface opposite of the camera, and has several M2 threaded inserts to attach it to the camera holder. SILICONES, INC. XP-565 is molded onto the rigid backing and coated with Print-On® Clear Silicone Ink mixed with a 13 micron non-leafing silver dollar aluminum pigment. Each camera is connected remotely to its camera driver board via a Arducam B0439 Sensor Extension Cable. The hand has two Arducam B0396 CamArray Hats allowing 8 synchronized camera modules to be interfaced to two Raspberry Pi 4. A single channel image is transmitted from each Raspberry Pi via ZeroMQ at 60Hz at 640x480 to the host computer.

V. EXPERIMENT SETUP

To validate the performance of the EyeSight Hand, we perform imitation learning on three tasks: cutting plasticine, picking and placing stacked plates, and opening bottles. We hope this will demonstrate the hand’s capabilities and the importance of tactile sensing in performing a diverse set of tasks, which require complex non-prehensile manipulation, forceful tool use, extrinsic force sensing, and precise dexterous manipulation. In this section, we describe the system used to collect data, provide a breakdown of each task, and discuss our imitation learning approach.

A. Teleoperation

Teleoperation is performed using the 3D Guidance trakStar tracking system. The operator wears a glove equipped with four magnetic sensor probes, three fastened near the finger nails of the thumb, index, and middle, while the other placed on the back of the palm. Each of these sensors provide a 6-DoF pose. To estimate the commanded joint angles, we perform kinematic retargeting via Task Space Vector (TSV) optimization [27, 14, 4]. A human hand model is fully determined by the four sensor poses, e.g., $\{H_1, H_2, \dots, H_4\}$. Accordingly, the URDF of the robotic hand specifies a set frames that are attached to the finger tips and the palm, e.g., $\{R_1, R_2, \dots, R_4\}$ which is parameterized by the joint angles of the robot $Q(t)$ at time t . For each set of coordinate frame we create a set of task space vectors, where each task space vector points from the origin of one coordinate system to another denoted as $r_{ij}(Q(t))$ and h_{ij} for the robot hand and human hand respectively in which $i = 1, 2, \dots, 4$, $j = 1, 2, \dots, 4$, and $i \neq j$. Then perform the following optimization routine:

$$\min_{Q(t)} \sum_{i \neq j} \|r_{ij}(Q(t)) - h_{ij}\|^2 + \alpha \|Q(t) - Q(t')\|^2, \quad (1)$$

in which t' is the previous time step and α is a weight to address the significance of the change of the robotic hand pose.

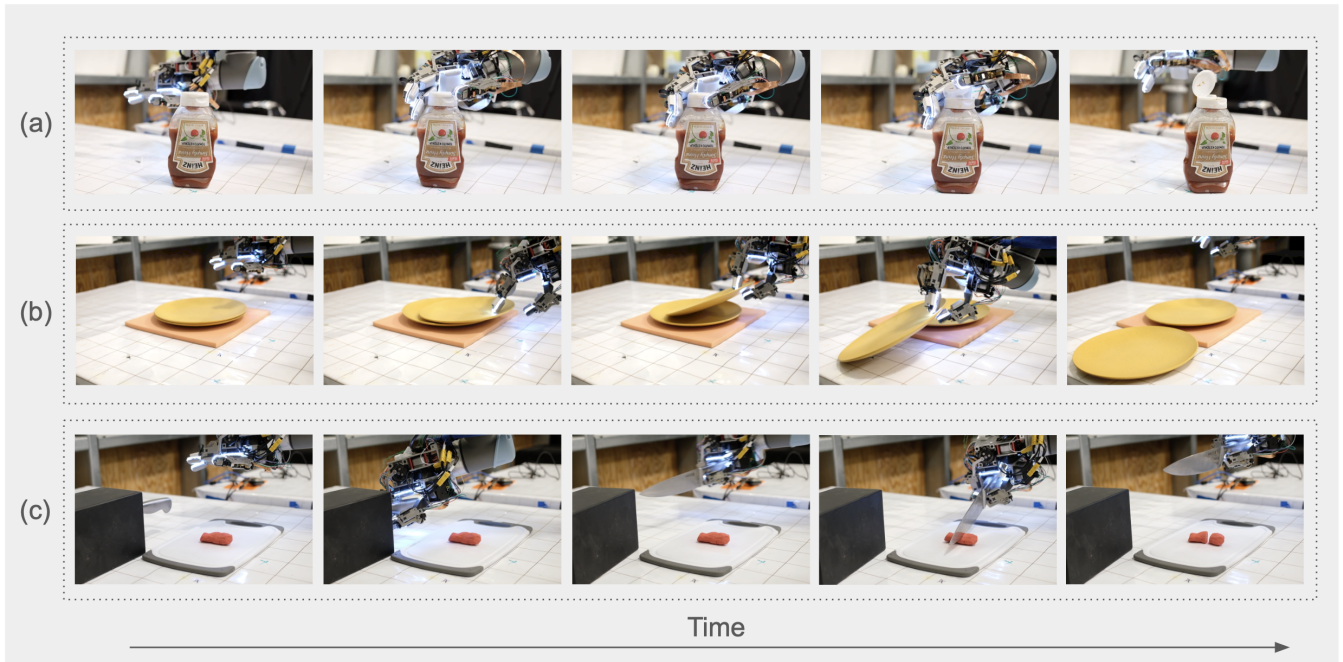


Fig. 5: Illustrations of the evaluation task. (a) The robot performs lid opening, where the robot must first approach the bottle, then constrain it with the index, middle finger, and palm, and then use the thumb to open the bottle. (b) The robot must approach a set of stack places, and then slide one over the other until it is able to grasp the plate. It then picks the plate and places next to the other plate. (c) The robot must remove a knife from a box, and use it to cut a plasticine roll fully.

In terms of implementation, Eq. 1 is solved using Sequential Least-Squares Quadratic Programming (SLSQP) [20, 17]. The optimization routine was written in JAX [7] and makes use of the Automatic Differentiation for rigid-body-dynamics AlgorithmMs library [2], to feed the forward and backwards solution of the objective to SLSQP in a timely manner. Finally, the result of the retargeting is used to directly command the joint angles of the robot hand. The teleoperation system operates at 125hz.

B. Task Specification

For each task we collect proprioceptive data from the EyeSight Hand and the UR5 robot, images from a global camera, images from a fisheye camera placed at the wrist of the hand, and images from the tactile sensors at 30Hz. For each task we collect 100 demonstrations.

1) *Bottle Opening*: For this task (Fig. 5(a)) a ketchup bottle is randomly placed within an area with slight perturbations to its initial orientation. The bottle is not fixed to the table so the robot must first approach the bottle and apply an appropriate constraint on the bottle so that it will not shift during manipulation, by applying downward force on the bottle’s body. The thumb tip then comes in contact with the lip of the bottle opening and forces it open. The robot must then fully swing the lid open. If the lid is not fully open it does not count as a success.

2) *Plate Pick and Place*: For this task (Fig. 5(b)) two stacked plates are placed on top of a hard foam block for slight friction. At the beginning of each episode, the robot must approach the stacked plates and make contact with the

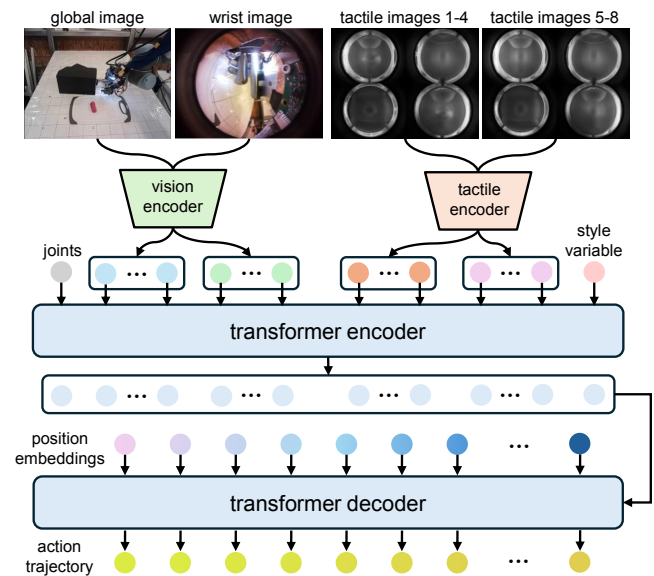


Fig. 6: Illustration of our input to the policy along with the architecture of the policy network.

top plate. The robot then gently slides the top plate over the bottom plate until it is in a graspable position. Too much force applied to the top plate will result in both plates sliding. The top plate is then lifted and gently placed next to the other plate. Initial positions of the stacked plates and foam block are slightly perturbed.

3) *Plasticine Cutting*: For this task (Fig. 5(c)) The robot must approach a knife stored in a box and pull it out of

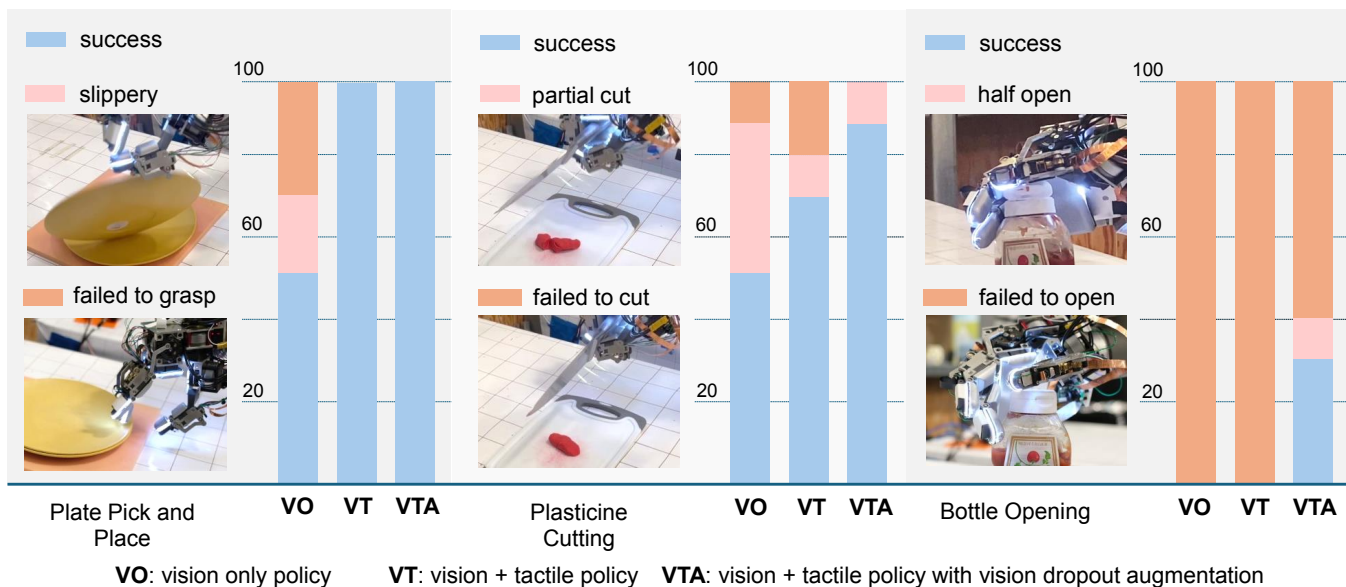


Fig. 7: Success rates of the various variants of our trained policies across different tasks, along with a detailed breakdown of the failure cases.

the box. The robot then goes and cuts a plasticine roll. The plasticine roll must be cut fully in half for the task to be considered successful. The knife box and cutting board are not fastened to the table, so excessive forces will cause them to slide. The initial position of the knife box is perturbed as well as the shape and position of the plasticine roll.

C. Imitation Learning

To learn dexterous manipulation, we adopt Action Chunking with Transformers (ACT) [48]. ACT is a generative model structured as a conditional variational autoencoder (CVAE), where the CVAE encoder learns a hidden style variable mapped to a diagonal Gaussian, and the CVAE decoder learns to predict the future action trajectory given the current observations and the style variable. This design accommodates the multi-modal nature of human-teleoperated action trajectories. For more details, we refer readers to [48]. In Fig. 6, we illustrate our input format and the architecture of the CVAE decoder, which serves as the policy network. The images from both the global camera and wrist camera are processed through a shared vision encoder to obtain their visual features. For the eight tactile images, we divide them into two subsets and concatenate every four images into a 2×2 super-image. The resulting two super-images are then processed through a shared tactile encoder. Both visual and tactile features, concatenated with the robot joint positions and the style variable, are fed through a transformer network to predict the future action trajectory in joint space.

We train three variants of ACT. The first, referred to as vision-only ACT, takes only visual images from the global and wrist cameras as input, without employing the tactile encoder branch during both training and testing phases. The second, referred to as vision-tactile ACT, incorporates both visual and tactile images. However, during our experiments,

we observed that the vision-tactile ACT sometimes ignored the tactile signal. Therefore, we propose a third version of ACT. In this version, during training, we introduce a random dropout for the input to the vision encoder—setting the entire images from the global and wrist cameras to zeros with a 30% chance. During testing, the input images are processed without dropout.

For our implementation, we resize all visual images and tactile super-images to 320×480 and adopt ResNet-18 [15] for both the vision and tactile encoders. The joint positions vector is 13-dimensional, consisting of 6 degrees of freedom (DoF) for the UR5 arm and 7 DoF for the hand. We predict the action trajectory for the next 20 timesteps, resulting in a 20×13 tensor. The prediction frequency is set to 15Hz. We train our policy on an Nvidia A6000 GPU with a batch size of 48. The learning rate is set to $1e-5$, and all policies are trained for 500 epochs. Following [12], we set the KL weight, hidden dimension, and feedforward dimension of ACT to 10, 1024, and 6400, respectively. During testing, our policy operates at 15Hz.

VI. EXPERIMENTAL RESULTS

We compared the results of policies trained with different strategies, rolling them out on the real robot across the three tasks introduced in V-B. Each experiment was conducted over 10 trials. In Fig. 7, we show the success rates of different policies on various tasks and a breakdown of the failure cases.

For the plate pick and place task, the vision-only policy achieved a mere 50% success rate. Incorporating tactile sensing increased the success rate to 100%. Without tactile sensing, there were two main failure modes: The first occurred when the robot attempted to slide the first plate but either failed to make contact or applied insufficient force,

leading to repeated unsuccessful sliding attempts without moving the plate. The second failure mode involved the robot grasping the plate without applying enough force, resulting in an unstable grasp and causing the plate to slip from its hand during movement.

For the plasticine cutting task, the vision-only policy also achieved a mere 50% success rate. Adding tactile sensing in a conventional manner improved the result to 70%. After adopting our vision dropout training strategy, the policy could achieve a success rate of 90%, perfectly grasping the knife and cutting the plasticine in half. Analyzing the failure cases, we found that without tactile sensing, the policy struggled to fully cut through the plasticine, as the visual difference between a partial and full cut is subtle in images. Meanwhile, in successful cases, the robot could apply excessive force, causing the cutting board to move with the knife. Incorporating tactile sensing significantly reduced the occurrences of partial cuts and prevented excessive force application during cutting, thus not moving the cutting board. The vision dropout training strategy further enhanced the policy's reliance on tactile sensing, improving the robustness of its execution.

The bottle opening task was particularly challenging. Both the vision-only policy and the vanilla vision-tactile policy failed in all trials to open the bottle, either attempting to open the lid without making contact or applying force in the wrong direction, causing the thumb fingertip to slip away from the lid. When adopting the vision dropout training strategy, the policy achieved a 30% success rate. The reason is that vision provides limited cues for the lid opening process, and the dropout strategy encourages the network to rely more on tactile sensing when predicting actions.

VII. CONCLUSION

In this work, we present the Eyesight Hand, a fully-actuated 7-DoF hand equipped with high-resolution tactile sensing and compliant actuation. Our design renders the hand suitable for the robot learning community due to its rich sensing capabilities, robustness to collisions, and morphology closely resembling that of a human hand. We develop an imitation learning baseline and demonstrate that the rich tactile sensing from our hand significantly enhances performance on several challenging real-world tasks. In the future, we aim to extend our hand with additional fingers and explore algorithms to better utilize the rich tactile sensing.

VIII. ACKNOWLEDGEMENT

The research was sponsored by the Army Research Office and was accomplished under ARO MURI Grant Number W911NF-23-1-0277. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government.

REFERENCES

- [1] Ilge Akkaya et al. "Solving rubik's cube with a robot hand". In: *arXiv preprint arXiv:1910.07113* (2019).
- [2] Automatic Differentiation for rigid-body-dynamics Algorithms. URL: <https://github.com/ami-iit/ADAM>.
- [3] OpenAI: Marcin Andrychowicz et al. "Learning dexterous in-hand manipulation". In: *The International Journal of Robotics Research* 39.1 (2020), pp. 3–20.
- [4] Dafni Antotsiou, Guillermo Garcia-Hernando, and Tae-Kyun Kim. "Task-oriented hand motion retargeting for dexterous manipulation imitation". In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. 2018, pp. 0–0.
- [5] Joseph T Belter, Jacob L Segil, and BS SM. "Mechanical design and performance specifications of anthropomorphic prosthetic hands: a review". In: *Journal of rehabilitation research and development* 50.5 (2013), p. 599.
- [6] Raunaq Bhirangi et al. *All the Feels: A dexterous hand with large-area tactile sensing*. 2024. arXiv: 2210.15658 [cs.RO].
- [7] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018. URL: <http://github.com/google/jax>.
- [8] Ilaria Cerulo et al. "Teleoperation of the SCHUNK S5FH under-actuated anthropomorphic hand using human hand motion tracking". In: *Robotics and Autonomous Systems* 89 (2017), pp. 75–84. ISSN: 0921-8890. DOI: <https://doi.org/10.1016/j.robot.2016.12.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0921889016302251>.
- [9] Tao Chen, Jie Xu, and Pulkit Agrawal. "A system for general in-hand object re-orientation". In: *Conference on Robot Learning*. PMLR. 2022, pp. 297–307.
- [10] Tao Chen et al. "Visual dexterity: In-hand reorientation of novel and complex object shapes". In: *Science Robotics* 8.84 (2023), eadc9244.
- [11] Elliott Donlon et al. "GelSlim: A High-Resolution, Compact, Robust, and Calibrated Tactile-sensing Finger". In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018, pp. 1927–1934.
- [12] Hongjie Fang et al. "Low-cost exoskeletons for learning whole-arm manipulation in the wild". In: *arXiv preprint arXiv:2309.14975* (2023).
- [13] Irmak Guzey et al. "Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play". In: *arXiv preprint arXiv:2303.12076* (2023).
- [14] Ankur Handa et al. "Dexpilot: Vision-based teleoperation of dexterous robotic hand-arm system". In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 9164–9170.
- [15] Kaiming He et al. *Deep Residual Learning for Image Recognition*. 2015. arXiv: 1512.03385 [cs.CV].
- [16] Wenlong Huang et al. "Generalization in dexterous manipulation via geometry-aware multi-task learning". In: *arXiv preprint arXiv:2111.03062* (2021).
- [17] Steven G. Johnson. *The NLOpt nonlinear-optimization package*. <https://github.com/stevengj/nlopt>. 2007.
- [18] Lynette A Jones and Susan J Lederman. *Human hand function*. Oxford university press, 2006.
- [19] Uikyum Kim et al. "Integrated linkage-driven dexterous anthropomorphic robotic hand". In: *Nature communications* 12.1 (2021), pp. 1–13.
- [20] Dieter Kraft. "Algorithm 733: TOMP—Fortran modules for optimal control calculations". In: *ACM Transactions on Mathematical Software* 20 (1994), pp. 262–281. DOI: 10.1145/192115.192124.

- [21] Vikash Kumar, Emanuel Todorov, and Sergey Levine. “Optimal control with learned local models: Application to dexterous manipulation”. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2016, pp. 378–383.
- [22] Sandra Q Liu and Edward H Adelson. “Gelsight fin ray: Incorporating tactile sensing into a soft compliant robotic gripper”. In: *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. IEEE. 2022, pp. 925–931.
- [23] Anusha Nagabandi et al. “Deep dynamics models for learning dexterous manipulation”. In: *Conference on Robot Learning*. PMLR. 2020, pp. 1101–1112.
- [24] Akhil Padmanabha et al. “Omnitact: A multi-directional high-resolution touch sensor”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 618–624.
- [25] Steffen Puhlmann, Jason Harris, and Oliver Brock. “RBO Hand 3: A Platform for Soft Dexterous Manipulation”. In: *IEEE Transactions on Robotics* 38.6 (Dec. 2022), pp. 3434–3449. ISSN: 1941-0468. DOI: [10.1109/tro.2022.3156806](https://doi.org/10.1109/tro.2022.3156806). URL: <http://dx.doi.org/10.1109/TRO.2022.3156806>.
- [26] Haozhi Qi et al. “General in-hand object rotation with vision and touch”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 2549–2564.
- [27] Yuzhe Qin et al. “Dexmv: Imitation learning for dexterous manipulation from human videos”. In: *European Conference on Computer Vision*. Springer. 2022, pp. 570–587.
- [28] Aravind Rajeswaran et al. “Learning complex dexterous manipulation with deep reinforcement learning and demonstrations”. In: *arXiv preprint arXiv:1709.10087* (2017).
- [29] Branden Romero, Filipe Veiga, and Edward Adelson. “Soft, Round, High Resolution Tactile Fingertip Sensors for Dexterous Robotic Manipulation”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020, pp. 4796–4802. DOI: [10.1109/ICRA40945.2020.9196909](https://doi.org/10.1109/ICRA40945.2020.9196909).
- [30] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. *LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning*. 2023. arXiv: [2309.06440](https://arxiv.org/abs/2309.06440) [cs.RO].
- [31] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. “Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning”. In: *arXiv preprint arXiv:2309.06440* (2023).
- [32] Yu She et al. “Exoskeleton-covered soft finger with vision-based proprioception and tactile sensing”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 10075–10081.
- [33] Zilin Si and Wenzhen Yuan. “Taxim: An example-based simulation model for gelsight tactile sensors”. In: *IEEE Robotics and Automation Letters* 7.2 (2022), pp. 2361–2368.
- [34] Zilin Si et al. *DIFFTACTILE: A Physics-based Differentiable Tactile Simulator for Contact-rich Robotic Manipulation*. 2024. arXiv: [2403.08716](https://arxiv.org/abs/2403.08716) [cs.RO].
- [35] Sudharshan Suresh et al. “Neural feels with neural fields: Visuo-tactile perception for in-hand manipulation”. In: *arXiv preprint arXiv:2312.13469* (2023).
- [36] Megha H. Tippur and Edward H. Adelson. “GelSight360: An Omnidirectional Camera-Based Tactile Sensor for Dexterous Robotic Manipulation”. In: *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. 2023, pp. 1–8. DOI: [10.1109/RoboSoft55895.2023.10122097](https://doi.org/10.1109/RoboSoft55895.2023.10122097).
- [37] Jun Wang et al. “CyberDemo: Augmenting Simulated Human Demonstration for Real-World Dexterous Manipulation”. In: *arXiv preprint arXiv:2402.14795* (2024).
- [38] Shaoxiong Wang et al. “GelSight Wedge: Measuring High-Resolution 3D Contact Geometry with a Compact Robot Finger”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 2021, pp. 6468–6475. DOI: [10.1109/ICRA48506.2021.9560783](https://doi.org/10.1109/ICRA48506.2021.9560783).
- [39] Achu Wilson et al. “Design of a Fully Actuated Robotic Hand With Multiple Gelsight Tactile Sensors”. In: *CoRR abs/2002.02474* (2020). arXiv: [2002.02474](https://arxiv.org/abs/2002.02474). URL: <https://arxiv.org/abs/2002.02474>.
- [40] Achu Wilson et al. “Design of a fully actuated robotic hand with multiple gelsight tactile sensors”. In: *arXiv preprint arXiv:2002.02474* (2020).
- [41] Manuel Wüthrich et al. “TriFinger: An Open-Source Robot for Learning Dexterity”. In: *CoRR abs/2008.03596* (2020). arXiv: [2008.03596](https://arxiv.org/abs/2008.03596). URL: <https://arxiv.org/abs/2008.03596>.
- [42] Hao Xiang et al. “Torque Estimation Base on Quasi-Direct Drive Actuators”. In: *2022 IEEE 5th International Electrical and Energy Conference (CIEEC)*. 2022, pp. 1637–1642. DOI: [10.1109/CIEEC54735.2022.9846110](https://doi.org/10.1109/CIEEC54735.2022.9846110).
- [43] Zhao-Heng Yin et al. “Rotating without seeing: Towards in-hand dexterity through touch”. In: *arXiv preprint arXiv:2303.10880* (2023).
- [44] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. “Gel-sight: High-resolution robot tactile sensors for estimating geometry and force”. In: *Sensors* 17.12 (2017), p. 2762.
- [45] Ying Yuan et al. “Robot synesthesia: In-hand manipulation with visuotactile sensing”. In: *arXiv preprint arXiv:2312.01853* (2023).
- [46] Jialiang Zhao and Edward H Adelson. “GelSight Svelte: A Human Finger-shaped Single-camera Tactile Robot Finger with Large Sensing Coverage and Proprioceptive Sensing”. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2023, pp. 8979–8984.
- [47] Jialiang Zhao and Edward H. Adelson. “GelSight Svelte: A Human Finger-Shaped Single-Camera Tactile Robot Finger with Large Sensing Coverage and Proprioceptive Sensing”. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2023, pp. 8979–8984. DOI: [10.1109/IROS55552.2023.10341646](https://doi.org/10.1109/IROS55552.2023.10341646).
- [48] Tony Z Zhao et al. “Learning fine-grained bimanual manipulation with low-cost hardware”. In: *arXiv preprint arXiv:2304.13705* (2023).
- [49] Jianshu Zhou et al. “A soft-robotic approach to anthropomorphic robotic hand dexterity”. In: *Ieee Access* 7 (2019), pp. 101483–101495.
- [50] Henry Zhu et al. “Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost”. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 3651–3657.