

AcTEExplore: Active Tactile Exploration on Unknown Objects

Amir-Hossein Shahidzadeh, Seong Jong Yoo, Pavan Mantripragada,
Chahat Deep Singh, Cornelia Fermüller, Yiannis Aloimonos

Abstract—Tactile exploration plays a crucial role in understanding object structures for fundamental robotics tasks such as grasping and manipulation. However, efficiently exploring such objects using tactile sensors is challenging, primarily due to the large-scale unknown environments and limited sensing coverage of these sensors. To this end, we present AcTEExplore, an active tactile exploration method driven by reinforcement learning for object reconstruction at scales that automatically explores the object surfaces in a limited number of steps. Through sufficient exploration, our algorithm incrementally collects tactile data and reconstructs 3D shapes of the objects as well, which can serve as a representation for higher-level downstream tasks. Our method achieves an average of 95.97% IoU coverage on unseen YCB objects while just being trained on primitive shapes.

I. INTRODUCTION

Human perception of the environment is a multifaceted process that involves multi-sensor modalities, including vision, audition, haptic, and *proprioception*. While deep learning has made significant progress in visual perception, conventional vision-only models have limitations compared to human perception. Humans excel at perceiving objects in challenging environments, utilizing their multi-sensor inputs [44], [12] such as the eyes for visual properties and the skin for tactile sensing which is essential to characterize physical properties such as texture, stiffness, temperature, and contour [62], [5], [28]. Thus, vision and tactile sensation have distinct roles in scene perception, each with unique requirements. Vision relies on direct line-of-sight unobstructed views, whereas tactile sensation only necessitates physical contact, enabling perception in challenging scenarios like occluded or dark environments when vision is limited. This distinction underscores the value of tactile sensing in scene perception, motivating the development of AcTEExplore. Our goal is to maximize contact with the object’s surface during exploration, thereby fully utilizing the benefits of tactile sensation that aligns with the capabilities of humans and other living beings to bridge the gap between machine and human perception.

The human skin, our largest organ, allows us to perceive contact with the external world. This has prompted behavioral studies [27], [6], [41] which investigate human manipulation skills, where the significance of tactile sensing becomes evident [66]. Similarly, the successful automation of robotic manipulation tasks heavily relies on their perceptual capabilities. Consequently, a multitude of tactile sensors

All authors are associated with the Perception and Robotics Group at the University of Maryland, College Park. The support by Brin Family Foundation, the Northrop Grumman Mission Systems University Research Program, ONR under grant award N00014-17-1-2622 and National Science Foundation under grant BCS 1824198 are gratefully acknowledged.

Supplementary material is available at <http://prg.cs.umd.edu/AcTEExplore>.




Fig. 1: **Reconstruction of a hammer.** (a) showcases the trajectory of the tactile sensor in 3D space. (b) depicts the intermediate tactile readings on the hammer’s surface, with the color gradient representing the passage of time. Following thorough tactile exploration, we achieve a complete object reconstruction (c), highlighting the effectiveness of our active strategy in exploring the entire reachable surface.

have been developed for robotic applications, encompassing optical-based sensors like DIGIT [32], GelSight [65] and Soft-bubble [1], which demonstrate remarkable proficiency in discerning skin deformation [59]. Conversely, bio-inspired electrode-based sensors such as SynTouch BioTac [60] necessitate extensive post-processing and finite-element modeling to accurately represent skin deformation, as studied in [46], [45].

Tactile sensors output a detailed local perspective of objects, giving them a unique role in surface exploration [58], [9], [52]. However, predicting future actions (moves) based on a single touch to explore an entire object’s shape is challenging. It involves two main issues: the agent getting stuck in cyclical movements or revisiting explored areas and the difficulty in inferring which action will maintain touch to avoid exploring empty space while consistently exploring the object. Balancing exploration and maintaining contact requires an active policy considering the tactile reading history and sensor trajectory. Our active policy facilitates the sensor to avoid *Non-Exploratory Scenarios*, enabling us to explore the objects within a limited number of actions. In **AcTEExplore**, we address the aforementioned challenges by formulating it as a Partially Observable Markov Decision Process (POMDP) where the policy is computed from only

a recent trajectory rather than the full interaction history.

To this end, we propose **AcTExplore**, an active method for tactile exploration that utilizes deep reinforcement learning, and the **core contributions** are given as follows: **(a) Exploring Object’s Surface** with minimal actions without limiting the approach to a specific distribution of objects. We achieve this by training the agent to learn dexterous movements from fundamental actions on primitive shapes (cube, sphere, etc.). Remarkably, as demonstrated in the experiments (Sec. IV), the learned behavior extends to unseen objects. **(b) Introducing Temporal Tactile Sensing** in the state representation (Sec. III-B) to enable Short-Term Memory (STM) on taxels (tactile receptors), inspired by various neurological and behavioral studies [33], [29]. **(c) Proposing a curiosity-driven Active Exploration Algorithm** for 3D reconstruction at scale that can be integrated into various high-level tasks in the future, such as Grasp Pose Refinement [18], Scene Perception [35], [56].

II. RELATED WORKS

Active SLAM involves traversing an unknown environment while simultaneously localizing and constructing a map [14], [8], [13], [15], [34]. In traditional SLAM, revisiting the marked areas is beneficial for correcting estimated localization errors [39], [51]. However, we aim to achieve an efficient exploration pipeline that minimizes revisits. Active SLAM is generally formulated as a Partially Observable Markov Decision Process (POMDP) [51], with various reward function formulations, including curiosity [8], coverage [16], or entropy-based [7]. Unlike the conventional active SLAM setting, which operates in unknown 2D spaces, our work focuses on exploring a 3D workspace with limited sensing space, leading to ambiguity [4].

Deep reinforcement learning in exploration. Recent advances in computing power and physics-based simulators have boosted research in virtual navigation and exploration. For instance, [43] trained an Asynchronous Advantage Actor-Critic (A3C) agent in a 3D maze, incorporating long short-term memory (LSTM) to provide the memory capability [54], [24]. Furthermore, [50] tackled the robot exploration problem using the D-optimality criterion as an intrinsic reward, significantly accelerating the training process. In our problem, we incorporate an exploration bonus as a reward function. This approach incentivizes the agent to explore undiscovered state and action pairs, leading to more sample efficient algorithms [3], [26].

Tactile applications. Tactile information plays a crucial role in human perception, encompassing tasks from object manipulation to emotional expressions. As a result, tactile sensors have been employed in various applications [36], [59], [48], [23], [22]. Especially when robots manipulate deformable objects, tactile sensors provide meaningful information that enhances system robustness alongside vision sensors [49]. Furthermore, tactile information has been used to estimate the pose of the objects [10], [4], [57], [67] or the relative pose of the gripper for object handling [31]. Similar to our work, tactile sensors have been employed to identify or reconstruct unknown objects [40], [37], [25], [47]. For instance, [63] designed a tactile object classification pipeline

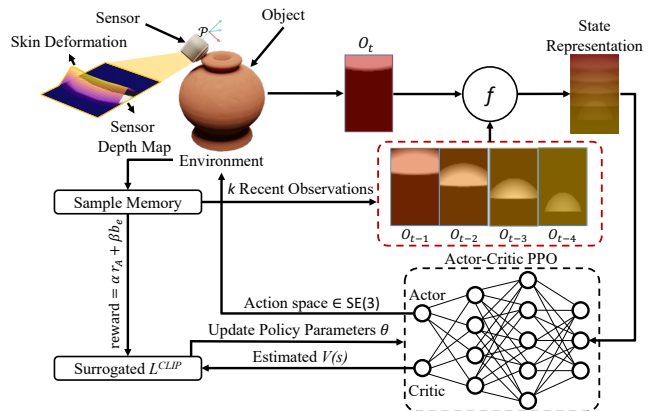


Fig. 2: **Overview.** This figure illustrates the key steps and components of AcTExplore in a scenario where the sensor moves upward along the jar’s edge. We employed Temporal Tactile Averaging for state representation f (Sec. III-B) to encode consecutive observations, enabling the perception of movement on the sensor vital for learning dexterous actions. We also incorporate an Upper Confidence Bound (UCB) exploration as a bonus to encourage effective exploration.

that actively collects tactile information while exploring the object. The closest previous studies that address the 3D reconstruction of unknown objects through information-theoretic exploration were [38], [20], [64]. Some of these were evaluated in simulation only, and some on simple objects very different from the real world YCB [11] objects, which present challenges in collision avoidance during planning when the objects’ geometry is unknown. Some other studies have focused on shape reconstruction, specifically handling missed segments individually [55], [19]. However, using a passive exploration algorithm, their primary focus was on shape completion. In contrast, our work addresses the challenges of 3D object active exploration in both simulation and real-world, facilitating the reconstruction process by exploring the object in limited trials.

III. METHOD

In AcTExplore, we consider a tactile sensor mounted on a robotic arm end-effector interacting with an unknown fixed 3D object. We have access to the end-effector pose $\mathcal{P}_t \in SE(3)$ from the forward kinematics (FK). The goal is to navigate an unknown object’s entire surface within a limited workspace to collect tactile data for the reasons mentioned in Sec. I. At time t , the model will utilize consecutive tactile data $\{O_t, O_{t-1}, \dots, O_{t-(k-1)}\}$ to generate exploratory action $a_t \in SE(3)$ based on k recent observations. The effectiveness of the Completeness process heavily relies on the robot’s exploration algorithm.

An exploration policy, denoted as π_θ (referred to as the explorer) determines the next action, $\arg \max_{a_t} \pi_\theta(a_t | s_t)$, based on state s_t to maximize the cumulative reward (Sec. III-D) which serves as an estimate of area coverage over an unknown object’s surface. Therefore, the problem formulation of AcTExplore encompasses four key components:

A. Observation Space

At each time step t , the sensor observes skin deformation as an image (Fig. 2) that can be converted into a depth map of skin [32], [59]. We'll denote this depth image as observation space $O_t \in R^{H \times W}$, corresponding to the sensor's deformation at pose \mathcal{P}_t .

B. State Representation

The state serves as the sole input to the explorer, so it has to be sufficiently informative, enabling the model to generate exploratory actions. Let s_t be the state input for π_θ at time t . Considering the possibility that multiple optimal actions correspond to the observation O_t at finger pose \mathcal{P}_t , it is advantageous to construct a state representation s_t that incorporates short-term memory. This can be achieved by using a sequence of k consecutive observations $\{O_t, O_{t-1}, \dots, O_{t-(k-1)}\}$ to encapsulate the complexity of the state, enabling the policy to generate the appropriate action for long-horizon exploration. Therefore, $s_t = f(O_t, O_{t-1}, \dots, O_{t-(k-1)})$, where f can be any function representing spatio-temporal information on the sensor like optical flow, which is however computationally costly to generate on the fly. This function aims to reduce the dimensionality of s_t while extracting critical features for the state representation [2], which are not feasible to infer from a single observation. This is particularly valuable for learning high-level, complex actions that require a wider view. We conduct extensive experiments investigating the effectiveness of various state representations in Sec. IV. In Fig. 2, we visualize the resulting tactile readings and our proposed functions f :

$$\text{TTS}(O_t, O_{t-1}, \dots, O_{t-(k-1)}) : O_t \parallel O_{t-1} \parallel \dots \parallel O_{t-(k-1)}$$

(Temporal Tactile Stacking)

$$\text{TTA}(O_t, O_{t-1}, \dots, O_{t-(k-1)}) : \sum_{i=0}^{k-1} \alpha_i O_{t-i}$$

(Temporal Tactile Averaging)

where α_i are decreasing and $\sum_{i=0}^{k-1} \alpha_i = 1$ so the most recent reading O_t is the most effective observation in s_t , and others will be averaged respectively.

C. Action Space

To efficiently explore the complex geometry of 3D objects, we enable the finger to move in a 6-degree-of-freedom (6DOF) action space denoted as $A \in SE(3)$. In this action space, we consider small incremental translations (x, y, z) and rotations (γ, θ, ψ) around the workspace frame, with the bottom of the finger as the reference point. The model selects one of the dimensions $(x, y, z, \gamma, \theta, \psi)$ and either increases or decreases its value by the specified step. This results in a total of 12 possible actions within the action space. This action space is chosen to facilitate the control of the sensor under kinematic and collision constraints of the arm in both simulation and real world.

While this action space can facilitate exploration and interaction, it can be enhanced by adding an action that allows the finger to return to the last touching location. This additional




Fig. 3: **Depth readings sliding over hammer** (a): Sensor aligned with object's surface, receiving more depth information and moving stably. (b): Misaligned rotation increases the probability of losing contact in future steps.

action serves the purpose of touch recovery and addressing Non-Exploratory scenarios where contact with the object may be temporarily lost. By including this *Touch Recovery action* (a_{TR}), the robot can reestablish contact and give up on the trajectories that aren't worthy of further exploration. It will also guarantee that the model never strays too far from the object, as it learns to perform this action after a certain number of steps without any touch. Therefore, the action space consists of $12 + 1$ actions in total.

D. Reward

We can break down the exploration objective into easier sub-goals. Specifically, in this work, we want to maximize objects' coverage through the tactile sensor's trajectory. To accomplish this, the reward function is divided into two components:

1) **Contact Area Reward** measures the contact area between the robot's finger and the object's surface. The rationale is that a larger contact area corresponds to more information being gathered from the object, as illustrated in Fig. 3. In other words, this reward will encourage actions that align the tactile sensing area of the finger with the object's surface.

2) **Exploration Bonus** is intended to encourage the agent further to explore the workspace from a global perspective. Drawing inspiration from [3], [26], we have introduced a memory mechanism to track all the agent's trajectories. This memory allows us to keep a count of the number of times the agent has performed a particular action a in a specific pose \mathcal{P} , denoted as $N(\mathcal{P}, a)$. Even so, using $N(\mathcal{P}, a)$ directly for exploration is not practical because most of the states would have $N(\mathcal{P}, a) = 0$, especially considering that the workspace is often continuous or high-dimensional. So instead, we define $\hat{N}(\mathcal{P}, a)$ as the the number of close poses (Sec. SI-C-1¹) in the trajectory history. By having access to $\hat{N}(\mathcal{P}, a)$ for each pose-action pair, we can incorporate a bonus term $b_e = \frac{1}{\sqrt{\hat{N}(\mathcal{P}, a)}}$ into the area reward r_A in case that $r_A > 0$

which indicates that the sensor is touching the object. The bonus term b_e is a term that incentivizes exploration and curiosity throughout the trajectory. By incorporating this bonus, actions that have been infrequently taken in the past on a specific \mathcal{P} are rewarded more prominently, encouraging the agent to explore less-visited state-action pairs. Furthermore, the agent is guided toward more comprehensive exploration and exhibits a tendency to venture into uncharted regions of the workspace. This aids in mitigating the impact of sub-optimal local optima and fosters a broader understanding of the environment, resulting in improved convergence and enhanced exploration behavior as discussed in [26].

3) **Penalties** The agent might learn two possible trivial local optima scenarios to maximize the area reward and exploration bonus without exploring the object in long-horizon. To address this issue, we'll define the necessary negative rewards: **Revisit Penalty** (P_{rev}): To prevent the agent from learning policies that involve revisiting recently visited poses, we'll construct a short-term memory $\mathcal{D} = \{\mathcal{P}_i\}_{i=t-m}^t$ of m recent interactions by time t . If an action leads the sensor to revisit a pose already present in its short-term history \mathcal{D} , a revisit penalty $P_{rev} < 0$ will be imposed to discourage such trajectories. **Touch Recovery Penalty** (P_{TR}): One possible scenario that may not be covered by P_{rev} is when the sensor moves freely in space without making contact with the object for more than m steps ($|\mathcal{D}|$), and then performs a touch recovery action, which has a positive area reward ($r_A > 0$). To prevent such scenarios, we introduce a negative reward P_{TR} each time the agent selects the touch recovery action (a_{TR}). However, despite this negative reward, the agent still tends to choose the touch recovery action due to the high value of $V(s_{t+1})$ associated with recovery actions. Additionally, we can control the number of actions without touch by adjusting P_{TR} .

By penalizing non-exploratory scenarios, the agent is incentivized to explore new areas, mitigating the risk of getting stuck in sub-optimal loops and performing dexterous actions on a long horizon (further discussion in Sec. SI-B-1¹).

With all of these considerations in mind, the reward function is formulated as follows:

$$r(s_t, a_t) = \begin{cases} \alpha r_A + \frac{\beta}{\sqrt{\hat{N}(\mathcal{P}_t, a_t)}}, & \text{if } r_A > 0 \text{ and } \mathcal{P}_{t+1} \notin \mathcal{D} \\ P_{rev}, & \text{if } r_A > 0 \text{ and } \mathcal{P}_{t+1} \in \mathcal{D} \\ P_{TR}, & \text{if } a_t = a_{TR} \text{ (touch recovery)} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where r_A represents the reward based on the area of contact between the sensor and the object's surface at time t , P_{rev} denotes the penalty term applied to actions leading to a previously visited pose in \mathcal{D} and $\hat{N}(\mathcal{P}_t, a_t)$ signifies number of times the agent has performed action a_t in pose \mathcal{P}_t by time t , which is used to calculate the exploration bonus term.

By combining these components within the reward function, we aim to achieve a balance between contact area maximization (i.e., *exploitation*), and avoidance of non-exploratory scenarios (i.e., *exploration*) which addresses the mentioned difficulties for training in unknown environments.

To compute variance-reduced advantage/value function estimators, AcTEmplore utilizes a modified Proximal Policy Optimization (PPO) algorithm by modeling the exploration objective as an intrinsic auxiliary reward while enriching the state with temporal representation. We summarize our overall method in Alg. 1. The environment is initialized with an object to be explored and a tactile sensor to move on the object's surface and collect observations actively. We define the exploration workspace such that the first touch with the object happens as the sensor moves toward the **workspace's center**. The agent continues to interact with the object to learn optimal actions through the multi-objective reward **AMB** (Sec. III-D) that estimates coverage that is not available during training on unknown objects.

Algorithm 1: ACTEXPLORE, PROCEDURE

```

for episode = 1, 2, ... do
   $\mathcal{D} \leftarrow$  List of size  $m$ 
   $\mathcal{P}_0 \leftarrow$  First touch pose
   $O_0 \leftarrow$  Tactile sensor reading at  $\mathcal{P}_0$ 
   $\hat{N}(p, a) \leftarrow 0$  for all  $(p, a) \in \mathcal{P} \times \mathcal{A}$ 
  for  $t = 0, 1, 2, \dots, T-1$  do
     $s_t \leftarrow f(\{O_i\}_{i=t-\min(k,t)}^t)$ 
     $a_t \leftarrow \arg \max_{a'} \pi_{\theta}(a' | s_t)$ 
     $\mathcal{P}_{t+1}, O_{t+1} \leftarrow \text{tactileSensor.step}(a_t)$ 
     $\hat{N}(\mathcal{P}_t, a_t) \leftarrow \hat{N}_t(\mathcal{P}_t, a_t) + 1$ 
     $r_A \leftarrow \text{nonZeroCount}(O_t) / \text{size}(O_t)$ 
     $r \leftarrow 0$ 
    if  $a_t = a_{TR}$  then
       $r \leftarrow P_{TR}$  // touch recovery
    else if  $\mathcal{P}_{t+1} \in \mathcal{D}$  then
       $r \leftarrow P_{rev}$  // revisit
    else if  $r_A > 0$  then
       $b_e \leftarrow \frac{1}{\sqrt{\hat{N}_t(\mathcal{P}_t, a_t)}}$ 
       $r \leftarrow \alpha r_A + \beta b_e$ 
       $\mathcal{D}.\text{add}(\mathcal{P}_{t+1})$ 
       $\delta_t \leftarrow r + \gamma V(s_{t+1}) - V(s_t)$ 
    end
    Compute advantages  $\hat{A}_{i \in [T]} : \sum_{j=i}^{T-1} \gamma^{j-i} \delta_{i+j}$ 
     $\theta \leftarrow \text{Optimize surrogate } L^{CLIP}(\theta, \hat{A}_{[T]})$ 
  end

```

IV. EXPERIMENTS

This section evaluates and analyzes our method AcTEmplore, with various rewards and states on zero-shot (*unseen*) objects. In addition, we validate our method on over 400 quantitative and qualitative experiments in reconstructing unknown objects with varying complexities. In our experiments, we use reconstruction accuracy as the metric for tactile exploration with a limited number of steps as it represents the AcTEmplore exploration potential.

A. Experimental Setup

Simulation. We employ TACTO [61], [21] to simulate tactile sensor skin deformation during object interactions and modified PPO from StableBaselines3 [53] in AcTEmplore.

¹Supplementary material: <http://prg.cs.umd.edu/AcTEmplore>




Fig. 4: **Training results.** For each row: The [top] compares state representations using AMB for the reward function, while the [bottom] showcases different reward settings using TTA for state representation. Note that episodes terminate when the agent surpasses 90% IoU, reaches the horizon steps, or reaches the workspace boundaries.

The TACTO simulator is calibrated with real-sensor data to ensure Sim-to-Real generalization. It generates depth map images from real-world signals, serving as our observation O . We train the agent only with primitive objects – sphere and cube – for 300K steps. These primitive objects are selected as they represent a broad range of object shapes, with the sphere having curvature and the cube having sharp edges, corners, and flat surfaces. To assess the model’s performance, we evaluate it on YCB objects that were not encountered during training time. This evaluation demonstrates the efficacy of training with primitives, which exhibit strong generalization capabilities for objects with realistic textures (Fig. S2¹). For the termination condition, each episode either spans 5000 steps (Sec. SI-A¹) or concludes once the Intersection over Union (IoU) metric exceeds 90%, or when the sensor leaves the workspace boundaries. This strategy is adopted to reduce the training time. In Tab. II, we show that these termination conditions do not limit the IoU performance during testing as our methods achieved over 90%.

Real-World System. We employ a UR10 arm to manipulate the 6D pose of the DIGIT (Fig. 6). This control is achieved by transforming changes in the DIGIT’s frame into a set of joint trajectories via inverse kinematics which are facilitated with `ur_rtde`. The resulting trajectories are executed only if free from self-collisions and within the defined workspace. When an invalid trajectory is generated, we select an alternative action based on the PPO’s advantage values.

Unlike simulations, where consecutive action executions while in contact with the object have minimal impact, our real-world implementation introduces significant shearing on the sensing surface. To ensure the safe execution of actions generated by our policy, we have adopted a strategy of lifting the DIGIT in the normal direction of the contact after each contact event. This strategy remains well-founded due to our policy’s consistent alignment of our sensor surface with the object’s surface and does not compromise its general applicability. Our method successfully transferred to real-world experiments without requiring further fine-tuning. Fig. 6 illustrates the effectiveness of our exploration policy on a drill in the real-world.

Baselines Configuration. To evaluate the efficacy of each component, we have established a collection of baselines for three different state rep. and reward functions in Tab. I.

TABLE I: Baseline Formulations. **TTA**: Temporal Tactile Averaging, **TTS**: Temporal Tactile Stacking (concatenation is denoted as $\|$), **TM**: binary Touch indicator ($\mathbb{I}(\cdot)$) + short Memory, **AM**: contact Area + short Memory, **AMB**: contact Area + short Memory + UCB Bonus

State	Depth	TTA	TTS
	O_t	$\sum_{i=0}^{k-1} \alpha_i O_{t-i}$	$O_t \ \dots \ O_{t-(k-1)}$
Reward	TM	AM	AMB
	$\mathbb{I}(O_t)$	$r_A(O_t)$	$\alpha r_A(O_t) + \frac{\beta}{\sqrt{N(P_{t,a_t})}}$

B. Analysis & Discussion

State Comparison. We compare different state representations using the same reward function (AMB), considering both representations with and without temporal information. This analysis highlights the influential impact of temporal

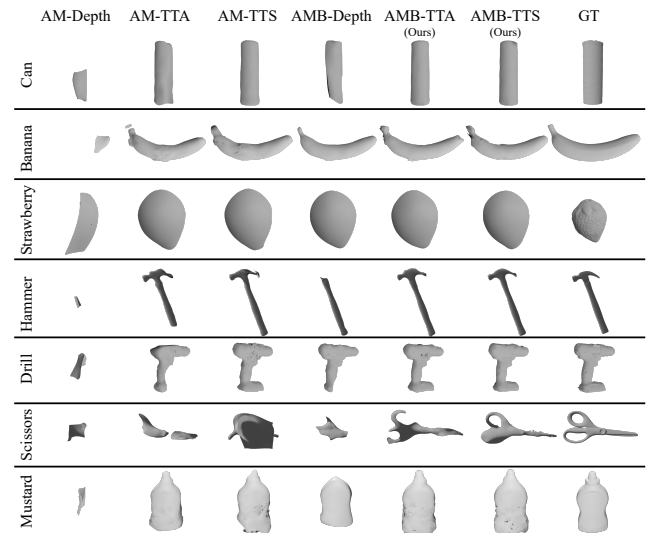


Fig. 5: **Qualitative results on unseen YCB objects** with different state and reward settings. We obtain point cloud data from active tactile exploration on the object’s surface. To generate a mesh from the collected point cloud, we apply Poisson surface reconstruction algorithm [30]. Further experiments are provided in supplementary materials.

TABLE II: **Quantitative results on unseen YCB objects:** The table presents IoU and Chamfer-L1 distance (cm) [42] between ground-truth and predicted meshes from methods in Tab. I. The surface area is listed below each object’s name as a severity metric. The details of metrics, confidence intervals, and step counts are given in the supplementary material¹.

Methods		Can (616 cm ²)	Banana (216 cm ²)	Strawberry (68 cm ²)	Hammer (410 cm ²)	Drill (591 cm ²)	Scissors (165.48 cm ²)	Mustard (454.54 cm ²)
		IoU \uparrow			(Chamfer- L_1 \downarrow)			
TM	depth	31.93 (2.66)	11.11 (7.52)	83.60 (0.44)	32.78 (1.86)	19.19 (4.1)	24.29 (8.15)	10.07 (4.07)
	TTA	17.60 (3.57)	6.03 (9.03)	41.0 (1.23)	14.85 (6.94)	28.15 (3.99)	14.17 (4.98)	19.94 (3.22)
	TTS	15.93 (5.22)	18.23 (5.48)	57.89 (0.88)	28.66 (2.47)	15.5 (3.97)	11.26(4.97)	14.55(2.95)
AM	depth	11.59 (5.49)	10.22 (6.84)	47.33 (1.16)	5.07 (7.69)	9.49 (4.03)	5.11(6.78)	11.04(5.16)
	TTA	72.70 (0.56)	97.70 (0.35)	100 (0.28)	79.80 (0.82)	57.58 (1.43)	41.77 (2.87)	71.72 (0.80)
	TTS	98.25 (0.22)	100 (0.34)	100 (0.31)	88.22 (0.44)	99.02 (0.37)	28.37 (2.38)	87.13 (0.59)
AMB	depth	41.45 (1.42)	98.64 (0.25)	100 (0.23)	61.42 (1.17)	79.68 (0.95)	31.99 (3.2)	65.74 (0.9)
	depth+LSTM	88.54 (0.3)	99.96 (0.28)	100 (0.24)	87.54 (0.49)	92.81 (0.36)	29.83 (0.58)	88.33 (0.36)
	TTA (ours)	89.6 (0.29)	100 (0.33)	100 (0.25)	98.22 (0.29)	98.85 (0.32)	67.02 (0.87)	95.91 (0.51)
	TTS (ours)	97.45 (0.20)	100 (0.3)	100 (0.25)	96.96 (0.28)	99.74 (0.31)	74.62 (0.61)	95.02 (0.49)

information on learning dexterous and high-level actions. As shown in Fig. 4, all state representations achieve a specified IoU during training. However, the state representations incorporating temporal information demonstrate higher stability, consistently reaching the 90% IoU objective after 200K steps. In contrast, the depth-only representation struggles to maintain the IoU objective and is outperformed by temporal representations in Tab. II. Furthermore, when considering the number of steps required to achieve the IoU objective, TTS training takes longer than TTA as $s_t^{TTS} \in R^{k \times H \times W}$ is k times bigger than $s_t^{TTA} \in R^{H \times W}$ which is averaging observations rather than stacking them. However, in our experiments in Fig. 5, we witnessed that both TTA and TTS are competitive, with TTS excelling on longer objects and TTA performing better on complex shapes.

Reward Comparison. In our pursuit of efficient exploration, we tried various reward functions mentioned in Tab. I. During training, we plotted the IoU and episode length until termination in Fig. 4. Notably, the AMB reward function outperformed the others, satisfying the IoU objective through encouraging exploration of less visited poses. In contrast, TM and AM cannot use environmental feedback as much as AMB can. This limitation arises from TM and AM’s deprivation of long-horizon history, which hampers their capacity to gather sufficient information through intrinsic rewards. As a result, AMB is better equipped to leverage environment feedback $\left(\frac{1}{N(P,a)}\right)$ effectively for improved exploration and sample efficiency. However, AM outperforms TM as it utilizes contact area information and can still align the sensor’s sensing area with the object surface to collect more information and maintain a reliable touch for future actions. Indeed, the disparity between TM and AM can also be understood as the distinction between using a touch sensor versus a tactile sensor for exploring an object.

Limitations and Future Work. The current formulation of our method has certain limitations. First, it assumes a moving sensor relative to a **fixed-pose** rigid object, necessitating a physically accurate simulator to narrow the sim2real gap for moving objects. Second, Although AcTEExplore is not restricted by object shape, it is designed to keep the sensor close to recent touching poses. This could pose challenges in environments with disconnected components. Workspace splitting can be a potential solution to address this problem.

Third, the sensor exhibits a small depth bias in the simulation resulting in larger reconstructions. While generally negligible, this bias becomes dominant when handling objects roughly the same size as the sensor, such as the strawberry shown in Fig. 5.

As a step towards benchmarking in tactile exploration, we have released our extensive explorations for YCB objects in Tab. S1¹ with a maximum of 5000 steps. While employing tactile sensors on multi-finger robotic hands may streamline the exploration process [58], there remains a promising direction for future research in modifying the POMDP that effectively handles collisions between sensors while maintaining object generalization.

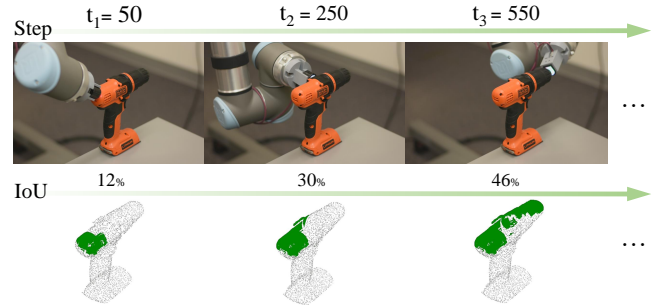


Fig. 6: **Real-World Exploration Execution.** Still frames from AcTEExplore’s exploration of a drill, starting from the rear and progressing towards the chuck. The second row shows the covered area per step, with IoU computed over the exploration workspace above the drill’s grip. t_i is the i -th step of the trajectory.

V. CONCLUSION

In this work, we introduced a novel reinforcement learning method using tactile sensing to explore unknown 3D objects actively. It addresses the need for an active exploration method to enable numerous works [18], [35], [56], [17] to become fully automated. AcTEExplore is not limited to specific shape distributions as it has only been trained on primitive shapes to learn fundamental movements by leveraging temporal tactile information and intrinsic exploration bonuses. We demonstrated this through our experiments with various shape complexities like a drill or a clay pot in both the real world and simulation.

REFERENCES

- [1] Alex Alspach, Kunimatsu Hashimoto, Naveen Kuppaswamy, and Russ Tedrake. Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation. pages 597–604, 04 2019.
- [2] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew, Jakob Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- [3] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, 2017.
- [4] Maria Bauza, Antonia Bronars, and Alberto Rodriguez. Tac2Pose: Tactile Object Pose Estimation from the First Touch. 2022. Publisher: arXiv Version Number: 2.
- [5] Wouter M. Bergmann Tiest and Astrid M. L. Kappers. The influence of visual and haptic material information on early grasping force. *Royal Society Open Science*, 6(3):181563, 2019.
- [6] Aude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364:eaat8414, 06 2019.
- [7] N. Botteghi, Beril Sirmacek, R. Schulte, M. Poel, and C. Brune. REINFORCEMENT LEARNING HELPS SLAM: LEARNING TO BUILD MAPS. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, 43, 2020.
- [8] Nicolò Botteghi, Rob Schulte, Beril Sirmacek, Mannes Poel, and Christoph Brune. Curiosity-Driven Reinforcement Learning Agent for Mapping Unknown Indoor Environments. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 1:129–136, 2021. Publisher: Copernicus GmbH.
- [9] Gabriele M. Caddeo, Andrea Maracani, Paolo D. Alfano, Nicola A. Piga, Lorenzo Rosasco, and Lorenzo Natale. Sim2real bilevel adaptation for object surface classification using vision-based tactile sensors, 2023.
- [10] Gabriele M. Caddeo, Nicola A. Piga, Fabrizio Bottarel, and Lorenzo Natale. Collision-aware in-hand 6d object pose estimation using multiple vision-based tactile sensors. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 719–725, 2023.
- [11] Berk Çalli, Aaron Walsman, Arjun Singh, Siddhartha S. Srinivasa, Pieter Abbeel, and Aaron M. Dollar. Benchmarking in manipulation research: The YCB object and model set and benchmarking protocols. *CoRR*, abs/1502.03143, 2015.
- [12] Céline Cappe, Gregor Thut, Vincenzo Romei, and Micah M. Murray. Selective integration of auditory-visual looming cues by humans. *Neuropsychologia*, 47(4):1045–1052, March 2009.
- [13] Devendra Singh Chaitan, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to Explore using Active Neural SLAM, April 2020. arXiv:2004.05155 [cs].
- [14] Fanfei Chen, Shi Bai, Tixiao Shan, and Brendan Englot. Self-learning exploration and mapping for mobile robots via deep reinforcement learning. In *Aiaa scitech 2019 forum*, page 0396, 2019.
- [15] Fanfei Chen, John D. Martin, Yewei Huang, Jinkun Wang, and Brendan Englot. Autonomous exploration under uncertainty via deep reinforcement learning on graphs. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6140–6147. IEEE, 2020.
- [16] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning Exploration Policies for Navigation, March 2019. arXiv:1903.01959 [cs].
- [17] Mauro Comi, Yijiong Lin, Alex Church, Alessio Tonioni, Laurence Aitchison, and Nathan F. Lepora. Touchsdf: A deepsf approach for 3d shape reconstruction using vision-based tactile sensing. *ArXiv*, abs/2311.12602, 2023.
- [18] Cristiana de Farias, Naresh Marturi, Rustam Stolkin, and Yasemin Bekiroglu. Simultaneous tactile exploration and grasp refinement for unknown objects. *CoRR*, abs/2103.00655, 2021.
- [19] Won Kyung Do and Monroe Kennedy. DenseTact: Optical Tactile Sensor for Dense Shape Reconstruction. *2022 International Conference on Robotics and Automation (ICRA)*, pages 6188–6194, May 2022.
- [20] Danny Driess, Daniel Hennes, and Marc Toussaint. Active multi-contact continuous tactile exploration with gaussian process differential entropy. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7844–7850, 2019.
- [21] Benjamin Ellenberger. Pybullet gymperium. <https://github.com/benelot/pybullet-gym>, 2018–2019.
- [22] Kanishka Ganguly, Pavan Mantripragada, Chethan M Parameshwara, Cornelia Fermüller, Nitin J Sanket, and Yiannis Aloimonos. Gradtac: Spatio-temporal gradient based tactile sensing. *Frontiers in Robotics and AI*, 9:898075, 2022.
- [23] Kanishka Ganguly, Behzad Sadrifaridpour, Krishna Bhavithavya Kidambi, Cornelia Fermüller, and Yiannis Aloimonos. Grasping in the dark: Compliant grasping using shadow dexterous hand and biotac tactile sensor.
- [24] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. IQA: visual question answering in interactive environments. *CoRR*, abs/1712.03316, 2017.
- [25] Francois Robert Hogan, Michael Jenkin, Sahand Rezaei-Shoshtari, Yogesh A. Girdhar, David Meger, and Gregory Dudek. Seeing through your skin: Recognizing objects with a novel visuotactile sensor. volume abs/2011.09552, 2020.
- [26] Chi Jin, Zeyuan Allen-Zhu, Sebastian Bubeck, and Michael I Jordan. Is q-learning provably efficient? In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [27] R.S. Johansson, U. Lundström, and R. Lundström. Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Research*, 244(1):17–25, 1982.
- [28] Astrid M. L. Kappers. Human perception of shape from touch. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581):3106–3114, 2011.
- [29] Tobias Katus and Søren K. Andersen. Chapter 21 - the role of spatial attention in tactile short-term memory. In Pierre Jolicœur, Christine Lefebvre, and Julio Martinez-Trujillo, editors, *Mechanisms of Sensory Working Memory*, pages 275–292. Academic Press, San Diego, 2015.
- [30] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, page 0, 2006.
- [31] Tarik Kelestemur, Robert Platt, and Taskin Padir. Tactile Pose Estimation and Policy Learning for Unknown Object Manipulation. 2022. Publisher: arXiv Version Number: 1.
- [32] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian H. Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, Dinesh Jayaraman, and Roberto Calandra. DIGIT: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *CoRR*, abs/2005.14679, 2020.
- [33] Rebecca Lawson, Alexandra M. Fernandes, Pedro B. Albuquerque, and Simon Lacey. Chapter 19 - remembering touch: Using interference tasks to study tactile and haptic memory. In Pierre Jolicœur, Christine Lefebvre, and Julio Martinez-Trujillo, editors, *Mechanisms of Sensory Working Memory*, pages 239–259. Academic Press, San Diego, 2015.
- [34] Huai-Jen Liang, Nitin J Sanket, Cornelia Fermüller, and Yiannis Aloimonos. SalientDso: Bringing attention to direct sparse odometry. *IEEE Transactions on Automation Science and Engineering*, 16(4):1619–1626, 2019.
- [35] Wenyu Liang, Qinyuan Ren, Xiaojiao Chen, Junli Gao, and Yan Wu. Dexterous manoeuvre through touch in a cluttered scene. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6308–6314, 2021.
- [36] Changyi Lin, Ziqi Lin, Shaoxiong Wang, and Huazhe Xu. DTact: A Vision-Based Tactile Sensor that Measures High-Resolution 3D Geometry Directly from Darkness. 2022. Publisher: arXiv Version Number: 1.
- [37] Justin Lin, Roberto Calandra, and Sergey Levine. Learning to Identify Object Instances by Touch: Tactile Recognition via Multimodal Matching. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3644–3650, May 2019. ISSN: 2577-087X.
- [38] Yujie Lu, Jianren Wang, and Vikash Kumar. Curiosity Driven Self-supervised Tactile Exploration of Unknown Objects. 2022. Publisher: arXiv Version Number: 1.
- [39] Andréa Macario Barros, Maugan Michel, Yoann Moline, Gwenolé Corre, and Frédéric Carrel. A comprehensive survey of visual slam algorithms. *Robotics*, 11(1), 2022.
- [40] Uriel Martinez-Hernandez, Nathan F. Lepora, and Tony J. Prescott. Active haptic shape recognition by intrinsic motivation with a robot hand. In *2015 IEEE World Haptics Conference (WHC)*, pages 299–304, 2015.
- [41] Luca Massari, Calogero Oddo, Edoardo Sinibaldi, Renaud Detry, Joseph Bowkett, and Kalind Carpenter. Tactile sensing and control of robotic manipulator integrating fiber bragg grating strain-sensor. *Frontiers in Neurobotics*, 13, 04 2019.

- [42] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [43] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J. Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharrshan Kumaran, and Raia Hadsell. Learning to navigate in complex environments. *CoRR*, abs/1611.03673, 2016.
- [44] Nicole E. Munoz and Daniel T. Blumstein. Multisensory perception in uncertain environments. *Behavioral Ecology*, 23(3):457–462, May 2012.
- [45] Yashraj Narang*, Balakumar Sundaralingam*, Miles Macklin, Arsalan Mousavian, and Dieter Fox. Sim-to-real for robotic tactile sensing via physics-based simulation and learned latent projections (*equal contribution). *IEEE Intl. Conf. on Robotics and Automation*, pages 6444–6451, 2021.
- [46] Yashraj S. Narang, Balakumar Sundaralingam, Karl Van Wyk, Arsalan Mousavian, and Dieter Fox. Interpreting and predicting tactile signals for the syntouch biotac. *CoRR*, abs/2101.05452, 2021.
- [47] Simon Ottenhaus, Lukas Kaul, Nikolaus Vahrenkamp, and Tamim Asfour. Active tactile exploration based on cost-aware information gain maximization. *International Journal of Humanoid Robotics*, 15:1850015, 02 2018.
- [48] K. Park, H. Yuk, M. Yang, J. Cho, H. Lee, and J. Kim. A biomimetic elastomeric robot skin using electrical impedance and acoustic tomography for tactile sensing. *Science Robotics*, 7(67):eabm7187, June 2022. Publisher: American Association for the Advancement of Science.
- [49] Leszek Pecyna, Siyuan Dong, and Shan Luo. Visual-Tactile Multimodality for Following Deformable Linear Objects Using Reinforcement Learning. *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3987–3994, October 2022. Conference Name: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) ISBN: 9781665479271 Place: Kyoto, Japan Publisher: IEEE.
- [50] Julio A. Placed and José A. Castellanos. A Deep Reinforcement Learning Approach for Active SLAM. *Applied Sciences*, 10(23):8386, January 2020. Number: 23 Publisher: Multidisciplinary Digital Publishing Institute.
- [51] Julio A. Placed, Jared Strader, Henry Carrillo, Nikolay Atanasov, Vadim Indelman, Luca Carlone, and José A. Castellanos. A Survey on Active Simultaneous Localization and Mapping: State of the Art and New Frontiers. *IEEE Transactions on Robotics*, pages 1–20, 2023. Conference Name: IEEE Transactions on Robotics.
- [52] Parth Potdar, David Hardman, Elijah Almanzor, and Fumiya Iida. High-speed tactile braille reading via biomimetic sliding interactions. *IEEE Robotics and Automation Letters*, 9(3):2614–2621, 2024.
- [53] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- [54] Dhruv Ramani. A short survey on memory based reinforcement learning. *ArXiv*, abs/1904.06736, 2019.
- [55] Edward J. Smith, Roberto Calandra, Adriana Romero, Georgia Gkioxari, David Meger, Jitendra Malik, and Michal Drozdal. 3d shape reconstruction from vision and touch. *CoRR*, abs/2007.03778, 2020.
- [56] S. Suresh, Z. Si, J. Mangelson, W. Yuan, and M. Kaess. ShapeMap 3-D: Efficient shape mapping through dense touch and vision. In *Proc. IEEE Intl. Conf. on Robotics and Automation, ICRA*, Philadelphia, PA, USA, May 2022.
- [57] Sudharshan Suresh, Maria Bauza, Kuan-Ting Yu, Joshua G. Mangelson, Alberto Rodriguez, and Michael Kaess. Tactile SLAM: Real-time inference of shape and pose from planar pushing. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11322–11328, May 2021. ISSN: 2577-087X.
- [58] Sudharshan Suresh, Haozhi Qi, Tingfan Wu, Taosha Fan, Luis Pineda, Mike Lambeta, Jitendra Malik, Mrinal Kalakrishnan, Roberto Calandra, Michael Kaess, Joseph Ortiz, and Mustafa Mukadam. Neural feels with neural fields: Visuo-tactile perception for in-hand manipulation. In *arXiv preprint arXiv:2312.1346*, December 2023.
- [59] Sudharshan Suresh, Zilin Si, Stuart Anderson, Michael Kaess, and Mustafa Mukadam. MidasTouch: Monte-Carlo inference over distributions across sliding touch. In *Proc. Conf. on Robot Learning, CoRL*, Auckland, NZ, December 2022.
- [60] SynTouch. Biotac product manual. Aug 2018.
- [61] Shaoxiong Wang, Mike Lambeta, Po-Wei Chou, and Roberto Calandra. TACTO: A fast, flexible and open-source simulator for high-resolution vision-based tactile sensors. *CoRR*, abs/2012.08456, 2020.
- [62] Xuelian Wei, Baocheng Wang, Zhiyi Wu, and Zhong Lin Wang. An open-environment tactile sensing system: Toward simple and efficient material identification. *Advanced Materials*, 34(29):2203073, 2022.
- [63] Jingxi Xu, Han Lin, Shuran Song, and Matei T. Ciocarlie. Tandem3d: Active tactile exploration for 3d object recognition, 2022.
- [64] Zhengkun Yi, Roberto Calandra, Filipe Veiga, Herke van Hoof, Tucker Hermans, Yilei Zhang, and Jan Peters. Active tactile object exploration with gaussian processes. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4925–4930, 2016.
- [65] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12), 2017.
- [66] 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, 2023.
- [67] Jialiang Zhao, Maria Bauzá, and Edward H. Adelson. Fingerslam: Closed-loop unknown object localization and reconstruction from visuo-tactile feedback, 2023.