AcTExplore: Active Tactile Exploration on Unknown Objects

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

Perception and Robotics Group
University of Maryland, College-Park

Introduction

Problem. Exploring objects with tactile sensors is non-trivial.

Why? Tactile sensors are like blind explorers exploring a vast landscape. Thorough exploring an object's surface with them is a formidable undertaking and creates exploration challenges: sensors might lose contact or become trapped in repetitive exploration trajectories.

Solution? We introduce a reinforcement learning approach for object reconstruction at scales that autonomously explores object surfaces in a limited number of steps. Through sufficient exploration, Our algorithm incrementally collects tactile data to reconstruct 3D shapes of objects. Achieving an average of **95.97%** IoU coverage on unseen YCB objects, using only primitive shapes (e.g. sphere, cube) for training.

Key Contributions

- Proposing an *Active Tactile Exploration Algorithm* for 3D reconstruction at scale using tactile sensing. The learned behavior extends to *unseen real-world objects*.
- Introducing a state representation, *Temporal Tactile Sensing*, to enable Short-Term Memory on tactile receptors, inspired by various neurological and behavioral studies.

Methodology

We formulate the exploration problem as a Partially Observable Markov Decision Process with the following setting:

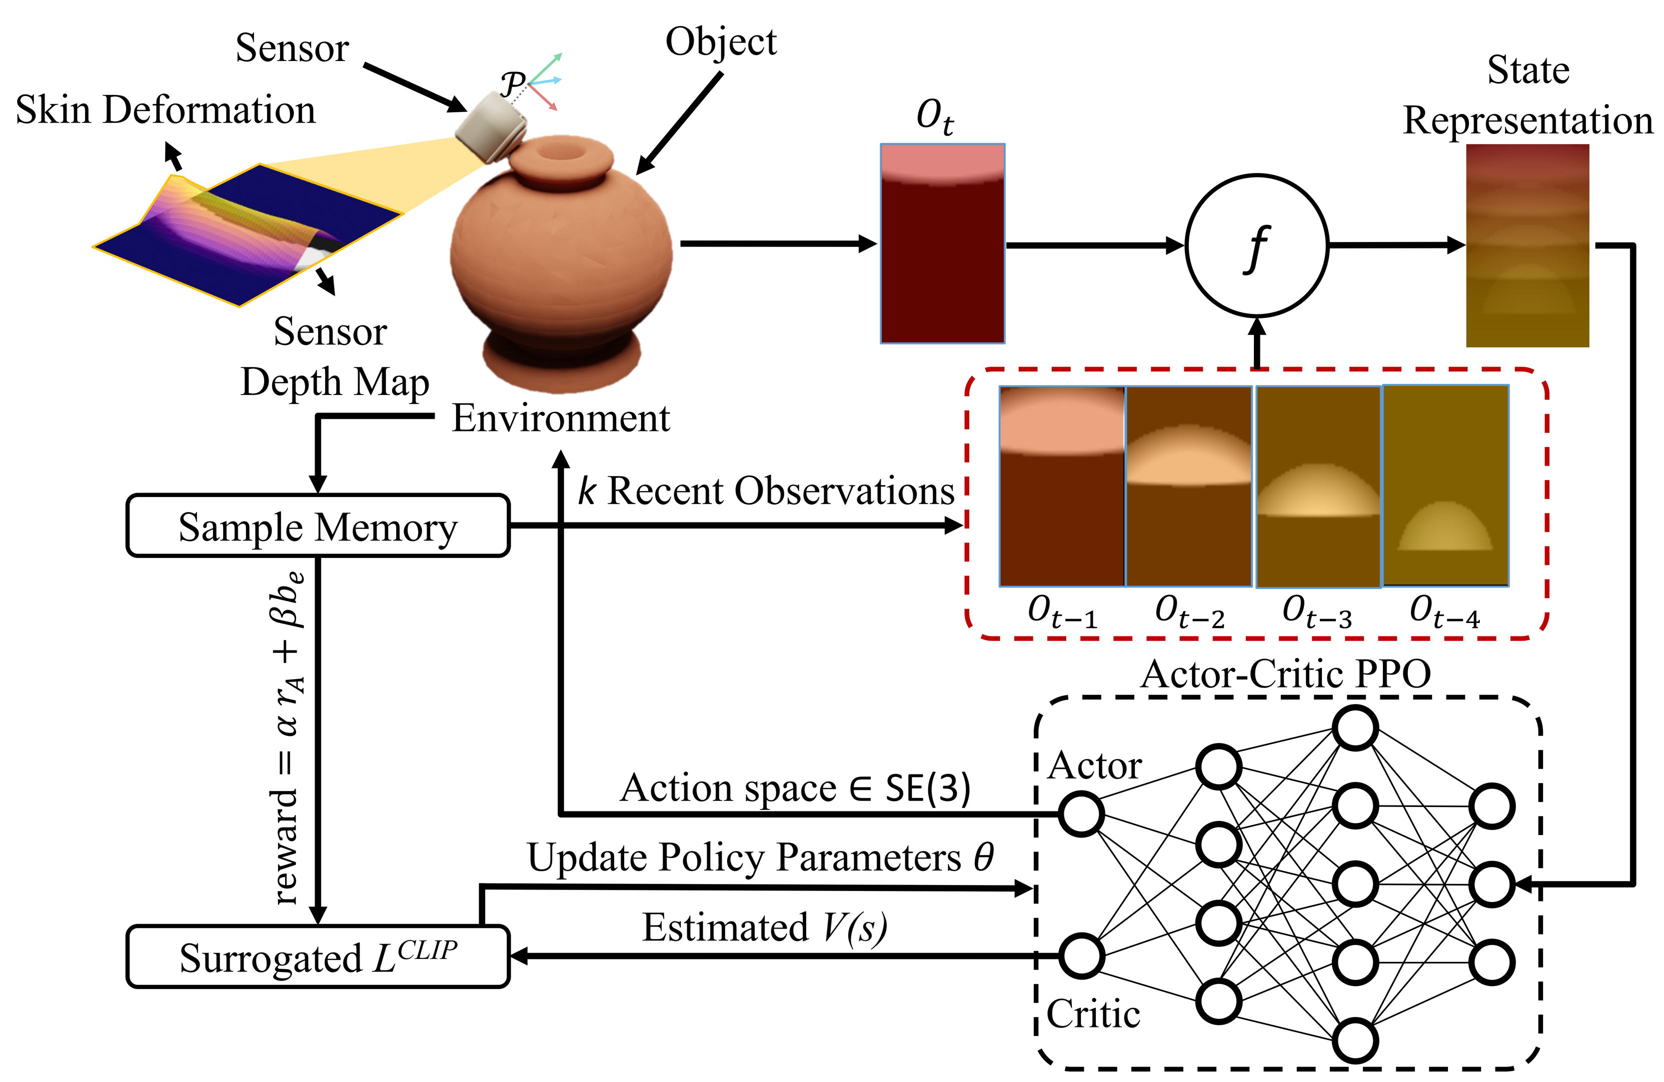

Observation Space: We consider depth image from the Tactile sensor as our observation.

$$O_t \in R^{H\times W}$$

State Representations: To enrich the state with spatio-temporal information, we introduced short-term memory in state representation.

$$\text{TTS}(O_t, O_{t-1}, ..., O_{t-(k-1)}) : O_t \,\|\, O_{t-1} \,\|\, \ldots \,\|\, O_{t-(k-1)} \qquad \textbf{(T}\text{emporal } \textbf{T}\text{actile } \textbf{S}\text{tacking)}$$

$$\text{TTA}(O_t, O_{t-1}, ..., O_{t-(k-1)}) : \sum_{i=0}^{k-1} \alpha_i O_{t-i} \qquad \textbf{(T}\text{emporal } \textbf{T}\text{actile } \textbf{A}\text{veraging)}$$

Action Space: We have a set of 12+1 distinct actions. The model selects one of the dimensions $(x, y, z, \gamma, \theta, \psi)$ and either increases or decreases its value by the specified step size. This results in a total of 12 possible actions. An additional action, denoted as a_{TR}, is included specifically to recover touch by moving the sensor back to the last touching pose.

Reward: Throughout the exploration we want to **encourage** the agent to maximize the contact area which aligns the sensor with object's surface by r_A. Moreover, the agent receives a bonus when performing new actions at touching poses to uncover new areas with $b_e = \frac{1}{\sqrt{\hat{N}(\mathcal{P}_t,a_t)}}$. Conversely, we want to **discourage** revisiting the same pose with $Prev$ and mitigate trivial local optima by penalizing Touch Recovery action a_{TR} with P_{TR}.

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

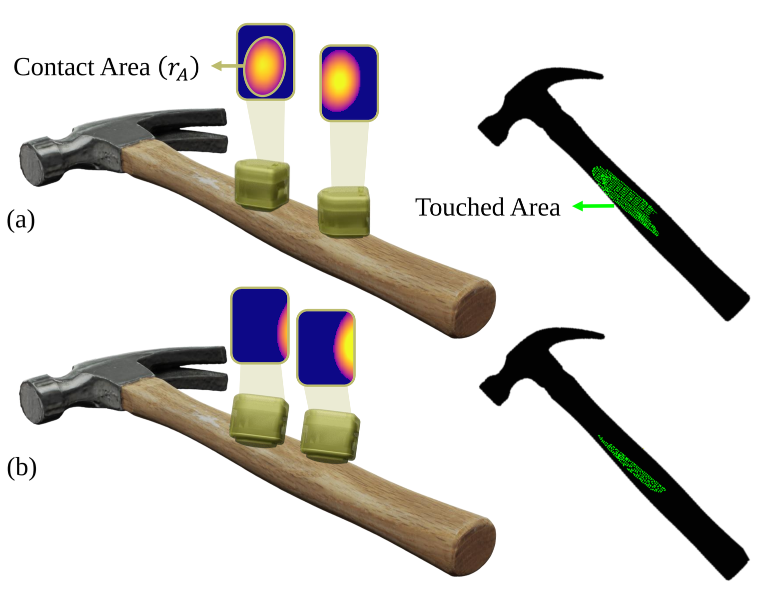

Experiments

To demonstrate the effectiveness of our proposed method, we ablate different state representations and rewards:

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

Also we have tested the algorithm on real-world unseen objects with different textures.

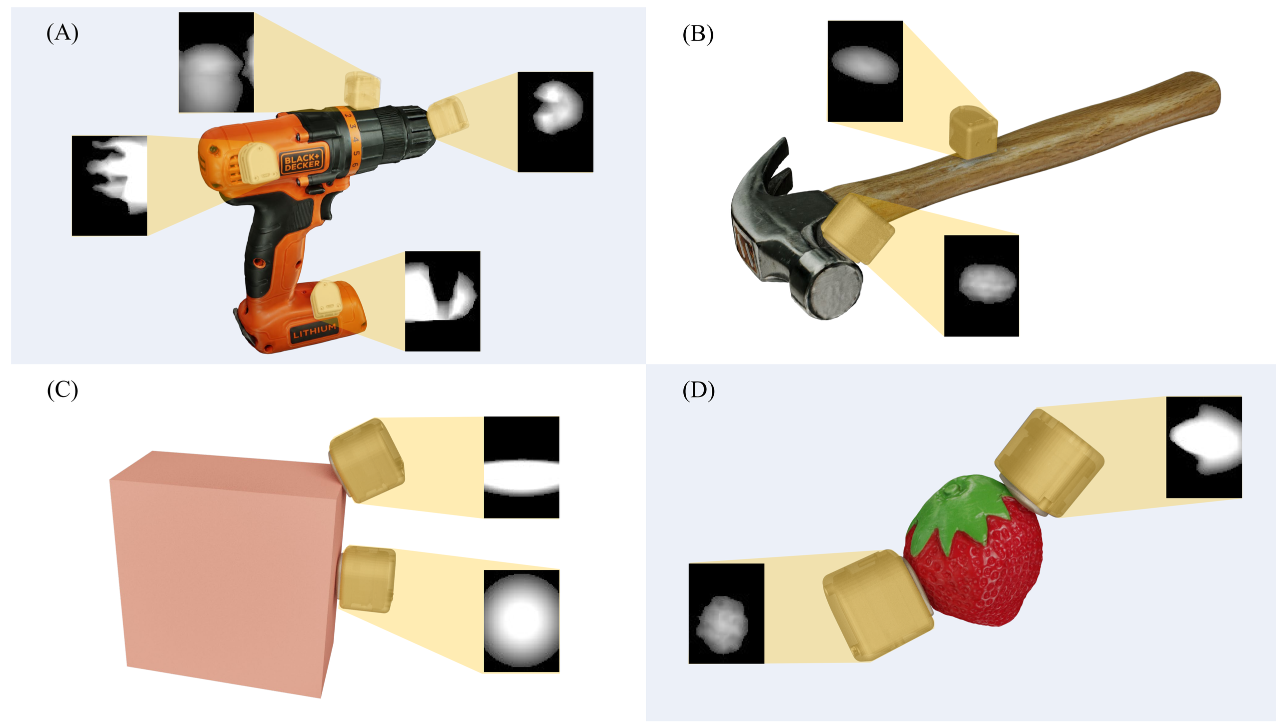

We then mounted the sensor on a robotic arm to manipulate the 6D-pose of the sensor and performed real-world execution of our exploration algorithm. Scan the QR code above to watch the demo.

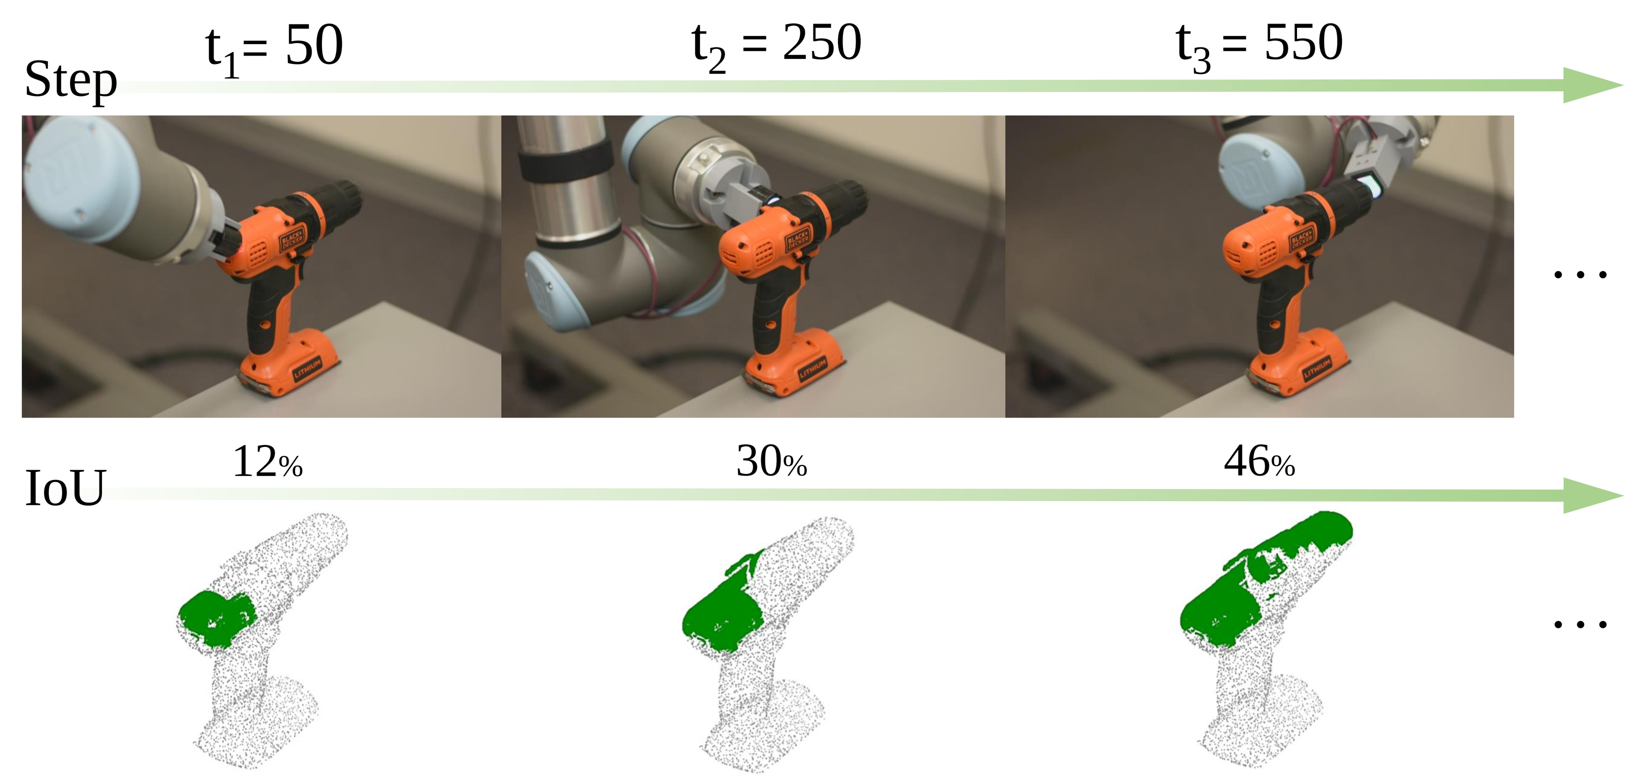

Results

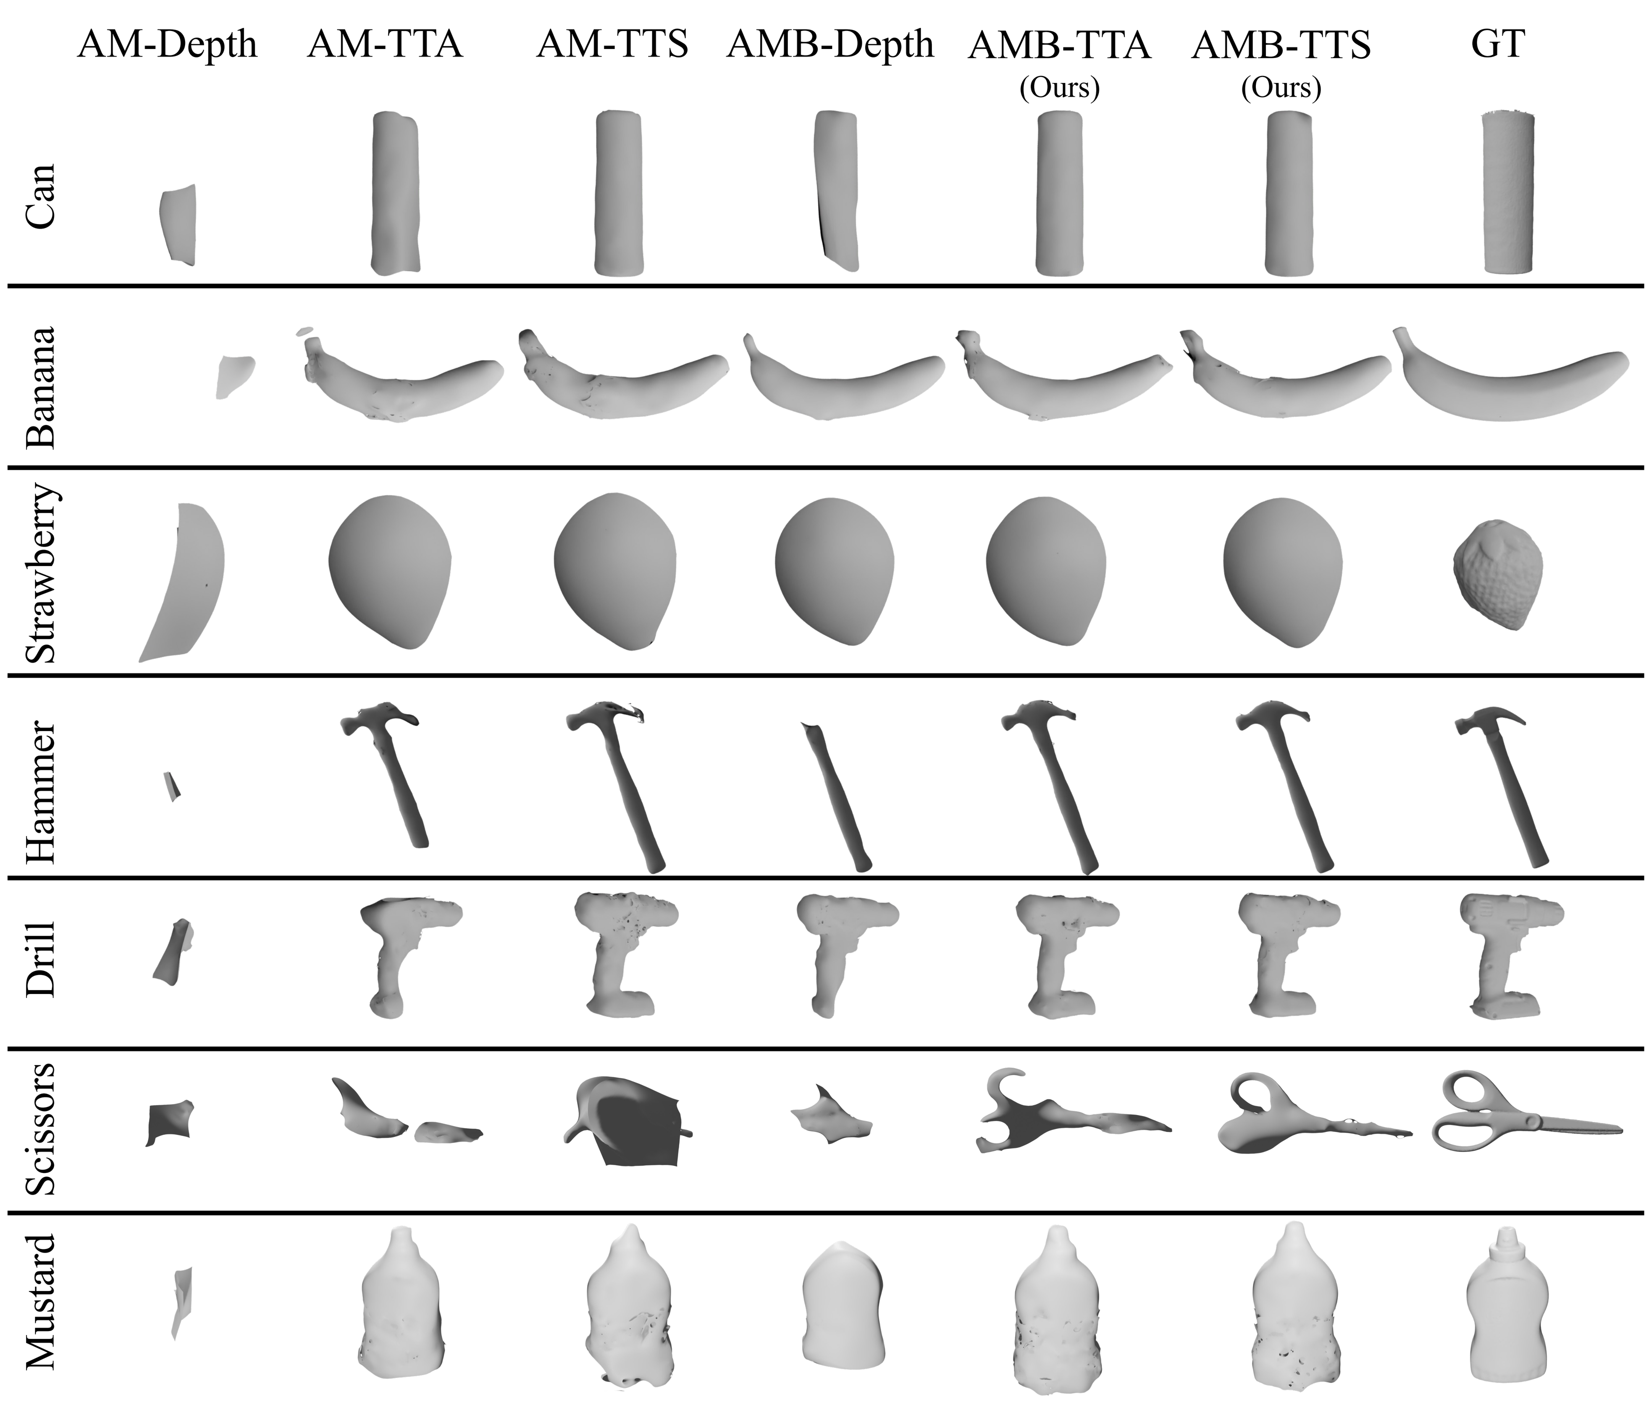

Summary and Future Works

AcTExplore addresses the need for an active exploration method to enable works such as grasp refinement and scene perception to become fully automated. AcTExplore 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 demonstrate this through our experiments with various shape complexities like a drill or a clay pot in both the real world and simulation.

amirsh@umd.edu