

Active End-Effector Pose Selection for Tactile Object Recognition through Monte Carlo Tree Search

Mabel M. Zhang, Nikolay Atanasov, and Kostas Daniilidis

Abstract—This paper considers the problem of active object recognition using touch only. The focus is on adaptively selecting a sequence of wrist poses that achieves accurate recognition by enclosure grasps. It seeks to minimize the number of touches and maximize recognition confidence. The actions are formulated as wrist poses relative to each other, making the algorithm independent of absolute workspace coordinates. The optimal sequence is approximated by Monte Carlo tree search. We demonstrate results in a physics engine and on a real robot. In the physics engine, most object instances were recognized in at most 16 grasps. On a real robot, our method recognized objects in 2–9 grasps and outperformed a greedy baseline.

I. INTRODUCTION

Tactile sensing for object recognition has been an area of research since the 1980s [1], [2], [3]. Major advances have been slow, partly due to the sparse nature of touch, which requires more sensing time for a large area coverage, plus motion planning and physical movement time. Additionally, manipulator hardware is expensive. In comparison, vision-based recognition has seen major improvement because of the rich data, rapid information gathering, and low cost.

However, scenarios exist where vision is unreliable, such as dark, dusty, smoky, or blurry underwater environments, transparent and reflective objects, occluded back sides, and objects in a bag. In these cases, tactile sensing is a better main modality. Furthermore, the ultimate goal of manipulation is to contact the object. Starting with contacts early on provides direct physical exteroception that vision cannot achieve. In fact, physical action is naturally integrated with perception in animals, who use various active tactile sensing organs [4]. Humans can recover shapes by touch alone.

While some disadvantages of tactile sensing can be compensated by better hardware, others can be compensated by efficient planning and exploitation of the limited input. In fact, Flanagan *et al.* [5] found that the key to sophisticated manipulation in the human hand lies more in the accurate prediction of motor commands and outcomes than in rapid sensing. This learning and prediction are the bases of active sensing. Active tactile sensing had early work in tandem with passive vision [6], [7], [8] and alone [9]. Active perception, as noted by Bajcsy [10], involves a selection strategy that trades off between task success and the energy to achieve it.

In this paper, we tackle touch-only object recognition by an autonomous active selection algorithm, which aims to

The authors are with the GRASP Laboratory, University of Pennsylvania, 3330 Walnut Street, Philadelphia, PA 19104, USA. {zmen@seas, atanasov@seas, kostas@cis}.upenn.edu. Grateful for support through the following grants: NSF-DGE-0966142 (IGERT), NSF-IIP-1439681 (IUCRC), ARL RCTA W911NF-10-2-0016, and a GSK grant.

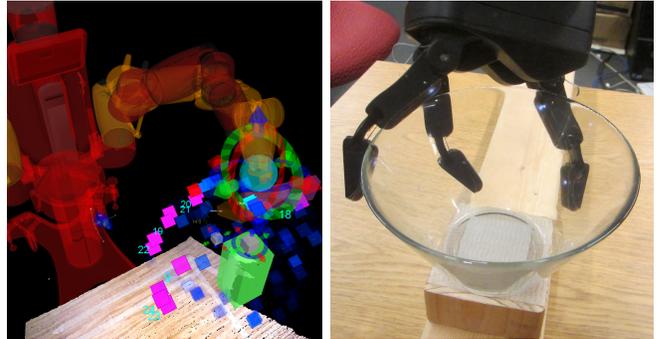


Fig. 1: Left: Experiment setup. Right: An adaptively selected pose.

select a minimum number of wrist poses for maximum recognition likelihood. We formulate the problem as a Markov Decision Process (MDP) and optimize for such a policy.

Our core idea is that consecutive tactile features observed on an object are related to the robot movements in between. Local features are not unique, repeating at symmetric parts and similar curvatures. Discretizing them across objects creates common features that can be modeled as a probability distribution, which we condition on observations and actions, independent of large state space dimensionality.

We focus on the active prediction problem and not the recognition. For recognition, we use an existing tactile object descriptor [11], the weakness in which is that recognition required hundreds of systematic touches, unrealistic in practice. Our goal is to eliminate that weakness by strategically selecting a small number of touches to observe the key features. We were able to decrease the number by a magnitude.

The novelty of our active approach has three parts. First, unlike typical active models, ours is independent of the state space, by conditioning on observations and actions. State space-dependent methods have search times proportional to state dimensionality, posing a limit on state definition. Second, unlike most active perception approaches, ours is not greedy. Third, we solve a high-level perception problem as opposed to a low-level sensor-focused haptics one. We target autonomous object-level recognition with cheap sparse pressure sensors, in contrast to most tactile recognition work, which are material-level recognition with expensive dense multi-modal sensors and predefined strokes. Our algorithmic abstraction is not limited to special sensors. We show successful prediction in a physics engine and on a real robot.

II. RELATED WORK

Tactile work has been done for reconstruction, localization, pose estimation, and recognition. Our work differs from haptics work on material recognition in three major ways.

First, we are solving a perception problem at the object level based on high-level geometry abstracted from sensor inputs, not at the material level that directly uses tactile vibrations. Second, we focus on active prediction of the most useful actions, whereas haptic recognition typically uses predefined motions. Third, we execute the actions autonomously.

Similar to our active end-effector pose selection, active viewpoints have been used to select camera poses to gather information, such as in [12]. Sensing with vision only is considered active perception. Our work is closer to interactive perception, which physically contacts the environment [13].

Early work have explored touch-only object recognition not involving active planning. Bajcsy [14] compared human haptic exploratory procedures (EPs) observed by Lederman and Klatzky [15] to robots, and Allen [3] extended them to a tactile robotic hand. Gaston [1], Grimson [2], and Siegel [16] used Interpretation Trees for recognition and pose estimation.

Active touch has been coupled with vision to various extents. Allen *et al.* [6], [8] used vision to guide active touch to invisible regions and explicitly fused the two for shape reconstruction. Stansfield [7] used an initial visual phase for rough object properties and a final haptic phase for detailed surface properties. Others explored solely using active touch. Schneiter [9] scheduled sensor motions based on [2] for recognition. Maekawa *et al.* [17] advanced through grid points for reconstruction as contacts were detected. Hsiao *et al.* [18] partitioned the workspace and represented each region as a state in a POMDP for optimal control policy.

Many recent active learning algorithms greedily maximize information gain (IG) [19], [20], [21], [22], [12] for property estimation or task success. Another recent development is adaptive submodularity, a diminishing return property [23], [24]. It was shown that entropy can be submodular and was used to greedily maximize IG for near-optimal touches [25].

Work most related to ours in active recognition are Pezzementi *et al.* [26] and Hausman *et al.* [27]. Both use tree search to select actions. However, Pezzementi’s tree was for motion planning, with nodes being collision-free configurations. Hausman’s tree nodes were entropy, which were minimized to find optimal poses to move the object into camera view. Our tree nodes are tactile observations, and we select end-effector poses to maximize recognition.

Different from greedy policies, a lookahead policy (*e.g.* tree search) explicitly optimizes cost and gain several steps ahead [28]. Its advantage is that it can avoid jumping on immediate high gains that are also extremely costly, and instead favor less costly actions that yield long-term gain.

Solving for lookahead policy directly is impractically costly, as every possible state in each step ahead needs to be considered. We tackle this in two ways. First, we use a Monte Carlo optimization method from reinforcement learning literature [29]. Second, instead of modeling the state space, we formulate a probability dependent only on the observations and actions. It is considerably lower dimensional and generalizes to any object descriptor and robot platform.

Monte Carlo tree search (MCTS) [30] has become popular for real-time decisions in AI. It is an online alternative to

dynamic programming and uses repeated simulations to construct a tree in a best-first order. Kocsic and Szepesvári [31] showed that tree policy using the UCT (Upper Confidence bounds applied to Trees) guarantees asymptotic optimality. Feldman and Domshlak [32] introduced BRUE, a purely exploring MCTS that guarantees exponential cost reduction. Silver and Veness [33] extended MCTS to partially-observable models. MCTS has been used for game solving [34] and belief-space planning in robotics [35], [36], [37], but has not been applied to manipulation.

III. PROBLEM FORMULATION

Our goal is to adaptively select a minimum sequence of end-effector poses to correctly recognize an object. The input is contact XYZ only, given by enclosure grasps, useful for sensing the volume and the global shape of an object [15].

1) *Recognition Descriptor*: We focus on optimizing the sequence of poses and use an existing tactile object descriptor [11] for recognition. We cap the sequence at $t = 1 : T$ poses. At time t , a grasp provides n contact points, resulting in $\binom{n}{3}$ triangles z_t [11]. Observed triangles $z_{1:t}$ are binned into a 3D histogram h_t . The three dimensions represent triangle parameters, *e.g.* two sides and an angle. The histogram is the object descriptor input to a classifier.

2) *Active Probability*: In between two consecutive observations z_t and z_{t+1} , the end-effector moves by some action a_{t+1} , which we model as $a \in SE(3)$, the translation and quaternion from the current wrist pose to a new one. As the hand moves, the previous ending pose becomes the next starting pose, hence removing the need for world frame coordinates for both the hand and the object pose. Let $c_m(a) \in [0, 1]$ denote the movement cost incurred by a .

To model the recursive chain of $z_t \rightarrow a_{t+1} \rightarrow z_{t+1} \rightarrow \dots$, we write the probability distribution $p(z_{t+1}|z_t, a_{t+1}, y)$. It is in terms of the next observation z_{t+1} , conditioned on the current observation z_t , the next action a_{t+1} that leads to z_{t+1} , and the unknown object class y .

3) *Training and Test*: During training (Alg. 1), two things are learned for each object: its histogram descriptor h and its $p(z_{t+1}|\cdot)$ distribution above. Training is done by moving the robot hand in a grid [11] around the object. Actions and observations are recorded to compute the two items. An action is defined between two wrist poses; n poses yield n^2 actions. Additionally, we train a support vector machine (SVM) classifier (Sec. V-A) on the descriptors. The SVM gives $p(y|h)$, the probability of class y given a histogram.

At test time, the robot chooses its next grasp a_{t+1} (Sec. IV-B) based on state $x_t = h_t$. Given the current histogram h_t , we can obtain the recognition probability $p(y|h_t)$.

Algorithm 1: Training stage

```

1 for each object do
2   define a grid of wrist poses P wrt object; execute P;
3   store triangle observations  $\{z\}$  from contacts;
4   store tallies of observations  $\{z\}$  per pose;
5   compute histogram descriptor  $h$ ;

```

Problem (Active Tactile Recognition). *Given an object with unknown class y , an initial information state x_0 , and a planning horizon of T steps, choose a control policy π to optimize the cost, which trades off between movement cost and misclassification probability:*

$$\min_{\pi} C_T(\pi) \triangleq \frac{\lambda}{T} \mathbb{E} \left[\sum_{t=0}^{T-1} c_m(\pi(x_t)) \right] + (1-\lambda) \mathbb{P}(\hat{y}_T \neq y) \quad (1)$$

where π maps current state x_t to next action a_{t+1} , and $\hat{y}_T = \arg \max_y p(y|h_T)$ is the maximum likelihood estimate of the object class. $\mathbb{P}(\hat{y}_T \neq y) = 1 - \max_y p(y|h_T)$ is the misclassification probability. $\lambda \in [0, 1]$ determines the relative importance of incurring movement cost (first term) to gather more information vs. making an incorrect recognition (second term).

IV. PROPOSED APPROACH

A. Markov Decision Process (MDP)

The problem can be represented by a finite-horizon MDP defined by $(\mathcal{X}, \mathcal{A}, \mathcal{T}, \mathcal{G}_t)$. \mathcal{X} is the state space. \mathcal{A} is a finite set of possible actions. The transition function

$$\mathcal{T}(x_t, a, x_{t+1}) \triangleq \sum_y p(z_{t+1} | z_t, a, y) p(y | h_t)$$

advances from state x_t to x_{t+1} given action a . Histogram $h_t \in \mathcal{X}$; z_t is determined by (x_t, a) ; and $h_{t+1} = (h_t, z_{t+1})$ initializes x_{t+1} .

$$\mathcal{G}_t(x_t, a, x_{t+1}) \triangleq \begin{cases} \frac{\lambda}{T} c_m(a), & 0 \leq t < T \\ (1-\lambda)(1 - \max_y p(y | x_t)), & t = T \end{cases}$$

is the stage cost. This corresponds to the two terms in Eqn. 1.

An MDP can be represented by a graph. Each state is a node, each action is an edge, and \mathcal{T} describes how to move from one node to another. \mathcal{G} is the cost associated with an edge. The graph is generated only at test time.

At the start of the process, a random action a_0 is selected. This generates observation z_0 , which initializes histogram h_0 at state x_0 . Then, we advance through the states by simulating possible actions and following the transitions \mathcal{T} to create nodes and edges (Sec. IV-B). Each node's z_t adds to the histogram. Over time, the histogram h_t is incrementally filled and resembles the true histogram from training, at which $p(y|h_t)$ would indicate a high probability.

1) *Relating Observations and Actions:* When we create a new node with z_t at test time, we do not make robot movements to observe an actual z_t , because moving after every edge would require hundreds of movements for the entire tree, making the search impractically slow. Instead, we rely on observations from training.

At the core of our approach is the relationship between observations and actions, modeled by $p(z_{t+1}|z_t, a_{t+1}, y)$, computed from training data (Sec. III-2). We trust this relationship to be reliable during training and carries over to test time, at which we directly sample this probability from training. Note that $p(z_{t+1}|\cdot)$ is independent of the state space \mathcal{X} , which is a probability of histograms $p(h_t)$,

high-dimensional (1000D) and exponential in search time. This independence and direct sampling from training allows $p(z_{t+1}|\cdot)$ to be computed quickly at test time (Sec. VI-A).

B. Monte Carlo Tree Search (MCTS)

This section describes how we generate a graph for the MDP at test time and select an optimal policy π from it. We represent the graph by a tree and use a Monte Carlo method. The reader should refer to [30] for an overview of MCTS and [29] for policy search. The accompanying video animates the concept. A simple example is shown in Fig. 2 and walked through in Sec. IV-B.3.

An optimal policy π outputs an optimal action sequence, which is defined as a path with maximum reward (or equivalently, minimum cost) from the root to a leaf. We seek a path that minimizes the objective cost in Eqn. 1.

After the root is created, Monte Carlo simulations select actions and follow \mathcal{T} to create new edges and nodes. Each simulation creates one new node. After a number of simulations, the tree is well populated, and the optimal path is selected. Each tree depth is a time step, with root at $t = 0$ and leaves at $t \leq T$, a defined horizon, or max tree depth.

1) *Choosing the Next Action a_{t+1} :* At time t , node x_t , the next action a_{t+1} is selected as follows. In an MDP that allows multiple actions per node (known as a multi-arm bandit problem [30]), the choice of an action faces an exploration-exploitation dilemma. Exploring new or less-seen actions generates unseen parts of the tree, making use of more training data. Re-visiting high-reward actions exploits branches that at the moment seem more likely to be optimal.

Balancing this dilemma ensures narrowing down the answer while keeping an open mind to see all of the tree. For contrast, greedy policies always exploit the highest-reward action and ignore the exploration half. We use the UCT [31] upper confidence bound to select actions to balance this dilemma. At a node x_t at depth t , the next action a_{t+1} is:

$$a_{t+1} = \arg \max_a \left((1 - C_{t_a}) + c \sqrt{\frac{2 \ln N}{N_a}} \right) \quad (2)$$

where the $C_{t_a} \in [0, 1]$ is $C_T(\pi)$ in Eqn. 1, computed in previous simulations and stored in node x_t . It is the cost of an available action edge a at the x_t (see backpropagation in Sec. IV-B.3). $1 - C_{t_a}$ is the reward. N is the number of times the node has been visited, and N_a is the number of times action a has been followed from the node.

The first term is exploitation; it favors actions with a high existing reward. The second term is exploration; it penalizes actions that have been followed many times. The two terms are balanced by weight c , picked by hand. The result is a well-explored bushy tree. In entirety, the bound selects an action a_{t+1} that minimizes cost C_{t_a} . Together with other actions on a path chosen this way from root to leaf, this minimizes the objective cost C_T in Eqn. 1.

2) *Inferring the Next Observation:* Given an action a_{t+1} , the next observation z_{t+1} is sampled from training data:

$$z_{t+1} \sim p(z_{t+1}|z_t, a_{t+1}) = \sum_y p(z_{t+1}|z_t, a_{t+1}, y) p(y|h_t) \quad (3)$$

This reflects the mapping of the MDP transition function $\mathcal{T}(x_t, a_{t+1}) \rightarrow x_{t+1}$. It describes how to move to the next node x_{t+1} , given the current node x_t and next action a_{t+1} . The class y is marginalized out, since the true y is unknown.

3) *Tree Search Procedure*: Now we put the pieces together and describe the procedure of each tree search simulation in Algs. 2, 3, 4. Alg. 2 outlines the top-level test stage procedure. Algs. 3 and 4 outline the tree search and tree policy. We will walk through a 5-node tree in Fig. 2, with horizon $T = 3$, the 5 nodes generated from 5 simulations.

Starting with an empty tree, some action is randomly selected and produces obs443 in Fig. 2 from training data. This initializes the root at $t = 0$, with 1 observation in histogram h_0 , which happens to be 0.89 distance from the closest object in training.

We will describe one full simulation. Each simulation starts at the root at depth $t = 0$ and must traverse a single path downward until the leaves at horizon depth $t = T$. Each depth contains actions a_t and nodes x_t . The intuition of a path from root to leaf in the real world is a sequence of T actions for the robot to execute.

In each simulation, one new node is created via the choice of a_{t+1} and z_{t+1} (Eqns. 2, 3), outlined in Alg. 4 treePolicy. This means early simulations cannot reach depth T via existing nodes, since the tree is still shallow. In Alg. 3, a recursive function treeSearch traverses the tree, incrementing in depth t (line 4). As long as a node exists (line 2), the tree policy is called (line 3) to continue down. When a desired node does not exist, it is created (line 8), which concludes the one node created in the current simulation, and this ends the tree policy. The rollout policy follows (line 9) and continues to depth T by randomly selecting a_{t+1} at each layer.

In Fig. 2, simulation 1, at $t = 0$, action $a_1 = p15$ is selected and produces $z_1 = obs3$. Since a node with obs3 at depth $t = 1$ does not yet exist, it is created, and this ends the tree policy. The rollout policy selects random a_2 and a_3 that produce temporary nodes x_2 and x_3 , not shown. The rollout policy operates on a temporary subtree that will be discarded at the end of the simulation. When the rollout policy reaches horizon T , the histogram h_T accumulated from observations $z_{0:T}$ on the path we took is fed to the classifier, which outputs the misclassification cost $\mathbb{P}(\hat{y}_T \neq y)$ in Eqn. 1.

We then trace the path back up to root and backpropagate this cost to store in each node on the path, as follows. At each depth t , reward $1 - C_{t_a}^a$ is updated for action a at node x_t . This is a standard discounted reward, $Q_a = Q_a + (r_a - Q_a)/N_a$, where Q_a is the node’s existing discounted reward, and r_a is the raw subtree reward (Alg. 3 lines 4–6).

In addition to backpropagation, the objective continues to be computed. At each depth t , edge a_t is accumulated to movement cost $c_m(a_t)$ in Eqn. 1. Intuitively, the reward at node x_t is the sum reward of actions on its current child path. When we arrive back at the root, the entire objective C_T has now been computed and stored to root under action p15 (Fig. 2). This concludes simulation 1. In the next 4 simulations, 4 more nodes are created and rewards computed similarly. The more simulations, the bushier the tree, and the

deeper the branches reach.

After many simulations, the tree search ends by extracting the optimal path (Alg. 2 line 10). Starting at the root, simply follow the highest-reward edges downward. This path defines an action sequence that minimizes C_T . The length of the sequence is $\leq T$, as some branches may not reach T , e.g. two right branches in Fig. 2. The optimal sequence is executed on a robot to obtain actual observations for recognition.

Algorithm 2: Test stage

```

1 object location given;
2 load training probabilities  $P_{train}$ ;
3 superimpose training poses  $\{p^o\}$  onto test object;
4 move robot to a pose  $p_0$  that contacts object;
5 close grippers; record observation  $z_0$ ; compute histogram  $h_0$ ;
6 for each tree do
7    $node_0 = \text{initNewTreeRoot}(h_0)$ ;
8   for each simulation do
9      $\text{treeSearch}(node_0, z_0)$ ;
10    actions = select max-reward root-to-leaf path;
11    for each action  $a_t = a_1 : a_T$  in actions do
12      move robot to  $a_t$ ; close grippers;
13      record observation  $z_t$ ; update histogram  $h_t$ ;
14     $z_0 = z_T$ ;  $h_0 = h_T$ ;
```

Algorithm 3: Tree search

```

1 function treeSearch ( $node_t, z_t$ );
2 if  $node_t$  exists in tree then
3    $a_{t+1}, z_{t+1} = \text{treePolicy}(node_t, z_t)$ ;
4   subtreeReward = treeSearch( $node_{t+1}, z_{t+1}$ );
5    $r = (1 - c_m(a_{t+1})) + \text{subtreeReward}$  // Eqn. 1
6    $node_t.\text{updateReward}(a_{t+1}, r)$ ;
7 else
8   create  $node_t$ ;
9    $r = \text{rolloutPolicy}(node_t, t)$  // Eqn. 1  $\mathbb{P}(\hat{y}_T \neq y)$ 
10 return r;
```

Algorithm 4: Tree policy

```

1 function treePolicy ( $node_t, z_t$ );
2  $a_{t+1} = \text{argmax}_a \text{UCT}(node_t.C_a, node_t.N_a)$  // Eqn. 2
3  $h_t = node_t.h$ ;
4 for each class  $y$  do
5    $p_y = p(z_{t+1}|z_t, a_{t+1}, y) = \text{sampleFrom } P_{train}(z_t, a_{t+1}, y)$ ;
6  $z_{t+1} = \text{marginalizeY}(\{p_y\})$  // Eqn. 3
7 return  $a_{t+1}, z_{t+1}$ ;
```

C. Implementation

The recognition of test objects is performed by alternating between MCTS and robot action execution. Note that we use the term iteration to refer to one tree search and action execution, e.g. two iterations means a tree search, an execution, a second tree search, and a second execution.

The object pose is assumed known and fixed. A first wrist pose is randomly selected from the training data, which store

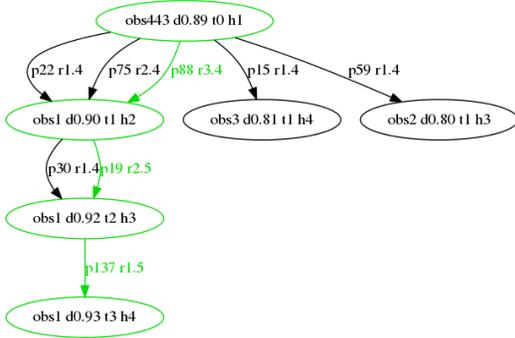


Fig. 2: A small example tree. Max-reward path highlighted. Node label is observation name, nearest neighbor distance, tree depth t , and number of items in histogram h_t . NN distance is inversely proportional to $p(y|h_t)$. Edge label is action name a_{t+1} and reward.

poses with respect to the object center. The observation z_0 is computed and initializes the root node of the first tree. The tree is generated and produces a sequence of relative wrist poses. This sequence is executed on the robot, with enclosure grasps onto the object at each wrist pose. Actual observations are taken, and a histogram is built and fed to the classifier. This completes one iteration. Then, the old tree is discarded, and the latest histogram initializes the root node of a new tree. MCTS is performed again on the new tree.

In order to generalize across objects, triangle observations z_t are discretized to their histogram bin centers. This is required to compute \mathcal{T} , which needs probabilities for the conditioned z_t for every object y . Otherwise, a triangle from one object might not exist in another to provide this probability. Histogram bin sizes are chosen as in [11].

For the movement cost $c_m(a_t)$, we computed the L_2 distance for translation and the angle for rotation, then normalized each to $[0, 1]$ and weighed both equally. We used $c = 1$ to weigh exploration and exploitation equally and $\lambda = 0.5$ to weigh movement and misprediction costs equally.

V. ANALYSIS USING A PHYSICS ENGINE

We validated the active recognition in a physics engine. The purpose is to analyze the method itself, without external constraints of a specific robot platform, such as joint and workspace limits. This lets us evaluate the core algorithm using an end-effector with no unreachable poses.

The tactile hardware is a RightHand Robotics ReFlex Beta Hand, which has 9 barometric pressure sensors on each of 3 fingers, 27 total. Each enclosure grasp typically gives non-zero values on 3–6 sensors. Fig. 3(a) shows objects used.

The XYZ positions of the non-zero sensors are used to compute the descriptor for recognition. We used the 3D histogram of triangles [11] as descriptor, as mentioned in Sec. III. A triangle requires three parameters to be uniquely described. The three side lengths of a triangle are denoted l_0, l_1, l_2 from large to small; similarly for the angles a_0, a_1, a_2 . The following results use the l_0, l_1, a_0 parameterization. Because the triangles are a relative measure, the descriptor is independent of object pose and movement.

For physics engine results, we built a stack in Gazebo for the hand’s tactile capabilities, including its guarded

enclosure, which closes all fingers and stops each finger when a sensor on the finger is in contact. Sensor values were simulated as Boolean contacts. To simulate wrist movement, we teleported the wrist to specified poses. This dramatically reduced the time required, which would otherwise involve motion planning and moving the joint of an arm.

We repeated the MCTS and action execution iterations until rewards in all simulations are depleted to 0, which typically takes 7–9 iterations. When all rewards are depleted, there is only one node in the tree; this happens because we do not allow repeated visits to nominal absolute poses, as they do not provide new observations.

A. Baseline Comparison

To illustrate the need for an active selection, we created a baseline. It uses the same training data as the tree search, except it selects poses to move to at uniform random. Fig 3(b) shows example poses chosen by the baseline and the tree policy for the cup; poses from the latter are more intuitive and capture the global shape better.

Fig. 3(a) and Table I show the 3D histograms and distances obtained by each method. Fig. 4(a) shows the progression of recognition through the iterations, in the form of distance to true class. We tried linear SVM, nearest neighbor (NN) with inner product distance, and NN with histogram intersection distance, of which the inner product performed the best. Tree policy performed better for all objects except teapot, which was correct in iterations 1–3 but diverged to mug in 4–9. This is reasonable as teapot, mug, and cup have similarities in the handle. Sphere was recognized by neither; the cause is evident in Fig. 3(a), as both were unable to capture the lowest bins in l_0 .

Fig. 5 shows wrist poses selected around the objects, and Fig. 6 shows contact points obtained in the physics engine by the two methods. Comparing Figs. 4(a), 5, and 6, even though the baseline sometimes recover a better object appearance, its recognition can still be wrong, e.g. bottle, mug. Tree policy recovered better contact cloud and recognition for cup, bowl, toilet paper. The only object that the baseline did better in all three was the teapot, most likely because tree policy tends to select poses at the bottom of the object, but the teapot’s top half provide identifying information.

Even though the baseline’s poses are more evenly distributed, they do not result in better recognition, other than teapot. Fig. 3(a) shows that baseline histograms are more distributed, whereas tree policy’s histogram bins are concentrated in the area lit up in the true histograms. Note that even though some baseline distances in Table I are closer, e.g. cup, toilet paper, its recognition in Fig. 4(b) is incorrect, meaning it is closer to some other class. This means that the tree policy correctly imposes a bias on the wrist poses selected - poses that result in high recognition certainty, as enforced by $\mathbb{P}(y_T \neq y)$ in the objective.

B. Recognition Performance

Fig. 4(b) shows recognition in the form of inner product distance to true class, for different iterations and simulations

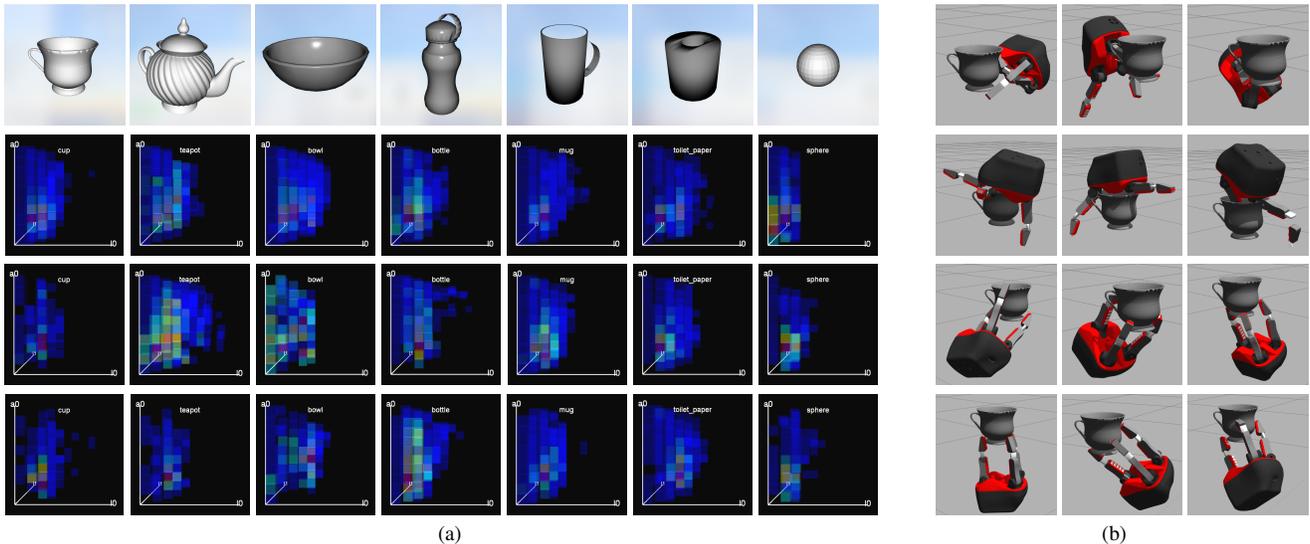


Fig. 3: (a). Synthetic meshes and their 3D histograms for visual comparison. From training (row 2), random baseline (row 3), tree policy (row 4). Note shape similarities between cup, teapot, and mug; bottle and mug; mug and toilet paper. Tree policy results resemble true histograms. Best viewed in color. (b). Wrist poses selected by random baseline (top 2 rows) and tree policy (bottom 2 rows) for the cup.

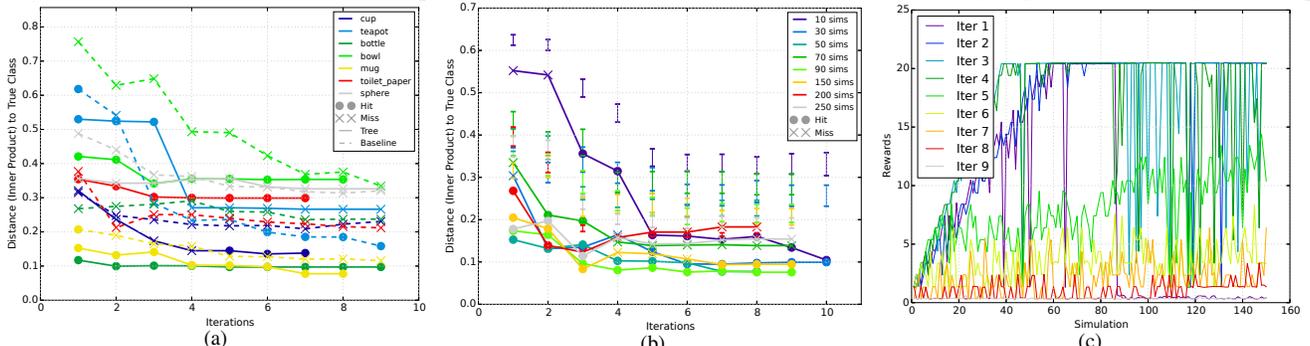


Fig. 4: (a). Physics engine results. Distance to true class vs. iteration, for baseline (dashed) and tree policy (solid). Dots and crosses show hits and misses. Each curve is an object. Tree policy has more hits and the closest distances. (b). Similar format, for different simulations for mug. Each curve is a simulation setting. Error bars are mean and variance in distances to all objects. (c). Rewards vs. simulations for mug; similar for all objects. Each curve is an iteration. Rewards diminish over iterations as unvisited high-probability poses are exhausted.

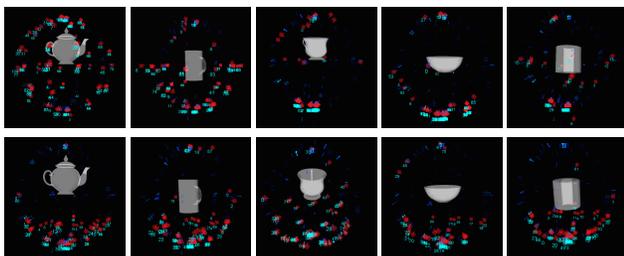


Fig. 5: Poses selected by random baseline (top) and tree (bottom), for teapot, mug, cup, bowl, and toilet paper.

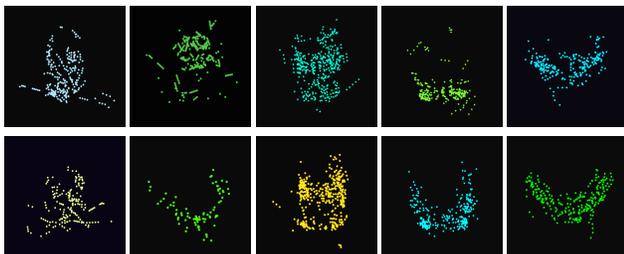


Fig. 6: Actual contacts obtained in physics engine from baseline (top) and tree (bottom), for bottle, teapot, mug, cup, and bowl.

	cup	teapot	bowl	bottle	mug	tlt ppr	sphere
Baseline	0.227	0.160	0.375	0.237	0.121	0.224	0.317
Tree	0.137	0.266	0.354	0.097	0.078	0.299	0.326

TABLE I: Test-time histogram distances $\in [0, 1]$ to ground truth histogram from training. Distances are from the last common iteration in Fig. 4(a). Columns correspond to Fig. 3(a).

for one object. Distances decrease as iterations increase, since the action executions provide increasingly descriptive histograms at the root of new trees.

The error bars show means and variances in distances to all objects. The difference between the error bars and distances to true class indicate that recognition converges early on, in as few as 2 iterations. That is 10–31 poses, average 25 poses across all simulation settings shown. The number of available actions at each node is 1063, pooled across all 7 objects’ trained poses. This shows that the tree successfully finds the few essential poses to fill the most discriminative features in the descriptor.

As simulations per tree increase, distances do not necessarily decrease, nor does recognition accuracy increase. 30–90 simulations perform better than 150–250 simulations. This can be due to our restriction of visiting a pose only

once. As the available poses are exhausted, the tree policy cannot find an edge at some nodes and has to switch to rollout policy, by design. This is further reflected in the rewards plot in Fig 4(c). Before 90 simulations, rewards steadily increase as the tree is more explored; after 90 simulations, deep troughs appear. The troughs are probably due to exhaustion of unvisited poses with high probability on a tree path. Similarly, rewards diminish as iterations increase, when fewer poses are available, eventually collapsing to a one-node tree in the last iteration where reward is 0.

Table II shows per-iteration predictions and distances for the cup. The cup starts as the third NN and moves up to first NN in iteration 5. It is often reasonably confused with a mug. The baseline always recognized the cup as a mug in all 9 iterations, for all 3 distance metrics.

Iter	1	2	3	4	5	6	7
Moves	14	14	17	18	7	5	2
Contacts	52	61	52	68	25	22	6
1st NN	teapot	mug	mug	mug	cup	cup	cup
1st dist	0.230	0.142	0.127	0.141	0.145	0.135	0.138
2nd NN	mug	cup	cup	cup	mug	mug	mug
2nd dist	0.284	0.237	0.174	0.145	0.156	0.158	0.158
3rd NN	cup	teapot	teapot	bottle	bottle	bottle	bottle
3rd dist	0.321	0.248	0.229	0.232	0.231	0.225	0.227

TABLE II: Tree result on cup

C. Number of Moves for Recognition

Table III shows the number of moves per iteration for all objects. Boldface shows the iteration in which recognition starts being correct, corresponding to Fig. 4(a). Teapot was correct in iterations 1–3 and diverged to mug. All other objects stayed correct. These are the upper bound moves for recognition, as we only ran recognition after each iteration, not after every move. Most objects were recognized within 16 moves, a large improvement over hundreds in [11].

Iteration	1	2	3	4	5	6	7	8	9
cup	14	14	17	18	7	5	2		
teapot	16*	13	16	15	8	5	4	1	1
bottle	16	12	18	17	8	5	4	2	1
bowl	15	14	16	16	7	4	3	1	
mug	14	14	15	14	11	6	4	1	
toilet paper	13	18	15	15	8	5	1		
sphere	16	13	17	14	10	6	5	3	1

TABLE III: Upper bound number of poses to recognize correctly

VI. REAL ROBOT EXPERIMENTS

On the real robot, we compare with a greedy baseline instead of random. Fig. 7(b) shows the experiment setup. We mounted the ReFlex Beta hand on a Baxter. An object is held fixed on a table. We trained 5 objects (Fig. 7(a)) on the real robot for active instance-based recognition. Note transparent objects pose significant challenge for vision systems.

The mug, bottle, jar, bowl, and glass were each trained with 34, 60, 39, 30, and 50 end-effector poses. With discretization (0.06 meters in translation, 0.05 in quaternion), this resulted in 138 possible actions at each MCTS node at test time. The goal is to recognize in considerably fewer poses. This would mean the active selection is able to select poses with discriminating features.

At test time, we ran MCTS (Sec. IV-B) to actively predict a sequence of end-effector poses. Then, the Baxter arm

autonomously moves to those poses, using motion planning in ROS Moveit for collision avoidance.

For a baseline, we compared with greedily selecting the immediate minimizer of the objective, *i.e.* zero step lookahead, equivalent to horizon=1. For tree policy, we used horizon=5. This means tree policy had up to 5 poses per iteration; greedy had 1. We used 20 simulations per iteration for both. Both were run until the recognition was correct for 3 consecutive iterations, some further until the distances leveled off. Results are in Table IV and Fig. 7(c). Example grasps selected by tree policy are in Fig. 8. A footage with per-move distances is in the accompanying video.

Some poses selected by either method were not successfully planned, due to joint limits and collisions in the workspace. Fig. 7(c) x-axis is the raw number of poses selected. Table IV shows the number of successful moves.

Tree policy recognized in significantly fewer iterations and shorter time in most cases. All objects were first recognized correctly in under 10 moves, significantly fewer than training. Greedy never recognized the glass correctly, always as mug. Tree policy recognized it twice in a row and then flip-flopped between glass and mug.

Object	# Iters until Correct				# Moves until Correct					
	T1	T3	G1	G3	T1	T3	T33	G1	G3	G33
jar	1	3	11	20	2	10	10	6	8	14
bottle	1	3	5	26	2	8	8	2	4	12
mug	4	6	3	9	9	12	12	1	3	3
bowl	3	-	9	-	5	8	-	3	5	-
glass	2	-	-	-	7	-	-	-	-	-

TABLE IV: Number of iterations and move it took to recognize correctly on real robot. T: tree policy, G: greedy. T1/G1: first time SVM recognizing correctly; T3/G3: SVM correct 3 times in a row; T33/G33: all 3 metrics correct 3 times in a row. - denotes never.



Fig. 8: Real robot actions selected by tree policy at test time.

A. Running Time

The running time for the tree search is directly proportional to horizon T . Each tree simulation takes 0.5 seconds for $T = 20$, 0.1 seconds for $T = 5$, and 0.02 seconds for $T = 1$. Times reported are on an Intel Xeon 3.6GHz quad-core desktop simultaneously running the rest of the experiment software. It can be improved by array access.

Even though the greedy approach ($T = 1$) took shorter time per iteration, it took many more iterations before correct recognition (Table IV). The reason is that $T = 1$ generates only one pose per iteration, and when the pose is unfeasible due to joint limits or collision, the iteration is wasted. Overall, the tree policy took significantly shorter time.

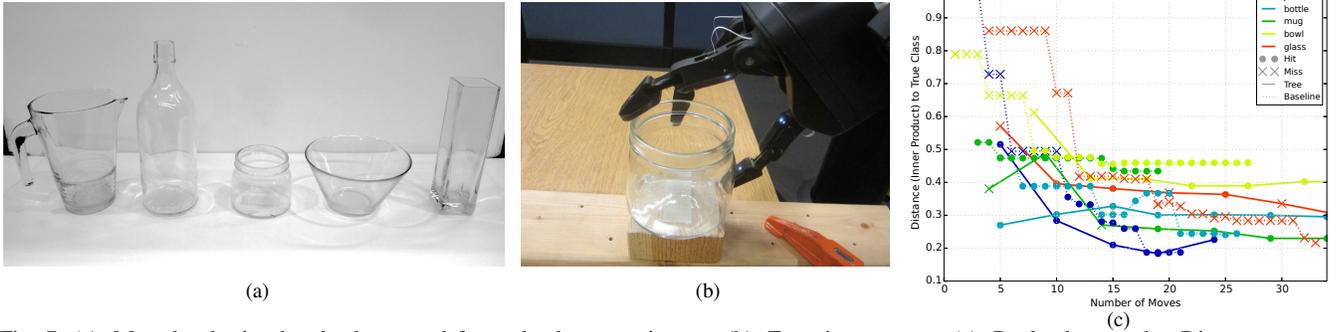


Fig. 7: (a). Mug, bottle, jar, bowl, glass, used for real robot experiments. (b). Experiment setup. (c). Real robot results. Distance to true class vs. number of poses, for baseline (dashed) and tree policy (solid). Dots and crosses show hits and misses. Each curve is an object.

VII. CONCLUSION

We described an algorithm for actively selecting a sequence of end-effector poses for the objective of confident object recognition. We formulated the problem as a MDP and associated tactile observations with relative wrist poses in training, which allows the next desired actions to be predicted by observations alone at test time. The method outperforms greedily selected poses in a physics engine and on a real robot.

An improvement to optimize recognition even more directly is to select actions that would produce the most salient features in the descriptor. Analysis methods exist for finding the most discriminative features in a classifier. The histogram descriptor makes this easy; each feature is simply a bin, which we already use as discretized observations z . To select the most salient action, simply select z_{t+1} that maximizes saliency in addition to recognition confidence.

REFERENCES

- [1] P. C. Gaston and T. Lozano-Perez, "Tactile recognition and localization using object models: The case of polyhedra on a plane," *PAMI*, 1984.
- [2] W. E. L. Grimson and T. Lozano-Perez, "Model-based recognition and localization from sparse range or tactile," *IJRR*, 1984.
- [3] P. K. Allen and P. Michelman, "Acquisition and interpretation of 3-D sensor data from touch," *TRO*, 1990.
- [4] T. J. Prescott, M. E. Diamond, and A. M. Wing, "Active tactile sensing," *Philosophical Transactions of the Royal Society B*, 2011.
- [5] J. R. Flanagan, M. C. Bowman, and R. S. Johansson, "Control strategies in object manipulation tasks," *Current Opinion in Neurobiology*, 2006.
- [6] P. Allen and R. Bajcsy, "Object recognition using vision and touch," *International Joint Conference on Artificial Intelligence*, 1985.
- [7] S. A. Stansfield, "A robotic perceptual system utilizing passive vision and active touch," *IJRR*, 1988.
- [8] P. K. Allen, "Integrating vision and touch for object recognition tasks," *IJRR*, 1988.
- [9] J. L. Schneider, "An objective tactile sensing strategy for object recognition and localization," in *ICRA*, 1986.
- [10] R. Bajcsy and M. Campos, "Active and exploratory perception," *CVGIP: Image Understanding*, 1992.
- [11] M. M. Zhang, M. Kennedy, M. A. Hsieh, and K. Daniilidis, "A triangle histogram for object classification by tactile sensing," in *IROS*, 2016.
- [12] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6D object pose and predicting next-best-view in the crowd," in *CVPR*, 2016.
- [13] J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, and G. S. Sukhatme, "Interactive Perception: Leveraging Action in Perception and Perception in Action," *arXiv:1604.03670*, 2016.
- [14] R. Bajcsy, S. J. Lederman, and R. L. Klatzky, "Object exploration in one and two fingered robots," in *ICRA*, 1987.
- [15] S. J. Lederman and R. L. Klatzky, "Hand movements: a window into haptic object recognition," *Cognitive Psychology*, 1987.
- [16] D. M. Siegel, "Finding the pose of an object in a hand," in *ICRA*, 1991.
- [17] H. Maekawa, K. Tanie, K. Komoriya, M. Kaneko, C. Horiguchi, and T. Sugawara, "Development of a finger-shaped tactile sensor and its evaluation by active touch," in *ICRA*, 1992.
- [18] K. Hsiao, L. Kaelbling, and T. Lozano-Perez, "Grasping POMDPs," in *ICRA*, 2007.
- [19] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *IROS*, 2009.
- [20] K. Hsiao, L. P. Kaelbling, and T. Lozano-Perez, "Task-driven tactile exploration," in *RSS*, 2010.
- [21] H. P. Saal, J.-A. Ting, , and S. Vijayakumar, "Active estimation of object dynamics parameters with tactile sensors," in *IROS*, 2010.
- [22] P. Hebert, T. Howard, N. Hudson, J. Ma, and J. W. Burdick, "The next best touch for model-based localization," in *ICRA*, 2013.
- [23] D. Golovin, A. Krause, and D. Ray, "Near-optimal bayesian active learning with noisy observations," in *NIPS*, 2010.
- [24] D. Golovin and A. Krause, "Adaptive submodularity: Theory and applications in active learning and stochastic optimization," *Journal of Artificial Intelligence Research (JAIR)*, 2011.
- [25] S. Javdani, M. Klingensmith, J. A. Bagnell, N. S. Pollard, and S. S. Srinivasa, "Efficient touch based localization through submodularity," in *ICRA*, 2013.
- [26] Z. Pezzementi, E. Plaku, C. Reyda, and G. Hager, "Tactile-object recognition from appearance information," *TRO*, 2011.
- [27] K. Hausman, C. Corcos, J. Miller, F. Sha, and G. S. Sukhatme, "Towards interactive object recognition," in *3rd Workshop on Robots in Clutter: Perception and Interaction in Clutter, IROS*, 2014.
- [28] W. B. Powell, *Approximate Dynamic Programming*. John Wiley & Sons, Inc, 2010, ch. 6.
- [29] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *Journal of Artificial Intelligence Research*, 1996.
- [30] C. Browne, E. Powley, D. Whitehouse, S. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of Monte Carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in Games*, 2012.
- [31] L. Kocsis and C. Szepesvári, "Bandit based Monte-Carlo Planning," in *European Conference on Machine Learning (ECML)*, 2006.
- [32] Z. Feldman and C. Domshlak, "Simple regret optimization in online planning for markov decision processes," *JAIR*, 2014.
- [33] D. Silver and J. Veness, "Monte-Carlo Planning in Large POMDPs," in *NIPS*, 2010.
- [34] X. Guo, S. Singh, H. Lee, R. Lewis, and X. Wang, "Deep learning for real-time Atari game play using offline Monte-Carlo tree search planning," in *NIPS*, 2014.
- [35] K. Hauser, "Randomized Belief-Space Replanning in Partially-Observable Continuous Spaces," in *Algorithmic Foundations of Robotics IX*, 2011.
- [36] J. Nguyen, N. Lawrance, and S. Sukkarieh, "Nonmyopic Planning for Long-term Information Gathering with an Aerial Glider," in *ICRA*, 2014.
- [37] M. Lauri and R. Ritala, "Stochastic Control for Maximizing Mutual Information in Active Sensing," in *ICRA Workshop on Robots in Homes and Industry*, 2014.