
Active 3D Shape Reconstruction from Vision and Touch

Edward J. Smith^{1,2,*}

David Meger²

Luis Pineda¹

Roberto Calandra¹

Jitendra Malik^{1,3}

Adriana Romero-Soriano^{1,2,†}

Michal Drozdal^{1,†}

¹ Facebook AI Research ² McGill University ³ University of California, Berkeley

Abstract

Humans build 3D understandings of the world through *active object exploration*, using jointly their senses of vision and touch. However, in 3D shape reconstruction, most recent progress has relied on *static datasets* of limited sensory data such as RGB images, depth maps or haptic readings, leaving the active exploration of the shape largely unexplored. In *active touch sensing* for 3D reconstruction, the goal is to actively select the tactile readings that maximize the improvement in shape reconstruction accuracy. However, the development of deep learning-based active touch models is largely limited by the lack of frameworks for shape exploration. In this paper, we focus on this problem and introduce a system composed of: 1) a haptic simulator leveraging high spatial resolution vision-based tactile sensors for active touching of 3D objects; 2) a mesh-based 3D shape reconstruction model that relies on tactile or visuotactile signals; and 3) a set of data-driven solutions with either tactile or visuotactile priors to guide the shape exploration. Our framework enables the development of the first fully data-driven solutions to active touch on top of learned models for object understanding. Our experiments show the benefits of such solutions in the task of 3D shape understanding where our models consistently outperform natural baselines. We provide our framework as a tool to foster future research in this direction.

1 Introduction

3D shape understanding is an active area of research, whose goal is to build 3D models of objects and environments from limited sensory data. It is commonly tackled by leveraging partial observations such as a single view RGB image [63, 71, 23, 46], multiple view RGB images [16, 25, 30, 32], depth maps [65, 78] or tactile readings [7, 51, 62, 47, 40]. Most of this research focuses on building shape reconstruction models from a *fixed set* of partial observations. However, this constraint is relaxed in the *active sensing* scenario, where additional observation can be acquired to improve the quality of the 3D reconstructions. In *active vision* [4], for instance, the objective can be to iteratively select camera perspectives from an object that result in the highest improvement in quality of the reconstruction [80] and only very recently the research community has started to leverage large scale datasets to learn exploration strategies that generalize to unseen objects [81, 5, 43, 53, 29, 82, 52, 6].

Human haptic exploration of objects both with and without the presence of vision has classically been analysed from a psychological perspective, where it was discovered that the developed tactile exploration strategies for object understanding were demonstrated to not only be ubiquitous, but also highly tailored to specific tasks [39, 35]. In spite of this, deep learning-based data-driven approaches to *active touch* for shape understanding are practically non-existent. Previous haptic exploration works consider objects

*Correspondence to: ejsmith@fb.com and edward.smith@mail.mcgill.ca

†Equal Contribution

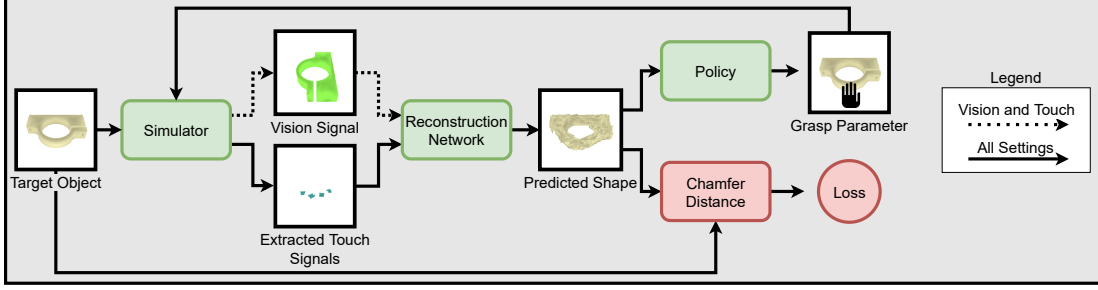



Figure 1: An overview of our active touch exploration framework. Given a 3D model, the simulator extracts touch and vision signals that are fed to the reconstruction model. The reconstruction model predicts a 3D shape that is used as an input to a policy model that decides where to touch next. The policies are trained to select grasps which minimize the Chamfer Distance.


independently, and build uncertainty estimates over point clouds, produced by densely touching the objects surface with point-based touch sensors [7, 77, 28, 18, 48]. These methods do not make use of learned object priors, and a large number of touches sampled on an object’s surface (over 100) is necessary to produce not only a prediction of the surface but also to drive exploration. However, accurate estimates of object shape have also been successfully produced with orders of magnitude fewer touch signals, by making use of high resolution tactile sensors such as [79], large datasets of *static* 3D shape data, and deep learning – see *e.g.* [58, 72]. Note that no prior work exists to learn touch exploration by leveraging large scale datasets. Moreover, no prior work explores active touch in the presence of visual inputs either (*e.g.* an RGB camera).

Combining the recent emergence of both data-driven reconstruction models from vision and touch systems [58, 72], and data-driven active vision approaches, we present a novel formulation for active touch exploration. Our formulation is designed to easily enable the use of vision signals to guide the touch exploration. First, we define a new problem setting over active touch for 3D shape reconstruction where touch exploration strategies can be learned over shape predictions from a learned reconstruction model with strong object priors. Second, we develop a simulator which allows for fast and realistic grasping of objects, and for extracting both vision and touch signals using a robotic hand augmented with high-resolution tactile sensors [79]. Third, we present a 3D reconstruction model from vision and touch which produces mesh-based predictions and achieves impressive performance in the single view image setting, both with and without the presence of touch signals. Fourth, we combine the simulator and reconstruction model to produce a tactile active sensing environment for training and evaluating touch exploration policies. The outline for this environment can be viewed in Figure 1.

Over the provided environment, we present a series of data-driven touch exploration models that take as an input a mesh-based shape reconstruction and decide the position of the next touch. By leveraging a large scale dataset of over 25k CAD models from the ABC dataset [36] together with our environment, the data-driven touch exploration models outperform baseline policies which fail to leverage learned patterns between object shape or the distribution of object shapes and optimal actions. We demonstrate our proposed data-driven solutions perform up to 18 % better than random baselines and lead to impressive object reconstructions relative to their input modality, such as those demonstrated in Figure 2. Our framework, training and evaluation setup, and trained models are publicly available on a GitHub repository to ensure and encourage reproducible experimental comparison³.



(a) Touch only exploration.



(b) Touch exploration with visual prior.

Figure 2: Target objects (top rows) and predicted 3D shapes (bottom rows) after 5 grasps have been selected.

³<https://github.com/facebookresearch/Active-3D-Vision-and-Touch>

2 Related Work

3D reconstruction. 3D shape reconstruction is a popular and well studied area of research with a multitude of approaches proposed across varying 3D representations [13, 67, 19, 54, 71, 23, 12, 75] and input modalities [27, 55, 73, 59]. While still a niche topic, a number of works have attempted to reconstruct 3D shape from only touch signals [7, 51, 62, 47, 40]. These works all assume point based touch signals which supply at most one point of contact of normal information. A limited number of methods have also been proposed for 3D reconstruction from both vision and touch signals [9, 26, 22, 72, 73, 58]. In contrast to these works which use static datasets, we focus on improving 3D reconstruction through tactile exploration. Shape from interaction methods have also been proposed for object reconstruction through hand interactions [44, 68, 49], though in contrast to our work an image of the interaction provides the additional information rather than touch readings.

Active sensing for reconstruction. Active vision aims to manipulate the viewpoint of a camera in order to maximize the information for a particular task [11, 80]. Similar to our setting, active vision can be useful in 3D shape reconstruction [66, 74, 17, 43, 14, 37], but where camera perspectives are planned instead of grasp locations. Only very recently have deep learning active vision approaches been proposed for 3D object reconstruction [81, 5, 43, 53]. Deep learning based active sensing for reconstruction has also recently emerged in the medical imaging domain, where the time spent performing MRI scans has been reduced by learning to select a small number of more informative frequencies over a pre-trained reconstruction model [29, 82, 52, 6].

Touch-based active exploration for shape understanding. Initial works tackling active tactile exploration focused on object recognition. [1–3]. While some of these methods do integrate both vision and touch for active shape understanding, they only focus on object recognition, operate over small datasets of already known and observed objects and do not leverage deep learning tools for learning patterns over large datasets of object shapes. Fleer and Moringen et. al. proposed to learn haptic exploration strategies using a reinforcement learning framework over a recurrent attention model, however these strategies were optimized for object classification and only trained over a dataset of 4 objects with a single, floating, depth-based tactile sensor array [20].

Several prior works have explored the problem of active acquisition of touch signals specifically for surface reconstruction. A common theme among most approaches is to predict object shape using Gaussian processes so that uncertainty estimates are naturally available to drive the selection of the next point to touch, which is performed using a single finger [28, 77, 18]. To improve the speed of shape estimation during tactile exploration Matsubara et. al. [42] consider both uncertainty and travel cost when selecting touches using graph-based path planning. Bierbaum et al. [8] instead drove exploration through a dynamic potential field approach made popular in robot navigation to produce point cloud predictions, and extracted touch signals using a comparatively unrealistic 5 finger robotic hand model in simulation. These methods all make use of point based tactile sensors and use deterministic strategies, which are tuned and evaluated over a very small number of objects. In contrast to these works, our approach is fully data driven, and as such, actions are selected using policies trained and evaluated over tens of thousands of objects. Moreover, predictions are made in mesh space fusing visual information with high resolution tactile signals extracted from objects in simulation using a realistic robot hand.

3 Active touch exploration

In our proposed active touch exploration problem, given a pre-trained shape reconstruction model over touch and optionally vision signals, the objective is to select the sequence of touch inputs which lead to the highest improvement in reconstruction accuracy. To tackle this problem, we define an active touch environment that contains a *simulator*, a *reconstruction model* (a pre-trained neural network), and a *loss function*. The simulator takes as input a 3D object shape O together with parameters describing a grasp, g , and outputs touch readings, t , of the 3D shape at the grasp locations along with an RGB image of the object, I . The reconstruction model is a neural network parametrized by ϕ , which takes an input X and produces the current 3D shape estimate \hat{O} as follows: $\hat{O} = f(X; \phi)$. In our setup, we investigate two reconstruction model variants which differ in their inputs: 1) the model only receives a set of touch readings, t , such that $X = t$, and 2) the model receives both a set of touch readings, t , and an RGB image rendering of the shape, I , such that $X = \{t, I\}$. The loss function takes as input the current belief of the object’s shape, \hat{O} , and the ground truth shape, O , and computes the distance between them: $d(O, \hat{O})$. Thus, active touch exploration can be formulated as sequentially selecting the optimal set of K grasp parameters $\{g_1, g_2, \dots, g_K\}$ that maximize the similarity between the ground truth shape O and the




Figure 3: Steps used in the tactile grasping simulator to produce simulated vision and touch signals.

reconstruction output after K grasps \hat{O}_K ⁴:

$$\arg \min_{g_1, g_2, \dots, g_K} d(O, \hat{O}_K), \quad (1)$$

where g_k determine the touch readings fed to the reconstruction network producing \hat{O}_k . We use the Chamfer distance (CD) [63] between the predicted and target surface as the distance metric in our active touch formulation. Further details with respect to the CD are provided in the supplemental materials. In the reminder of this section, we provide details about the touch simulator and the reconstruction model.

3.1 Active touch simulator

Conceptually, our simulator can be described in the five steps depicted in Figure 3. First, an object is loaded onto the environment. Second, an action space is defined around the 3D object by uniformly placing 50 points on a sphere centered at the center of the loaded object. Third, to choose a grasp, one of the points is selected and a 4-digit robot hand is placed such that its 3rd digit lies on the point and the hand’s palm lies tangent to the sphere. Fourth, the hand is moved towards the center of the object until it comes in contact with it. Last, the fingers of the hand are closed until they reach the maximum joint angle, or are halted by the contact with the object. As a result, the simulator produces 4 touch readings (one from each finger of the hand) and one RGB image of the object. Note that each action is defined by its position index on the sphere of 50 actions. This parameterization is selected specifically as it does not require any prior knowledge of the object, other than its center, and in simulation it consistently leads to successful interactions between the hand’s touch sensors and the object’s surface.

In our simulator, all steps are performed in python across the robotics simulator PyBullet [15], the rendering tool Pyrender [41], and PyTorch [50]. For a given grasp and object, the object is loaded into PyBullet [15], along with a Wonik’s Allegro Hand [56] equipped with vision-based touch sensors [38] on each of its fingers, and then the point in space corresponding to the action to be performed is selected and the grasping procedure is performed using PyBullet’s physics simulator. Pose information from the produced grasps is then extracted and used by Pyrender to render both a depth map of the object from the perspective of each sensor and an RGB image of the object from a fixed perspective. The depth maps are then converted into simulated touch signals using the method described in [58]. All steps in this procedure are performed in parallel or using GPU accelerated computing, and as a result across the 50 grasping options of 100 randomly chosen objects, simulated grasps and touch signals are produced in ~ 0.0317 seconds each on a Tesla V100 GPU with 16 CPU cores. Our simulator supports two modes of tactile exploration *grasping* and *poking*. In the grasping scenario, the hand is performing a full grasp of an object using all four fingers. While, in the poking scenario, only the index finger of the hand is used for touch sensing. Further details on this simulated environment are provided the supplemental materials.

3.2 Shape reconstruction

We take a chart-based approach to reconstruction [24], beginning from [58], which used charts to fuse vision and touch signals for shape prediction, and extending it to effectively leverage touch positional information while handling increasing number of touches, and to efficiently predict the object shape from the touch readings. In particular, shape is predicted in the mesh representation by repeatedly deforming a collection of independent mesh surface elements – *i.e.* charts – using a graph convolutional network (GCN). The full pipeline for this reconstruction method is highlighted in Figure 4.

Predicting local shape from touch readings. Available touch signals are passed through a touch convolutional neural network (CNN), which takes a set of touch readings as input, and produces a set of mesh surface elements, referred to as *touch charts*, representing the surface of the touched object where the touches occurred. We train the touch CNN to directly minimize the CD [63] between the predicted touch charts and the local surface at the corresponding touch site. By contrast, [58] predicts local point clouds which are then converted to charts via iterative optimization.

⁴We use a subscript K to indicate that the reconstruction has been obtained after observing K touch readings.

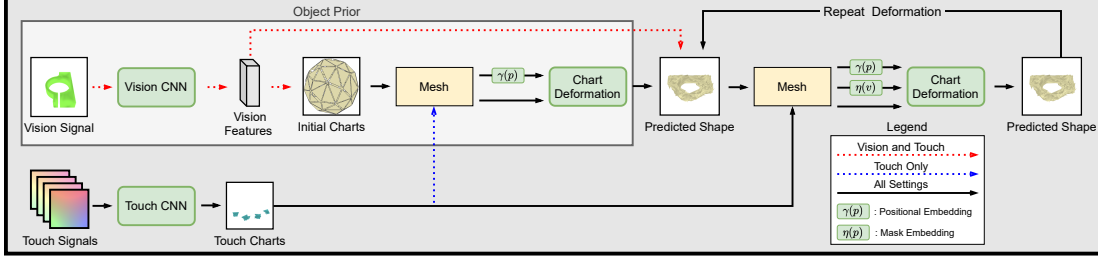


Figure 4: Our pipeline to 3D object reconstruction from vision and touch.

Chart-based mesh representation and vertex features. A mesh in the shape of a sphere is initially created from a large collection of mesh surface elements, known as *charts*, where each vertex is defined by its location. We call these mesh surface elements *vision charts* in the presence of vision signals, and *touch charts* in their absence. When leveraging touch signals, the mesh of initial charts is augmented with a set of N additional *touch charts*, which may be initialized with the previous touch CNN vertex location predictions or uninitialized – *i.e.* all vertex locations are set to zero – where N denotes the maximum number of touch charts expected in the active exploration process. Unlike [58], we encode the vertex locations by leveraging positional embeddings [45, 57, 70] to capture higher frequency shape information. When required, a mask embedding is appended to each vertex in the mesh to indicate if it originates from an uninitialized touch chart, a predicted touch chart, or a vision chart. With this setup, a variable number of touch charts can be expected by the subsequent deformation network. In addition, if a vision signal is available, the image is passed through a standard CNN and the extracted vision features are projected onto the vertices of the mesh using perceptual feature pooling [71, 61, 58]. Thus, the resulting mesh is composed of touch and eventually vision charts, whose vertex features include the above-described positional vertex embedding, mask embedding when needed, and if available, the corresponding projected visual features. The mesh connectivity enabling communication among vertices and charts is defined following [58].

Mesh deformation process. The chart-based mesh representation is fed to a mesh deformation model composed of two GCNs, by contrast [58] used a single GCN to handle the mesh deformation process. Decoupling the iterative GCN of [58] is crucial to ensure the use of vertex positional information by the model. Since the mesh deformation process of any object starts with the same sphere, parameter sharing across iterations results in the model ignoring the positional vertex information in the predictions of the initial shape belief, hindering the use of the added touch charts in subsequent iterations. In our model, the first GCN aims to learn an object prior, which can be learnt either from vision signals (vision prior) or from touch signals (touch prior). Note that the expected input of the first GCN is static *w.r.t* the position of vision and uninitialized touch charts. The second GCN takes the object belief resulting from the first GCN and refines it through a 2-step deformation process, in which each step recomputes the projection of the image features onto the mesh. Since the input to the second GCN is expected to be different in each case, we can leverage parameter sharing in its 2-step deformation process. Note that touch signals are only included in the second GCN when leveraging vision signals; however, they are included in the first GCN in the touch-only setting when available.

Training. The parameters of the whole reconstruction pipeline, including both GCNs and the vision CNN, are jointly optimized to minimize the CD between the predicted and target surface [63, 61]. With this setup, a potential vision signal and a variable number of touch signals can be leveraged to produce a surface prediction in a single model pass. Further details as well as a comparison highlighting the superiority of our shape reconstruction pipeline *w.r.t* [58] can be found in the supplemental material.

4 Data-driven policies for touch exploration

Our touch exploration framework takes advantage of the reconstruction model introduced in Section 3.2 to predict 3D shapes in the mesh space, and defines policies to select the positions of the next touches to acquire in order to maximize the similarity between predicted and target shape. While useful for graphics applications and efficient for representing surfaces, meshes are difficult to process and computationally heavy to compare. To combat these issues, we propose to use mesh embeddings of reduced dimensionality to facilitate the learning of our policies. The mesh embedding is extracted from the bottleneck of a mesh autoencoder which is trained offline from the shape predictions to produce a learned embedding space.

We also use the mesh embeddings to allow for efficient distance metric computation over predicted shapes – *i.e.* Euclidean distance in the embedding space.

Autoencoder for shape embedding. The encoder, e , takes as an input a surface mesh, and produces a mesh embedding. Following our shape reconstruction model, we use positional embeddings [45, 57, 70] to represent the vertices in the mesh. The mesh is then passed through a series Zero-Neighbor GCN layers [61, 34] to update the vertex features, followed by a channel-wise max pooling operation across vertices to produce a latent encoding. The decoder takes the resulting latent encoding and follows the FoldingNet [76] architecture to yield a point cloud with 2,024 points recovering the object shape. The autoencoder is trained by minimizing the CD between the input mesh and the predicted point cloud.

In the remainder of this section, we outline the object-specific and dataset-specific policies selected for the purpose of comparison. In object specific policies the current object shape is considered when deciding which action to perform. In dataset-specific policies the full training set of objects is leveraged to determine the optimal fixed trajectory to perform across all test set objects.

4.1 Object-specific policies

Nearest Neighbor (NN). We compute *myopic greedy* trajectories for all objects in the training set, and save the object reconstruction at each time step along with the action – grasping parameters – leading to its best immediate improvement. Then, when evaluating, we search our record of training reconstructions and their corresponding actions to find the most similar reconstruction to the current object belief and simply copy the grasping parameters from our record.

To allow for easy and efficient comparisons, we leverage the learned mesh encoder, e , to perform similarity search with the ℓ_2 distance.

Supervised (Sup.). In this model we attempt to learn the improvement each action will provide for a given object using regression. An independent network h_i is trained for each time step i to predict the relative improvement which will result from taking each action. Each network is comprised of set of fully connected layers with ReLU activations and takes as input the set of already performed grasp parameters, and the embedding of the current and initial shape reconstruction produced by the pre-trained autoencoder. The networks are trained sequentially, such that the i -th network learns to predict the relative improvement of each possible i -th action after performing the actions predicted by the networks of the previous time steps. When evaluating the performance of the supervised approach at time step i , h_i is used to determine which action will lead to highest improvement, and this action, g_i , is selected:

$$g_i = \arg \max_{\mathcal{G} \setminus \{g_0, \dots, g_{i-1}\}} h_i(\{g_0, \dots, g_{i-1}\}, e(\hat{O}_0), e(\hat{O}_i)), \quad (2)$$

where \mathcal{G} represents the set of all grasps.

Double Deep Q-Networks (DDQN). In the last object-specific model we leverage the discrete deep RL method Double Deep Q-Networks (DDQN) [69]. In our case, the value network takes as input the set of actions already performed and the embedded current reconstruction of the object, and predicts a value for every possible action. We propose two value network architectures. In the first, referred to as DDQN_m, the value network takes as input the mesh reconstruction of an object, where an embedding of the performed actions is appended to every vertex’s feature vector. The network architecture is identical to that of encoder, e , and produces a small shape embedding, which is then passed through a few fully connected layers to predict a value for every action. In the second, referred to as DDQN_i, the current reconstruction is passed through the pre-trained mesh encoder, e , producing a shape embedding which is concatenated with the embedding of the actions performed, and fed through a few fully connected layers to predict a value for every action. Note that the first setup benefits from a complete understanding of the current shape belief, as the mesh is fed to the value network. The second setup benefits from the pre-computed shape embeddings, which already contain the rich information necessary for reasoning over the object and allows for a simplistic network design which has already been demonstrated to perform well in deep reinforcement learning (RL) settings [69, 21]. In both cases, the action selected is the one which the value network Q predicts has the highest value:

$$g_i = \arg \max_{\mathcal{G} \setminus \{g_0, \dots, g_{i-1}\}} Q(\{g_0, \dots, g_{i-1}\}, \hat{O}_i). \quad (3)$$

4.2 Dataset-specific policies

Most frequent best action (MFBA).

This policy selects the first action by computing the performance of all actions over all objects in the training set, and then chooses the most common best action. For the second action, the performance of performing the first fixed action followed by all remaining actions is computed, and the most common best performing second action is chosen. This is repeated until a full trajectory is obtained. Then, when evaluating, this trajectory is selected every time, regardless of the object reconstruction.

Lowest error best action (LEBA). This policy is in effect identical to the MFBA except that the action which leads to the greatest average error improvement is selected and fixed at every time step rather than the most common best action.

5 Experiments

In this section, we validate the reconstruction and autoencoder models. We then compare object-specific and dataset-specific policies to several baselines. Additional experimental details can be found in the supplementary material.

5.1 Experimental setup

5.2 Shape reconstruction from static vision and touch signals

Dataset of CAD models. The dataset used is made up of 40,000 objects sampled from the ABC dataset [36, 58], a CAD model dataset of approximately one million objects. This dataset poses a much harder generalization challenge than other 3D object datasets, such as [10] due to its highly variable object shapes, and lack of clearly defined classes over which biases can be learned. The geometry of these objects were decimated such that all objects possess approximately 500 vertices. Those objects which could not be reduced to this size due to geometric constraints and those which possessed multiple disconnected parts were automatically removed, leading to set of 26,545 usable object models. This set of objects was split into 5 sets; 3 training sets ⁵ of size 7,700 object each, a validation set comprised 2,000 objects, and a test set of size 1,000.

Baselines and Oracle.

(1) *Random* baseline. As a naive baseline, a random policy is considered, which selects for every time step and object one of the available actions uniformly at random. This is the standard baseline for any exploration algorithm. (2) *Even* baseline. As a second naive baseline, we consider a policy which randomly selects an evenly spaced set of 5 actions over the sphere of possible actions. This baseline is chosen as it will result in uniform coverage of the target object, which is intuitively a useful and strong strategy for object understanding in our task. (3) *Oracle*. As a near-optimal target for the performance of our policies a *myopic* oracle policy is considered. In this policy, for a given object and time step the action which resulted in the best improvement is selected.

This policy possesses unfair hindsight, which is not accessible to all others, and so should be seen as an upper-bound point of comparison. The true optimal policy cannot be computed in a reasonable time frame, however due to the diminishing return of rewards from actions in this framework, the provided myopic oracle policy represents a close approximation.

Experimental scenarios. We examine the performance of our exploration framework and models across 4 learning settings: (1) *poking only*, T_P , where only touch from the third finger of the hand is leveraged; (2) *grasping only*, T_G , where only touch signals from all hand sensors are used during shape exploration; (3) *poking with vision*, $V \& T_P$, an extension of T_P that includes a visual input signal; and (4) *grasping with vision*, $V \& T_G$, an extension of T_G that leverages a visual signal.

We evaluate the performance of our proposed reconstruction method in the target domain of 3D reconstruction from vision and touch, and compare it to the current state-of-the-art [58]. While Wang et. al.

⁵Three training sets are used to mitigate the progressive overfitting which might occur from sequentially training the elements of our pipeline (reconstruction, autoencoder and policy networks) on the same data.

Model	Grasp #	
	0	1
T_G [58]	25.586 \pm 0.069	9.016 \pm 0.358
T_G [ours]	24.864 \pm 0.266	8.220 \pm 0.389
$V \& T_G$ [58]	2.653 \pm 0.022	2.637 \pm 0.042
$V \& T_G$ [ours]	2.538 \pm 0.098	2.486 \pm 0.102

Table 1: Comparison between our reconstruction model and state-of-art in terms of CD.




Figure 5: Two objects from the test set, along with their four nearest neighbors in the test set measured in the latent space of our trained autoencoder (V&T_G setting).

Input	Oracle	Baselines		Dataset specific		Object-specific			
		Random	Even	MFBA	LEBA	NN	DDQN _m	DDQN _i	Sup
T _P	19.35 ±0.00	36.38 ±0.29	33.25 ±0.48	32.40 ±1.04	29.85 ±0.39	33.46 ±0.51	32.41 ±0.40	31.10 ±0.34	31.21 ±0.67
T _G	16.38 ±0.00	25.83 ±0.14	24.53 ±0.27	23.46 ±0.07	23.04 ±0.09	24.34 ±0.29	23.92 ±0.14	23.84 ±0.23	23.70 ±0.27
V&T _P	78.95 ±0.00	94.56 ±0.34	93.95 ±0.29	93.59 ±0.32	92.36 ±0.25	91.79 ±0.15	93.75 ±0.48	92.62 ±0.30	93.12 ±0.38
V&T _G	77.18 ±0.00	90.65 ±0.34	90.29 ±0.32	89.39 ±0.11	89.31 ±0.25	88.53 ±0.25	90.07 ±0.51	89.32 ±0.17	89.46 ±0.23

Table 2: Comparison of active touch exploration strategies. Numbers represent a ratio between CD after 5 actions and initial CD (lower is better).

[72] do consider both vision and touch for 3D reconstruction, touch is not fused directly for prediction but rather used for shape refinement in sim2real transfer. The results of this experiment can be seen in Table 1 where the 4 highest performing models have been selected from the validation set while doing hyper-parameter search, and mean and variance numbers across these 4 models on the test set are shown. From these results it is clear that the proposed method outperforms the baseline comparison in all settings, and validates its model choices in our target multi-modal domain.

5.3 Shape Autoencoder

For each of the four learning settings an autoencoder is trained by leveraging the output of the reconstruction model. We qualitatively validate the shape embedding learnt by our autoencoder by visualizing object shapes and their nearest neighbors in the learnt embedding space. Figure 5 depicts 2 random objects sampled from the test set, along with the 4 other objects closest to their latent encoding in the test set for the V&T_G setting. Moreover, in the supplemental materials we highlight the average CD between the input and output meshes across 5 grasps, in our 4 learning settings, and observe that they are low relative to the error of the corresponding reconstruction models. The visual similarity of objects to their closest neighbors in the latent space along with the relatively low CD achieved demonstrates that the learned latent encodings possess important shape information which may be leveraged in the proposed active exploration policies.

5.4 Active Touch Exploration Results

We examine the performance of all described active touch exploration strategies over 5 grasps and show the results of all policies over the 4 learning settings in Table 2. For the DDQN_m, DDQN_i, and Supervised policies, the 5 highest performing models have been selected from the validation set while doing hyper-parameter search, and mean and variance numbers across these 5 models on the test set are reported. For the Random, and Even results, the strategies are repeated 5 times over the test set. For the MFBA, LEBA, and NN results, 5 random subsets of 40% of the training data are made to produce different fixed trajectories, or training set latent distributions. Figure 6 highlights the distributions of action selected by each strategy in the T_G and V&T_G settings. Here, the points of all actions on the sphere are transformed into their corresponding UV coordinates in an image, and the intensity value for every pixel corresponding to an action is set to its relative frequency computed over the test set. The visible area of the sphere of actions from the camera’s perspective is highlighted in orange, and non-visible in blue. Figure 2 highlights object predictions after 5 grasps with action selected using the DDQN_i model in the T_G and V&T_G settings.

From Table 2, we can see the clear, and expected trend of the Oracle strategy being the best performing and the Random baseline being the worst, followed by the Even baseline which is slightly better. In all cases,




Figure 6: Distribution of selected actions (greyscale encoded) for all policies in the T_G and $V\&T_G$ settings, with visible area of the sphere of actions from the camera highlighted in orange.




Figure 7: Action selection for the Oracle, LEBA, NN, and Supervised strategies in the $V\&T_G$ setting, where the arrows indicate the direction the hand moves towards the object for each selected action.

LEBA performs better than MFBA, most probably due its selection strategy being directly in line with the target objective. In both touch only settings, the LEBA policy performs best. For the learned policies, even after 4 grasps with a full hand of sensors, enough information about the current shape must not be available to properly learn the best action to perform. This is supported by the distributions in top line of Figure 6 where object-specific policies seem to employ close to a fixed strategy. In contrast to LEBA, these policies cannot leverage the full training dataset simultaneously and so their fixed strategy performs worse. From this we can see that even in the absence of meaningful shape information, leveraging the dataset of objects provides significantly better action selections leading to improved reconstruction accuracy.

In the vision and touch settings ($V\&T_P$ and $V\&T_G$), we see a reverse trend, with the NN policies performing best in both grasping settings. From this we can see that in the presence of better shape understanding (due to additional vision input), more successful action selection can be discriminated per object. This is supported in Figure 6 where we see that with vision priors, far more variable actions are performed, meaning that the object shape is now being appropriately considered. Moreover, in row 2 of Figure 6 a trend can be observed where some policies avoid actions where the touched surface is more likely to be visible from the camera. In the $V\&T_G$ setting, the NN, $DDQN_M$, and supervised policies select visible actions only 40.20%, 45.91% and 43.98% of the time respectively. Compared to the random policy which selects these actions 48.25 % of the time, and the oracle which selects them 43.80% of the time, this indicates that these policies have learned the intuitive strategy of avoiding sampling areas of the surface which have already been observed. In Figure 7 the different action selection strategies employed by various policies and the Oracle are shown for the $V\&T_G$ over 2 randomly sampled objects in the test set. Here the green arrows indicated from which direction the hand moves towards the objects for all 5 action selected. In both cases we see the object specific strategies tend to be fairly uniform around the sides, possibly to correctly identify the unseen dimensions of the objects.

6 Discussion

Limitations. There exist some limitations which should be highlighted. First, the reconstruction method aims to exclusively minimize the CD, which leads to poor visual object quality in the mesh representation [23, 71]. While attractiveness regularizers are available, we opted to focus exclusively on accuracy to make clear distinctions in improvement from information across different touch options. This is at the cost of visual quality. Second, the chosen shape agnostic grasp parameterization, where the hand always moves towards the center of the object leads touch sites biased towards the center of objects which possess dramatically different dimensional scales. An example of this can be seen the the first object of Figure 7, where because the object is long and thin, all touches will lie in the center of the object, as highlighted by the direction of the arrows. Finally, our environment requires full 3D shape supervision for training, which while easily available in simulation, limits its application to real world scenarios.

Societal Impact. Our work contributes towards improved understanding of the three dimensional world in which we all live. In particular, we provide new framework to study touch perception in a simulated environments. Thus, we advance the understanding of the importance of haptic information in the task of active 3D understanding, especially when used in tandem with visual signals. We envision our contributions to be relevant for real world robot-object manipulation. However, the improved active 3D

object understanding could have positive impact in fields beyond robotics, such as automation and virtual reality. Failures of these models could arise if not enough testing is performed prior to the deployment of the automation tools. To mitigate these risks, we encourage further investigation focusing on active 3D reconstruction system generalization limits both in the simulated and real-world scenarios.

Conclusions. In this paper, we explored the problem of data-driven active touch for 3D object reconstruction from vision and touch. We introduced a tactile-grasping simulator which allows for the efficient production of vision and touch signals from selected grasp parameters, and built a new 3D reconstruction method from vision and touch which achieves impressive performance both with and without haptic inputs. Over these elements and a large dataset of simulated objects, we built an active touch exploration environment which allows for the training and testing of active touch policies for 3D shape reconstruction. Finally, we produced an array of data-driven active touch policies which we compared to a set of baselines. The benefit of leveraging data for active touch is then highlighted by the superior reconstruction results of learned policies both in the presence of poor and rich shape information. In the presence of only touch information, the most successful exploration strategies learn a deterministic trajectory over the training data to employ indiscriminately over test objects even in the presence of shape information, indicating that either not enough information is present or that this information cannot be learned over with the current models. In the vision and touch settings the most fruitful strategies learn to select grasps based on the current belief of the objects’ shape, and experiments also indicate that learned models tend to favour grasps which occur on occluded areas of the objects’ surface.

References

- [1] Peter Allen. Surface descriptions from vision and touch. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 394–397. IEEE, 1984.
- [2] Peter K Allen. Integrating vision and touch for object recognition tasks. *The International Journal of Robotics Research*, 7(6):15–33, 1988.
- [3] Peter K Allen and Paul Michelman. Acquisition and interpretation of 3-d sensor data from touch. 1990.
- [4] Y. Aloimonos, I. Weiss, and A. Bandyopadhyay. Active vision. *International Journal of Computer Vision*, 1:333–356, 1988.
- [5] Kumar Ashutosh, Saurabh Kumar, and Subhasis Chaudhuri. 3d-nvs: A 3d supervision approach for next view selection. *arXiv preprint arXiv:2012.01743*, 2020.
- [6] Tim Bakker, Herke van Hoof, and Max Welling. Experimental design for mri by greedy policy search. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18954–18966. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/daed210307f1dbc6f1dd9551408d999f-Paper.pdf>.
- [7] A. Bierbaum, I. Gubarev, and R. Dillmann. Robust shape recovery for sparse contact location and normal data from haptic exploration. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3200–3205, 2008.
- [8] Alexander Bierbaum, Matthias Rambow, Tamim Asfour, and Rudiger Dillmann. A potential field approach to dexterous tactile exploration of unknown objects. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 360–366. IEEE, 2008.
- [9] Mårten Björkman, Yasemin Bekiroglu, Virgile Högman, and Danica Kragic. Enhancing visual perception of shape through tactile glances. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 11 2013. doi: 10.1109/IROS.2013.6696808.
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [11] Shengyong Chen, Youfu Li, and Ngai Ming Kwok. Active vision in robotic systems: A survey of recent developments. *The International Journal of Robotics Research*, 30(11):1343–1377, 2011.
- [12] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6970–6981, 2020.
- [13] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 628–644. Springer, 2016.
- [14] Cl Connolly. The determination of next best views. In *Proceedings. 1985 IEEE international conference on robotics and automation*, volume 2, pages 432–435. IEEE, 1985.
- [15] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. *GitHub repository*, 2016.
- [16] A. Dame, V. A. Prisacariu, C. Y. Ren, and I. Reid. Dense reconstruction using 3d object shape priors. In *2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1288–1295, 2013.
- [17] Jeffrey Delmerico, Stefan Isler, Reza Sabzevari, and Davide Scaramuzza. A comparison of volumetric information gain metrics for active 3d object reconstruction. *Autonomous Robots*, 42(2):197–208, 2018.

- [18] Danny Driess, Peter Englert, and Marc Toussaint. Active learning with query paths for tactile object shape exploration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [19] Haoqiang Fan, Hao Su, and Leonidas Guibas. A point set generation network for 3d object reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 38, 2017.
- [20] Sascha Fleer, Alexandra Moringen, Roberta L Klatzky, and Helge Ritter. Learning efficient haptic shape exploration with a rigid tactile sensor array. *PloS one*, 15(1):e0226880, 2020.
- [21] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596. PMLR, 2018.
- [22] Gabriela Zarzar Gandler, Carl Henrik Ek, Mårten Björkman, Rustam Stolkin, and Yasemin Bekiroglu. Object shape estimation and modeling, based on sparse gaussian process implicit surfaces, combining visual data and tactile exploration. *Robotics and Autonomous Systems*, 126:103433, 2020. ISSN 0921-8890. doi: <https://doi.org/10.1016/j.robot.2020.103433>. URL <http://www.sciencedirect.com/science/article/pii/S0921889019303495>.
- [23] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [24] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 216–224, 2018.
- [25] C. Häne, N. Savinov, and M. Pollefeys. Class specific 3d object shape priors using surface normals. In *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 652–659, 2014.
- [26] Jarmo Ilonen, Jeannette Bohg, and Ville Kyrki. Three-dimensional object reconstruction of symmetric objects by fusing visual and tactile sensing. *The International Journal of Robotics Research*, 33(2):321–341, 2014.
- [27] Eldar Insafutdinov and Alexey Dosovitskiy. Unsupervised learning of shape and pose with differentiable point clouds. *arXiv preprint arXiv:1810.09381*, 2018.
- [28] N. Jamali, C. Ciliberto, L. Rosasco, and L. Natale. Active perception: Building objects’ models using tactile exploration. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 179–185, Nov 2016. doi: 10.1109/HUMANOIDS.2016.7803275.
- [29] Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019.
- [30] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, pages 365–376. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6640-learning-a-multi-view-stereo-machine.pdf>.
- [31] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. *arXiv preprint arXiv:1711.07566*, 2017.
- [32] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, and Peter Henry. End-to-end learning of geometry and context for deep stereo regression. In *IEEE International Conference on Computer Vision (ICCV)*, pages 66–75. IEEE Computer Society, 2017. doi: 10.1109/ICCV.2017.17. URL <https://doi.org/10.1109/ICCV.2017.17>.
- [33] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [34] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- [35] Roberta L Klatzky, Susan J Lederman, and Dana E Matula. Haptic exploration in the presence of vision. *Journal of Experimental Psychology: Human Perception and Performance*, 19(4):726, 1993.
- [36] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [37] Adarsh Kowdle, Yao-Jen Chang, Andrew Gallagher, and Tsuhan Chen. Active learning for piecewise planar 3d reconstruction. In *CVPR 2011*, pages 929–936. IEEE, 2011.
- [38] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian 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. *IEEE Robotics and Automation Letters (RA-L)*, 5(3):3838–3845, 2020. doi: 10.1109/LRA.2020.2977257.
- [39] Susan J Lederman and Roberta L Klatzky. Extracting object properties through haptic exploration. *Acta psychologica*, 84(1):29–40, 1993.
- [40] Shan Luo, Joao Bimbo, Ravinder Dahiya, and Hongbin Liu. Robotic Tactile Perception of Object Properties: A Review. *arXiv e-prints*, art. arXiv:1711.03810, November 2017.
- [41] Matthew Matl. Pyrender. <https://github.com/mmatl/pyrender>, 2019.
- [42] Takamitsu Matsubara and Kotaro Shibata. Active tactile exploration with uncertainty and travel cost for fast shape estimation of unknown objects. *Robotics and Autonomous Systems*, 91:314–326, 2017.
- [43] Miguel Mendoza, J Irving Vasquez-Gomez, Hind Taud, L Enrique Sucar, and Carolina Reta. Supervised learning of the next-best-view for 3d object reconstruction. *Pattern Recognition Letters*, 133:224–231, 2020.
- [44] Damien Michel, Xenophon Zabulis, and Antonis A Argyros. Shape from interaction. *Machine Vision and Applications*, 25(4):1077–1087, 2014.
- [45] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020.
- [46] J Krishna Murthy, GV Sai Krishna, Falak Chhaya, and K Madhava Krishna. Reconstructing vehicles from a single image: Shape priors for road scene understanding. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 724–731. IEEE, 2017.
- [47] Simon Ottenhaus, Martin Miller, David Schiebener, Nikolaus Vahrenkamp, and Tamim Asfour. Local implicit surface estimation for haptic exploration. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 850–856, 11 2016. doi: 10.1109/HUMANOIDS.2016.7803372.
- [48] 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(01):1850015, 2018.
- [49] Paschalis Panteleris, Nikolaos Kyriazis, and AA Argyros. Recovering 3d models of manipulated objects through 3d tracking of hand-object interaction. In *IEEE International Conference on Computer Vision Workshops (OUI 2015-ICCVW 2015)*, Santiago, Chile, 2015.

- [50] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [51] Z. Pezzementi, C. Reyda, and G. D. Hager. Object mapping, recognition, and localization from tactile geometry. In *2011 IEEE International Conference on Robotics and Automation*, pages 5942–5948, 2011.
- [52] Luis Pineda, Sumana Basu, Adriana Romero, Roberto Calandra, and Michal Drozdal. Active mr k-space sampling with reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 23–33. Springer, 2020.
- [53] SG Potapova, AV Artemov, SV Sviridov, DA Musatkina, DN Zorin, and EV Burnaev. Next best view planning via reinforcement learning for scanning of arbitrary 3d shapes. *Journal of Communications Technology and Electronics*, 65(12):1484–1490, 2020.
- [54] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1(2):4, 2017.
- [55] Jason Rock, Tanmay Gupta, Justin Thorsen, JunYoung Gwak, Daeyun Shin, and Derek Hoiem. Completing 3d object shape from one depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2484–2493, 2015.
- [56] SimLab. Allegro hand overview, 2016. URL http://wiki.wonikrobotics.com/AllegroHandWiki/index.php/Allegro_Hand_Overview. [Online; accessed 25-May-2020].
- [57] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.
- [58] Edward Smith, Roberto Calandra, Adriana Romero, Georgia Gkioxari, David Meger, Jitendra Malik, and Michal Drozdal. 3d shape reconstruction from vision and touch. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14193–14206. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/a3842ed7b3d0fe3ac263bcabd2999790-Paper.pdf>.
- [59] Edward J Smith and David Meger. Improved adversarial systems for 3d object generation and reconstruction. In *Conference on Robot Learning (CoRL)*, pages 87–96, 2017.
- [60] Edward J. Smith, Scott Fujimoto, and David Meger. Multi-view silhouette and depth decomposition for high resolution 3d object representation. In *Advances in Neural Information Processing Systems*, pages 6479–6489, 2018.
- [61] Edward J. Smith, Scott Fujimoto, Adriana Romero, and David Meger. Geometrics: Exploiting geometric structure for graph-encoded objects. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, volume 97 of *Proceedings of Machine Learning Research*, pages 5866–5876. PMLR, 2019. URL <http://proceedings.mlr.press/v97/smith19a.html>.
- [62] Nicolas Sommer, Miao Li, and Aude Billard. Bimanual compliant tactile exploration for grasping unknown objects. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 6400–6407, 09 2014. doi: 10.1109/ICRA.2014.6907804.
- [63] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B. Tenenbaum, and William T. Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. *CoRR*, abs/1804.04610, 2018. URL <http://arxiv.org/abs/1804.04610>.

- [64] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [65] Julián Tachella, Yoann Altmann, Nicolas Mellado, Aongus McCarthy, Rachael Tobin, Gerald S Buller, Jean-Yves Tournet, and Stephen McLaughlin. Real-time 3d reconstruction from single-photon lidar data using plug-and-play point cloud denoisers. *Nature communications*, 10(1):1–6, 2019.
- [66] Glenn H Tarbox and Susan N Gottschlich. Planning for complete sensor coverage in inspection. *Computer Vision and Image Understanding*, 61(1):84–111, 1995.
- [67] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 209–217. IEEE, 2017.
- [68] Dimitrios Tzionas and Juergen Gall. 3d object reconstruction from hand-object interactions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 729–737, 2015.
- [69] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [71] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. *arXiv preprint arXiv:1804.01654*, 2018.
- [72] S. Wang, J. Wu, X. Sun, W. Yuan, W. T. Freeman, J. B. Tenenbaum, and E. H. Adelson. 3d shape perception from monocular vision, touch, and shape priors. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1606–1613, Oct 2018. doi: 10.1109/IROS.2018.8593430.
- [73] David Watkins-Valls, Jacob Varley, and Peter Allen. Multi-modal geometric learning for grasping and manipulation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7339–7345. IEEE, 2019.
- [74] Chenming Wu, Rui Zeng, Jia Pan, Charlie CL Wang, and Yong-Jin Liu. Plant phenotyping by deep-learning-based planner for multi-robots. *IEEE Robotics and Automation Letters*, 4(4):3113–3120, 2019.
- [75] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *Advances in Neural Information Processing Systems 32*, pages 492–502. Curran Associates, Inc., 2019.
- [76] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 206–215, 2018.
- [77] Zhengkun Yi, Roberto Calandra, Filipe Fernandes Veiga, Herke van Hoof, Tucker Hermans, Yilei Zhang, and Jan Peters. Active tactile object exploration with Gaussian processes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4925–4930, 2016. doi: 10.1109/IROS.2016.7759723.
- [78] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*, pages 728–737, 2018.
- [79] Wenzhen Yuan, Siyuan Dong, and Edward Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, Nov 2017. ISSN 1424-8220. doi: 10.3390/s17122762. URL <http://dx.doi.org/10.3390/s17122762>.
- [80] Rui Zeng, Yuhui Wen, Wang Zhao, and Yong-Jin Liu. View planning in robot active vision: A survey of systems, algorithms, and applications. *Computational Visual Media*, pages 1–21, 2020.

- [81] Rui Zeng, Wang Zhao, and Yong-Jin Liu. Pc-nbv: A point cloud based deep network for efficient next best view planning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7050–7057. IEEE, 2020.
- [82] Zizhao Zhang, Adriana Romero, Matthew J Muckley, Pascal Vincent, Lin Yang, and Michal Drozdal. Reducing uncertainty in undersampled mri reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2049–2058, 2019.

Appendix

In this appendix, we supply additional information, experiments, and results. This includes a more in depth review of the simulator (Section A), the full definition of the Chamfer Distance (Section B), the **network architecture and training details** (Section C) for all experiments, additional results and experiments (Section D), **licenses** (Section E) for assets and frameworks used, and descriptions of **compute** (Section F) used for all tasks. Sections containing elements from the checklist are boldfaced.

A Simulator

In the following section details with respect to the tactile grasping environment will be provided, including the manner in which grasps on an object are performed, and an explanation of how the touch signals are simulated.

A.1 Grasping

A batch of PyBullet simulator [15] threads are kept constantly available throughout training. Each thread possesses the Allegro hand [56] pre-loaded. When a batch of objects needs to be grasped with a set of chosen actions, each object is loaded into one of the threads such that the object’s center is at position $[0, 0, 0.6]$, where the max size of the object across any dimension is 0.5. In a given thread and with a given action, a fixed sphere of points with radius 1 is imagined around the object, and the chosen action is used to pick its corresponding point. The hand is then loaded in an open position facing the object with the third digit placed at the point, and the palm of the hand tangent to the imagined sphere. A line from the point to the center of the object is then imagined, and the closest intersection of this line with the convex hull of the object is computed. The hand is moved along the imagined line up to the intersection, and then reoriented so that the palm of the hand is tangent to the surface of the object at the intersection site. The grasp is then performed by increasing all of the hand’s joint angles in 5 maximal joint updates. Here the sensors on the fingers will either touch the object in some way, be impeded from performing the touch by some geometry in the hand or object, or miss the object entirely and the finger will close to the palm. The pose information of the hand for each thread is then exported from the thread.

A.2 Simulating Touch

A Pyrender [41] scene is kept constantly available throughout training to allow for touch signals to be simulated. The object is first loaded into the scene in the same position and orientation as in the PyBullet thread. For each touch sensor in the hand, the pose information of its finger when the grasp was performed is used to place a simulated perspective camera in its same position and orientation. A depth image is then produced from this camera. This depth image is used to produce a simulated touch signal, through the same method as in [58].

Vision signals from each object are also produced using Pyrender. To improve the simulation time, this is performed offline, and stored images are loaded from memory when needed. To generate this vision signal (image), the object is given a random colour texture and placed alone in an empty scene with 4 point lights placed in fixed positions around it. The produced images are of size $256 \times 256 \times 3$.

A.3 Simulation Time

For the complexity of the required simulation task, the time to simulate all 4 touches for a single grasp is relatively quick. This due to a number of optimization choices which include:

- All operations are performed within the same python instance.
- GPU acceleration is leveraged for simulating the touch when possible, such as when rendering with Pyrender, and performing heavy mathematical operations in PyTorch [50].
- The meshes for all objects possess on the order of 500 vertices. As a result, the time required for operations such as loading the mesh into both simulators, performing the physics simulation in PyBullet, and rendering images in Pyrender is minimal.
- The number of PyBullet physics simulation steps required to perform the grasp is only five, and all other computations for the grasps – e.g. the hand placement relative to the object – are computed out of simulation.
- Objects remain loaded within Pybullet and Pyrender between grasps in the same trajectory.

As a result of these optimizations, the average time to extract the 4 simulated touch signals from a single grasp across all grasp options and 100 random objects from the third training set is ~ 0.0317 seconds on a machine with 1 Tesla V100 GPU and with 16 CPU cores.

B Chamfer Distance

The Chamfer distance (CD) is used extensively to compare predictions for training and evaluation in this paper. The Chamfer distance takes as input two set of point clouds and computes the average minimal distance between points in each set. As our prediction are in the mesh representation, we first uniformly sample points on the surface of each mesh, O and \hat{O} , to produce point clouds S and \hat{S} . We then compute the difference:

$$CD(O, \hat{O}) = \sum_{p_1 \in S} \min_{p_2 \in \hat{S}} \|p_1 - p_2\|_2^2 + \sum_{p_2 \in \hat{S}} \min_{p_1 \in S} \|p_2 - p_1\|_2^2. \quad (4)$$

For all experiments, we sample 30,000 points from each object, and average over 3 computations of the CD. This is over more points than normally used, and usually no averaging is performed. These changes were employed to address the stochasticity induced by converting our meshes to point clouds, and ensure that changes in CD can be attributed to different touches being performed rather than variance in the approximation.

C Network and Training details

In the following section, the network architectures and training details will be described for all learned models.

C.1 Touch CNN Model

A convolutional neural network (CNN) is trained to convert an input touch signal into a chart representing the local surface of the object where the touch occurred. A chart here is a mesh sheet, meant to represent a small subsection of a full mesh surface. These produced charts are referred to as *touch* charts. To do this, the touch signal is passed through a series of CNN layers, reshaped into a vector, and passed through a series of fully connected layers to predict the vertex positions of a small fix size mesh with 25 vertices. The meshes is then rotated and translated with the known pose information of the finger which performed the touch. The parameters of this network are trained to minimize CD between the produced mesh and the ground truth surface of the object at the location of the touch. The same trained network is used across all active touch settings as the objective of the network does not change between them. The touch signal input image is of size $121 \times 121 \times 3$. The network architecture is provided in Table 7. The network was trained using the Adam optimizer [33] with a learning rate of 0.001 for 300 epochs at a batch size of 256. Hyper-parameter search was performed in a grid search over learning rate $\{0.0001, 0.001\}$ and batch size $\{128, 256\}$. The performance of the models was evaluated on the validation set every epoch, and the best performing model across these evaluations was selected. The objects used in this training were from the first training set, and in each batch, a set of random touches are sampled from these objects. With this trained network, a touch chart can be produced for any given touch in a grasp, representing local surface where the touch takes place.

C.2 Reconstruction Model

The model takes as input a set of touch signals and a vision signal, all in the form of RGB images and produces from them a collection of charts which have been deformed and arranged to make up the full surface of the predicted shape. Initially, all touch signals are converted into touch charts using the pre-trained touch CNN.

In the vision and touch setting, a collection of 76 vision charts are formed in the shape of a sphere and the vertices on the borders of these charts share edges with each other to allow communication between them. Recall that the iterative mesh deformation process is performed by two sets of networks which do not share parameters. In the first part of the deformation process, we pass the image signal through a VGG-like CNN with network architecture described in Table 8 to extract image features. We then use perceptual feature pooling [71, 61] to project the extracted image features onto the vertices of our mesh (composed of the above-described charts). The feature vectors this process produces are of size 118. This mesh is then passed through a graph convolutional network with Zero-Neighbor GCN layers [61] with network architecture described in Table 9 and the output of this is added to the original mesh to predict a new location for every vertex. A set of K empty touch charts which are initially given position $(0, 0, 0)$ for all vertices are then appended to the mesh of vision charts, and then every predicted touch chart is used to

replace one of the empty touch charts each. All touch charts have one of its center vertices connected by an edge to every vision chart’s border vertices to allow communication between vision and touch charts. A mask embedding of size 118 is produced indicating if each vertex in this graph is from a vision chart, empty touch chart, or predicted touch chart. The vertex positions of each vertex are concatenated with a Nerf positional embedding [45] of length 10 to produce feature vectors which are then passed through 3 fully connected layers with ReLU activations to grow their shape to size 118. In the second part of the deformation process, we again pass the image signal through a VGG-like CNN with network architecture described in Table 8 to extract image features which are projected onto the vertices of the mesh using perceptual feature pooling. The features from the positional embedding, mask embedding and image are then added together and passed again through a graph convolutional network with Zero-Neighbor GCN layers with network architecture described in Table 9. The output of this network is added to the input mesh to predict a new location for every vertex in the vision charts. This process is repeated once more. More precisely, the resulting mesh is passed through the same GCN network to produce a final update to the positions of all vertex charts. The combination of vision and touch charts then makes up the final prediction of shape. The same GCN, CNN, positional embedding, and mask embedding parameters are used in the second and third deformation steps.

In the touch only setting, a collection of 76 vision charts (here referred to as touch charts) is formed in the shape of a sphere and the vertices on the borders of these charts share edges with each other to allow communication between them. An additional set of K empty touch charts, which are initially given position $(0, 0, 0)$ for all vertices, are then appended to the initial mesh of charts. Every predicted touch chart is used to replace one of the empty touch charts in this additional set of charts. All predicted touch charts have one of its center vertices connected by an edge to every other chart’s border vertices to allow communications. A mask embedding of size 50 is produced indicating if each vertex in this mesh is from a chart in the initial sphere, empty touch chart, or predicted touch chart. The vertex positions of each vertex are concatenated with a Nerf positional embedding [45] of length 10 to produce a feature vectors which are then passed through 3 fully connected layers with ReLU activations to shrink their size to 50. The features from the positional embedding and mask embedding are then added together and passed through a graph convolutional network with Zero-Neighbor GCN layers with network architecture described in Table 9. The output of this network is added to the original mesh to predict a new location for every vertex. The process of computing these features and passing the resulting mesh through a GCN network to produce an update to the positions of all vertex charts is repeated twice more. The combination of charts then makes up the final prediction of shape. The same GCN parameters are used in the second and third deformation steps, and the positional embedding and mask embedding parameters are used in all three deformations.

In all settings, the reconstruction models were trained to minimize the CD between predicted and target meshes using the Adam optimizer [33] with a learning rate of 0.0001 for 1,000 epochs at a batch size of 12 and with patience of 70. In the poking setting $K = 5$, and in the grasping setting $K = 20$. A grid search was performed over hyper-parameters including the number of GCN layers $\{8, 10, 12, 15\}$, the hidden dimension size in the GCN layers $\{300, 350, 400\}$, and the percentage of vertex features shared in every Zero-Neighbor GCN layer $\{33\%, 35\%\}$. The performance of the models was evaluated on the validation set every epoch, and the best performing model across these evaluations was selected. The objects used in this training were from the first training set, and for each instance in the batch a random number of grasps between 0 and 5 is sampled.

C.3 Autoencoder Models

Recall that the autoencoder creates latent embeddings of the predicted shapes produced using the trained reconstructed models. The autoencoder takes as input a mesh in the form of a graph, passes it through a GCN with architecture described in Table 9, concatenates the max value from each vertex feature position across all vertices to produce a features vector, and passes it through 4 fully connected layers with ReLU activations to produce a latent embedding of size 200. For the decoder, the latent embedding is passed through a FoldingNet [76] decoder to produce the predicted point cloud of 2024 points. This setup is trained to minimize the CD between the predicted point cloud and the input mesh. In all settings, the models were trained using the Adam optimizer [33] with a learning rate of 0.0001 for 1000 epochs at a batch size of 12 and with patience of 70. A grid search was performed over hyper-parameters including the number of GCN layers $\{8, 10, 12, 15\}$, the hidden dimension size in the GCN layers $\{300, 350, 400\}$, and the percentage of vertex features shared in every Zero-Neighbor GCN layer $\{8\%, 15\%\}$. The performance of the models was evaluated on the validation set every epoch, and the best performing model across these

evaluations was selected in each setting. The objects used in this training were from the second training set, and for each instance in the batch a random number of grasps between 0 and 5 is sampled.

C.4 DDQN Policies

We make use of the DDQN [69], a standard and highly successful deep reinforcement learning algorithm for solving MDPs over discrete action spaces. In this method, the policy is defined by greedily selecting the action which maximizes a learned value function, at every time step. The value function predicts the value of any given action in the current state, where the value is defined as the expected future cumulative reward. The value function is typically implemented as a deep neural network and is trained to minimize the temporal difference error [64] over sampled data from a replay buffer of previously experienced trajectories in the environment. A full explanation of this method can be found in [69]. In the DDQN_m setting, the action values are predicted directly from the predicted mesh shape. Here, the set of actions performed is described as a k -hot mask which is passed through 3 fully connected layers with ReLU activations to produce an action embedding of size 100. A mask embedding of size 100 is produced indicating if each vertex in the mesh is from a vision chart (charts in the initial sphere), empty touch chart, or predicted touch chart. The vertex positions of each vertex are concatenated with a Nerf positional embedding [45] of length 10 to produce feature vectors which are then passed through 3 fully connected layers with ReLU activations to grow their shape to size 100. The action, positional and mask embeddings are then concatenated together and passed through a GCN with network architecture described in Table 9 to produce a feature vector of size 50 at every vertex. The max value across all vertices is then computed to produce a vector of size 50 representing the value of each action. In this setting, the models were trained using the standard DDQN framework as described in [69] using the Adam optimizer [33] with a learning rate of 0.001 for 500,000 episodes with a network update batch size of 128. Validation was regularly performed over the validation set to identify the best models. A grid search was performed over hyper-parameters epsilon decay $\{0.999993, 0.999996\}$, discount factor $\{0.9, .99\}$, hidden GCN dimension size $\{100, 200\}$, memory capacity $\{300, 000, 100, 000\}$, and normalization of the reward by the CD of the initial object belief or the current.

For DDQN_i, the action value is predicted over the latent embedding of the predicted mesh. Here, the set of actions performed is again described as a k -hot mask passed through 3 fully connected layers with ReLU activations to produce an action embedding of size 50. This is concatenated with both the latent embeddings of the current mesh and the initial mesh prediction, and then passed through L fully connected layers with hidden dimension H and ReLU activation to produce a vector of size 50 representing the value of every action. In this setting, the models were trained using the standard DDQN framework as described by [69] using the Adam optimizer [33] with a learning rate of 0.001 for 500k episodes with a network update batch size of 12. Validation was constantly performed over the validation set to identify the best models. A grid search was performed over hyper-parameters epsilon decay $\{0.999993, 0.999996\}$, discount factor $\{0.9, .99\}$, number of hidden dimensions $H \{35, 75, 100, 200\}$, number of hidden layers $L \{2, 3\}$, memory capacity $\{300, 000, 100, 000\}$, and normalization by the CD of the initial object belief or the current.

C.5 Supervised Policy

For the supervised learning policies, an individual network is trained to learn the value of actions at each of the 5 time steps. The first network is trained to predict the improvement induced by performing every action based on the current belief of the object given that no actions have been performed yet. Here, the improvement of an action is defined as $\frac{CD(\hat{O}_{k+1}, O)}{CD(\hat{O}_k, O)}$, where $k = 0$. Once this network is trained, we move to training network for the second step. Here, for every object in a batch, action one is first performed based on selecting the action with the the highest predicted value from the first network. the second network is trained to predict the improvement induced by performing every action based on the current belief of the object given that one action has been selected from the first network and performed. This continues until all 5 networks have been trained. When testing this policy the action at time step t is selected by passing the current object belief through the t -th network and taking the action which value is maximized. In all networks, the set of actions performed is described as a k -hot mask which is passed through 2 fully connected layers with ReLU activations and hidden dimension of 100 to produce an action embedding of size 50. This is concatenated with both the latent embedding of the current mesh and the initial mesh prediction and then passed through L fully connected layers with hidden dimension H and ReLU activations to produce a vector of size 50 representing the value of every action. Each network is trained to minimize the mean squared error between the predicted action improvements and the true improvements using the Adam optimizer [33] for 300 epochs with patience 20 and batch size 64. The

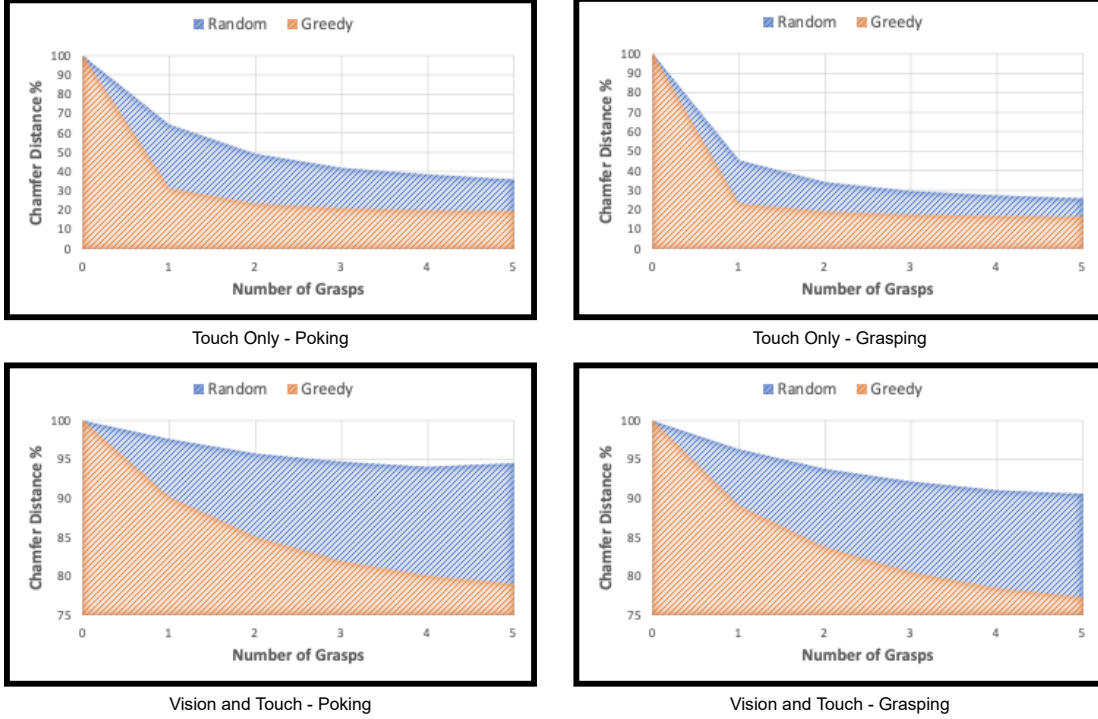


Figure 8: Graphs demonstrating the average reconstruction accuracy of the trained model in each of the 4 learning settings, across different number of grasps and the greedy oracle and random baseline

performance of the models was evaluated on the validation set every epoch, and the best performing model across these evaluations was selected in each setting. A grid search was performed over hyper-parameters: number of fully connected layers $L \in \{2, 3, 4\}$, number of hidden dimensions $H \in \{50, 100, 200\}$ and learning rate $\{0.001, 0.0003\}$.

D Additional Results

In this section, additional experiments and results are provided.

D.1 Mesh Reconstruction

The performance of the best trained reconstruction models on the test set when randomly picking actions across the 4 active touch settings is shown in Table 3. The performance of the best trained models on the test set when greedily picking the best action across the 4 active touch settings is shown in Table 4. Figure 8 shows the change in relative reconstruction accuracy from 0 to 5 touches when picking actions using the greedy and random policies. The large difference in reconstruction performance between the random and greedy policies highlights the need for learned policies which select more informative grasps.

We compare the performance of the proposed model to other state of the art single image 3D object reconstruction models on the 3D Warehouse Dataset [10] using the exact training and evaluation setup described in [23]. The results of this experiment can be seen in Table 5. Here, $F1^{k*\tau}$ is the harmonic mean of the percentage of predicted points with distance at most $k * \tau$ from any ground truth points and the percentage of ground truth points with distance at most $k * \tau$ from any predicted point. The proposed method nearly matches the performance of the best performing method, Mesh-RCNN [23], and notably performs significantly better than the only previous method built for leveraging vision and touch [58].

D.2 Autoencoder

In Table 6, the chosen autoencoder models' reconstruction Chamfer distances on the test set across all 4 settings are shown. In Figure 9, two random objects are shown in each learning setting along with the 4 closest objects to them in the respective learned latent space of objects. The visual similarity of objects to their closest neighbors in the latent space along with the relatively low CD achieved demonstrates that the

Touches	0	1	2	3	4	5
T _P	100	64.28	49.01	41.95	38.31	35.91
T _G	100	45.26	34.10	29.55	27.31	25.76
V&T _P	100	97.66	95.76	94.71	94.09	94.51
V&T _G	100	96.37	93.78	92.14	91.06	90.60

Table 3: Mesh reconstruction results across all 4 learning settings with actions chosen using the random policy. The units displayed are the percentage of Chamfer distance relative to the Chamfer distance of the initial object belief.

	CD(↓)	F1 ^τ (↑)	F1 ^{2τ} (↑)
N3MR [31]	2.629	3.80	47.72
3D-R2N2 [13]	1.445	39.01	54.62
PSG [19]	0.593	48.58	69.78
MVD [60]	-	66.39	-
GEOMETRICS [61]	-	67.37	-
Pixel2Mesh [71]	0.463	67.89	79.88
MeshRCNN [23] (Pretty)	0.391	69.83	81.76
VT3D [58]	0.369	69.52	82.33
MeshRCNN [23] (Best)	0.306	74.84	85.75
Ours	0.346	73.58	84.78

Table 5: Single image 3D shape reconstructing results on the 3D Warehouse Dataset using the evaluation from [23] and [71].

Touches	0	1	2	3	4	5
T _P	100	31.40	23.45	20.87	19.87	19.35
T _G	100	23.09	18.99	17.49	16.76	16.38
V&T _P	100	90.19	85.00	81.90	80.03	78.95
V&T _G	100	89.03	83.71	80.51	78.46	77.18

Table 4: Mesh reconstruction results across all 4 learning settings with actions chosen using the greedy policy. The units displayed are the percentage of Chamfer distance relative to the Chamfer distance of the initial object belief.

Grasps	0	1	2	3	4	5
T _P	0.334	0.435	0.436	0.435	0.438	0.445
T _G	0.405	0.514	0.488	0.470	0.462	0.459
V&T _P	0.516	0.516	0.516	0.516	0.517	0.517
V&T _G	0.477	0.477	0.477	0.477	0.477	0.477

Table 6: Autoencoder average Chamfer distance scores across the 4 active learning settings and 5 grasps.

learned latent encodings possess important shape information which may be leveraged in the proposed active exploration policies.

D.3 Policies

Figure 10 highlights the distributions of action selected by each strategy. Here, the points of all actions on the sphere are transformed into their corresponding UV coordinates in an image, and the intensity value for every pixel corresponding to an action is set to its relative frequency computed over the test set. The visible area of the sphere of actions from the camera’s perspective is highlighted in orange, and non-visible in blue. Figure 11 displays shape reconstructions after 5 grasps resulting from the DDQN_I policy. In Figure 12, the different action selection strategies employed by various policies and the Oracle are shown for 2 randomly sampled objects in the test set.




Figure 9: Objects from the test set, along with their four nearest neighbors in the test set measured in the latent space of our trained autoencoder in the 4 learning settings.




Figure 10: Distribution of selected actions (greyscale encoded) for all policies in all settings, with visible area of the sphere of actions from the camera highlighted in orange.

E Licences

All licensed software and assets, along with their licenses are:

1. **PyBullet:** MIT License
<https://github.com/bulletphysics/bullet3>
2. **PyRender:** MIT License
<https://github.com/mmatl/pyrender>
3. **Pytorch3D:** BSD 3-Clause License
<https://github.com/facebookresearch/pytorch3d>
4. **ABC Dataset:** MIT License <https://github.com/deep-geometry/abc-dataset>
5. **Wonik Allegro Hand:** BSD 2-Clause "Simplified" License
https://github.com/simlabrobotics/allegro_hand_ros_catkin

F Compute

All described actions were performed on machines with either 1 or 2 Tesla V100 GPUs and with 16 CPU cores.




Figure 11: Target objects (top rows) and predicted 3D shapes (bottom rows) after 5 grasps have been selected following the DDQN₁ policy in all settings.

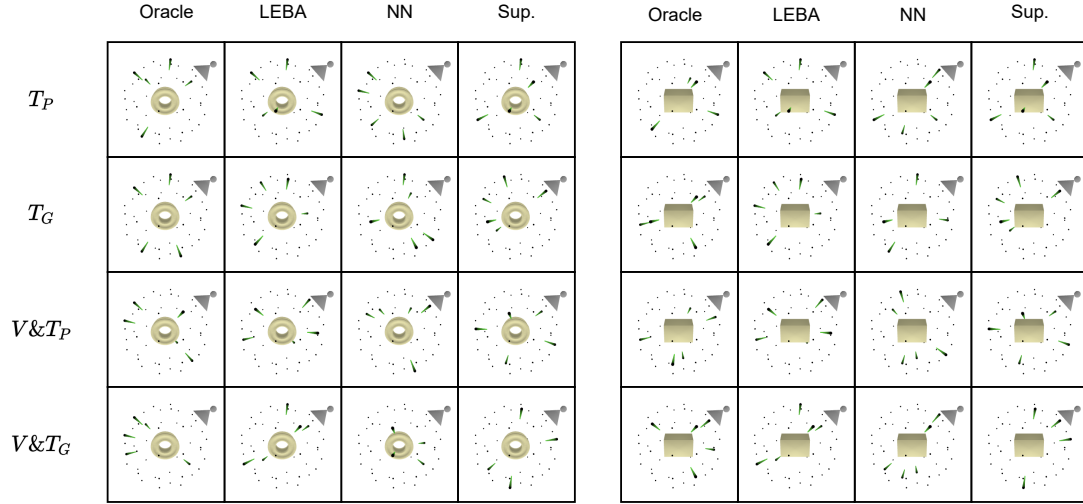


Figure 12: Action selection for the Oracle, LEBA, NN, and Supervised strategies, where the arrows indicate the direction the hand moves towards the object for each selected action.

Index	Input	Operation	Output Shape
(1)	Input	Conv (5×5) + BN + ReLU	$8 \times 61 \times 61$
(2)	(1)	Conv (5×5) + BN + ReLU	$8 \times 61 \times 61$
(3)	(2)	Conv (5×5) + BN + ReLU	$8 \times 61 \times 61$
(4)	(3)	Conv (5×5) + BN + ReLU	$8 \times 61 \times 61$
(5)	(4)	Conv (5×5) + BN + ReLU	$16 \times 31 \times 31$
(6)	(5)	Conv (5×5) + BN + ReLU	$16 \times 31 \times 31$
(7)	(6)	Conv (5×5) + BN + ReLU	$16 \times 31 \times 31$
(8)	(7)	Conv (5×5) + BN + ReLU	$16 \times 31 \times 31$
(9)	(8)	Conv (5×5) + BN + ReLU	$32 \times 16 \times 16$
(10)	(8)	Conv (5×5) + BN + ReLU	$32 \times 16 \times 16$
(11)	(10)	Conv (5×5) + BN + ReLU	$32 \times 16 \times 16$
(12)	(11)	Conv (5×5) + BN + ReLU	$32 \times 16 \times 16$
(13)	(12)	Conv (5×5) + BN + ReLU	$64 \times 8 \times 8$
(14)	(13)	Conv (5×5) + BN + ReLU	$64 \times 8 \times 8$
(15)	(14)	Conv (5×5) + BN + ReLU	$64 \times 8 \times 8$
(16)	(15)	Conv (5×5) + BN + ReLU	$64 \times 8 \times 8$
(17)	(16)	Conv (5×5) + BN + ReLU	$128 \times 4 \times 4$
(18)	(17)	Conv (5×5) + BN + ReLU	$128 \times 4 \times 4$
(19)	(18)	Conv (5×5) + BN + ReLU	$128 \times 4 \times 4$
(20)	(19)	Conv (5×5) + BN + ReLU	$128 \times 4 \times 4$
(21)	(20)	FC + ReLU	2048
(22)	(21)	FC + ReLU	1024
(23)	(22)	FC + ReLU	512
(24)	(23)	FC + ReLU	256
(25)	(24)	FC + ReLU	128
(26)	(25)	FC + ReLU	75
(27)	(26)	FC	25×3

Table 7: Architecture for CNN used to convert touch signals into touch charts.

Index	Input	Operation	Output Shape
(1)	Input	Conv (5×5) + BN + ReLU	$6 \times 256 \times 256$
(2)	(1)	Conv (5×5) + BN + ReLU	$6 \times 254 \times 254$
(3)	(2)	Conv (5×5) + BN + ReLU	$16 \times 126 \times 126$
(4)	(3)	Conv (5×5) + BN + ReLU	$16 \times 124 \times 124$
(5)	(4)	Conv (5×5) + BN + ReLU	$6 \times 122 \times 122$
(6)	(5)	Conv (5×5) + BN + ReLU	$16 \times 120 \times 120$
(7)	(6)	Conv (5×5) + BN + ReLU	$32 \times 59 \times 59$
(8)	(7)	Conv (5×5) + BN + ReLU	$32 \times 57 \times 57$
(9)	(8)	Conv (5×5) + BN + ReLU	$32 \times 55 \times 55$
(10)	(8)	Conv (5×5) + BN + ReLU	$32 \times 53 \times 53$
(11)	(10)	Conv (5×5) + BN + ReLU	$64 \times 26 \times 26$
(12)	(11)	Conv (5×5) + BN + ReLU	$64 \times 24 \times 24$
(13)	(12)	Conv (5×5) + BN + ReLU	$64 \times 22 \times 22$
(14)	(13)	Conv (5×5) + BN + ReLU	$64 \times 20 \times 20$

Table 8: Architecture for perceptual feature pooling in vision and touch setting where features from layers 2, 6, 10, and 14 are extracted.

Index	Input	Operation	Output Shape
(1)	Input	ZN-GCN Layer (C)	$ V \times H$
(2)	(1)	ZN-GCN Layer(C)	$ V \times H$
....			
K-1	(K-2)	ZN-GCN Layer (C)	$ V \times H$
K	(K-1)	GCN Layer	$ V \times O$

Table 9: Architecture for deforming charts positions where H is the chosen hidden dimension size , K is the chosen number of layers, C the percentage of vertex features shared between neighboring vertices in the ZN-GCN layer, and O is the output vertex feature vector size.