

Act the Part: Learning Interaction Strategies for Articulated Object Part Discovery

Samir Yitzhak Gadre¹ Kiana Ehsani² Shuran Song¹
¹ Columbia University ² Allen Institute for AI
<https://atp.cs.columbia.edu>

Abstract

People often use physical intuition when manipulating articulated objects, irrespective of object semantics. Motivated by this observation, we identify an important embodied task where an agent must play with objects to recover their parts. To this end, we introduce Act the Part (AtP) to learn how to interact with articulated objects to discover and segment their pieces. By coupling action selection and motion segmentation, AtP is able to isolate structures to make perceptual part recovery possible without semantic labels. Our experiments show AtP learns efficient strategies for part discovery, can generalize to unseen categories, and is capable of conditional reasoning for the task. Although trained in simulation, we show convincing transfer to real world data with no fine-tuning. A summary video, interactive demo, and code will be available at <https://atp.cs.columbia.edu>.

1. Introduction

How do people and animals make sense of the physical world? Studies from cognitive science indicate observing the consequences of one’s actions plays a crucial role [17, 38, 3]. Gibson’s influential work on affordances argues visual objects ground action possibilities [14]. Work from Tucker *et al.* goes further, suggesting what one sees affects what one does [44]. These findings establish a plausible biological link between seeing and doing. However, in an age of data-driven computer vision, static image and video datasets [40, 24, 2] have taken center stage.

In this paper, we aim to elucidate connections between perception and interaction by investigating articulated object part discovery and segmentation. In this task, an agent must recover part masks by choosing strategic interactions over a few timesteps. We do not assume dense part labels or known kinematic structure [1, 23]. We also do not interact randomly [33]. Rather, we *learn* an agent capable of *holding* and *pushing*, allowing us to relax the assumption that objects are fixed to a ground plane [28]. Our task and

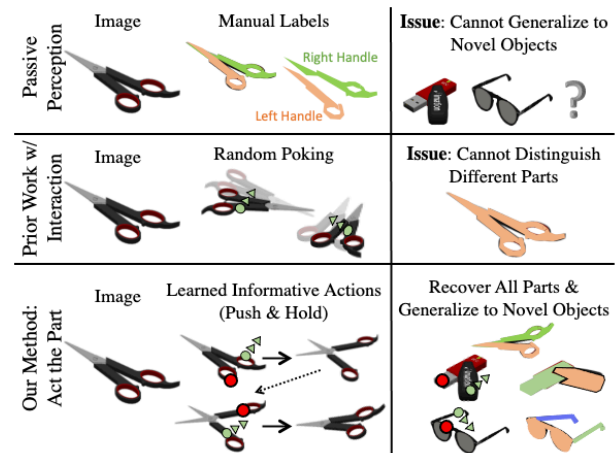


Figure 1. **Interaction for Part Discovery.** Passive part segmentation algorithms require detailed annotation and cannot generalize to new categories. While motion can help discover new objects, prior work cannot infer actions for understanding individual parts. Our work, Act the Part, learns interaction strategies that expose parts and generalize to unseen categories.

approach novelty are highlighted in Fig. 1.

Segmentation from strong supervision and random interaction is widely studied; however, creating *informative* motion to enable category level generalization while relaxing supervision is less explored in the community. We identify the following hurdles, which make this direction salient and difficult. Motion cannot be assumed in a scene as objects seldom move spontaneously. Even with agent interaction, not all actions create perceivable motion to give insight about articulation. Actions might activate only a small number of parts, so diversity of action and aggregation of potentially noisy perceptual discoveries is necessary. Generalization of interaction and perception to unseen categories without retraining or fine-tuning is also desirable. These facets are often overlooked in prior work but are at the heart of this paper.

To address these challenges, we introduce **Act the Part (AtP)**, which takes visual observations, interacts intelligently, and outputs part masks. Our key insight is to couple

action selection and segmentation inference. Given an RGB input image and the part segmentation belief, our interaction network reasons about where to hold and push to move undiscovered parts. By reasoning about changes in visual observations, our perception algorithm is able to discover new parts, keep track of existing ones, and update the part segmentation belief.

We evaluate our approach on eight object categories from the PartNet-Mobility dataset [9, 29, 49] and a ninth multilink category, which we configure with three links. Our experiments suggest: (1) AtP learns effective interaction strategies to isolate part motion, which makes articulated object part discovery and segmentation possible. (2) Our method generalizes to unseen object instances and categories with *different* numbers of parts and joints. (3) Our model is capable of interpretable conditional reasoning for the task—inferring where and how to push given arbitrary hold locations.

We also demonstrate transfer to real images of unseen categories (without fine-tuning) and introduce a toolkit to make PartNet-Mobility more suitable for future research.

2. Related Work

Our approach builds on existing work in interactive perception [6], where visual tasks are solved using agent intervention. We also position our work alongside existing methods in articulated object understanding.

Interactive Perception for Rigid Objects. Instance segmentation of rigid objects from interaction is well studied [13, 5, 46, 32, 7, 33, 12]. Similar work infers physical properties [35, 52] and scene dynamics [30, 50, 45]. These approaches typically employ heuristic or random actions. In contrast, we *learn* to act to expose articulation.

For learning interaction strategies, Lohmann *et al.* [26] learn to interact with rigid objects to estimate their segmentation masks and physical properties. Yang *et al.* [54] learn to navigate to recover amodal masks. These algorithms do not change object internal states in structured ways for articulated object part discovery.

There is also work that leverages multimodal tactile and force inputs [10]. Inspired by this work, we explore using touch feedback in our learning loop. However, we assume only binary signals (e.g., the presence of shear force), which is easier to obtain in real world settings.

Passive Perception for Object Structures. Existing work extracts parts from pairs of images [53, 51], point clouds [55] or videos [41, 27, 25]. In these settings, agents do not have control over camera or scene motion. While the assumption that structures move spontaneously is valid for robot arms or human limbs, the premise breaks down when considering inanimate objects. Even when motion exists, it is not guaranteed to give insight about articulation. We address these issues by learning how to create informative

motion to find and extract parts.

Other work tackles part segmentation from a single image [47, 43, 19, 22, 1, 23] or point clouds [36, 37, 48, 18]. These algorithms are trained with full supervision (e.g., pixel labels) or assume strong category-level priors (e.g., known kinematics or single networks per category). In contrast, our approach uses flow and touch feedback as supervision and makes no class specific assumptions. As a result, we are able to learn *a single model* for *all* our object categories, which encompass diverse kinematic structures.

Interactive Perception for Articulated Objects. In traditional pipelines, agents are carefully programmed to execute informative actions to facilitate visual feature tracking [42, 21, 34]. Other classical approaches improve on action selection for downstream perception [4, 31, 15]. However, these methods assume known object structure, which is used to design heuristics. In contrast, we employ a framework, which allows learning actions directly from pixels without known object models.

Recently, Mo *et al.* [28] present a learnable framework to estimate action affordances on articulated objects from a single RGB image or point cloud. However, they do not consider using their learned interactions for multistep part discovery and segmentation.

3. Approach

Our goal is to learn how to interact with articulated objects to discover and segment parts without semantic supervision. This poses many technical challenges: (1) With repetitive actions, an agent may not explore all parts. (2) Actions resulting in rigid transformations are undesirable. (3) Erroneous segmentation makes tracking parts over time difficult. To begin exploring these complexities, we consider articulated objects in table-top environments.

First, we formally define the task and environment details (Sec. 3.1). We then explain the three components of our approach: an interaction network (Sec. 3.2) to determine what actions to take, a part network (Sec. 3.3) to recover masks from motion, and a history aggregation algorithm (Sec. 3.4) to keep track of discovered parts. Finally, we explain the reward formulation (Sec. 3.5) and combine our modules to present the full pipeline (Sec. 3.6), Act the Part (AtP). Our approach is summarized in Fig. 2.

3.1. Problem Formulation

General Setting. Let \mathcal{O} denote a set of articulated objects, each with $n \leq N$ parts. At each timestep t , an agent gets an observation $I_t \in \mathbb{R}^{H \times W \times C}$, and executes an action $a_t \in \mathcal{A}$ on an object $o \in \mathcal{O}$, where \mathcal{A} is the set of all possible actions. Additional sensor readings $s_t \in \mathbb{R}^l$ complement visual perception. The action results in the next observation I_{t+1} . Given the sequence of T observations, sensor readings, and actions, the goal is to infer part mask




Figure 2. **Model overview.** (a) The interaction network computes hold and push from an image observation and current part memory. The physics simulator gives the action effects yielding the next observation. (b) The part network takes the action and image observations to infer the motion masks for the part that moved, one aligned to the current frame and one to the next frame. (c) The history aggregation module incorporates newly discovered parts and updates existing parts in the memory.

$M_T \in \{1, 2, \dots, N + 1\}^{H \times W}$, where each pixel is assigned a value corresponding to N part labels or background.

Task Details. We consider \mathcal{O} to be a set of common household objects with $n \leq 3$ parts, $T = 5$, $W, H = 90$, and $C = 3$ (RGB). All objects have revolute joints and no fixed base link. Each $a \in \mathcal{A}$ represents a tuple: an image pixel to hold, another pixel to push, and one of eight push directions. The directions are discretized every 45° and are parallel to the ground plane. We take $s_t \in \{0, 1\}^3$, representing binary signals for detecting contact on the hold and push grippers and a binary shear force reading on the hold gripper to emulate touch.

Environment Details. To enable large-scale training and ground truth part segmentation (for benchmarking only), we use a simulated environment. However, we also show our model generalizes to real-world images without fine-tuning. Our simulation environment is built using PyBullet [11] with Partnet-Mobility [9, 29, 49] style dataset assets.

Our environment supports two generic actions. First, a hold action parameterized by its location and implemented as a fixed point constraint between the gripper and a part. Second, a push action parameterized by the location and the direction of the applied force. Actions are easily extensible to facilitate future 2D and 3D object interaction research.

3.2. Learning to Act to Discover Parts

Given a visual observation of an object, we want to create motion by interacting to expose articulation. We give the agent two sub-actions every timestep: hold and push. The action space directly motivates network and reward design.

Conditional Bimanual Action Inference. The interaction task reduces to finding pixels to hold and push and determining the push direction. To decrease the search space, we discretize the push direction into eight options (45° apart). We consider a constant magnitude push force parallel to the ground plane. We condition the predicted push location and direction on the predicted hold location. This allows us to

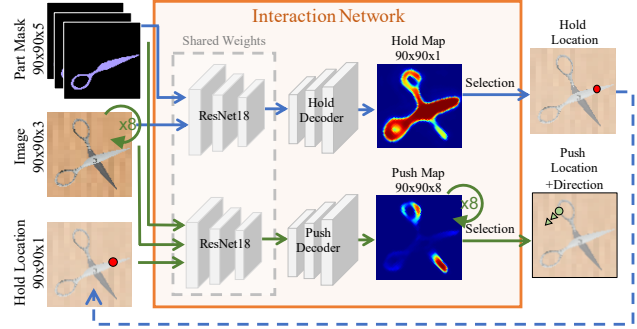


Figure 3. **Interaction network.** Given an image and the current belief of part segmentation, our network predicts a hold and a push conditioned on the hold.

synchronize sub-actions without considering all pairs.

Interaction Network. At every timestep, we predict one step pixel-wise reward for holding and pushing at the spatial resolution of the input image, similar to Zeng *et al.* [56]. As shown in Fig. 3, we use a shared ResNet18 [16] with two residual decoder heads wired with U-Net [39] skip connections. At each timestep t , we have a current belief about the part segmentation. This is represented as a part memory $V_t \in \{0, 1\}^{H \times W \times N}$, where each channel encodes a different part mask. Given an image I_t and V_t , the network predicts a hold reward map $H_t \in [0, 1]^{H \times W}$, where each entry estimates the reward for holding that pixel. We uniformly sample one of the top $k = 100$ pixels from H_t as the hold location. Sampling encourages optimization over the top k actions, which we notice is necessary for the model to learn effective strategies.

Since we wish to infer pushing based on holding, we encode the hold as a 2D Gaussian h_t centered at the hold location with standard deviation of one pixel [20]. In doing so, we can pass the hold location in a manner that preserves its spatial relationship to I_t and V_t . To predict the push reward maps, we pass eight rotations of I_t , V_t , and h_t —every 45° —through the push network. The rotations allow the network to reason implicitly about pushing in all eight directions, while reasoning explicitly only about pushing right [56]. We consider the output map with the largest reward, whose index encodes the push direction, and sample uniformly from the top $k = 100$ actions to choose the push location. An emergent property of our network is conditional reasoning, where hold locations can come from anywhere and the network still reasons about a synchronized push. We demonstrate this capability on real world data in our experiments (Sec. 4.3).

During training, we rollout the current interaction network for seven timesteps for each training instance, which gives the network more opportunities to learn from its mistakes. The images for the last 10 iterations of rollouts are saved in a training buffer. For each image we also buffer




Figure 4. **Part network.** Given a pair of observations and the action that caused the change, this network predicts motion masks aligned to each observation.

optical flow, touch sensors binary output, and intermediate predictions, which are used to define the reward target (Sec. 3.5). We use pixel-wise binary cross entropy loss to learn the hold and push reward maps. Interaction network roll-outs are also a data generation process. All pairs of frames, executed actions, and optical flow ground truth are saved as a dataset to fit the part network.

3.3. Learning to Discover Parts from Action

After an action is executed, we wish to recover the moved part. To do so, we create a part network to predict two masks for the pixels that moved—one mask M_t aligned to the current frame and the other M_{t+1} to the next frame.

Part Network. Our part network (Fig. 4) takes the observations before and after the interaction. Additionally we pass in the hold location h_t and a spatial encoding p_t of the push location and direction. p_t has a 2D Gaussian centered at the push pixel, analogous to h_t . To encode direction, we add Gaussians of smaller mean value in the direction of the push, forming a trail. The network is comprised of a shared encoder with two decoder heads to predict M_t and M_{t+1} . Using consistent forward and backward flow collected during interaction network training, we threshold at zero to acquire target motion masks. We supervise predictions using binary cross-entropy loss.

3.4. History Aggregation

We introduce a history aggregation algorithm to updated part memory V , based on predicted M_t and M_{t+1} . Our algorithm classifies the type of update into four categories: (1) no movement, (2) finding a new part, (3) moving an existing part, (4) entangling parts. These labels are used to decide how to update V and influence the reward (Sec. 3.5).

New Part. If M_t does not overlap significantly with any channels in V , it is likely to be a new part. A free channel c is assigned: $V^c \leftarrow M_{t+1}$. If there is significant overlap between M_t and a mask V^c , relative only to the area of M_t , there is indication two parts are assigned to V^c that must be split: $V^c \leftarrow V^c - (V^c \cap M_t)$ and $V^{c+1} \leftarrow M_{t+1}$. Finding a new part is the most desirable case.

Optical Flow	Touch Sensor	Part Memory	Hold Reward	Push Reward
x	1/0	-	N/A	0
✓	1	New part	1	1
✓	1	Existing part	.5	.5
✓	0	-	0	N/A
✓	1	Entangled part	0	N/A

Table 1. **Reward Calculation.** N/A indicates no backpropagation due to insufficient information. For more details refer to Appx. C.

Existing Part. If there is significant overlap between M_t and a mask V^c , relative to the areas of both M_t and V^c , we execute the update: $V^c \leftarrow M_{t+1}$. This case is less desirable than discovering a new part.

Entangled Parts. If there is significant overlap between M_t and a mask V^c , relative to the area of only V^c , it suggests our action is entangling movement of more than one part. During training; $V^c \leftarrow M_{t+1}$. During testing, we use Iterative Closest Point (ICP) to get the correspondences between V^c and M_{t+1} , yielding $T \in SE(2)$, to execute the updates: $V^c \leftarrow (T \circ V^c) \cap M_{t+1}$, then $V^{c+1} \leftarrow M_{t+1} - V^c$. Entangled part actions are the least desirable, as reflected in our reward described next.

For more details on handling edge cases (e.g., all channels being filled at allocation time), refer to Appx. B.

3.5. Reward

During training, reward for the interaction network is determined from the optical flow, touch feedback, and history aggregation case. The reward conditions and values are shown in Tab. 1.

As presented, reward is sparse; however, we leverage touch and flow to make the reward more dense. If the touch sensor feels no force but flow exists, we know the agent should not hold or push in areas of no flow, which should correspond to the background. We can safely supervise with reward 0 for all such pixels for both hold and push reward maps. If the touch sensor feels a force, flow exists, and we have moved a new or existing part, then we can make the push reward dense. We compute the L2-norm of the flow field and normalize by the max value. If we moved a new part, these values give a dense target for the push map prediction. If we moved an existing part, we scale the target dense push values by the *existing part* reward of 0.5. For more details please see Appx. C.

3.6. Putting Everything Together: Act the Part

We begin by training the interaction network using motion masks from the thresholded flow for history aggregation. We then train our part network using the entire dataset of interactions to learn to infer motion masks. At inference, we first predict and execute an action. We infer motion masks and run history aggregation to output a segmentation mask at every timestep. Further architecture and training details are provided in Appx. D and E.

4. Evaluation

Five Act the Part (AtP) models, trained with different seeds, are evaluated on 20 unseen instances from four seen categories (scissors, knife, USB, safe) and 87 instances from five unseen categories (pliers, microwave, lighter, eyeglasses, and multilink). The multilink objects have three links in a chain similar to eyeglasses. Critically, all train instances have *two links*; however, during testing, we evaluate on objects with *two* and *three links*. See Appx. A.1 for information about the number of instances per category.

To initialize instances, we uniformly sample start position, orientation, joint angles, and scale. Dataset, test initialization, and pre-trained models will be released for reproducibility and benchmarking.

4.1. Metrics and Points of Comparison

For each test data point, we allow the agent to interact with the object five times. We collect three perceptual metrics to evaluate performance on part discovery and segmentation. Two additional metrics measure effectiveness of the actions for part discovery. Let \mathcal{G} , \mathcal{H} denote the sets of ground truth and predicted binary part masks respectively.

Average Percentage Error (APE). To measure errors in number of parts discovered, we compute $|\mathcal{G} - \mathcal{H}|/|\mathcal{G}|$.

Part-aware Intersection over Union (IoU). We use Hungarian matching to solve for the maximal IoU bipartite match between \mathcal{G} and \mathcal{H} . Unmatched parts get IoU of 0. Final IoU is determined by summing part IoUs and dividing by $\max(|\mathcal{G}|, |\mathcal{H}|)$. The metric penalizes both errors in mask prediction and failure to discover masks (e.g., if one of two parts is discovered, maximum IoU is 50%).

Part-aware Hausdorff distance @ 95% (d_{H95}). We notice IoU is sensitive for thin structures. For example, a small pixel shift in a thin rod can lead to IoU of 0. To provide a better metric for these structures, we measure d_{H95} , which is a part-aware variant of a common metric in medical image segmentation [8]. The directed Hausdorff distance @ 95% between some masks $G \in \mathcal{G}$ and $H \in \mathcal{H}$ is $d_{H95}^d(G, H) := P_{95} \min_{h \in H} \|g - h\|_2$ where P_{95} gives the 95-th percentile value over pixel distances.

The metric is robust to a small number of outliers, which would otherwise dominate. The symmetric measure is given as $d_{H95}(G, H) := \max(d_{H95}^d(G, H), d_{H95}^d(H, G))$. We use Hungarian matching to find minimal d_{H95} bipartite matches between \mathcal{G} and \mathcal{H} . If $|\mathcal{G}| \neq |\mathcal{H}|$, we compute the distance of unmatched parts against a matrix of ones at the image resolution. Distances are summed and normalized by $\max(|\mathcal{G}|, |\mathcal{H}|)$.

Effective Steps. A step is effective if the hold is on an object link, the push is on another link, and the action creates motion.

Optimal Steps. An interaction is optimal if it is effective

and a new part is discovered. If all the parts have already been discovered, moving a *single* existing part in the interaction is not penalized.

We compute the average of perceptual metrics for each category at every timestep over five models trained with different random seeds. Hence IoU, APE, and d_{H95} yield mIoU, MAPE, and \bar{d}_{H95} . For evaluation in Tab. 2, we consider metrics after the fifth timestep. Efficient and optimal step scores are averaged for each category over all timesteps (in contrast to being considered only at the fifth timestep).

Baselines and Ablations. We compare the AtP framework trained with and without touch reward, [Ours-Touch] and [Ours-NoTouch] respectively, with the following alternative approaches to study the efficacy of our interaction network. All methods use the same part network trained from the full AtP rollouts:

- Act-Random: hold and push locations and the push direction are uniformly sampled from the action space.
- Act-NoHold: The agent applies a single push action every step. Single modes of interaction are widely used in interactive segmentation algorithms [13, 5, 46, 32, 33, 12]; however, this ablation differs from these works as push is learned for the part segmentation task.
- Act-NoPart: The interaction network does not take the part memory and considers each moved part as a new part for reward calculation.

For the modified reward used to train the above networks see Appx. C. We also design three oracle algorithms using simulation state to provide performance upper bounds:

- GT-Single: GT object mask as output. This gives an upper bound for part-unaware algorithms that segment objects from background via interaction.
- GT-Act: Optimal step based on ground truth state, but use of AtP part network for the mask inference. This is conceptually similar to [34], which uses expert actions for part segmentation.
- GT-Act-Mot: Optimal step based on ground truth state with motion masks from the ground truth flow.

4.2. Benchmark Results

To validate the major design decisions, we run a series of quantitative experiments in simulation to compare different algorithms. We also provide qualitative results in Fig. 5. In Sec. 4.3, we evaluate our model on real world data.

Does Interaction Help Part Discovery? First we want to validate if AtP learns effective interaction strategies for part discovery by accumulating information over time. To evaluate, we plot the part mIoU w.r.t. interaction steps in Fig. 6. As expected, the upper bounds peaks at 2 and 3 steps for pliers and multilink, respectively. While other algorithms' performance saturate quickly with one or two interactions, [Ours-Touch] and [Ours-NoTouch] are able to




Figure 5. **Qualitative Results.** Our network learns a policy to interact with unseen objects and categories. While it is only trained on objects with two part, it also learns to reason about three part objects. Due to space limitation, only three interaction steps are shown in this figure. For more results, please refer to our project website: <https://atp.cs.columbia.edu>.

improve with more interactions. These plots indicate that while the learned interaction strategies may not be optimal (compared to upper bounds using ground truth state), they are informative for discovering new parts of the object and self-correct errors over time. Results from other categories

are presented in Appx. G, where we see all AtP curves approach the upper bounds.

Does Part Prediction Help with Action Selection? Our interaction network takes the current belief of the part segmentation as input and obtains reward for new part discov-

Method	Seen Categories (Novel Instances)				Unseen Categories					
	Scissors ↓/↓/↑	Knife ↓/↓/↑	USB ↓/↓/↑	Safe ↓/↓/↑	Pliers ↓/↓/↑	Microwave ↓/↓/↑	Lighter ↓/↓/↑	Eyeglasses ↓/↓/↑	Multilink ↓/↓/↑	
GT-Single	0.5 / 36.2 / 27.8	0.5 / 42.5 / 32.7	0.5 / 31.7 / 29.5	0.5 / 39.4 / 44.3	0.5 / 47.3 / 28.1	0.5 / 38.9 / 43.3	0.5 / 38.2 / 34.1	0.66 / 56.2 / 18.4	0.66 / 51.5 / 14.2	
GT-Act	0.01 / 4.3 / 78.0	0.02 / 4.5 / 81.6	0.0 / 5.7 / 82.7	0.02 / 2.1 / 89.7	0.01 / 4.3 / 78.4	0.03 / 2.4 / 87.8	0.0 / 3.3 / 88.0	0.03 / 7.4 / 64.6	0.10 / 7.2 / 75.3	
GT-Act+Mot	0.0 / 1.6 / 88.4	0.0 / 0.9 / 92.9	0.0 / 2.4 / 91.5	0.0 / 0.6 / 91.7	0.0 / 0.9 / 92.7	0.0 / 0.4 / 94.2	0.0 / 0.6 / 94.8	0.0 / 4.2 / 82.5	0.0 / 5.2 / 86.9	
Act-Random	0.62 / 36.2 / 22.0	0.63 / 41.6 / 24.1	0.47 / 26.8 / 33.1	0.62 / 40.8 / 32.5	0.56 / 39.9 / 25.0	0.58 / 37.9 / 36.0	0.62 / 40.4 / 24.5	0.70 / 53.5 / 12.5	0.78 / 52.6 / 10.7	
Act-NoHold	0.46 / 34.4 / 28.5	0.43 / 35.5 / 35.2	0.40 / 30.0 / 33.2	0.38 / 32.1 / 42.7	0.41 / 35.4 / 30.5	0.41 / 31.3 / 40.7	0.45 / 39.6 / 34.3	0.40 / 41.6 / 19.6	0.53 / 48.0 / 19.9	
Act-NoPart	0.25 / 15.3 / 53.8	0.29 / 20.4 / 51.6	0.44 / 19.5 / 47.6	0.37 / 18.3 / 49.1	0.34 / 20.6 / 48.5	0.43 / 19.7 / 46.3	0.49 / 26.4 / 42.0	0.33 / 27.1 / 34.4	0.40 / 31.0 / 39.7	
Ours-NoTouch	0.19 / 13.1 / 58.0	0.16 / 13.2 / 66.2	0.14 / 10.4 / 68.6	0.15 / 9.9 / 75.0	0.15 / 12.5 / 59.8	0.14 / 9.2 / 74.1	0.22 / 14.5 / 64.4	0.28 / 26.2 / 37.9	0.25 / 24.6 / 46.6	
Ours-Touch	0.10 / 8.5 / 65.6	0.16 / 12.2 / 65.9	0.09 / 8.3 / 75.3	0.17 / 10.1 / 74.2	0.13 / 9.7 / 64.9	0.14 / 8.3 / 75.4	0.25 / 15.1 / 62.8	0.24 / 21.8 / 43.0	0.22 / 20.0 / 54.7	

Table 2. **Perception performance.** MAPE [frac] / \bar{d}_{H95} [pixels] / mIoU [%]. Image resolution is 90×90 . Numbers are evaluated after the fifth interaction. Numbers are averaged over five models trained with different seeds.




Figure 6. **IoU w.r.t. Interaction Steps.** Results on two unseen object categories show our methods (pink and brown) approach the oracle baseline over time.




Figure 7. **Effective and Optimal Steps.** Our method learns an efficient policy that chooses optimal steps (i.e., actions that discover new parts) more frequently than other approaches.

ery. We hope this design would encourage the algorithm to focus on selecting actions that provide information gain (e.g., push new parts to discover them). To validate this design, we compare AtP to an ablated version, [Act-NoPart], which is not mask-aware. Interestingly, this model performs efficient actions at roughly the same rate as [Ours-Touch] (Fig. 7); however, [Ours-Touch] is better at finding optimal steps (resulting in new part discovery). Histograms for all other categories are presented in Appx. G and corroborate these findings. This result is also supported in Tab. 2, which shows degradation on all perceptual metrics when part-awareness is not exploited.

Is Holding Necessary? In contrast to a vast majority of prior work that use simple push actions, our algorithm uses bimanual actions for object interaction (i.e., simultaneous hold and push). Our hypothesis is that such actions give the system a better chance at disentangling motion between

different moving parts and therefore aid part discovery. To validate this hypothesis, we compare our algorithm with an agent that performs only push actions [Act-NoHold]. The result in Tab. 2 shows that without the hold action the system performance is much worse at part segmentation. [Act-NoHold] has trouble discovering more than one object part, since the whole object is likely to be moved during interaction. Furthermore, this result suggests more complex perceptual modules are necessary to get push-only policies to achieve competitive performance at this task. While this is an interesting direction, disentangling the motion of many moving parts is non-trivial and out of scope for this paper.

Does Touch Feedback Help? In this experiment, we want to evaluate the effect of touch feedback. Looking at Tab. 2, we see that [Ours-Touch] outperforms [Ours-NoTouch] in most categories. A similar trend is noticeable when looking at action performance in Figs. 6 and 7. We conjecture this is due to the benefit of using touch signal to define more specific reward cases, which is ultimately reflected in the final system performance. However, we are still able to learn helpful interaction strategies even without touch.

Generalization to Unseen Objects and Categories. Our algorithm does not make category-level assumptions, therefore the same policy and perception model should work for unseen object categories with different kinematic structures. More specifically, we wish to probe generalization capabilities of our model to unseen instances from seen categories and novel categories.

The algorithm’s generalizability is supported by results in Tab. 2, where mIoU, MAPE, and \bar{d}_{H95} are comparable for seen and unseen categories. Performance on eyeglasses is slightly worse, however, still impressive as our model is only trained on instances with two links. Furthermore, for eyeglasses, MAPE value falls under 0.33, suggesting the model finds the three parts in most cases. IoU performance on the multilink category is better than on eyeglasses; however, MAPE is comparable, suggesting that eyeglasses are particularly challenging for reasons other than having three links. These results support that our method learns to interact intelligently and reason about motion in spite of differing shape, texture, or structure in the test objects.




Figure 8. **Conditional Action on Real Images.** (a) Varying the hold location, we observe the model is able to reason where to push right. (b) Fixing the hold location, we observe the model reasons about a good direction to push (i.e., top left).

4.3. Real World Results

In these experiments, we want to validate [Ours-Touch] performance on real world data. Since our algorithm does not need prior knowledge of objects or special sensory input during inference, we *directly* test our learned model on real world RGB images of unseen categories taken by smartphone cameras. To build a pipeline that demonstrates the viability of our model on real world data, a camera is positioned over an articulated object and an image is captured. Our trained model runs interaction inference, predicts hold and push actions, and provides a human operator with instructions on what to execute. A next frame image is sent back to the model, at which point it runs the part network, history aggregation, and another round of interaction inference. More details and a discussion about experimental limitations can be found in Appx. F.

Conditional Action Reasoning. We visualize the conditional action inference result from the interaction network on real world images. Fig. 8 shows two types of visualizations. In example (a), we pick various hold positions and analyze the “push right” reward prediction maps (recall: pushing is conditioned on holding). We notice that the affordance prediction switches between the links depending on the hold location, which indicates the network’s understanding about the object structure. When hold is placed in free space or between the links, the push reward predictions are not confident about pushing anywhere. These results suggest that our model is able to disentangle push predictions from its own hold predictions, thereby demonstrating a form of conditional reasoning.

In example (b), we further probe the model’s reasoning about the push direction by visualizing different push maps for the same holding position. Among all directions, the network infers the highest score on the top-left rotation, which would push the scissors open. The result suggests that the algorithm is able to pick a push direction that would lead to informative motion, when reasoning over many options.

Interaction Experiment. Next, we evaluate both perception and interaction networks together with the real world




Figure 9. **Interaction Experiment.** Images were captured and sent to AtP, where hold and push are predicted. A human agent executes actions following the instructions. From pairs of images, masks are recovered and aggregated by AtP, and the next instruction is given. This is especially interesting because the lighting, shadows, and artifacts of the images taken with a phone are different from our simulated environment (i.e., sim2real gap).

physical interactions. To validate performance independent of robot execution accuracy, a human is instructed to execute the actions. Fig. 9 shows the predicted actions, affordances and final object part masks discovered by the algorithm. Without any fine-tuning, the algorithm shows promising results on inferring interaction strategies and reasoning about the observed motion for part discovery. Please refer to Appx. G for more real world experiment results and failure case analysis.

5. Conclusion and Future Work

We present Act the Part (AtP) to take visual observations of articulated objects, interact strategically, and output part segmentation masks. Our experiments suggest: (1) AtP is able to learn efficient strategies to isolate and discover parts. (2) AtP generalizes to novel categories of objects with unknown and unseen number of links—in simulation and the real world. (3) Our model demonstrates conditional reasoning about how to push based on arbitrary hold locations. We see broad scope for future work including extensions to 3D part segmentation and singular frameworks for rigid, articulated, and deformable object understanding. We hope this paper will inspire others in the vision and robotics communities to investigate perception and interaction in tandem.

Acknowledgements. Thank you Shubham Agrawal, Jessie Chapman, Cheng Chi, the Gadres, Bilkit Githinji, Huy Ha, Gabriel Ilharco Magalhães, Sarah Pratt, Fiadh Sheeran, Mitchell Wortsman, and Zhenjia Xu for valuable conversations. This work was supported in part by the Amazon Research Award and NSF CMMI-2037101.

References

- [1] Ben Abbatematteo, Stefanie Tellex, and George Konidaris. Learning to generalize kinematic models to novel objects. *CoRL*, 2019. 1, 2
- [2] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. 1
- [3] Renée Baillargeon. Infants’ physical world. *Current Directions in Psychological Science*, 13(3), 2004. 1
- [4] Patrick R. Barragan, Leslie Kaelbling, and Tomás Lozano-Pérez. Interactive bayesian identification of kinematic mechanisms. *ICRA*, 2014. 2
- [5] Mårten Björkman and Danica Kragic. Active 3d scene segmentation and detection of unknown objects. *ICRA*, 2010. 2, 5
- [6] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *T-RO*, 2017. 2
- [7] Arunkumar Byravan and Dieter Fox. Se3-nets: Learning rigid body motion using deep neural networks. *ICRA*, 2017. 2
- [8] Vikram Chalana and Yongmin Kim. A methodology for evaluation of boundary detection algorithms on medical images. *IEEE Transactions on Medical Imaging*, 16(5):642–652, 1997. 5
- [9] 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. 2, 3, 11
- [10] Vivian Chu, Ian McMahan, Lorenzo Riano, Craig G. McDonald, Qin He, Jorge Martinez Perez-Tejada, Michael Arigo, Trevor Darrell, and Katherine J. Kuchenbecker. Robotic learning of haptic adjectives through physical interaction. *Robotics and Autonomous Systems*, 63, 2015. 2
- [11] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. 2016. 3, 11
- [12] Andreas Eitel, Nico Hauff, and Wolfram Burgard. Self-supervised transfer learning for instance segmentation through physical interaction. *IROS*, 2019. 2, 5
- [13] Paul M. Fitzpatrick. First contact: an active vision approach to segmentation. *IROS*, 2003. 2, 5
- [14] James J Gibson. The ecological approach to visual perception. *Psychology Press*, 1979. 1
- [15] Karol Hausman, Scott Niekum, Sarah Osentoski, and Gaurav S. Sukhatme. Active articulation model estimation through interactive perception. *ICRA*, 2015. 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 3, 12
- [17] Richard Held and Alan Hein. Movement-produced stimulation in the development of visually guided behavior. *Journal of comparative and physiological psychology*, 56, 1963. 1
- [18] Jiahui Huang, He Wang, Tolga Birdal, Minhyuk Sung, Federica Arrigoni, Shi-Min Hu, and Leonidas Guibas. Multibodysync: Multi-body segmentation and motion estimation via 3d scan synchronization. *arXiv preprint arXiv:2101.06605*, 2021. 2
- [19] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. *CVPR*, 2019. 2
- [20] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. *NeurIPS*, 2018. 3
- [21] Dov Katz, Moslem Kazemi, J. Andrew Bagnell, and Anthony Stentz. Interactive segmentation, tracking, and kinematic modeling of unknown 3d articulated objects. *ICRA*, 2013. 2
- [22] Timothy E. Lee, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Oliver Kroemer, Dieter Fox, and Stan Birchfield. Camera-to-robot pose estimation from a single image. *ICRA*, 2020. 2
- [23] Xiaolong Li, He Wang, Li Yi, Leonidas Guibas, A. Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. *CVPR*, 2020. 1, 2
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft coco: Common objects in context. *ECCV*, 2014. 1
- [25] Qihao Liu, Weichao Qiu, Weyao Wang, Gregory D. Hager, and Alan L. Yuille. Nothing but geometric constraints: A model-free method for articulated object pose estimation. *arXiv preprint arXiv:2012.00088*, 2020. 2
- [26] Martin Lohmann, Jordi Salvador, Aniruddha Kembhavi, and Roozbeh Mottaghi. Learning about objects by learning to interact with them. *NeurIPS*, 2020. 2
- [27] Roberto Martín Martín and Oliver Brock. Online interactive perception of articulated objects with multi-level recursive estimation based on task-specific priors. *IROS*, 2014. 2
- [28] Kaichun Mo, Leonidas Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. *arXiv preprint arXiv:2101.02692*, 2021. 1, 2
- [29] Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. *CVPR*, 2019. 2, 3, 11
- [30] Iman Nematollahi, Oier Mees, Lukas Hermann, and Wolfram Burgard. Hindsight for foresight: Unsupervised structured dynamics models from physical interaction. *IROS*, 2020. 2
- [31] Stefan Otte, Johannes Kulick, Marc Toussaint, and Oliver Brock. Entropy-based strategies for physical exploration of the environment’s degrees of freedom. *IROS*, 2014. 2
- [32] Joni Pajarinen and Ville Kyrki. Decision making under uncertain segmentations. *ICRA*, 2015. 2, 5
- [33] Deepak Pathak, Yide Shentu, Dian Chen, Pulkit Agrawal, Trevor Darrell, Sergey Levine, and Jitendra Malik. Learning instance segmentation by interaction. *CVPRW*, 2018. 1, 2, 5

- [34] Sudeep Pillai, Matthew Walter, and Seth Teller. Learning articulated motions from visual demonstration. *RSS*, 2014. 2, 5
- [35] Lerrel Pinto, Dhiraj Gandhi, Yuanfeng Han, Yong-Lae Park, and Abhinav Gupta. The curious robot: Learning visual representations via physical interactions. *ECCV*, 2016. 2
- [36] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CVPR*, 2017. 2
- [37] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017. 2
- [38] Gwendolyn E. Roberson, Mark T. Wallace, and James A. Schirillo. The sensorimotor contingency of multisensory localization correlates with the conscious percept of spatial unity. *Behavioral and Brain Sciences*, 24(5), 2001. 1
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *MICCAI*, pages 234–241, 2015. 3, 13
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 2015. 1
- [41] Tanner Schmidt, Richard A. Newcombe, and Dieter Fox. Dart: Dense articulated real-time tracking. *RSS*, 2014. 2
- [42] Jürgen Sturm, Cyrill Stachniss, and Wolfram Burgard. A probabilistic framework for learning kinematic models of articulated objects. *JAIR*, 2011. 2
- [43] Stavros Tsogkas, Iasonas Kokkinos, George Papandreou, and Andrea Vedaldi. Semantic part segmentation with deep learning. *arXiv preprint arXiv:1505.02438*, 2015. 2
- [44] Mike Tucker and Rob Ellis. On the relation between seen objects and components of potential actions. *Journal of experimental psychology. Human perception and performance*, 24, 1998. 1
- [45] Hsiao-Yu Fish Tung, Zhou Xian, Mihir Prabhudesai, Shomit Lal, and Katerina Fragkiadaki. 3d-oes: Viewpoint-invariant object-factorized environment simulators. *CoRL*, 2020. 2
- [46] Herke van Hoof, Oliver Kroemer, and Jan Peters. Probabilistic segmentation and targeted exploration of objects in cluttered environments. *T-RO*, 2014. 2, 5
- [47] Jianyu Wang and Alan Yuille. Semantic part segmentation using compositional model combining shape and appearance. *CVPR*, 2015. 2
- [48] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *TOG*, 2019. 2
- [49] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. *CVPR*, 2020. 2, 3, 11
- [50] Zhenjia Xu, Zhanpeng He, Jiajun Wu, and Shuran Song. Learning 3d dynamic scene representations for robot manipulation. *CoRL*, 2020. 2
- [51] Zhenjia Xu, Zhijian Liu, Chen Sun, Kevin Murphy, William Freeman, Joshua Tenenbaum, and Jiajun Wu. Unsupervised discovery of parts, structure, and dynamics. *ICLR*, 2019. 2
- [52] Zhenjia Xu, Jiajun Wu, Andy Zeng, Joshua B Tenenbaum, and Shuran Song. Densephysnet: Learning dense physical object representations via multi-step dynamic interactions. *RSS*, 2019. 2
- [53] Jingyu Yan and Marc Pollefeys. A general framework for motion segmentation: Independent, articulated, rigid, non-rigid, degenerate and non-degenerate. *ECCV*, 2006. 2
- [54] Jianwei Yang, Zhile Ren, Mingze Xu, Xinlei Chen, David J. Crandall, D. Parikh, and Dhruv Batra. Embodied amodal recognition: Learning to move to perceive objects. *ICCV*, 2019. 2
- [55] Li Yi, Haibin Huang, Difan Liu, Evangelos Kalogerakis, Hao Su, and Leonidas Guibas. Deep part induction from articulated object pairs. *TOG*, 37(6), 2019. 2
- [56] Andy Zeng, Shuran Song, Stefan Welker, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. Learning synergies between pushing and grasping with self-supervised deep reinforcement learning. *IROS*, 2018. 3