

Understanding 3D Object Interaction from a Single Image

Shengyi Qian[†]
[†]University of Michigan
 syqian@umich.edu

David F. Fouhey^{†‡}
[‡]New York University
 david.fouhey@nyu.edu

<https://jasonqsy.github.io/3DOI/>

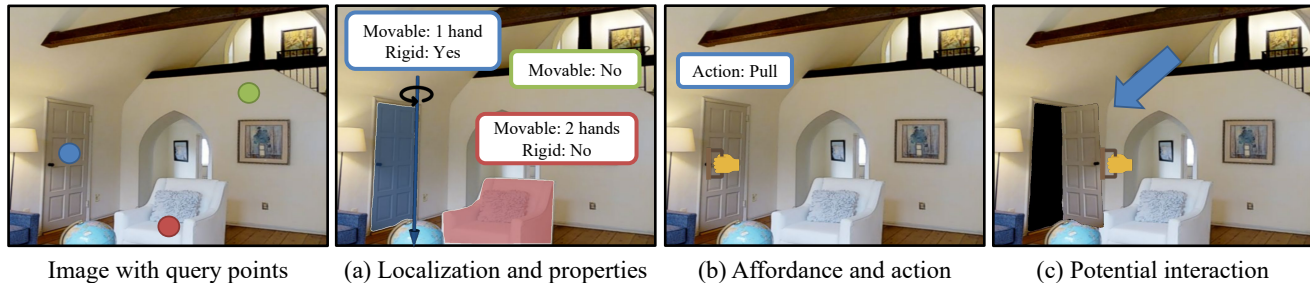


Figure 1. Given a single image and a set of query points \bullet \circ \bullet , our approach predicts: (a) whether the object at the location can be moved \updownarrow , its rigidity \updownarrow and articulation class \updownarrow , and location \updownarrow ; (b) an affordance \updownarrow and action \updownarrow ; and (c) potential 3D interaction for articulated objects. This ability can assist intelligent agents to better manipulate objects or explore the 3D scene.

Abstract

Humans can easily understand a single image as depicting multiple potential objects permitting interaction. We use this skill to plan our interactions with the world and accelerate understanding new objects without engaging in interaction. In this paper, we would like to endow machines with the similar ability, so that intelligent agents can better explore the 3D scene or manipulate objects. Our approach is a transformer-based model that predicts the 3D location, physical properties and affordance of objects. To power this model, we collect a dataset with Internet videos, egocentric videos and indoor images to train and validate our approach. Our model yields strong performance on our data, and generalizes well to robotics data.

1. Introduction

What can you do in Figure 1? This single RGB image conveys a rich, interactive 3D world where you can interact with many objects. For instance, if you grab the chair with two hands, you can move it as a rigid object; the pillow can be picked up freely and squished; and door can be moved, but only rotated. This ability to recognize and interpret potential affordances in scenes helps humans plan our interactions and more quickly learn to interact with objects. The goal of this work is to give the same ability to computers.

Obtaining such an understanding of potential interac-

tions from a single 3D image is beyond the current state of the art in scene understanding because it spans multiple disparate subfields of computer vision. For instance, single image 3D has made substantial progress [50, 47, 68, 18], but primarily focuses on the scene *as it exists*, as opposed to *as it could be*. There has been an increasing interest in understanding articulation [34, 65, 49], but these works primarily focus on articulation *as it occurs* in a 3D model or carefully collected demonstrations, instead of *as it could occur*. Finally, while there is long-standing work on enabling robots to learn interaction and potential interaction points [48, 56], these works focus primarily on evaluation in primarily the same environment (*e.g.* the lab) and do not focus on applying the understanding in entirely new environments.

We propose to bootstrap this interactive understanding by developing (1) a problem formulation, (2) a rich dataset of annotations on challenging images, and (3) a transformer-based approach. We frame the problem of recognizing the articulation as a prediction-at-a-query-location problem: given an image and 2D location, our method aims to answer “what can I do here?” in the style of classic point-and-click games like *Myst*. We frame “what can I do here” via a set of common questions: whether the object can be moved, its extent when moved and location in 3D, rigidity, whether there are constraints on its motion, as well as estimates of how one would interact the object. To maximize the potential for downstream transfer, our questions are chosen to be generic rather than specific to particular hands or end-effectors: knowing where to act or the degrees

of freedom of an object may accelerate reinforcement learning even if one must still learn end-effector-specific skills.

In order to tackle the task, we introduce a transformer-based model. Our approach, described in Section 5 builds on a detection backbone such as Segment-Anything [30] in order to build on the advances and expertise of object detection. We extend the backbone with additional heads that predict each of our “what I can I do here” tasks, and which can be trained end-to-end. As an advantage of our formulation, we can train the system on sparse annotations; we believe this will be helpful for eventually converting our direct supervision to supervision via video.

Powering our approach is a new dataset, described in Section 4, which we name the 3D Object Interaction dataset (3DOI). In order to maximize the likelihood of generalizing to new environments, the underlying data comes from diverse sources, namely Internet and egocentric videos as well as 3D renderings of scene layouts. We provide annotations of our tasks on this data and, due to the source of the data, we also naturally obtain 3D supervision in the form of depth and normals. In total, the dataset has over 50K objects across 10K images, as well as over 31K annotations of non-interactable objects (e.g., floor, wall).

Our experiments in Section 6 test how well our approach recognizes potential interaction, testing on both unseen data in 3DOI as well as robotics data. We compare with a number of alternatives, including generalizing from data of demonstrations [49, 46] and synthetic data [65], as well alternate network designs. Our approach outperforms these models and shows strong generalization to the robotics dataset WHIRL [1].

To summarize, we see our primary contributions as: (1) the novel task of detecting 3D object interactions from a single RGB image; (2) 3D Object Interaction dataset, which is the first large-scale dataset containing objects that can be interacted and their corresponding locations, affordance and physical properties; (3) A transformer-based model to tackle this problem, which has strong performance on the 3DOI dataset and robotics data.

2. Related Works

Our paper proposes to extract 3D object interaction from a single image. This problem lies at the intersection of 3D vision, object detection, human-object interaction and scene understanding. It is also closely related to downstream robotics applications.

Interactive scene understanding. Recently, the computer vision community is increasingly interested in understanding 3D dynamics of objects. It is motivated by human-object interaction [4, 17, 53], although humans do not need to be present in our setting. Researchers try to understand the 3D shapes, axes, movable parts and affordance on synthetic data [44, 65, 45, 26, 63, 34, 61, 63],

videos [49, 21, 19, 46, 41, 33] or point clouds [27, 25]. Our work is mainly related to [49, 21, 19] since they work on real images, but is different from them on two aspects: (1) they need video or multi-view inputs, but our input is only a single image; (2) their approaches recover objects which are being interacted, while our approach understands potential interactions before any interactions happen. Finally, OPD [26, 57] tackles a similar problem for articulated objects, but ours also work for non-articulated objects.

Object detection. The training anchor-based object detection pipeline basically follows the pipeline of Mask R-CNN [23, 31, 51, 29]. As the development of transformer-based models goes, DETR [3], AnchorDETR [62] and MaskFormer [6] approach object detection as a direct set prediction problem. Recently, Kirillov *et al.* proposes Segment Anything Model [30], which predicts object masks from input prompts such as points or boxes. Our network needs to be built on decoder-based backbones [3, 6, 30], and we choose SAM [30] due to its state-of-the-art performance.

Single image 3D. Since our problem requires us recover 3D object interaction instead of 2D from a single image, it is also related to single image 3D. In the recent few years, researchers have developed many different approaches to recover 3D from a single image, including depth [68, 50, 36, 5, 14], surface normals [59, 15], 3D planes [40, 39, 28] and shapes [8, 42, 18, 47]. Our work is built upon their works. Especially, our architecture is motivated by DPT [50] which trains ViT for both segmentation and depth estimation.

Robotics manipulation. Manipulation of objects is a long-term goal of robotics. Researchers have developed various solutions for different kinds of objects in different scenes, ranging from articulated objects [56, 48, 9, 11, 64, 20] to deformable objects [66, 67, 60, 7]. While manipulation is not the goal of our paper, understanding objects and the environment in 3D is typically an important part of a manipulation pipeline. Our paper mainly improves the perception part, which can potentially improve manipulation. Therefore, we also test our approach on robotics data [1], to show it can generalize.

3. Overview

Given a single image, our goal is to be able to answer “What could I do here?” with the object at a query point. We introduce annotations in Section 4 as well as a method for the task in Section 5. Before we do so, we present a unified explanation for the questions we answer as well as the rationale for choosing these questions. We group our questions into six property types, some of which are further subdivided. Not all objects support all questions: objects that cannot be moved, for instance, do not have other properties and objects that can be freely moved do not have rotation axes. We further note that some objects defy these




Figure 2. Example annotations of our 3DOI dataset. Our images come from Internet videos [49], egocentric videos [10] and renderings of 3D dataset [13]. ● is the query point, and ▼ is the affordance.

properties – ball joints, for example, permit a 2D subspace of motion – our goal is to identify a large subspace of potential interactions.

Movable 🖐️ The most important subdivision is whether the object at the query point can be moved. This follows work in both 3D scene understanding [55] and human-object interaction [53] that subdivide objects into how movable they are. We group objects into three categories based on how easily the object can be moved: (1) *fixtures* which effectively cannot be moved, such as walls and floor; (2) *one hand* objects that can be moved with a single hand, such as a water bottle or cabinet door; (3) *two hand* objects that require two hands to move, such as a large TV. We frame the task as three-way classification.

Localization 📍 Understanding the extent of an object is important, and so we localize the object in the world. Since our objects consist of a wide variety of categories, we frame localization as 2D instance segmentation as in [23, 3], as well as a depthmap to localize the object in 3D [50, 68]. These properties can be estimated for most objects.

Rigidity 🛠️ To understand action, one primary distinction is rigid-vs-non-rigid since rigid objects are subject to substantially simpler rules of motion [35]. We therefore classify whether the object is rigid or not.

Articulation 📐 Most rigid objects can further decomposed as permitting freeform, rotational / revolute, or translation / prismatic motion [56]. Each of these requires different end-effector interactions to effectively interact with. We frame the articulation category as a three-way classification problem, and recognizing the rotation axis as a line prediction problem following [49].

Action 🖐️ We also want to understand what the potential action could be to interact with the object. Here we focus on three types of actions: pull, push or other.

Affordance 📍 Finally, we want to know where we should interact with the object. For example, we need to manipulate the handle if we want to open a door. We predict a probability map which is over the location of the affordance.

4. 3D Object Interaction Dataset

One critical component of our contribution is accurate annotations of object interactions, as there is no publicly available data. In this paper, we introduce 3D Object Interaction dataset (3DOI), which is the first dataset. We picked data that can be easily integrated with 3D, including a 3D dataset, so that we have accurate 3D ground truth to train our approach. Examples of our data are shown in Figure 2.

Images. Our goal is to pick up diverse images representing real-world scenarios. In particular, we want our images contain a lot of everyday objects we can interact with. Therefore, we sample 10K images from a collection of publicly available datasets: (1) Articulation [49] comes from third-person Creative Commons Internet videos. Typically, a video clip contains humans manipulating articulated objects in households. We randomly sample 3K images from the articulation dataset; (2) EpicKitchen [10] contains egocentric videos making foods in kitchen environments. We sample 2K images from EpicKitchen; (3) Taskonomy [69] is an indoor 3D dataset with real 2D image and corresponding 3D ground truth. We use the renderings by Omnidata [13]. We sample 5k images from the taskonomy split of Omnidata starter dataset. Overall, there are 10K images.

Annotation. With a collection of images with potential objects we can interact, we then turn to manual annotation. For a single image, we select around 5 interactable query points, including both large and small objects. For each query point, we annotate: (*Movable* 🖐️) one hand, two hand, or fixture. (*Localization* 📍) The bounding box and mask of the part this point belonging to. (*Rigidity* 🛠️) Rigid, or nonrigid. (*Articulation* 📐) Rotation, translation or freeform. We also annotate their rotation axes. (*Action* 🖐️) Pull, push or others. (*Affordance* 📍) A keypoint which indicates where we should interact with the object. At the same time, our taskonomy [69] images come with 3D ground truth, including depth and surface normals. We also annotate 31K query points of fixtures. Finally, we split 10K images into a train/val/test set of 8k/1k/1k split, respectively.

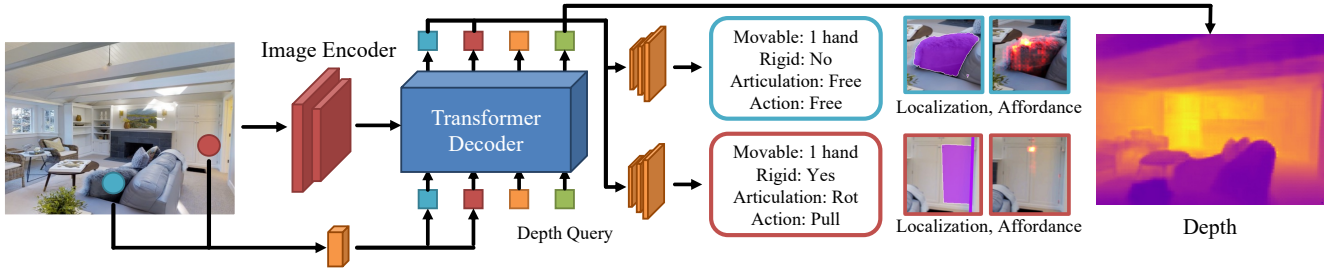


Figure 3. Overview of our approach. The inputs of our network is a single image and a set of query points \bullet . For each query point, it predicts the potential 3D interaction, in terms of movable ✍ , location 📍 , rigidity 🔪 , articulation 🔧 , action 👉 and affordance 👉 . In addition, the input of transformer decoder includes a learnable depth query, which estimates the dense depth to recover 3D object interaction for articulated objects.

Availability and Ethics. Our images come from three publicly available datasets. Taskonomy does not contain any humans. The video articulation dataset comes from Creative Commons Internet videos. We do not foresee any ethical issues in our dataset.

5. Approach

We now introduce a model which can take an image and a set of query point and answer all of questions we asked in Section 3, including movable, localization, rigidity, articulation, action and affordance. A brief overview of our approach is shown in Figure 3.

Since our inputs include a set of query points and our outputs include both bounding boxes and segmentation masks, we mainly extend SAM [30] to build our model. Compared with traditional detection pipeline such as Mask R-CNN [23], we can use a query point to naturally guide SAM to detect the corresponding object. Mask R-CNN generates thousands of anchors for each image, which is challenging to find the correct matching. However, we also compare with alternative network architectures in our experiments for completeness. We find they can also work despite being worse than SAM. For simplicity, we assume there is only a single query point. But our model can accept hundreds of query points at a time.

5.1. Backbone

The goal of our backbone is to map an image I and a query point $[x, y]$ to a pooled feature $h = f(I; [x, y])$. Full details are in the supplemental.

Image Encoder. Our image encoder is a MAE [22] pretrained Vision Transformer (ViT) [12], following SAM [30]. They map a single image I to the memory of the transformer decoder.

Query Point Encoder. We transfer the query point $[x, y]$ to positional encodings [58], which is then feed into the transformer decoder. We use the embedding k to guide the transformer to produce the feature h for different query points.

Transformer Decoder. The decoder accepts inputs of the

memory from the encoder, and an embedding k of the query point. It produces a embedding h for each query point, and we use it to predict all the properties, like a ROI feature.

5.2. Prediction Heads

We now describe how to map from the pooled feature h to the features. Each prediction is done by a separate head that handles each output type.

Movable ✍ We add a linear layer and map the hidden embedding h to the prediction of movable. We use the standard cross entropy loss to train it.

Localization 📍 We follow SAM standard practice to predict segmentation masks. We predict segmentation masks using mask decoder and train them using focal loss [37] and DICE [43] as loss functions. For depth, we have a separate depth transformer decoder with a corresponding learnable depth query. We train depth using scale- and shift-invariant L1 loss and gradient-matching loss following [68, 50, 36]. The shift and scale are normalized per image.

Rigidity 🔪 Similar to movable, we add a linear layer to predict whether the object is rigid or not. We train the linear layer using a standard binary cross entropy loss.

Articulation 🔧 We first add a linear layer to predict whether the interactive object is rotation, translation or freeform, and we use the standard cross entropy loss to train it. For the rotation axis, we follow [49, 70] to represent an axis as a 2D line (θ, r) . Any points on this line satisfy $x \cos(\theta) + y \sin(\theta) = r$ where θ represents the angle and r represents the distance from the object center to the line. In training, we represent the 2D line as $(\sin 2\theta, \cos 2\theta, r)$, so that the axis angle is in a continuous space [71]. We use a 3-layer MLP to predict the axis, similar to bounding boxes as both tasks require localization. We use L1 loss to train it.

Action 👉 Similar to movable, we add a linear layer to predict what the potential action is to interact with the object. We train the linear layer using a standard binary cross entropy loss.

Affordance 👉 Our prediction of affordance is a probability map, while our annotation is a single keypoint. How-

ever, affordance can have multiple solutions. Therefore, we transform the annotation of affordance to a 2D gaussian bump [32] and train the network using a binary focal loss [37]. We set the weight of positive examples to be 0.95 and that of negative ones to be 0.05 to balance positives and negatives, as there are more negatives than positives.

Our total loss is a weighted linear combination of all losses mentioned above. Details are in supplemental.

5.3. Implementation Details

Full architectural details of our approach are in the supplemental. In practice, we use three different transformer decoders for mask, depth and affordance. The image encoder, query point encoder and mask decoder are pretrained on SAM [30]. Other parts, including affordance head and depth head, are trained from scratch. We use an AdamW optimizer of the learning rate 10^{-4} , and train our model for 200 epochs.

6. Experiments

We have introduced an approach that can localize and predict the properties of the moving part from an image. In the experiments, we aim to answer the following questions: (1) how well can one localize and predict the properties of the moving part from an image; (2) how well do alternative approaches to the problem do? We evaluate our approach on our 3DOI dataset and test the generalization to robotics data WHIRL [1].

6.1. Experimental Setup

We first describe the setup of our experiments. Our method aims to look at a single RGB image and infer information about the moving part given a keypoint. We therefore evaluate our approach on two challenging datasets, using metrics that capture various aspects.

Datasets. We train and validate our approach on two datasets: 3DOI dataset (described in Section 4), and the WHIRL dataset [1]. WHIRL [1] is a robotics dataset including every-day objects and settings, for example drawers, dishwashers, fridges in different kitchens, doors to various cabinets. We use WHIRL to validate the generalization of our approach and downstream applications in the robotics settings. We split the first frame of all WHIRL videos and annotate them using the same pipeline as our datasets. Typically, humans are not present in the first frame and it’s before any manipulation.

Metrics. We report standard practices of evaluation for all of our predictions. For all metrics, the higher the better. These metrics are detailed as follows:

- Movable 🚪, Rigidity 🛠️, and Action 🖐️: We report accuracy as these are multiple choice questions.

- Localization 📏: We report Intersection-over-Union (IoU) for our predictions of bounding boxes and masks [38]. We report threshold accuracy for depth [14].

- Articulation 🗑️: We report accuracy for articulation type. The rotation axis is a 2D line. Therefore, we report EA-Score between the prediction and the ground truth, following [49, 70]. EA-Score [70] is a score in $[0, 1]$ to measure the angle and euclidean distance between two lines.

- Affordance 🖐️: It’s a probability map and we report the histogram intersection (or SIM) following [46, 2, 33, 41].

Baselines. We compare our approach with a series of baselines, to evaluate how well alternative approaches work on our problem. We first evaluate 3DADN [49], SAPIEN [65], and InteractionHotspots [46] using their pretrained checkpoints, to test how well learning from videos or synthetic data works on our problem. We then train two query-point-based model, ResNet MLP [24] and COHESIV [54], to test how well alternative network architectures work on our problem. The details are introduced as follows.

- (3DADN [49]): 3DADN detects articulated objects which humans are interacting with, extending Mask R-CNN [23]. It is trained on Internet videos. We drop the temporal optimization part since we work on a single image. For each image, it can detect articulated objects, as well as the type (translation or rotation), bounding boxes, masks and axes. Since the inputs of 3DADN do not include a query point, we compare the predicted bounding boxes and the ground truth to find the matching detection, and evaluate other metrics. We lower the detection threshold to 0.05 to ensure we have enough detections to match our ground truth.

- (SAPIEN [65]): The training frames of 3DADN [49] typically have human activities. However, our dataset does not require humans to be present, which may lead to generalization issues. Alternatively, we are interested in whether we can just learn the skill from synthetic data. We train 3DADN [49] on renderings of synthetic objects generated by SAPIEN. SAPIEN is a simulator which contains a large scale set of articulated objects. We use the renderings provided by 3DADN and the same evaluation strategies.

- (InteractionHotspots [46]): While 3DADN and SAPIEN can detect articulated objects as well as their axes, they cannot tell the affordance. InteractionHotspots learns affordance from watching OPRA [16] or Epic-Kitchen [10] videos. Since InteractionHotspots cannot detect objects, we apply a center crop of the input image based on the query point, and resize it to the standard input shape (224, 224). We use the model trained on Epic-Kitchen as it transfers better than OPRA.

Additionally, we want to test alternative network architectures trained on our 3DOI dataset. We use the same loss as ours to train it on 3DOI, to ensure fair comparison.

- (ResNet MLP [24]): ResNet MLP uses a ResNet-50 encoder to extract features from input images. We then sample

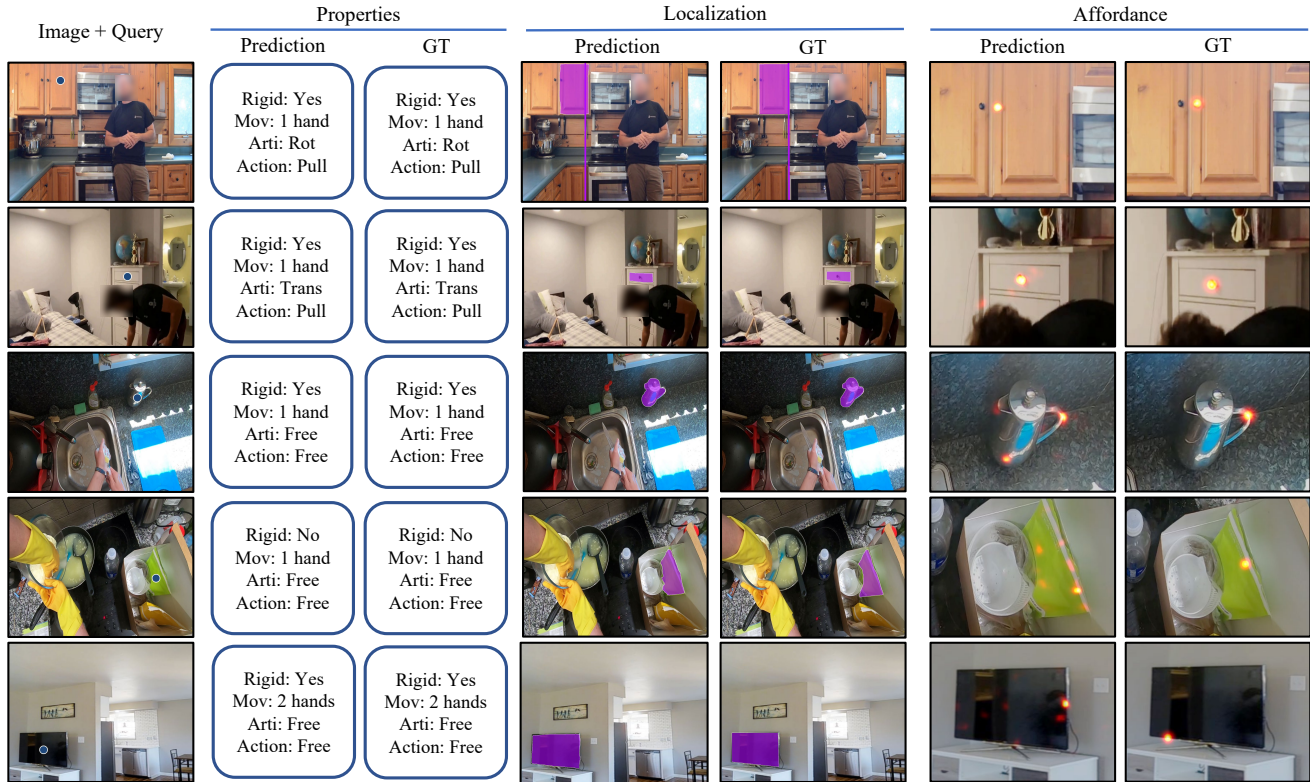


Figure 4. Results on our 3DOI dataset. ● indicates the query point. (Row 1, 2) Our approach can correctly recognize articulated objects, as well as its type (rotation or translation), axes, and affordance. (Row 3, 4) Our approach can recognize rigid and nonrigid objects in egocentric video. (Row 5) Our approach can recognize objects need to be moved by two hands, such as a TV. We note that the affordance of these objects have multiple solutions. Affordance is zoomed manually for better visualization. Affordance colormap: min max.

the corresponding spatial features from the feature map using the 2D coordinates of keypoints. We train ResNet MLP on all tasks except mask, affordance and depth, as these tasks requires dense predictions for each pixel. Adding a separate decoder to ResNet makes it a UNet-like architecture [52], which is beyond the scope of ResNet.

- (COHESIV [54]): We also pick another model COHESIV, which designed for the prediction-at-a-query-location problem. Given an input image and corresponding hand location as a query, COHESIV predicts the segmentation of hands and hand-held objects. We adopt the network, as it produces a feature map of queries. we sample an embedding from the feature map according to the query point, concatenate it with image features, and produce multiple outputs.

6.2. Results

First, we show qualitative results in Figure 4. For articulated objects (drawers, cabinets, etc.), our approach can recognize its location, kinematic model (rotation or translation), axes and handle. It can also recognize rigid or non-rigid objects, as well as light or heavy ones. It works on both third-person images or egocentric videos. And all of these are achieved in a single model. For articulated ob-




Figure 5. Prediction of 3D potential interaction of articulated objects. ● indicates the query point. In prediction 1, 2, and 3, we rotate the object along its rotation axis, or translate the object along its normal direction.

jects, we utilize the outputs and further show their potential 3D interaction in Figure 5. Full details in supplemental.

We then compare our approach with a set of baselines. The quantitative results are reported in Table 1. 3DADN [49] is much worse than our approach, since it can only detect objects which are being articulated. It fails to detect objects humans are not interacting. Instead, our ap-

Table 1. Quantitative results on our 3DOI dataset. Cat. means category. We report accuracy for all category classification, including movable, rigid, articulation and action. We report mean IoU for box and mask, EA-Score for articulation axis, and SIM for affordance. For all metrics, the higher the better.







Methods	Movable 	Localization 		Rigidity 	Articulation 		Action 	Affordance 
	Cat.	Box	Mask	Cat.	Cat.	Axis	Cat.	Probability
3DADN [49]	-	8.53	6.45	-	44.3	5.63	-	-
SAPIEN [65]	-	5.94	4.57	-	41.6	1.79	-	-
InteractionHotspots [46]	-	-	-	-	-	-	-	0.047
ResNet MLP [24]	72.5	21.4	-	81.9	51.9	68.3	58.8	-
COHESIV [54]	71.5	28.3	35.2	81.2	68.0	67.2	71.5	0.013
Ours	85.3	69.9	77.1	90.1	89.4	80.3	89.7	0.167



Figure 6. Comparison of 3DADN [49], SAPIEN [65] and our approach. \bullet indicates the query point. 3DADN has a strong performance when humans are present. However, it has difficulty detecting objects without human activities. SAPIEN does not generalize well to real images. However, it is sometimes better than 3DADN when humans are not present.

proach can detect any objects can be interacted, regardless of human activities. SAPIEN is worse than 3DADN, which suggests learning from synthetic objects has a huge domain gap. This is consistent with the observation of 3DADN. Visual comparisons are shown in Figure 6.

We compare our prediction of the affordance map with InteractionHotspots [46]. Our approach outperforms InteractionHotspots significantly, with a 3.5x improvement. A visual comparison is shown in Figure 7. While InteractionHotspots predicts a cloud-like probability map, our approach is typically very confident about its prediction. However, the overall performance is relatively low, mainly due to ambiguity of affordance on deformable objects.

To explore alternative network architectures, we compare our approach with ResNet MLP [24] and COHESIV [54], which are trained on our data with the same loss functions. ResNet MLP is reasonable on movable, rigidity, and action. It is especially bad on bounding box localization, which is why we typically rely on a detection pipeline such as Mask R-CNN [23]. COHESIV learns reasonable bounding boxes and masks, which is a huge improvement over ResNet MLP. The performance of movable drops compared with ResNet MLP, while that of kinematic and action improves. Overall, our approach outperforms both ResNet MLP and COHESIV, mainly due to the introduction of transformers.

Finally, we evaluate depth on our data. Having state-of-the-art depth estimation is orthogonal to our goal, since

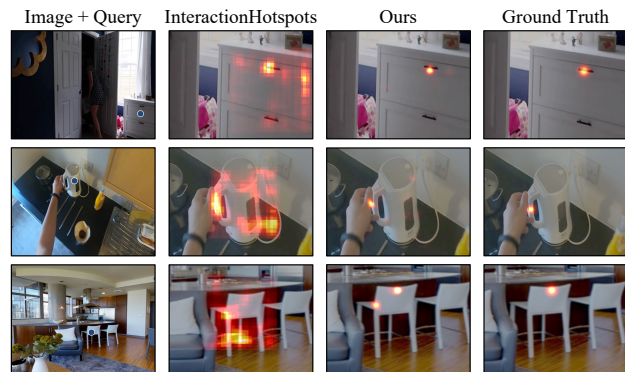



Figure 7. Comparison of InteractionHotspots [46] and our approach. \bullet indicates the query point. We find InteractionHotspots typically makes a cloud like probability map on our data. Our model is very confident about its prediction, while there can be multiple solutions. Prediction and GT are zoomed manually for better visualization. Affordance colormap: min  max.







we only need reasonable depth to localize objects in 3D and render potential 3D interactions. In fact, state-of-the-art depth estimation models are trained on over ten datasets and one million images [50, 68, 13], while our dataset only has 5K images with depth ground truth. We just report the evaluation of depth estimation, in order to show our model has learned reasonable depth. On our data, 96.7% pixels are within the 1.25 threshold, 99.3% pixels are within the 1.25² threshold.

6.3. Generalization Results

To test whether our approach and models trained on our 3DOI dataset can generalize, we further evaluate our approach on WHIRL [1], a robotics dataset manipulating every-day objects. Since WHIRL is a small-scale dataset, we test our model on WHIRL without finetuning. Our results are shown in Figure 8. For both articulated objects and deformable objects, our approach can successfully recover its kinematic model, location and affordance.

We also quantitatively evaluate our approach on WHIRL. We report our results in Table 2. Similar to our 3DOI dataset, our approach outperforms 3DADN [49], SAPIEN [65] and InteractionHotspots [46] significantly.

Table 2. Quantitative results on robotics data [1]. Cat. means category. We report accuracy for all category classification, including movable, rigid, articulation and action. We report mean IoU for the boxes and masks, EA-Score for articulation axis, and SIM for affordance probability map. For all metrics, the higher the better.

Methods	Movable 	Localization 		Rigidity 	Articulation 		Action 	Affordance 
	Cat.	Box	Mask	Cat.	Cat.	Axis	Cat.	Probability
3DADN [49]	-	13.8	10.1	-	53.3	4.03	-	-
SAPIEN [65]	-	9.14	6.15	-	51.1	0.0	-	-
InteractionHotspots [46]	-	-	-	-	-	-	-	0.045
ResNet MLP [24]	88.8	14.1	-	80.0	51.1	57.1	51.1	-
COHESIV [54]	86.7	37.1	38.7	82.2	73.3	66.1	73.3	0.015
Ours	91.1	68.7	70.2	95.6	80.0	68.5	84.4	0.148

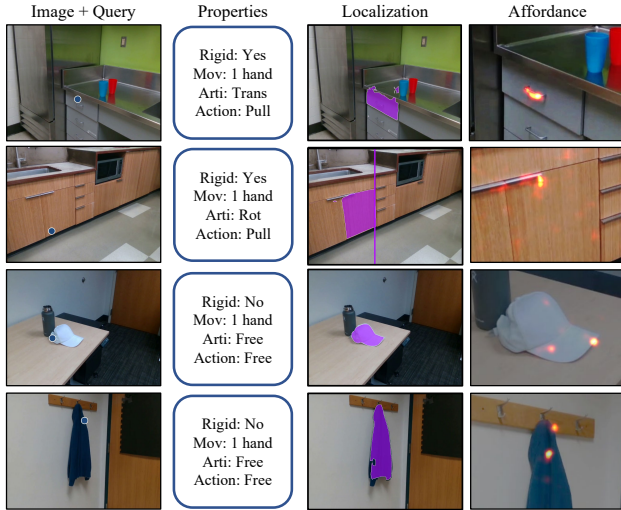



Figure 8. Results on robotics data [1]. ● indicates the query point. Without finetuning, our approach generalizes well to robotics data, which indicates its potential to help intelligent agents to better manipulate objects. Row 1 and 2 are articulated objects. Row 3 and Row 4 are deformable objects. Affordance is zoomed manually for better visualization. Affordance colormap: min  max.

The performance gap is even larger. We believe it is because humans are not present in most images of the dataset.

We compare our approach with ResNet MLP [24] and COHESIV [53], which are also trained on our 3DOI dataset. Our model outperforms both ResNet MLP and COHESIV consistently. The improvement on dense predictions (Localization and Affordance) is significant, due to the design of mask decoder. The improvement on other properties is relatively small. It illustrates models trained on our 3DOI dataset generalize well to robotics data, regardless of network architectures.

6.4. Limitations and Failure Modes

We finally discuss our limitations and failure modes. In Figure 9, we show some predictions are hard to make from visual cues: Some articulated objects are symmetric and humans rely on common sense to guess its rotation axis. There are also hard examples when predicting the rigidity and movable. Finally, we only annotate a single keypoint

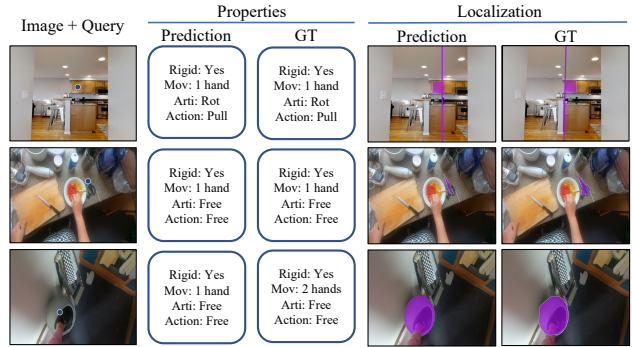


Figure 9. Typical failure modes of our approach. ● indicates the query point. **Row 1:** Our predicted rotation axis is on the wrong side when the objects look symmetric. **Row 2:** Our predicted mask is partial when the scissors are occluded. **Row 3:** Our model thinks the trash bin can be picked up by 1 hand, potentially since its material looks plastic.

for each object instance as affordance. But some objects may have multiple keypoints as affordance.

7. Conclusion

We have presented a novel task of predicting 3D object interactions from a single RGB image. To solve the task, we collected the 3D Object Interaction dataset, and proposed a transformer-based model which predicts the potential interactions of any objects according to query points. Our experiments show that our approach outperforms existing approaches on our data and generalizes well to robotics data.

Our approach can have positive impacts by helping build smart robots that are able to understand the 3D scene and manipulate everyday objects. On the other hand, our approach may be useful for surveillance activities.

Acknowledgments This work was supported by the DARPA Machine Common Sense Program. This material is based upon work supported by the National Science Foundation under Grant No. 2142529. We thank Shikhar Bahl and Deepak Pathak for their help with WHIRL data, Georgia Gkioxari for her help with the figure, and Tiange Luo, Ang Cao, Cheng Chi, Yixuan Wang, Mohamed El Banani, Linyi Jin, Nilesh Kulkarni, Chris Rockwell, Dandan Shan, Siyi Chen for helpful discussions.

References

- [1] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. 2022. [2](#), [5](#), [7](#), [8](#)
- [2] Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? *IEEE transactions on pattern analysis and machine intelligence*, 41(3):740–757, 2018. [5](#)
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. [2](#), [3](#)
- [4] Yu-Wei Chao, Zhan Wang, Yugeng He, Jiaxuan Wang, and Jia Deng. Hico: A benchmark for recognizing human-object interactions in images. In *ICCV*, 2015. [2](#)
- [5] Weifeng Chen, Shengyi Qian, and Jia Deng. Learning single-image depth from videos using quality assessment networks. In *CVPR*, 2019. [2](#)
- [6] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021. [2](#)
- [7] Cheng Chi and Dmitry Berenson. Occlusion-robust deformable object tracking without physics simulation. In *IROS*, 2019. [2](#)
- [8] 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 *ECCV*, 2016. [2](#)
- [9] Cristina Garcia Cifuentes, Jan Issac, Manuel Wüthrich, Stefan Schaal, and Jeannette Bohg. Probabilistic articulated real-time tracking for robot manipulation. *IEEE Robotics and Automation Letters*, 2(2):577–584, 2016. [2](#)
- [10] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision (IJCV)*, 130:33–55, 2022. [3](#), [5](#)
- [11] Karthik Desingh, Shiyang Lu, Anthony Opipari, and Odest Chadwicke Jenkins. Factored pose estimation of articulated objects using efficient nonparametric belief propagation. In *ICRA*, 2019. [2](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [4](#)
- [13] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *ICCV*, 2021. [3](#), [7](#)
- [14] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015. [2](#), [5](#)
- [15] Rui Fan, Hengli Wang, Bohuan Xue, Huaiyang Huang, Yuan Wang, Ming Liu, and Ioannis Pitas. Three-filters-to-normal: An accurate and ultrafast surface normal estimator. *IEEE Robotics and Automation Letters*, 6(3):5405–5412, 2021. [2](#)
- [16] Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, and Joseph J Lim. Demo2vec: Reasoning object affordances from online videos. In *CVPR*, 2018. [5](#)
- [17] Georgia Gkioxari, Ross Girshick, Piotr Dollár, and Kaiming He. Detecting and recognizing human-object interactions. In *CVPR*, 2018. [2](#)
- [18] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *ICCV*, 2019. [1](#), [2](#)
- [19] Mohit Goyal, Sahil Modi, Rishabh Goyal, and Saurabh Gupta. Human hands as probes for interactive object understanding. In *CVPR*, 2022. [2](#)
- [20] Arjun Gupta, Max E Shepherd, and Saurabh Gupta. Predicting motion plans for articulating everyday objects. In *ICRA*, 2023. [2](#)
- [21] Sanjay Haresh, Xiaohao Sun, Hanxiao Jiang, Angel X Chang, and Manolis Savva. Articulated 3d human-object interactions from rgb videos: An empirical analysis of approaches and challenges. *arXiv preprint arXiv:2209.05612*, 2022. [2](#)
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. [4](#)
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. [2](#), [3](#), [4](#), [5](#), [7](#)
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [5](#), [7](#), [8](#)
- [25] Cheng-Chun Hsu, Zhenyu Jiang, and Yuke Zhu. Ditto in the house: Building articulation models of indoor scenes through interactive perception. In *ICRA*, 2023. [2](#)
- [26] Hanxiao Jiang, Yongsan Mao, Manolis Savva, and Angel X Chang. Opd: Single-view 3d openable part detection. In *ECCV*, 2022. [2](#)
- [27] Zhenyu Jiang, Cheng-Chun Hsu, and Yuke Zhu. Ditto: Building digital twins of articulated objects from interaction. In *CVPR*, 2022. [2](#)
- [28] Linyi Jin, Shengyi Qian, Andrew Owens, and David F. Fouhey. Planar surface reconstruction from sparse views. In *ICCV*, 2021. [2](#)
- [29] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *CVPR*, 2019. [2](#)
- [30] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *ICCV*, 2023. [2](#), [4](#), [5](#)
- [31] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *CVPR*, 2020. [2](#)
- [32] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 2018. [5](#)
- [33] Gen Li, Varun Jampani, Deqing Sun, and Laura Sevilla-Lara. Locate: Localize and transfer object parts for weakly supervised affordance grounding. In *CVPR*, 2023. [2](#), [5](#)
- [34] Xiaolong Li, He Wang, Li Yi, Leonidas J Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *CVPR*, 2020. [1](#), [2](#)

- [35] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*, 2018. 3
- [36] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018. 2, 4
- [37] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 4, 5
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 5
- [39] Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, and Jan Kautz. PlaneRCNN: 3D plane detection and reconstruction from a single image. In *CVPR*, 2019. 2
- [40] Chen Liu, Jimei Yang, Duygu Ceylan, Ersin Yumer, and Yasutaka Furukawa. Planenet: Piece-wise planar reconstruction from a single rgb image. In *CVPR*, 2018. 2
- [41] Hongchen Luo, Wei Zhai, Jing Zhang, Yang Cao, and Dacheng Tao. Learning affordance grounding from exocentric images. In *CVPR*, 2022. 2, 5
- [42] Tiange Luo, Honglak Lee, and Justin Johnson. Neural shape compiler: A unified framework for transforming between text, point cloud, and program. 2022. 2
- [43] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, 2016. 4
- [44] Kaichun Mo, Leonidas Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *ICCV*, 2021. 2
- [45] Jiteng Mu, Weichao Qiu, Adam Kortylewski, Alan Yuille, Nuno Vasconcelos, and Xiaolong Wang. A-sdf: Learning disentangled signed distance functions for articulated shape representation. In *ICCV*, 2021. 2
- [46] Tushar Nagarajan, Christoph Feichtenhofer, and Kristen Grauman. Grounded human-object interaction hotspots from video. In *ICCV*, 2019. 2, 5, 7, 8
- [47] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *CVPR*, 2020. 1, 2
- [48] Sudeep Pillai, Matthew R Walter, and Seth Teller. Learning articulated motions from visual demonstration. In *RSS*, 2014. 1, 2
- [49] Shengyi Qian, Linyi Jin, Chris Rockwell, Siyi Chen, and David F. Fouhey. Understanding 3d object articulation in internet videos. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [50] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 1, 2, 3, 4, 7
- [51] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 2
- [52] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 6
- [53] Dandan Shan, Jiaqi Geng, Michelle Shu, and David Fouhey. Understanding human hands in contact at internet scale. In *CVPR*, 2020. 2, 3, 8
- [54] Dandan Shan, Richard E.L. Higgins, and David F. Fouhey. COHESIV: Contrastive object and hand embedding segmentation in video. In *NeurIPS*, 2021. 5, 6, 7, 8
- [55] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012. 3
- [56] Jürgen Sturm, Cyrill Stachniss, and Wolfram Burgard. A probabilistic framework for learning kinematic models of articulated objects. *Journal of Artificial Intelligence Research*, 41:477–526, 2011. 1, 2, 3
- [57] Xiaohao Sun, Hanxiao Jiang, Manolis Savva, and Angel Xuan Chang. Opdmulti: Openable part detection for multiple objects. *arXiv preprint arXiv:2303.14087*, 2023. 2
- [58] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020. 4
- [59] Xiaolong Wang, David F. Fouhey, and Abhinav Gupta. Designing deep networks for surface normal estimation. In *CVPR*, 2015. 2
- [60] Yixuan Wang, Dale McConachie, and Dmitry Berenson. Tracking partially-occluded deformable objects while enforcing geometric constraints. In *ICRA*, 2021. 2
- [61] Yian Wang, Ruihai Wu, Kaichun Mo, Jiaqi Ke, Qingnan Fan, Leonidas J Guibas, and Hao Dong. Adaafford: Learning to adapt manipulation affordance for 3d articulated objects via few-shot interactions. In *ECCV*, 2022. 2
- [62] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based object detection. *arXiv preprint arXiv:2109.07107*, 3(6), 2021. 2
- [63] Fangyin Wei, Rohan Chabra, Lingni Ma, Christoph Lassner, Michael Zollhöfer, Szymon Rusinkiewicz, Chris Sweeney, Richard Newcombe, and Mira Slavcheva. Self-supervised neural articulated shape and appearance models. In *CVPR*, 2022. 2
- [64] Ruihai Wu, Yan Zhao, Kaichun Mo, Zizheng Guo, Yian Wang, Tianhao Wu, Qingnan Fan, Xuelin Chen, Leonidas Guibas, and Hao Dong. Vat-mart: Learning visual action trajectory proposals for manipulating 3d articulated objects. *arXiv preprint arXiv:2106.14440*, 2021. 2
- [65] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *CVPR*, 2020. 1, 2, 5, 7, 8
- [66] Zhenjia Xu, Cheng Chi, Benjamin Burchfiel, Eric Cousineau, Siyuan Feng, and Shuran Song. Dextairity: Deformable manipulation can be a breeze. *arXiv preprint arXiv:2203.01197*, 2022. 2

- [67] Fengyu Yang, Chenyang Ma, Jiacheng Zhang, Jing Zhu, Wenzhen Yuan, and Andrew Owens. Touch and go: Learning from human-collected vision and touch. In *NeurIPS*, 2022. 2
- [68] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *CVPR*, 2021. 1, 2, 3, 4, 7
- [69] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. 3
- [70] Kai Zhao, Qi Han, Chang-Bin Zhang, Jun Xu, and Ming-Ming Cheng. Deep hough transform for semantic line detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 4, 5
- [71] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. 4