

Generating 6DoF Object Manipulation Trajectories from Action Description in Egocentric Vision

Tomoya Yoshida¹ Shuhe Kurita² Taichi Nishimura³ Shinsuke Mori¹

¹Kyoto University ²National Institute of Informatics ³Sony Interactive Entertainment



Figure 1. **6DoF object manipulation trajectory.** This task aims to generate a sequence of 6DoF object poses from an action description and an initial state comprising the visual input and the object’s initial pose.

Abstract

Learning to use tools or objects in common scenes, particularly handling them in various ways as instructed, is a key challenge for developing interactive robots. Training models to generate such manipulation trajectories requires a large and diverse collection of detailed manipulation demonstrations for various objects, which is nearly unfeasible to gather at scale. In this paper, we propose a framework that leverages large-scale ego- and exo-centric video datasets — constructed globally with substantial effort — of Exo-Ego4D to extract diverse manipulation trajectories at scale. From these extracted trajectories with the associated textual action description, we develop trajectory generation models based on visual and point cloud-based language models. In the recently proposed egocentric vision-based in-a-quality trajectory dataset of HOT3D, we confirmed that our models successfully generate valid object trajectories, establishing a training dataset and baseline models for the novel task of generating 6DoF manipulation trajectories from action descriptions in egocentric vision. Our dataset and code: <https://biscue5.github.io/egoscaler-project-page/>.

This work is also done in Research and Development Center for Large Language Models, National Institute of Informatics (NII LLMC).

1. Introduction

Use of tools and manipulation of objects at their will is an essential ability for developing robots that assist human activities. Especially having knowledge of manipulating numerous tools or objects in our work places and houses following the requests from humans is the next key technology for developing interactive robots that support us in daily scenes. One way to achieve this is to preliminarily learn various ways of manipulating objects from the human demonstration, as the recent robotic studies suggest that human demonstration acts as quite effective supervision for object manipulation trajectory generation [15, 24, 104]. However, this human demonstration-driven approach evokes the well-known limitation of data scarcity: the object manipulation trajectories from real human demonstrations are quite informative yet extremely limited and almost prohibitive for covering numerous objects we use in our daily lives.

One promising solution for this is to make use of videos of daily working scenes for extracting human demonstrations of various object manipulations. Egocentric videos are becoming popular in virtual and augmented reality (VR/AR) [6] and even recording our daily life activities [21, 63] in both industry and research fields. For tracking human motions, HOT3D [3] provides quality 3D poses and models of hands and objects from precise lab recordings of human activity. Unfortunately, the precise record-

ing of human motion trajectories is still considerably costly, and hence such data is still limited for developing models that generate object manipulation trajectories whilst general purpose egocentric recordings are becoming available at scale as of the Ego4D [29] and Ego-Exo4D [30] datasets. Hereby, we assume extracting human motions for manipulating objects from egocentric vision datasets at scale is the key break-through for developing interactive robots.

In this paper, we introduce a framework to extract an object manipulation trajectory, a sequence of 6DoF object positions and rotations, from egocentric videos. Our framework consists of the following: determination of temporal action span, extracting objects with the open-vocabulary segmentation model [74] and tracking them with the dense 3D point tracker [93], determining the camera coordinates, and obtaining the object rotation by the singular value decomposition from the translation between scenes. Applying this framework to egocentric videos from Ego-Exo4D [30], we construct EgoTraj with 28,497 object manipulation trajectories and corresponding action descriptions. Sec. 2.1 presents the dataset statistics from previous studies and ours. Our dataset emerges as the largest, containing 28,497 action descriptions and trajectories. While existing datasets such as H2O [46], EgoPAT3D [53], and HOT3D [3] rely on lab-recordings with multiple cameras or depth cameras, our dataset is extracted from ordinary egocentric videos with a single camera without pre-recorded camera coordinates or depth, and encompasses an entirely different order of magnitude in the variety of action verbs and objects from real-world scenarios. As our method does not rely on pre-recorded camera coordinates, our method is applicable for most of egocentric video datasets, covering diverse object manipulations from egovision.

With the extracted trajectories of EgoTraj, we develop a branch of models to generate object manipulations from the textual action description. The overview of the task, including extracted 6DoF object trajectory, is presented in Fig. 1. In experiments, we develop object manipulation trajectory generation models based on visual and point cloud-based language models with discretized trajectory tokens [9, 15, 44, 108]. Our experimental results on the HOT3D dataset demonstrate that our models successfully generate valid object manipulation trajectories. Furthermore, we conduct the experiments of generating action descriptions using our object trajectories, suggesting that utilizing trajectories enhances the generation performance, particularly for the similarity of verbs in the action descriptions. To the best of our knowledge, this is the first attempt to extract large-scale 6DoF object manipulation trajectories from egocentric videos without pre-recorded camera coordinates and develop visual and point cloud-based language models for generating object manipulation trajectories from textual action descriptions.

Dataset	#AD	#Verb	#Object	#Trajectories
EgoPAT3D [†] [53]	-	26	11	11,141
H2O [†] [46]	-	12	8	13,653
HOT3D [3]	986	30	33	986
EgoTraj (Ours)	28,497	228	4,158	28,497

Table 1. Dataset statistics from previous studies and ours. The “#AD” column denotes the number of action descriptions. “#Trajectories” of [†] is reported in [5]. Note that our dataset is constructed through a machine learning-based approach, while others are through precise annotations.

2. Related Work

2.1. Object Interaction Understanding

Understanding object interaction involves two key factors of object affordance and manipulation motion. Koppula and Saxena [45] addressed human action anticipation by using object affordances to predict future motion trajectories. They demonstrated that understanding object affordance and predicting motion trajectories is beneficial for reactive robotic response. Subsequently, this field has been extensively explored in the contexts of object affordance learning [19] and object interaction prediction [29].

Object Affordance Learning. Object affordance learning aims to localize objects [14, 60, 77] or their specific parts [31, 48, 49, 61, 66] necessary for performing certain actions. This task has gained attention for applications in robotic manipulation [2, 40, 65]. However, most studies focus on grasping affordance, limiting various manipulation understandings. More recently, hand-object interaction generation [8, 11, 16, 98, 101] has been proposed, addressing both affordance learning and manipulation understanding in 3D space. This task builds upon richly annotated video datasets [46, 59] that include hand poses, object meshes, and scene point clouds.

Object Interaction Prediction. Object interaction prediction includes various tasks such as interaction anticipation [25, 41, 57, 57, 64], next active object detection [7, 26, 45, 57, 71, 83, 96], and hand trajectory forecasting [58, 61, 62, 103]. These tasks have potential applications such as VR/AR [27], human-robot interaction [22, 81]. However, most studies addressed these tasks in 2D space, limiting their applications. To mitigate this limitation, recent work has extended these tasks into 3D space [5, 53]. For instance, Bao *et al.* [5] addressed 3D hand trajectory forecasting by proposing an uncertainty-aware state space transformer, extending conventional 2D trajectory forecasting task to 3D. Despite these advancements, addressing diverse manipulation motions remains challenging due to the high cost of annotations in 3D space. In contrast, our work proposes a framework to extract 6DoF object manipulation trajectories from egocentric videos, enabling the handling of diverse manipulation motions. From

these extracted trajectories with the associated action descriptions, we address text-guided object manipulation trajectory generation.

2.2. Egocentric Vision Dataset

Egocentric videos encompass diverse object manipulation and hand activities, serving as potential resources for object interaction learning. A wide range of datasets and benchmarks have been proposed [17, 18, 29, 30, 51, 78]. For example, Ego-Exo4D [30] is a large-scale egocentric-exocentric dataset, providing rich multi-modal resources including gaze, point clouds, and texts. Based on these large-scale datasets, multiple detailed tasks have been developed [33, 58, 72, 75, 79, 85]. Additionally, several datasets have also been developed for specific purposes, including hand-object-interaction and VR tasks [3, 32, 46, 59, 90]. HOT3D [3] is an egocentric dataset for 3D hand and object tracking. Trajectories are captured using optical markers and OptiTrack cameras, resulting in precise hand poses and 6DoF object trajectories. While several other datasets [32, 46, 59] also provide similar hand and object poses, HOT3D covers a wider range of trajectories involving diverse objects and action scenarios. To this end, we utilize 6DoF object manipulation trajectories from HOT3D to validate model performance.

2.3. Object Pose Estimation / Tracking

Estimating object pose from visual sensors, termed object pose estimation, is a crucial topic for robotic perception [20]. Object pose tracking is an extended task of object pose estimation, aiming to efficiently track temporal object poses [52, 91]. Methods in both tasks often rely on objects' CAD models [35, 52, 69] or several reference images [34, 36, 80, 91]. Although pose estimation from visual input is similar to our task, we aim to understand and generate pose transformations during manipulation. Consequently, our task does not require any 3D object models or reference images of objects.

2.4. Language Models for Visual / Robotic Tasks

Recent advances in large auto-regressive models, such as large language models (LLMs) and vision-language models (VLMs), have established these models as powerful tools in computer vision [38, 50, 55, 56, 82, 95, 100, 107] and robotics [9, 44, 108]. For instance, PointLLM [94], a 3D-LLM for object point cloud understanding, has shown strong performance in 3D object captioning and has also demonstrated capability in scene-level point cloud understanding. Moreover, by embedding spatio-temporal information into language space, several studies have achieved comparable or superior performance to conventional approaches in image, video, and point cloud understanding benchmarks, such as object detection [12, 88], object track-

ing [87], temporal action localization [39], and 2D/3D visual grounding [88, 107]. Inspired by these studies, we develop trajectory generation models based on VLMs, formalizing our task as next token prediction task.

3. Dataset Construction Framework

We construct a large-scale training dataset containing 6DoF object manipulation trajectories along with images, depth maps, and action descriptions from an egocentric vision dataset of Ego-Exo4D [30]. Due to the cost and scalability, we propose an automated approach here. Note that our framework does not require any camera extrinsic parameters in advance. Additionally, we utilize precise 6DoF object manipulation trajectories extracted from HOT3D [3] dataset for the evaluation set.

3.1. Training Data Creation

As shown in Fig. 2, our pipeline consists of four stages. First, given an egocentric video, we determine the start and end timestamps of the action and identify the manipulated object within the scene. Second, we extract the position sequence of the manipulated object using an open-vocabulary segmentation model [74] and a dense 3D point tracker [93]. Third, we project the sequence into the camera coordinate system of the first frame using point cloud registration. Fourth, we extract a rotation sequence by computing the transformation between the two object point clouds using singular value decomposition.

Temporal Action Localization. Based on the original annotated timestamps, we first split the raw egocentric video into multiple clips. Each clip spans a four-second interval centered on each timestamp. Next, we sample image frames within each clip and assign them sequential indices. We input these indexed frames along with the corresponding action description, originally annotated in Ego-Exo4D, into OpenAI GPT-4o [1] to determine the action's start and end timestamps: t_{start} and t_{end} . We denote the frames at t_{start} and t_{end} as f_{start} and f_{end} , respectively. Moreover, we use the model to extract the manipulated object name and determine whether the object is rigid or not. We filter out non-rigid objects as the transformation of the non-rigid objects is left for future work. We provide the system prompt in the Appendix.

Position Sequence Extraction. We first use an open-vocabulary object segmentation model [74] to identify the manipulated object in the first frame f_{start} . This model takes f_{start} and the object name as inputs, and outputs a 2D segmentation map. Next, we employ an off-the-shelf monocular depth estimation model [97] and SpaTracker [93]. SpaTracker is a state-of-the-art dense point tracking model that tracks any 2D points in 3D space. Based on a sequence of frames between t_{start} and t_{end} , we obtain depth maps for

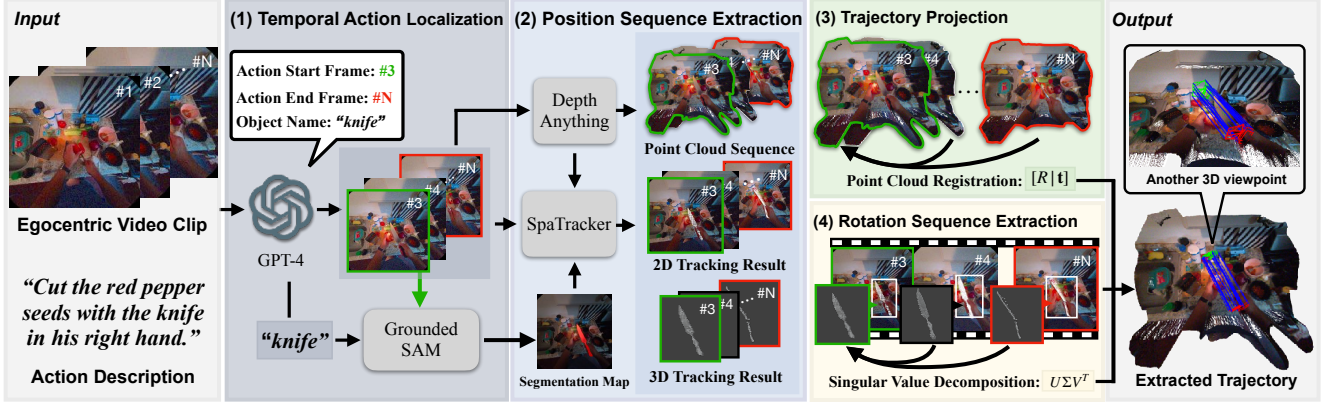


Figure 2. **Trajectory extraction from egocentric videos.** Four steps of (1) temporal action localization, (2) position sequence extraction, (3) trajectory projection, and (4) rotation sequence extraction. Our project page includes visualization of resulting trajectories with videos.

each frame using a depth estimation model. For SpaTracker, we input the sequence of frames, the corresponding depth maps, and the segmentation map into the model. The model then outputs 3D tracking results of the segmented object.

Trajectory Projection. Motions in egocentric videos can be separated into the object motions and the camera motions by the camera wearer. Due to the camera motions, tracking results are not aligned with the camera coordinate system of the start frame f_{start} . However, the camera motion between two consecutive frames can be assumed to be small and is represented by a projection matrix. To obtain the projection matrix between two consecutive frames, we use Depth Anything [97] to estimate depth, convert each RGB-D image frame into point clouds, and then apply point cloud registration to determine the projection matrix of the camera extrinsic parameters. We successfully obtain the sequence of the camera extrinsic parameters from each frame to f_{start} by multiplying these projection matrices. Recently, DUST3R [89] can jointly perform camera pose estimation and point cloud construction, internally using point cloud registration of RoMa [10] and gradient-based optimization. However, we notice DUST3R can squeeze detailed 3D motions of objects or hands in egocentric views, leaving this direction for future work.

Each 2D image frame is first converted into a point cloud using an estimated depth map. Let (i, j) denote the image pixel coordinates and d_{ij} the corresponding depth, then

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = d_{ij} K^{-1} \begin{bmatrix} i \\ j \\ 1 \end{bmatrix}, \quad (1)$$

where K is the camera’s intrinsic parameters. Next, we compute the normals and Fast Point Feature Histograms (FPFH) [76] for each point cloud to perform registration. For initial alignment, we use RANSAC-based [23] feature matching between two point clouds using these FPFH features. We further refine the alignment using the colored it-

erative closest point algorithm [68]. With the above operations, we obtain camera extrinsic parameters for each frame. The 3D tracking results are then projected to the camera coordinate system of f_{start} by multiplying the coordinates at each timestep by the corresponding projection matrix.

Rotation Sequence Extraction. As shown in Fig. 2, SpaTracker effectively handles self-occlusion, enabling the computation of temporal object pose transformations. Given the tracking result of the manipulated object, the object pose is computed using singular value decomposition (SVD). First, we set the initial pose of the manipulated object to an identity transformation: $[0, 0, 0]^T$. We then compute the covariance matrix H between the object point cloud at the initial timestamp and each subsequent timestamp. Applying singular value decomposition to H gives $H = U\Sigma V^T$. Finally, we compute the rotation matrix $R = VU^T$, where $R \in \mathbb{R}^{3 \times 3}$, and convert R into a rotation vector $r \in \mathbb{R}^3$. By concatenating the position and rotation sequences, we obtain the 6DoF object manipulation trajectory. Additionally, we calculate the minimum 3D bounding box from the object point cloud at the initial timestamp.

3.2. Evaluation Data

For evaluation, we utilize the existing 6DoF hand and object tracking dataset of HOT3D [3]. HOT3D is an egocentric view dataset for 3D hand and object tracking. The dataset is constructed using optical markers and multiple infrared OptiTrack cameras, yielding precise hand and object 6DoF information. The dataset is recorded with Project Aria glasses [21] and Quest 3 [63]. Unfortunately, the dataset does not include temporal action descriptions, resulting in a misalignment with our task setting. To mitigate this difference, we split the raw videos and determine the action start and end timestamps with action descriptions as described in Sec. 3.1. Additionally, we obtain depth maps using the depth estimation model and align the trajectories with the estimated depth maps for models that require depth or point

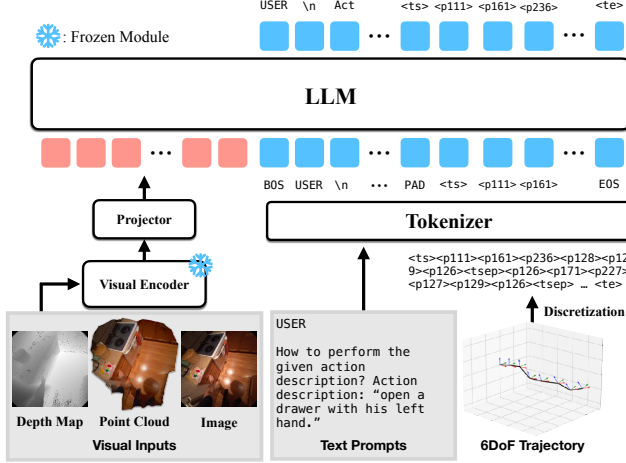


Figure 3. **Overview of model architecture.** Our model architecture utilizes visual and point cloud-based language models as backbones and extends them by incorporating extended vocabularies for trajectory tokenization.

clouds. Further details are provided in the Appendix.

4. 6DoF Object Manipulation Generation

Task Setting. Our task is to generate a sequence of 6DoF object poses for object manipulation, based on an action description and an initial state comprising a visual input and the initial pose of the object. Each pose in this sequence is defined by the position of the object’s centroid and its rotation in 3D space.

Model. Considering recent advancements in multi-modal language models [12, 39, 87, 88], we develop object manipulation trajectory generation models based on visual and point cloud-based language models (VLMs), formalizing our task as next token prediction task. This is achieved by incorporating an extended vocabulary for trajectory tokenization into the VLMs. An overview of the model architecture is depicted in Figure 3.

Discretization of object poses. Inspired by robotic foundation models [9, 15, 44, 108], we discretize each continuous dimension into 256 bins. For each variable, we map the target variable to one of the 256 bins, where the bins are uniformly distributed within the bounds of each variable. Using this discretization, we obtain an $N \times 6$ array of discrete integers from a trajectory, where each point $(x, y, z, \text{roll}, \text{pitch}, \text{yaw}) \in [0, 255]^6$. We reserve $N = 256$ special tokens as trajectory tokens.

Backbones. As our framework is compatible with any vision-language models (VLMs), we employ several VLMs as the backbone architecture.

- **BLIP-2:** BLIP-2 [50] is a VLM for image captioning and visual question answering. The model uses a pre-trained vision transformer from CLIP [73] and OPT

(2.7B/6.7B) [102] as its language decoder.

- **VILA:** VILA [54] is a state-of-the-art model for downstream vision-and-language tasks. It uses a similar architecture to BLIP-2 and is capable of reasoning over multiple images.
- **PointLLM:** PointLLM [94] is a 3D-LLM designed for 3D vision tasks, such as 3D question answering and 3D grounding. The model uses Point-BERT [99] from ULIP-2 [95] as a colored point cloud encoder and LLaMA [84] with the Vicuna (7B/13B) [105] checkpoints as the language decoder.
- **MiniGPT-3D:** MiniGPT-3D [82] is a light-weight 3D-LLM for downstream 3D vision understanding tasks. This model uses a similar architecture to PointLLM and incorporates an extended Mixture of Experts [42, 43].

As existing VLMs do not utilize depth information, we extend BLIP-2 [50] and VILA [54] by incorporating an additional encoder for depth maps. To evaluate the effectiveness of depth information, we use a pre-trained image encoder as the depth encoder for each model. We concatenate the encoded depth features with image features and provide them as inputs to LLM. For the objective function, we employ cross-entropy loss across all backbone models.

5. Experiments

5.1. Experimental Setting

Dataset. We apply our framework (Sec. 3.1) to Ego-Exo4D [30] dataset, resulting in collecting 28,497 samples. We split them into 27,788 training samples and 709 validation samples. Additionally, we use data extracted from HOT3D [3] as the test split, resulting in 986 test samples. We sample image frames at 20 fps from each video clip, resulting in the average extracted trajectory length of 33.17 frames. For model training and evaluation, trajectories longer than 20 frames are cropped to 20 frames in length due to the token sequence length limit, resulting in the average length being 15.39 frames.

Baselines. We employ a traditional transformer architecture with a multi-layer perceptron head (Seq2Seq) that recursively predicts the next pose of trajectories [86, 106]. This architecture is often utilized for vehicle and pedestrian trajectory forecasting such as [28]. Further details are in the Appendix. We also apply a 3D hand trajectory forecasting model of USST [5]. USST is an uncertainty-aware state space Transformer model that utilizes the attention mechanism and aleatoric uncertainty. It generates future hand positions in 3D space from a past temporal sequence of the RGB frames and 3D hand positions. USST is different from our VLM-based models in several major points. First, *our models use action descriptions* while USST does not. Instead of action descriptions, USST depends on the past visual features. Second, USST can predict hand positions

Model	Input	3DoF Training				6DoF Training				
		3D pos.		2D pos.		3D pos.		2D pos.		3D rot.
		ADE	FDE	ADE	FDE	ADE	FDE	ADE	FDE	GD
Seq2Seq [86]	Image	0.269	0.475	0.215	0.372	0.374	0.625	0.325	0.528	0.559
BLIP-2 (2.7B) [50]	Image	0.278	<u>0.464</u>	0.218	<u>0.346</u>	0.280	<u>0.465</u>	0.218	0.344	0.545
BLIP-2 (6.7B) [50]	Image	0.286	0.487	0.224	0.365	0.286	0.475	0.219	0.349	0.543
VILA (3B) [54]	Image	0.500	0.662	0.428	0.546	0.477	0.619	0.390	0.489	0.826
VILA (8B) [54]	Image	0.316	0.512	0.249	0.388	0.293	0.478	0.225	0.347	0.545
Seq2Seq [86]	Image + Depth	0.302	0.541	0.250	0.438	0.341	0.558	0.280	0.452	0.590
BLIP-2 (2.7B) [50]	Image + Depth	0.288	0.481	0.227	0.363	<u>0.275</u>	0.458	<u>0.211</u>	<u>0.332</u>	0.543
BLIP-2 (6.7B) [50]	Image + Depth	0.283	0.477	0.218	0.351	<u>0.282</u>	0.469	<u>0.219</u>	<u>0.345</u>	<u>0.542</u>
VILA (3B) [54]	Image + Depth	0.301	0.492	0.229	0.355	0.294	0.480	0.223	0.344	0.541
VILA (8B) [54]	Image + Depth	0.298	0.494	0.234	0.368	0.318	0.513	0.253	0.391	0.563
MiniGPT-3D (2.7B) [82]	Point cloud	0.299	0.487	0.236	0.368	0.281	0.467	0.218	0.342	0.544
PointLLM (7B) [94]	Point cloud	<u>0.274</u>	0.459	0.210	0.328	0.271	0.458	0.208	0.327	0.541

Table 2. **Comparison of 3DoF and 6DoF object manipulation trajectory generation.** The “Input” column indicates the visual modality input for each model. Lower values are preferable for all metrics. Best and secondary results are viewed in **bold** and underline.

	Pre-trained	EgoTraj	3D pos.		2D pos.	
			ADE	FDE	ADE	FDE
USST [5]		✓	0.319	0.464	0.225	0.315
USST [5]	✓		0.303	0.414	0.243	0.310
USST [5]	✓	✓	0.254	0.448	0.186	0.307

Table 3. **Comparison with existing 3DoF trajectory forecasting methods.** “Pre-trained” means the model is trained with densely recorded datasets of H2O [46] and EgoPAT3D [53].

considering the center of the hand but not rotations, while *ours* can predict both the positions and rotations of the objects. As USST is the hand position prediction model, we use object poses instead of hand poses for our task.

Evaluation Metrics. We evaluate model performance using the following three metrics.

- **Average Displacement Error (ADE):** ADE is the l_2 distance of points in each time step between the predicted and ground-truth trajectories.
- **Final Displacement Error (FDE):** FDE is the l_2 distance of the final points between the predicted and ground-truth trajectories.
- **Geodesic Distance (GD):** GD [92] is the angular difference between two rotation matrices as:

$$d(R^*, R) = \arccos\left(\frac{\text{tr}(R^{*T}R) - 1}{2}\right),$$

where R^* and R denote the ground-truth and predicted 3D rotations, respectively. tr denotes the trace of matrix.

When generating trajectories, some VLM models can halt further generation when they predict a stop token. While this feature allows models to generate arbitrary lengths of trajectories that are not possible for the existing models as of USST [5], this causes a problem in the evaluation due to the varying trajectory length. Therefore, we apply padding or cutting up for the final pose of a generated trajectory to fit the length of the ground-truth for the evaluation.

5.2. Results

Tab. 2 presents the results of text-guided 3DoF and 6DoF object manipulation trajectory generation in HOT3D. VLM-based models outperform conventional Seq2Seq methods. Specifically, when comparing PointLLM to the Seq2Seq (*v*) model, PointLLM achieves a reduction of 0.103 m in ADE (3D). These results demonstrate the effectiveness of utilizing VLMs for object manipulation trajectory generation.

Analysis of VLM-based models reveals two key insights. First, when comparing model performance across different visual modalities, the utilization of point clouds significantly contributes to performance improvement. Additionally, the use of depth maps in 2D-VLMs tends to result in better performance. These results suggest that incorporating spatial information such as depth and point clouds boosts model performance. Second, although general scaling laws suggest that larger LLMs achieve better performance, this trend does not consistently apply to trajectory prediction. For instance, BLIP-2 (2.7B) outperforms VILA (8B). These results suggest that improving model performance for our task requires not only increasing model size but also considering the specific characteristics of each LLM.

Comparison with Existing Trajectory Set. Tab. 3 presents the performance of the USST models [5]. The USST model that is trained with our extracted trajectories from scratch performed on-par with the USST model that is pre-trained with densely recorded datasets [46, 53]. When we further fine-tuned the pre-trained USST model with our dataset, we observed the performance gain in many metrics, confirming the effectiveness of the automatically extracted trajectories.

Probabilistic Sampling. There are possible multiple trajectories that match the given action description [13, 47]. It is difficult for existing models to generate multiple possible trajectories as USST does not rely on action descrip-


“Pick up the cellphone from the table with the right hand.”



“Stir the contents of the bowl with the wooden spoon using the right hand.”



“Pick up the bamboo plate from the table with both hands.”



Initial state

Generation

Figure 4. Qualitative results of PointLLM [94]. Generated trajectories are illustrated using 3D bounding boxes for visualization.

Sampling	3D pos.		2D pos.		3D rot.
	ADE	FDE	ADE	FDE	GD
1	0.271	0.458	0.208	0.327	0.541
3	0.236	0.401	0.178	0.280	0.543
10	0.212	0.363	0.158	0.252	0.540

Table 4. Performance comparison in probabilistic sampling for PointLLM [94]. “Sampling” column indicates the number of samples generated.

tions. Fortunately, our model uses the language modeling approach that allows generating multiple different trajectories using nucleus sampling [37]. We employ PointLLM as it demonstrated the best performance among the baselines for this purpose. For each instance, we report the lowest ADE (3D) score from the sampled trajectories. Tab. 4 shows that increasing the number of samples improves model performance. These results imply that our model successfully generates multiple trajectory candidates to perform the given action descriptions. Meanwhile, the performance improvement for GD is smaller than other positional metrics, indicating the difficulty in generating accurate rotational movements.

Data Scalability. As our framework automatically generates manipulation trajectories, we investigate the impact of dataset scale on model performance. We utilized training data at ratios of 1.0, 0.5, 0.1, 0.05, and 0.01 of the full dataset to assess performance across different dataset scales. Fig. 5 presents the performance comparison under varying dataset scales. We use PointLLM for this purpose




Figure 5. Comparison of performance across different dataset scales for PointLLM [94].

as it achieved the highest performance in Tab. 2. The results, especially in ADE and FDE, indicate that model performance improves as the dataset size increases, demonstrating the effectiveness of our dataset creation framework.

Qualitative Results. Fig. 4 illustrates the generated trajectories of PointLLM along with the changes of the 6DoF object pose from the initial pose. It is interesting that regardless of the difference of the objects in HOT3D from those in the Ego-Exo4D videos, our method successfully generates understandable object manipulation trajectories depending on action descriptions and their verbs. For example, the case in the middle row is with the verb “stir” and it yields a different trajectory from other examples.

Action Description Change. We survey the impact of the action descriptions on the generation of trajectories, as it is highly expected that the generated trajectories are modified depending on the action descriptions. Fig. 6 presents two generated trajectories from the trained PointLLM, confirming two different movements depending on the action de-

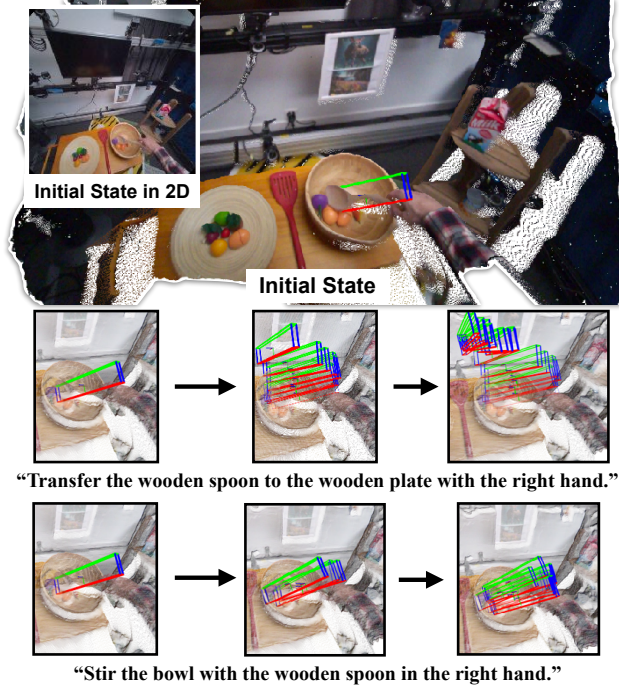


Figure 6. **PointLLM [94] results with different action descriptions.** The generations are visualized with 3D bounding boxes.

	Trajectory	Sim_v	METEOR	BLEU-4
BLIP-2 [50]		58.59	23.76	11.93
	✓	67.58	24.19	12.48
PointLLM [94]		43.05	17.42	3.25
	✓	53.33	17.56	4.44

Table 5. **Effect of trajectory information on image and 3D captioning performance** The “Trajectory” column indicates whether trajectory information is utilized by the model (✓). Higher values indicate better performance across all metrics.

scriptions: the wooden spoon is pulled away from the plate in the top row while it stays on the bowl and is being shaken on the bottom row. Further visualizations including demonstration videos in our project page.

5.3. Action Description Generation

Finally, we conduct experiments of generating action descriptions with BLIP-2 and PointLLM. Indeed, current VLMs are so effective that they can learn to generate at first-look plausible action descriptions from solely 2D image 3D scenery inputs. Therefore, we test two conditions of with and without trajectories. In addition to the standard metrics of METEOR [4] and BLEU-4 [67] for action description evaluation, we also report the similarity of the verbs between the generated and annotated trajectories, Sim_v , because they reflect the differences in the actions by camera wearers in egovision. Sim_v is computed as the cosine similarity of word embeddings [70]. Tab. 5 shows the results from BLIP-2 [50] and PointLLM [94]. The results indi-

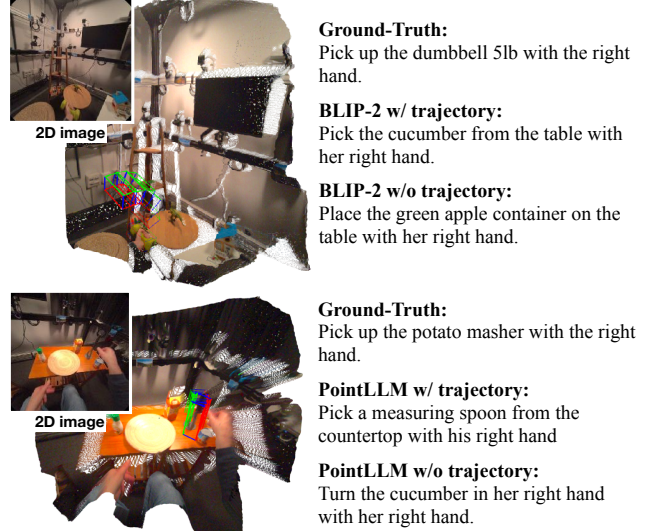


Figure 7. **Qualitative results in image and 3D captioning task.** The top and bottom figures depict generated captions of BLIP-2 [50] and PointLLM [94], respectively. “w/” and “w/o” indicate whether models utilize trajectory information.

cate that utilizing trajectories boosts performance, particularly verb similarity, by nearly 10% in both image and 3D captioning, confirming the effectiveness of our framework for extracting accurate manipulation trajectories that reflect human motion. Fig. 7 illustrates the qualitative results of image and 3D captioning, indicating that object motion trajectories enhance verb similarity in both models.

6. Conclusion

We developed a scalable framework to extract 6DoF object manipulation trajectories from egocentric videos. Using these trajectories, we created generation models for object manipulation trajectories based on visual and point cloud-based language models. Our experimental results on HOT3D dataset demonstrate that our models learn to generate object trajectories from action descriptions, providing fair baseline models for the new task. Moreover, utilizing trajectories for visual captioning tasks reveals that our dataset creation framework successfully extracts meaningful trajectories from egocentric videos. For future work, we construct a massive-scale dataset by adapting the proposed approach to other egocentric video datasets and investigate its effectiveness in robotic manipulation.

Limitations. The current framework has two limitations. First, because our framework defines an object pose as a 6DoF bounding box, it cannot be applied to deformable objects such as clothes. Second, we observed failure cases in our dataset due to object segmentation and point-cloud registration. Segmentation can fail when similar objects are present, while registration can fail when camera pose changes abruptly. Further details are in the Appendix.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP22K17983, JP22KK0184 and JST CRONOS JP-MJCS24K6.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [2] Shikhar Bahl, Russell Mendonca, Lili Chen, Unnat Jain, and Deepak Pathak. Affordances from human videos as a versatile representation for robotics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [3] Prithviraj Banerjee, Sindi Shkodrani, Pierre Moulon, Shreyas Hampali, Fan Zhang, Jade Fountain, Edward Miller, Selen Basol, Richard Newcombe, Robert Wang, Jakob Julian Engel, and Tomas Hodan. Introducing hot3d: An egocentric dataset for 3d hand and object tracking. *arXiv preprint arXiv:2406.09598*, 2024. 1, 2, 3, 4, 5
- [4] Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 2005. 8
- [5] Wentao Bao, Lele Chen, Libing Zeng, Zhong Li, Yi Xu, Junsong Yuan, and Yu Kong. Uncertainty-aware state space transformer for egocentric 3d hand trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 5, 6
- [6] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, and Elad Shulman. ARKitscenes - a diverse real-world dataset for 3d indoor scene understanding using mobile RGB-d data. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 1
- [7] Gedas Bertasius, Hyun Soo Park, Stella X. Yu, and Jianbo Shi. First-person action-object detection with egonet. In *Proceedings of Robotics: Science and Systems (RSS)*, 2017. 2
- [8] Samarth Brahmabhatt, Cusuh Ham, Charles C. Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [9] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, et al. RT-1: Robotics Transformer for Real-World Control at Scale. In *Proceedings of Robotics: Science and Systems*, 2023. 2, 3, 5
- [10] Romain Brégier. Deep regression on manifolds: A 3d rotation case study. In *2021 International Conference on 3D Vision (3DV)*, 2021. 4
- [11] Junuk Cha, Jihyeon Kim, Jae Shin Yoon, and Seungryul Baek. Text2hoi: Text-guided 3d motion generation for hand-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [12] Ting Chen, Saurabh Saxena, Lala Li, David J. Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. In *10th International Conference on Learning Representations, ICLR*, 2022. 3, 5
- [13] Chiho Choi and Behzad Dariush. Looking to relations for future trajectory forecast. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 6
- [14] Ching-Yao Chuang, Jiaman Li, Antonio Torralba, and Sanja Fidler. Learning to act properly: Predicting and explaining affordances from images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [15] Open X-Embodiment Collaboration. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024. 1, 2, 5
- [16] Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer, and Gregory Rogez. Ganhand: Predicting human grasp affordances in multi-object scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [17] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [18] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, 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*, 130(1):33–55, 2022. 3
- [19] Thanh-Toan Do, Anh Nguyen, and Ian Reid. Affordancenet: An end-to-end deep learning approach for object affordance detection. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 2
- [20] Guoguang Du, Kai Wang, Shiguo Lian, and Kaiyong Zhao. Vision-based robotic grasping from object localization, object pose estimation to grasp estimation for parallel grippers: a review. *Artificial Intelligence Review*, 54(3):1677–1734, 2021. 3
- [21] Jakob Engel, Kiran Somasundaram, Michael Goesele, Albert Sun, Alexander Gamino, Andrew Turner, Arjang Talattof, Arnie Yuan, Bilal Souti, Brigid Meredith, et al. Project aria: A new tool for egocentric multi-modal ai research. *arXiv preprint arXiv:2308.13561*, 2023. 1, 4
- [22] Irving Fang, Yuzhong Chen, Yifan Wang, Jianghan Zhang, Qiushi Zhang, Jiali Xu, Xibo He, Weibo Gao, Hao Su, Yiming Li, and Chen Feng. Egopat3dv2: Predicting 3d action

target from 2d egocentric vision for human-robot interaction. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024. 2

- [23] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 1981. 4
- [24] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile ALOHA: Learning bimanual mobile manipulation using low-cost whole-body teleoperation. In *8th Annual Conference on Robot Learning*, 2024. 1
- [25] Antonino Furnari and Giovanni Maria Farinella. What would you expect? anticipating egocentric actions with rolling-unrolling lstms and modality attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2
- [26] Antonino Furnari, Sebastiano Battiato, Kristen Grauman, and Giovanni Maria Farinella. Next-active-object prediction from egocentric videos. *Journal of Visual Communication and Image Representation*, 49:401–411, 2017. 2
- [27] Nisal Menuka Gamage, Deepana Ishtaweera, Martin Weigel, and Anusha Withana. So predictable! continuous 3d hand trajectory prediction in virtual reality. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, 2021. 2
- [28] Francesco Giuliari, Irtiza Hasan, Marco Cristani, and Fabio Galasso. Transformer networks for trajectory forecasting. In *25th International Conference on Pattern Recognition (ICPR)*, 2021. 5
- [29] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [30] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. *arXiv preprint arXiv:2311.18259*, 2023. 2, 3, 5
- [31] Andrew Guo, Bowen Wen, Jianhe Yuan, Jonathan Tremblay, Stephen Tyree, Jeffrey Smith, and Stan Birchfield. Handal: A dataset of real-world manipulable object categories with pose annotations, affordances, and reconstructions. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023. 2
- [32] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [33] Shengyu Hao, Wenhao Chai, Zhonghan Zhao, Meiqi Sun, Wendi Hu, Jieyang Zhou, Yixian Zhao, Qi Li, Yizhou Wang, Xi Li, and Gaoang Wang. Ego3dt: Tracking every 3d object in ego-centric videos. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024. 3
- [34] Xingyi He, Jiaming Sun, Yuang Wang, Di Huang, Hujun Bao, and Xiaowei Zhou. Onepose++: Keypoint-free one-shot object pose estimation without cad models. In *Advances in Neural Information Processing Systems*, 2022. 3
- [35] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [36] Yisheng He, Yao Wang, Haoqiang Fan, Jian Sun, and Qifeng Chen. Fs6d: Few-shot 6d pose estimation of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [37] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *8th International Conference on Learning Representations, ICLR*, 2020. 7
- [38] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. In *Advances in Neural Information Processing Systems*, 2023. 3
- [39] De-An Huang, Shijia Liao, Subhashree Radhakrishnan, Hongxu Yin, Pavlo Molchanov, Zhiding Yu, and Jan Kautz. Lita: Language instructed temporal-localization assistant. *arXiv preprint arXiv:2403.19046*, 2024. 3, 5
- [40] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *7th Annual Conference on Robot Learning*, 2023. 2
- [41] Tatsuya Ishibashi, Kosuke Ono, Noriyuki Kugo, and Yuji Sato. Technical report for ego4d long term action anticipation challenge 2023. *arXiv preprint arXiv:2307.01467*, 2023. 2
- [42] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991. 5
- [43] M.I. Jordan and R.A. Jacobs. Hierarchical mixtures of experts and the em algorithm. In *Proceedings of 1993 International Conference on Neural Networks (IJCNN)*, 1993. 5
- [44] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024. 2, 3, 5
- [45] Hema S. Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):14–29, 2016. 2
- [46] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 6
- [47] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with inter-

- acting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [48] Gen Li, Varun Jampani, Deqing Sun, and Laura Sevilla-Lara. Locate: Localize and transfer object parts for weakly supervised affordance grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [49] Gen Li, Nikolaos Tsagkas, Jifei Song, Ruairidh Mon-Williams, Sethu Vijayakumar, Kun Shao, and Laura Sevilla-Lara. Learning precise affordances from egocentric videos for robotic manipulation. *arXiv preprint arXiv:2408.10123*, 2024. 2
- [50] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. 3, 5, 6, 8
- [51] Yin Li, Miao Liu, and James M. Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [52] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [53] Yiming Li, Ziang Cao, Andrew Liang, Benjamin Liang, Luoyao Chen, Hang Zhao, and Chen Feng. Egocentric prediction of action target in 3d. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 6
- [54] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 5, 6
- [55] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, 2023. 3
- [56] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [57] Miao Liu, Siyu Tang, Yin Li, and James M. Rehg. Forecasting human-object interaction: Joint prediction of motor attention and actions in first person video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [58] Shaowei Liu, Subarna Tripathi, Somdeb Majumdar, and Xiaolong Wang. Joint hand motion and interaction hotspots prediction from egocentric videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [59] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [60] Liangsheng Lu, Wei Zhai, Hongchen Luo, Yu Kang, and Yang Cao. Phrase-based affordance detection via cyclic bilateral interaction. *IEEE Transactions on Artificial Intelligence*, 4(5):1186–1198, 2023. 2
- [61] Junyi Ma, Xieyuanli Chen, Wentao Bao, Jingyi Xu, and Hesheng Wang. Madiff: Motion-aware mamba diffusion models for hand trajectory prediction on egocentric videos. *arXiv preprint arXiv:2409.02638*, 2024. 2
- [62] Junyi Ma, Jingyi Xu, Xieyuanli Chen, and Hesheng Wang. Diff-ip2d: Diffusion-based hand-object interaction prediction on egocentric videos. *arXiv preprint arXiv:2405.04370*, 2024. 2
- [63] Meta. Quest 3, 2023. 1, 4
- [64] Antoine Miech, Ivan Laptev, Josef Sivic, Heng Wang, Lorenzo Torresani, and Du Tran. Leveraging the present to anticipate the future in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019. 2
- [65] Soroush Nasiriany, Sean Kirmani, Tianli Ding, Laura Smith, Yuke Zhu, Danny Driess, Dorsa Sadigh, and Ted Xiao. Rt-affordance: Affordances are versatile intermediate representations for robot manipulation. *arXiv preprint arXiv:2411.02704*, 2024. 2
- [66] Toan Nguyen, Minh Nhat Vu, An Vuong, Dzung Nguyen, Thieu Vo, Ngan Le, and Anh Nguyen. Open-vocabulary affordance detection in 3d point clouds. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023. 2
- [67] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 2002. 8
- [68] Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Colored point cloud registration revisited. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017. 4
- [69] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [70] Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014. 8
- [71] Hamed Pirsiavash and Deva Ramanan. Detecting activities of daily living in first-person camera views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 2
- [72] Chiara Plizzari, Shubham Goel, Toby Perrett, Jacob Chalk, Angjoo Kanazawa, and Dima Damen. Spatial cognition from egocentric video: Out of sight, not out of mind. *arXiv preprint arXiv:2404.05072*, 2024. 3
- [73] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 2021. 5

- [74] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 2, 3
- [75] Ivan Rodin, Antonino Furnari, Kyle Min, Subarna Tripathi, and Giovanni Maria Farinella. Action scene graphs for long-form understanding of egocentric videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [76] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2009. 4
- [77] Johann Sawatzky, Yaser Souri, Christian Grund, and Jürgen Gall. What object should i use? - task driven object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [78] Fadime Sener, Dibyadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [79] Yale Song, Eugene Byrne, Tushar Nagarajan, Huiyu Wang, Miguel Martin, and Lorenzo Torresani. Ego4d goal-step: Toward hierarchical understanding of procedural activities. In *Advances in Neural Information Processing Systems*, 2023. 3
- [80] Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and Xiaowei Zhou. Onepose: One-shot object pose estimation without cad models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [81] Yoshiyuki Tanaka. Robot-aided rehabilitation methodology for enhancing movement smoothness by using a human hand trajectory generation model with task-related constraints. *Journal of Human-Robot Interaction*, 4(3): 101–119, 2015. 2
- [82] Yuan Tang, Xu Han, Xianzhi Li, Qiao Yu, Yixue Hao, Long Hu, and Min Chen. Minigpt-3d: Efficiently aligning 3d point clouds with large language models using 2d priors. *arXiv preprint arXiv:2405.01413*, 2024. 3, 5, 6
- [83] Sanket Thakur, Cigdem Beyan, Pietro Morerio, Vittorio Murino, and Alessio Del Bue. Leveraging next-active objects for context-aware anticipation in egocentric videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024. 2
- [84] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 5
- [85] Vadim Tschernezki, Ahmad Darkhalil, Zhifan Zhu, David Fouhey, Iro Laina, Diane Larlus, Dima Damen, and Andrea Vedaldi. Epic fields: Marrying 3d geometry and video understanding. In *Advances in Neural Information Processing Systems*, 2023. 3
- [86] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017. 5, 6
- [87] Junke Wang, Dongdong Chen, Chong Luo, Bo He, Lu Yuan, Zuxuan Wu, and Yu-Gang Jiang. Omnivid: A generative framework for universal video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 5
- [88] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *Proceedings of the 39th International Conference on Machine Learning*, 2022. 3, 5
- [89] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 4
- [90] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Fruejri, Neel Joshi, and Marc Pollefeys. Holoassist: an egocentric human interaction dataset for interactive ai assistants in the real world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [91] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [92] Paul Wohlhart and Vincent Lepetit. Learning descriptors for object recognition and 3d pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 6
- [93] Yuxi Xiao, Qianqian Wang, Shangzhan Zhang, Nan Xue, Sida Peng, Yujun Shen, and Xiaowei Zhou. Spatialtracker: Tracking any 2d pixels in 3d space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3
- [94] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023. 3, 5, 6, 7, 8
- [95] Le Xue, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip-2: Towards scalable multimodal pre-training for 3d understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 5
- [96] Dejie Yang and Yang Liu. Active object detection with knowledge aggregation and distillation from large models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [97] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Un-

- leashing the power of large-scale unlabeled data. *arXiv preprint arXiv:2401.10891*, 2024. [3](#), [4](#)
- [98] Yufei Ye, Abhinav Gupta, and Shubham Tulsiani. What’s in your hands? 3d reconstruction of generic objects in hands. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#)
 - [99] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [5](#)
 - [100] Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2023. [3](#)
 - [101] Mengqi Zhang, Yang Fu, Zheng Ding, Sifei Liu, Zhuowen Tu, and Xiaolong Wang. Hoidiffusion: Generating realistic 3d hand-object interaction data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [2](#)
 - [102] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022. [5](#)
 - [103] Zichen Zhang, Hongchen Luo, Wei Zhai, Yang Cao, and Yu Kang. Pear: Phrase-based hand-object interaction anticipation. *arXiv preprint arXiv:2407.21510*, 2024. [2](#)
 - [104] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Proceedings of Robotics: Science and Systems*, 2023. [1](#)
 - [105] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, 2023. [5](#)
 - [106] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [5](#)
 - [107] Ziyu Zhu, Xiaojian Ma, Yixin Chen, Zhidong Deng, Siyuan Huang, and Qing Li. 3d-vista: Pre-trained transformer for 3d vision and text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [3](#)
 - [108] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of The 7th Conference on Robot Learning*, 2023. [2](#), [3](#), [5](#)