

# EgoAgent: A Joint Predictive Agent Model in Egocentric Worlds

Lu Chen<sup>1\*</sup> Yizhou Wang<sup>2\*</sup> Shixiang Tang<sup>2†</sup> Qianhong Ma<sup>3</sup> Tong He<sup>4</sup>  
 Wanli Ouyang<sup>2</sup> Xiaowei Zhou<sup>1</sup> Hujun Bao<sup>1</sup> Sida Peng<sup>1†</sup>

<sup>1</sup>State Key Lab of CAD&CG, Zhejiang University <sup>2</sup>The Chinese University of Hong Kong  
<sup>3</sup>Shanghai Jiao Tong University <sup>4</sup>Shanghai Artificial Intelligence Laboratory

\*Equal contribution <sup>†</sup>Corresponding author

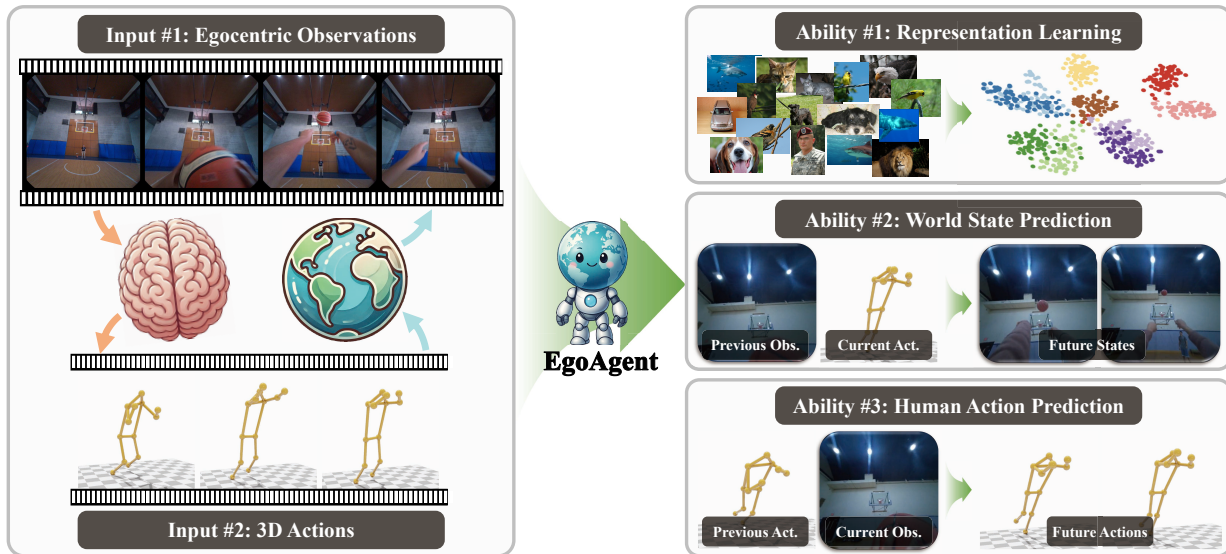


Figure 1. (a) Humans learn from the world by continuously perceiving the egocentric world, predicting future states, and taking actions to achieve their goals. (b) Inspired by the human learning process, we propose **EgoAgent**, a joint predictive agent model in egocentric worlds that can learn to represent, predict, and act from egocentric videos and synchronized 3D human body motions.

## Abstract

Learning an agent model that behaves like humans—capable of jointly perceiving the environment, predicting the future, and taking actions from a first-person perspective—is a fundamental challenge in computer vision. Existing methods typically train separate models for these abilities, which fail to capture their intrinsic relationships and prevent them from learning from each other. Inspired by how humans learn through the perception-action loop, we propose *EgoAgent*, a unified agent model that simultaneously learns to represent, predict, and act within a single transformer. *EgoAgent* explicitly models the causal and temporal dependencies among these abilities by formulating the task as an interleaved sequence of states and actions. It further introduces a joint embedding-action-prediction architecture with temporally asymmetric predictor and observer branches, enabling synergistic optimization across all three capabilities. Comprehensive evaluations of *EgoAgent* on representative tasks such as image classification, egocen-

tric future state prediction, and 3D human motion prediction demonstrate the superiority of our method. The code and trained models will be publicly available at <https://github.com/zju3dv/EgoAgent>.

## 1. Introduction

“Cognition is embodied; it arises from bodily interactions with the world,” as stated by Linda B. Smith<sup>1</sup> in [43]. This perspective emphasizes that cognition emerges from continuous, dynamic interactions between the brain, body, and environment, rather than relying on symbolic systems detached from real-world experiences [8, 40, 42]. Through these ongoing interactions, humans naturally acquire foundational abilities such as *visual perception*, *predictive un-*

<sup>1</sup>Linda Brawn Smith is a cognitive scientist recognized for her work in developmental psychology and cognitive science, with over 53,000 citations. She won the David E. Rumelhart Prize for theoretical contributions to cognitive science and is a member of both the National Academy of Sciences and the American Academy of Arts and Science.

*derstanding of world dynamics, and action anticipation.* Replicating these abilities in artificial systems remains one major challenge in AI research, as highlighted by Yann LeCun [30]. Motivated by this human learning process, we propose an agent model that jointly learns these abilities from egocentric RGB observations and their responsive actions. Such abilities are essential for applications in robotics, augmented reality, and gaming, where agents must understand what is happening in the world, anticipate what will happen next, and decide what actions to take.

Existing approaches typically address these capabilities in isolated tasks: (1) Visual representation learning [10, 27, 46], which encodes high-level representations of human observations of the world; (2) World models [37, 51], which learn predictive representations of world state transitions conditioned on actions; (3) Action prediction [9, 36], which forecasts future human actions based on prior motion. However, research in cognitive science, particularly the famous *Common Coding Theory* [8, 39], claims that perception and action are not separate processes but deeply intertwined and mutually reinforcing, operating within a shared representational space. Therefore, modeling these abilities within a unified framework is not only mutually beneficial but also represents a more principled approach for agent models to achieve human-like understanding in egocentric worlds.

While combining these abilities offers clear benefits, integrating them into one model remains challenging. Human interaction with the world involves a continuous loop of perceiving egocentric observations and taking 3D actions, creating complex interdependencies between observations and actions in time and causality. A successful model must therefore learn from this interleaved and tightly-coupled structure. The core difficulty, which remains an open question, lies in designing learning frameworks and supervision signals to capture these intricate connections.

In response, we propose *EgoAgent*, a joint predictive agent model that simultaneously learns to represent world observations, predict future states, and act based on learned representations (Fig. 1). *EgoAgent* introduces two key innovations: (1) Formulating human-world interactions as interleaved sequences of “state-action-state-action” tokens processed by the causal attention mechanism. This design explicitly models both how observations trigger actions and how those actions, in turn, influence the subsequent states of the world. (2) A joint embedding-action-prediction (JEAP) architecture with temporally asymmetric predictor-observer branches. The predictor leverages past information to forecast future human actions and world states, while the observer extracts the target world states from raw observations. By minimizing the prediction error between these two branches in a shared semantic feature space, JEAP is tailored to learn both world dynamics and human policy based on high-level visual representations, thereby integrating the

three tasks into a cohesive learning framework.

Extensive experiments show that *EgoAgent* effectively handles perception, prediction, and action tasks, even surpassing state-of-the-art methods across individual benchmarks. Specifically, it outperforms the leading egocentric video pretraining method [44] by **+1.32%** Top1 accuracy on ImageNet-1K [15] for image classification and **+3.90%** success rate of moving cube in the TriFinger simulator [48].<sup>2</sup> On Ego-Exo4D [20], it achieves **+16.28%** Top1 accuracy and **+16.95%** mAP for future state prediction, while also exceeding 3D human motion prediction methods by improving MPJPE by **-0.82** at a 30 fps prediction rate.

Our contributions are as follows: (1) We develop the first agent model that can simultaneously learn to represent egocentric observations, predict future states, and generate 3D actions in a unified representational space. (2) We propose a JEAP architecture with temporally asymmetric branches to learn both semantic understanding and temporal transitions of the dynamic multi-modal egocentric world. (3) We evaluate our *EgoAgent* on egocentric video-action datasets and show its superior abilities in visual representation, world states prediction, and 3D human motion prediction tasks.

## 2. Related Work

**World Models and Agent Models.** Ha and Schmidhuber [23] defined the world model as a predictive representation of the environment, while the agent model utilizes this learned representation to make decisions and take actions. Yann LeCun [30] further envisioned that an agent model can generate actions either reactively or through iterative optimization using the world model. Recently, *generative world models* [2, 49, 51] have demonstrated success in autonomous driving [47], robotics [22], and game control [5], mostly employing an encoder-decoder architecture to learn transitions between video frames. GAIA-1 [28] maps multi-modal driving signals into discrete tokens, and processes them using an autoregressive transformer. In contrast, *JEPA world models* [3, 7, 18] focus on learning state transitions through the joint embedding predictive architecture (JEPA) [30] without generating videos explicitly. MC-JEPA [6] incorporates self-supervised learning with optical flow estimation to learn content features and motion dynamics. In this work, we extend the JEPA world model into a reactive agent model, which not only predicts world states based on human actions but also predicts actions based on learned world state transitions.

**Egocentric Visual Representation Learning.** Traditional visual representation learning methods are trained on large-scale image datasets, such as ImageNet [15], using either

<sup>2</sup>Image classification serves as a standard benchmark for visual perception, while TriFinger evaluates the transferability of learned representations to egocentric manipulation with a three-finger robot. Together, these tasks assess a model’s ability to learn meaningful visual representations.

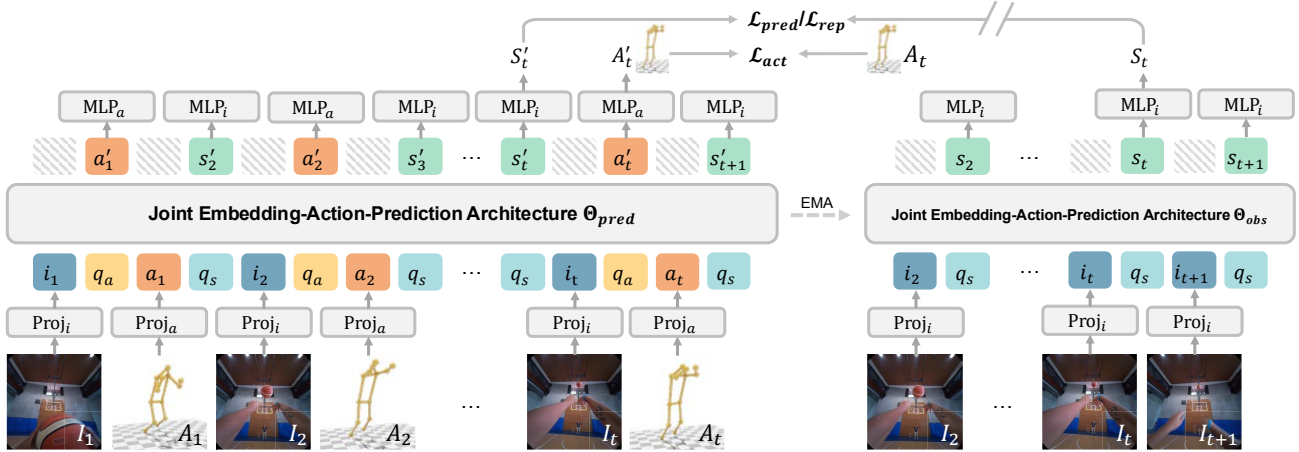


Figure 2. The overall framework of EgoAgent. EgoAgent adopts the Joint Embedding-Action-Prediction Architecture to project input egocentric video frames  $I_t$  and 3D human poses  $A_t$  into image tokens  $i_t$  and action tokens  $a_t$ . Learnable action query token  $q_a$  and future state query token  $q_s$  are appended after the image and action tokens to stimulate EgoAgent’s predictor  $\Theta_{pred}$  to predict the next action  $A'_t$  and world state  $S'_{t+1}$ . The action prediction loss  $\mathcal{L}_{act}$  is used to supervise the predicted actions. Since EgoAgent predicts world states  $S'_{t+1}$  in a semantic feature space, the target world state  $S_{t+1}$  is obtained from a momentum observer  $\Theta_{obs}$ . Features from the predictor and observer are aligned through a state prediction loss  $\mathcal{L}_{pred}$ . When 3D poses are not paired with video frames, EgoAgent randomly crops different views from the video and learns representative features only with image tokens  $i_t$  and state query tokens  $q_s$  by a self-supervised representation loss  $\mathcal{L}_{rep}$ . For easier understanding, we only illustrate the supervisions on predicting the state  $S_t$  and action  $A_t$ .

supervised [25, 29, 46] or self-supervised methods [10, 12, 26]. However, these datasets primarily consist of third-person perspectives, which do not fully align with how humans perceive the world. To bridge this gap, researchers have turned to egocentric videos, which offer a more natural and immersive viewpoint for learning visual representations. Models like R3M [38] and VIP [32] leverage the temporal relationships in egocentric videos to learn generalizable representations from Ego4D [19] for robotics applications. Recently, DoRA [44] learns object-level representations by tracking movements on an egocentric long video dataset called WalkingTours. In this paper, we propose an agent model that not only learns robust representations from egocentric videos but also predicts future states and actions.

**3D Human Motion Prediction.** Human motion prediction models [31, 45, 52] aim to forecast future body movements given historical poses. Aksan et al. [1] proposed a transformer-based architecture to capture the spatial-temporal dependencies of human motion over short and long horizons. Cui et al. [14] represent the human skeleton as a dynamic graph to adaptively learn joint connection strengths for better prediction accuracy. Other works [33, 53] incorporate physics-based priors to ensure the generated motions adhere to physical principles. Recent research has integrated multimodal cues such as textual instructions [35], eye gaze [54], and 3D objects [50] as conditions for human motion prediction. Notably, Cao et al. [9] proposed a three-stage framework that consists of goal prediction, path planning, and pose finalization, which effectively enhances scene-aware motion prediction.

### 3. Method

Our goal is to develop an agent model that can *represent world states, predict future states, and generate future actions* by simulating how humans interact with the world through egocentric vision and 3D skeletal motion. The overview of our framework is presented in Fig. 2. We first introduce our joint embedding-action-prediction architecture in Sec. 3.1. We then introduce how EgoAgent can be trained in Sec. 3.2. Finally, we demonstrate the application of EgoAgent in the three tasks in Sec. 3.3.

#### 3.1. Joint Embedding-Action-Prediction Architecture for EgoAgent

Given egocentric observations  $V = \{I_1, I_2, \dots, I_t\}$  and human actions  $A = \{A_1, A_2, \dots, A_t\}$ , our agent model aims to predict the future world state  $S_{t+1}$  associated with the next observation  $I_{t+1}$  while simultaneously generate plausible human action  $A_{t+1}$ . The world states are represented as semantic embeddings, whereas human actions are 3D poses.

**Interleaved joint prediction.** EgoAgent begins by encoding the egocentric video  $V$  and human action sequence  $A$  into semantic embeddings. For video frames, we follow the approach of ViT [16], applying a convolutional layer to each video frame  $I_t$  to produce a feature map. This map is then subdivided and flattened to obtain the image feature  $i_t$ . For human actions, we utilize a convolutional layer followed by layer normalization (LN) and a Gaussian error linear unit (GeLU) activation to map the 3D body pose  $A_t$  into action feature  $a_t$ . These feature vectors are then processed by a causal transformer network.

As illustrated in Fig. 2, at each time step  $t$ , we assemble a structured sequence of interleaved tokens, including the image token  $i_t$ , an action query token  $q_a$ , the action token  $a_t$ , and a future state query token  $q_s$ . The query tokens are learnable embeddings that serve as explicit prompts, directing the model to predict either the next action or a future world state. Notably, we insert  $q_a$  right after the image tokens and  $q_s$  after the action tokens to model the causal and temporally interlaced nature of perceiving observations and taking actions. Specifically, enabled by transformer’s causal attention mechanism,  $q_a$  incorporates all preceding image tokens  $i_{[0:t]}$  and action tokens  $a_{[0:t-1]}$  to predict  $a'_t$ . Similarly,  $q_s$  integrates  $i_{[0:t]}$  and the newly appended  $a_{[0:t]}$  to predict  $s'_{t+1}$ . We utilize separate MLPs to map these embeddings to 3D actions  $A'_t$  and world states  $S'_{t+1}$ .

**JEPA with temporally asymmetric branches.** Unlike actions, which can be directly obtained and supervised, world states must be learned from observations while also be used as supervision signals for state prediction. To achieve this, we propose a temporally asymmetric predictor-observer architecture based on JEPA [18, 30], a framework that combines representation learning with predictive learning. In addition to the predictor branch executing the interleaved joint prediction, we introduce an observer branch that only processes image inputs by (1) extracting current-frame features for self-supervised representation learning and (2) providing next-frame features for future state supervision alongside the predictor. Specifically, the observer branch takes image tokens  $i_t$  and state query tokens  $q_s$  as input, and outputs current state embeddings  $s_t$ , which are then fed into the same MLP for obtaining world states  $S_t$  from image embeddings in the predictor. A key advantage of our query-based design is that it effectively decouples the shared state/representation components from the predictor’s action-specific ones. This prevents conflicting gradient updates, allowing the observer’s parameters to be updated as an exponential moving average (EMA) of the predictor’s, despite the observer not processing the action modality.

**Learning in semantic feature space.** Notably, unlike previous methods [4, 51] that rely on pretrained reconstruction-based tokenizers, such as VQGAN [17], to convert images into discrete tokens, we employ learnable convolutional layers to project images into continuous semantic embeddings. Discrete tokens prioritize pixel-level reconstruction over high-level semantic information, whereas humans make predictions based on abstract concepts rather than pixels. Therefore, our approach aligns more closely with human reasoning and significantly enhances the model’s performance in future state prediction (see Sec. 4.2).

### 3.2. Training EgoAgent

**Base model.** We employ InternLM, an open-source LLM [41], as the foundational architecture for EgoAgent,

adapting it to the JEPA architecture mentioned in Sec. 3.1. This is motivated by two key considerations: (1) Our objective of near-future world state and action prediction naturally aligns with the next-token prediction mechanism inherent in LLMs. (2) The structured yet flexible architecture of LLMs provides a scalable foundation for our specific tasks. We do NOT initialize the LLM with any pretrained weights, as we consider that visual perception and prediction abilities can be learned without language priors [4].

**Supervisions for predicting world states and actions.** EgoAgent consists of two branches: a predictor branch, which predicts future world states  $S'_{t+1}$  and actions  $A'_t$ , and an asymmetric observer branch, which extracts the future world states  $S_{t+1}$  using only the egocentric image  $I_{t+1}$ . Given the input token sequence at time step  $t$ , the loss functions for future state and action prediction are defined as:

$$\mathcal{L}_{act}(t) = \mathcal{L}_1(A'_t, A_t), \quad (1)$$

$$\mathcal{L}_{pred}(t) = \mathcal{L}_{dino}(S'_{t+1}, sg[S_{t+1}]), \quad (2)$$

where  $sg[\cdot]$  denotes the stop-gradient operation,  $\mathcal{L}_1$  and  $\mathcal{L}_{dino}$  represent the L1 loss and the DINO loss [10], respectively. Following common practices in self-supervised learning [10, 26],  $sg[\cdot]$  is applied to block the gradients from back-propagating to the observer branch. The weights of the observer branch, including the MLPs, are updated at each training iteration using an EMA of the predictor’s.

**Self-supervision for learning powerful representations.** When humans learn to interact with the environment, they first develop an understanding of the observed world and objects within, which then aids them in predicting future states of the world and making appropriate responses. Inspired by this process, we introduce an additional self-supervised learning loss on EgoAgent to facilitate learning representative features from egocentric videos from scratch:

$$\mathcal{L}_{rep}(t) = \mathcal{L}_{dino}(\Theta_{pred}(I_t^{v1}), \Theta_{obs}(I_t^{v2})). \quad (3)$$

Here,  $\Theta_{pred}$  and  $\Theta_{obs}$  denote the predictor and observer networks of EgoAgent, respectively,  $I_t^{v1}$  and  $I_t^{v2}$  denote the two different views derived from the egocentric image  $I_t$ .

**Overall objective function.** Given an input interleaved token sequence containing  $t$  time steps, we have the overall objective function  $\mathcal{L}$  defined as:

$$\mathcal{L} = \frac{1}{t} \sum_{k=0}^t (\lambda_{rep} \mathcal{L}_{rep} + \lambda_{pred} \mathcal{L}_{pred} + \lambda_{act} \mathcal{L}_{act}), \quad (4)$$

where  $\lambda_{rep}$ ,  $\lambda_{pred}$ , and  $\lambda_{act}$  are the corresponding loss weights for representing the world, predicting future world states and predicting actions, respectively.

### 3.3. Task-Specific Inference

After training, EgoAgent inherently acquires the abilities to solve the three tasks designed to mirror human cognition:

**World state prediction.** Given a sequence of egocentric video frames and human motion, EgoAgent is capable of predicting the future world state  $S_{t+1}$  by  $S_{t+1} = \Theta(I_1, A_1, \dots, I_t, A_t)$ , where  $\Theta$  denotes the EgoAgent model. Once predicted, the future world image can then be retrieved by measuring similarity within the learned feature space, enabling applications that require foresight into upcoming visual states.

**Human action prediction.** From an egocentric view and informed by previous body motions, EgoAgent can generate plausible future action  $A_t$  represented as 3D human motion by  $A_t = \Theta(I_1, A_1, \dots, A_{t-1}, I_t)$ . This predictive capability of human motion is essential for applications in humanoid robotics, virtual reality, and interactive gaming, where anticipating human-like movements is crucial.

**Visual representation.** EgoAgent learns world states as feature representations, allowing it to perceive meaningful semantic features from input images by  $S_t = \Theta(I_t)$ . These features can be directly applied to representation tasks such as image classification and visuomotor policy learning.

## 4. Experiments

### 4.1. Experimental Setups

**Datasets.** We train EgoAgent on two egocentric datasets, WalkingTours (WT) [44] and Ego-Exo4D [20]. WT is a video-only dataset with approximately 1.5M high-resolution frames. Ego-Exo4D has 221.26 hours of egocentric videos along with 376K manually labeled and 9.2M automatically generated 3D body poses. To eliminate the camera projection difference between Ego-Exo4D and WT, we undistort the raw fisheye sensor data in Ego-Exo4D into a pinhole camera model. We apply a sliding window filter of 20 frames to extract continuous pose sequences, resulting in 1,410,119 clips of synchronized egocentric videos and 3D body poses. Based on the EgoPose split in Ego-Exo4D, we reserve the validation set of Ego-Exo4D-v2<sup>3</sup> for evaluation and use the remaining clips for training, yielding 1,378,672 training clips and 31,447 evaluation clips.

**Implementation Details.** To alleviate the training burden, we sample one image frame every five frames while retaining all 3D body poses. Specifically, for each clip with 20 frames, we divide it into  $T = 4$  time steps and sample the first image frame at each step. This results in every time step consisting of one image frame paired with five frames of 3D body poses. For self-supervised learning on egocentric videos, We follow [44] to avoid cropping noisy positive pairs and adopt two global crops, six local crops, and the same augmentations consistent with DINO [10].

Following the configurations in [24], we train EgoAgent with two model sizes, 300M and 1B, on the pro-

<sup>3</sup>The whole Ego-Exo4D dataset is divided into Ego-Exo4D-v1 and Ego-Exo4D-v2 on the official website (<https://ego-exo4d-data.org>).

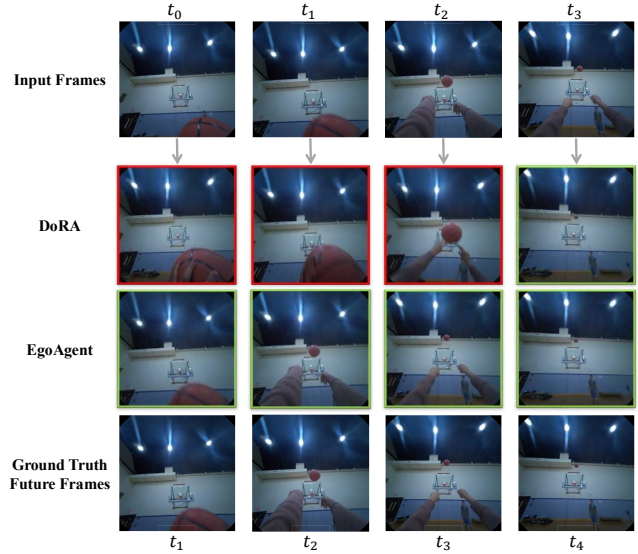


Figure 3. Retrieval results for egocentric world state prediction. Images correctly retrieved are highlighted with green borders.

cessed dataset for 72,000 iterations with a linear warm-up of 1800 iterations. We set the base learning rate to  $6 \times 10^{-4}$  with a cosine decay scheduler. The loss weights are set as  $\lambda_{rep} = 2, \lambda_{pred} = 1, \lambda_{act} = 3$ , respectively. All models are trained from scratch using FP16 to speed up. The training takes 25/60 hours in total with a batch size of 1920 on 32/48 NVIDIA A100 GPUs for the 300M/1B model, respectively.

### 4.2. World State Prediction

To evaluate the effectiveness of EgoAgent in predicting future states of the world, we perform an egocentric video prediction task with feature retrieval. Specifically, given the egocentric video frames and 3D human poses at a certain time step, we build a query set that stores the EgoAgent-predicted features of the next frame and a gallery set that contains the features directly extracted from all video frames. At time step  $t$ , if the predicted world state  $S'_{t+1}$  in the query set can correctly retrieve the corresponding  $S_{t+1}$  in the gallery set, we treat it as a successful prediction. For traditional representation models, since there are no action conditions, the output state  $S'_{t+1}$  would be exactly the state at time step  $t$ , *i.e.*,  $S_t$ . To avoid retrieving this exact same state, when querying with input image  $I_t$ , we ignore the scene state  $S_t$  in the gallery set. Note that  $S_t$  is NOT removed when testing EgoAgent. Following common practice in retrieval tasks, we use Top1 accuracy and mean average precision (mAP) as metrics.

As shown in Tab. 1, EgoAgent outperforms traditional representation models that trained with time-contrastive loss [32, 38, 44] or view-contrastive loss [10] on egocentric videos by a large margin. Specifically, EgoAgent-300M outperforms DoRA by **+12.86%** in Top1 accuracy and **+13.05%** in mAP, indicating that its predicted fea-

Table 1. Performance comparison across three tasks. (1) For world state prediction, we report the feature retrieval accuracy of the predicted next-frame features by the averaged Top1 accuracy and mAP on total timesteps  $T = 4$ . (2) For human action prediction, we report the 3D motion prediction results on Ego-Exo4D by MPJPE (cm) and MPJVE (cm/s) with time gaps of 1/30 and 1/10 seconds. (3) For visual representation, we report the  $k$ -NN evaluation results on ImageNet-100 and ImageNet-1K by Top1 and Top5 accuracy (%). We evaluate with the officially released checkpoints of representation models trained on egocentric video datasets and train all motion prediction models from scratch. \* denotes the model trained by ourselves for a fair comparison, ↓ indicates lower is better.

Method	Training Dataset	World State Prediction		3D Human Motion Prediction				Visual Representation			
		Ego-Exo4D		Ego-Exo4D (30fps)		Ego-Exo4D (10fps)		ImageNet-100		ImageNet-1K	
		Top1 Acc.	mAP	MPJPE↓	MPJVE↓	MPJPE↓	MPJVE↓	Top1	Top5	Top1	Top5
VIP [32]	Ego4D	1.67	6.10	×	×	×	×	8.04	21.14	1.59	1.59
R3M [38]	Ego4D	24.42	37.26	×	×	×	×	4.82	14.42	0.77	0.77
DINO [10]	WT <sub>Venice</sub>	28.24	43.42	×	×	×	×	40.12	63.74	22.18	35.85
DoRA [44]	WT	30.15	45.01	×	×	×	×	55.08	78.06	34.52	52.50
DoRA* [44]	WT+Ego-Exo4D	29.78	44.34	×	×	×	×	51.74	74.14	31.46	47.82
Diffusion Policy-C [13]	Ego-Exo4D	×	×	28.07	206.96	27.95	114.75	×	×	×	×
Diffusion Policy-T [13]	Ego-Exo4D	×	×	25.92	353.24	25.85	148.82	×	×	×	×
HumanMAC [11]	Ego-Exo4D	×	×	19.21	94.22	17.68	77.43	×	×	×	×
siMLPe [21]	Ego-Exo4D	×	×	13.33	81.94	12.20	60.65	×	×	×	×
EgoAgent-300M	WT+Ego-Exo4D	43.01	58.06	12.92	82.18	11.89	59.96	55.14	76.56	34.65	51.42
EgoAgent-1B	WT+Ego-Exo4D	<b>46.43</b>	<b>61.96</b>	<b>12.51</b>	<b>81.45</b>	<b>11.65</b>	<b>58.99</b>	<b>56.48</b>	<b>78.12</b>	<b>35.84</b>	<b>53.03</b>




Figure 4. Retrieval results for egocentric world state prediction from the same input observation with diverse human poses.

tures are beyond the input images. In contrast, traditional models rely solely on semantic similarity, struggling to distinguish similar frames in a video. Furthermore, scaling up the model enhances performance, as EgoAgent-1B achieves further performance gains upon EgoAgent-300M with **+3.42%** Top1 accuracy and **+3.90%** mAP. Note that when trained on our combined dataset, DoRA exhibits performance degradation in both world state prediction and visual representation tasks. We attribute this to its tracker’s inability to track abrupt object motions in Ego-Exo4D. For further analysis, please refer to Sec. 4.4. Fig. 3 visualizes the retrieval results of predicted features, where DoRA fails to retrieve the correct future states when humans have large movements, while EgoAgent successfully predicts the future features that retrieve the correct next-frame images.

To further evaluate whether EgoAgent accurately models the causal relationship between actions and future world states, we conduct a qualitative analysis by conditioning next-state predictions on various poses with the same initial observation, including dribbling, raising, shooting, and passing the ball. Fig. 4 shows the retrieved images, where all predicted features successfully retrieve images that accurately reflect the intended motion dynamics. This demon-

strates EgoAgent’s ability to understand both the semantics of actions and their consequences in the egocentric world.

### 4.3. 3D Human Motion Prediction

We compare EgoAgent with the state-of-the-art video-based motion generation model, Diffusion Policy [13] with convolution (C) and transformer (T) architectures, and unconditional motion prediction models [11, 21] using the EgoPose from Ego-Exo4D [20]. Each testing clip contains 5 frames of egocentric images and 20 frames of 3D human motion. We set the prediction target as the last 15 frames of 3D motion. The images and first 5 frames of 3D motion are adopted as input for Diffusion Policy and EgoAgent, while only the first 5 poses are used as input for unconditional motion prediction models. Following Ego-Exo4D [20], the mean per-joint position error (MPJPE) in centimeters (cm) and the mean per-joint velocity error (MPJVE) in centimeters per second (cm/s) are adopted as evaluation criteria.

Quantitatively, as shown in Tab. 1, EgoAgent achieves state-of-the-art performance with the lowest MPJPE and competitive MPJVE. Specifically, EgoAgent-300M improves siMLPe [21] by **-0.41** and **-0.31** MPJPE on the evaluation set of 30 fps and 10 fps, respectively. When scaling up to 1B, the prediction errors are further reduced. These results demonstrate that EgoAgent can predict future actions accurately given egocentric observations and past actions.

We also evaluate the motion prediction results qualitatively. As illustrated in Fig. 5, EgoAgent produces relatively small errors compared to the ground truth. Compared to Diffusion Policy [13], EgoAgent achieves higher accuracy on non-visible joints. Specifically, in the early frames ( $t_1$  and  $t_2$ ), the missing right hand in the image leads to large prediction errors in Diffusion Policy, whereas EgoAgent accurately estimates hand positions by leveraging prior

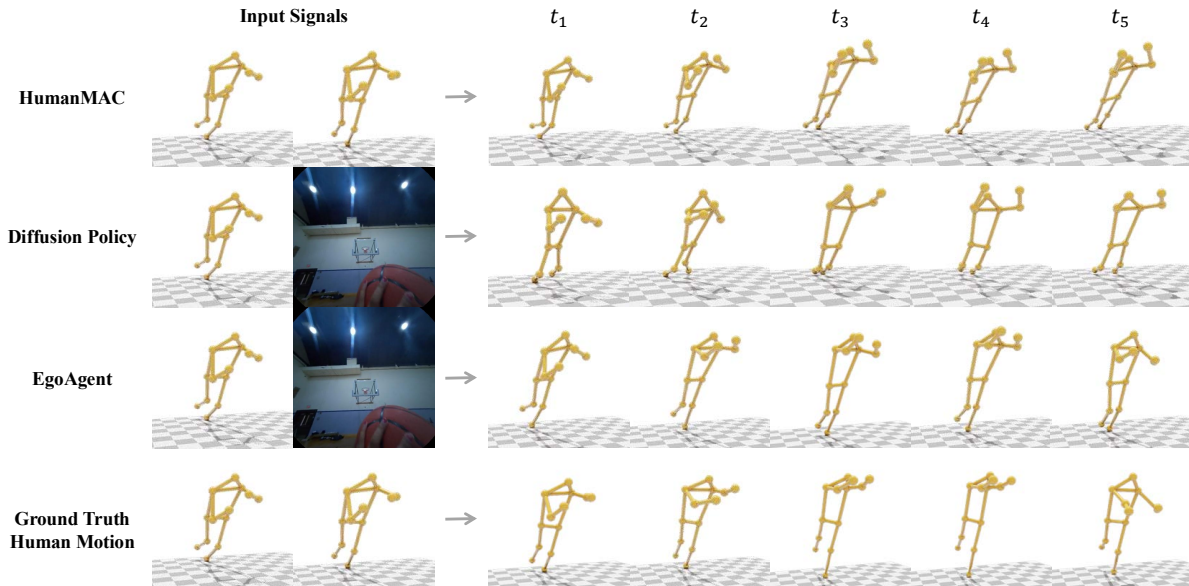


Figure 5. Visualizations of the 3D human motion prediction task. After observing the egocentric frames, EgoAgent can generate accurate human skeletons even though most of the body joints are not visible in the input frame.

motion cues. Compared to HumanMAC [11], EgoAgent shows smaller accelerated errors. The predicted skeletons in HumanMAC gradually lean toward the floor ( $t_{[3:5]}$ ) after several frames, whereas EgoAgent integrates visible body part information from egocentric observations to correct accelerated errors, leading to more stable motion predictions.

#### 4.4. Visual Representation

We examine EgoAgent’s visual perception ability by two aspects: (1) image classification accuracy by  $k$ -NN on ImageNet-100 and ImageNet-1K [15], which is the widely-used golden standard for evaluating visual representations, and (2) egocentric manipulation success rate in an embodied simulator called TriFinger [48], which measures the transferability of learned visual representations to visual-motor control. For both tasks, the models are frozen and used to extract visual features. For image classification, we utilize a  $k$ -nearest neighbor classifier with  $k = 20$ .

As shown in Tab. 1, EgoAgent achieves the best representation performance. Specifically, EgoAgent-1B outperforms the officially released DoRA by **+1.40%** and **+1.32%** Top1 accuracy on ImageNet-100 and ImageNet-1K, respectively, demonstrating that learning to predict world states and actions can facilitate the task of representing the world. For a fair comparison, we train DoRA on the same combined dataset (WT and Ego-Exo4D), which instead results in worse performance than training on WT alone. We suggest this is because, unlike WT videos, which feature stable motion, Ego-Exo4D includes fewer objects but more dynamic movements (objects frequently move out of sight). As a result, the tracker in DoRA struggles to maintain consistent object tracking, leading to degraded performance. In

Table 2. Success rate (%) on the TriFinger [48] benchmark, where each model’s pretrained visual representation is fixed, and additional linear layers are trained as the policy network.

Methods	Training Dataset	Reach Cube	Move Cube
DINO [10]	WT <sub>Venice</sub>	78.03	47.42
DoRA [44]	WT	82.40	48.13
DoRA [44]	WT+Ego-Exo4D	81.55	45.41
EgoAgent-300M	WT+Ego-Exo4D	82.61	54.21
EgoAgent-1B	WT+Ego-Exo4D	<b>85.72</b>	<b>57.66</b>

contrast, EgoAgent learns robust representation by modeling the temporal causality of state features and actions.

The TriFinger benchmark [48] involves a three-finger robot performing two tasks (reach cube and move cube) from an egocentric view. We freeze the pretrained models and use a 3-layer MLP as the policy network, training each task with 100 demonstrations following [34]. As shown in Tab. 2, EgoAgent achieves the highest success rates on both tasks, outperforming the egocentric representation models DoRA [44] with increases of **+3.32%** and **+3.9%** respectively. This result shows that by incorporating human action prediction into the learning process, EgoAgent demonstrates the ability to learn more effective representations that benefit both image classification and embodied manipulation tasks. This highlights the potential of leveraging human-centric motion data to bridge the gap between visual understanding and actionable policy learning.

#### 4.5. Ablation Study

To evaluate the effectiveness of EgoAgent’s specialized designs, we conduct several ablation studies using EgoAgent-300M with a short learning schedule of 14,400 iterations.

Table 3. Ablation studies: removing one of the three tasks during training (b-d), training on a single task (e-g), and using pretrained image latents from VQGAN’s reconstruction-based latent space (h-i). All variants are trained with 14,400 iterations.

Method	Training Task			World State Prediction		Human Motion Prediction		Visual Representation	
	Pred.	Act.	Rep.	Ego-Exo4D		Ego-Exo4D (@30fps)		ImgNet-100	ImgNet-1K
				Top1 Acc.	mAP	MPJPE↓	MPJVE↓	Top1 Acc.	Top1 Acc.
(a) Baseline	✓	✓	✓	<b>37.77</b>	<b>53.72</b>	14.49	88.61	<b>41.64</b>	<b>22.28</b>
(b) w/o $\mathcal{L}_{pred}$		✓	✓	✗	✗	14.70	89.62	39.12	20.97
(c) w/o $\mathcal{L}_{act}$	✓		✓	34.86	49.13	✗	✗	39.92	21.31
(d) w/o $\mathcal{L}_{rep}$	✓	✓		25.90	40.17	14.49	88.82	✗	✗
(e) only $\mathcal{L}_{pred}$	✓			33.23	47.28	✗	✗	✗	✗
(f) only $\mathcal{L}_{act}$		✓		✗	✗	14.32	91.13	✗	✗
(g) only $\mathcal{L}_{rep}$			✓	✗	✗	✗	✗	40.80	21.56
(h) Pixel-level latent	✓	✓		20.62	37.12	<b>13.57</b>	<b>84.86</b>	1.00	0.10
(i) Pixel-level latent	✓	✓	✓	15.63	23.85	16.25	100.49	31.20	13.22

### Joint learning of representation, prediction, and action.

EgoAgent demonstrates promising performance across the three tasks. To understand how these tasks interact during training, we conduct *leave-one-out* experiments and *one-task-only* experiments, as shown in Tab. 3 (a-g). We observe several key findings: (1) Removing any one task degrades the performance of the other two, highlighting the complementary roles of these tasks in learning. Specifically, removing the representation loss causes the largest drop in state prediction (-11.87% in Top1 accuracy and -13.55% in mAP), while removing the prediction loss significantly weakens representation learning, reducing Top1 accuracy by -2.52% on ImageNet-100 and -1.31% on ImageNet-1K. This suggests that state prediction and visual representation are closely interconnected. For motion prediction, excluding the prediction loss leads to the most significant decline, increasing MPJVE by 1.01. (2) Training on individual tasks yields worse performances on each task than learning all three tasks together, except for MPJPE in motion prediction. We suspect that when trained only on motion prediction, the model focuses excessively on minimizing joint position errors. However, without context from the other tasks, it struggles to maintain smooth, consistent motion across frames, resulting in a higher MPJVE of 2.52.

**Analysis of task dependencies.** We also identify several interesting dependencies among the three tasks: (1) *Representation* serves as the foundation for both future state and action prediction as it captures the semantic information of the current world. (2) *Prediction and action* jointly provide effective supervision for representation learning by modeling world dynamics and predicting actions, improving representation by 1.16% in Top1 accuracy. However, training on prediction or action alone cannot enhance representation, as actions without accurate next-state predictions or next-state predictions conditioned on incorrect actions may provide noisy supervision. (3) *Representation and action* guide the model to integrate learned representations with action dynamics, leading to more accurate predictions of fu-

ture states. Meanwhile, *representation and prediction* work together to help the model understand how actions influence state transitions, refining its ability to generate actions.

**Training in semantic feature space.** To evaluate the effectiveness of training agent models with a high-level semantic feature space, we use a pretrained VQGAN [17], which is designed to reconstruct pixel-level details of images, as the image tokenizer. The results are shown in Tab. 3 (a,h,i). Training EgoAgent in VQGAN’s pretrained pixel-level latent space results in significant performance drops in world state prediction and nearly nullifies the ability for visual representation. While motion prediction improves, likely due to the unified task formulation, the overall performance degradation makes this latent space unsuitable. Forcing representation learning with this pixel-level latent space not only yields modest representation performance but also significantly weakens the other two tasks compared to learning in continuous semantic features. This suggests that high-level semantics are more effective than low-level details for understanding world states and distinguishing different contents in adjacent time steps.

## 5. Conclusion and Future Work

In this paper, we presented EgoAgent, a joint predictive agent model that simultaneously learns to represent world observations, predict future world states, and generate 3D human motions in egocentric worlds. EgoAgent integrates these three abilities within a joint embedding-action-prediction architecture with temporally asymmetric predictor-observer branches. Through extensive experiments, we demonstrated that EgoAgent outperforms existing state-of-the-art methods, with the joint learning of perception, prediction, and action tasks mutually reinforcing each other, leading to improved performance across all three capabilities. Future work will explore incorporating finer-grained hand representations for precise tasks such as object manipulation and integrating long-term memory mechanisms for tasks requiring extended temporal dependencies.

**Acknowledgments:** This work was partially supported by the Zhejiang Provincial Natural Science Foundation of China (No. LR25F020003), NSFC (No. 62402427, No. U24B20154), Zhejiang University Education Foundation Qizhen Scholar Foundation, and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University. We also acknowledge the support from the JC STEM Lab of AI for Science and Engineering, funded by the Hong Kong Jockey Club Charities Trust, and the Research Grants Council of Hong Kong (No. CUHK14213224).

## References

- [1] Emre Aksan, Manuel Kaufmann, Peng Cao, and Otmar Hilliges. A spatio-temporal transformer for 3D human motion prediction. In *IEEE International Conference on 3D Vision*, pages 565–574, 2021. 3
- [2] Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos J Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. *Advances in Neural Information Processing Systems*, 37: 58757–58791, 2024. 2
- [3] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023. 2
- [4] Yutong Bai, Xinyang Geng, Karttikeya Mangalam, Amir Bar, Alan L Yuille, Trevor Darrell, Jitendra Malik, and Alexei A Efros. Sequential modeling enables scalable learning for large vision models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22861–22872, 2024. 4
- [5] Chris Bamford and Simon M Lucas. Neural game engine: Accurate learning of generalizable forward models from pixels. In *IEEE Conference on Games*, pages 81–88, 2020. 2
- [6] Adrien Bardes, Jean Ponce, and Yann LeCun. MC-JEPA: A joint-embedding predictive architecture for self-supervised learning of motion and content features. *arXiv preprint arXiv:2307.12698*, 2023. 2
- [7] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mido Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *Transactions on Machine Learning Research*, 2024. 2
- [8] Lawrence W Barsalou. Grounded cognition. *Annual Review of Psychology*, 59(1):617–645, 2008. 1, 2
- [9] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *European Conference on Computer Vision*, pages 387–404, 2020. 2, 3
- [10] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 2, 3, 4, 5, 6, 7
- [11] Ling-Hao Chen, Jiawei Zhang, Yewen Li, Yiren Pang, Xiaobo Xia, and Tongliang Liu. Humanmac: Masked motion completion for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9544–9555, 2023. 6, 7
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607, 2020. 3
- [13] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 2023. 6
- [14] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3D human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6519–6527, 2020. 3
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. 2, 7
- [16] 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. *International Conference on Learning Representations*, 2021. 3
- [17] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. 4, 8
- [18] Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and leveraging world models in visual representation learning. *arXiv preprint arXiv:2403.00504*, 2024. 2, 4
- [19] 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*, pages 18995–19012, 2022. 3
- [20] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19383–19400, 2024. 2, 5, 6
- [21] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to MLP: A simple baseline for human motion prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4809–4819, 2023. 6
- [22] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems*, 31, 2018. 2

- [23] David Ha and Jürgen Schmidhuber. World models. In *Thirty-second Annual Conference on Neural Information Processing Systems*, 2018. 2
- [24] Zhiwei Hao, Jianyuan Guo, Chengcheng Wang, Yehui Tang, Han Wu, Han Hu, Kai Han, and Chang Xu. Data-efficient large vision models through sequential autoregression. In *International Conference on Machine Learning*, pages 17572–17596, 2024. 5
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 3
- [26] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 3, 4
- [27] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 2
- [28] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. GAIA-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023. 2
- [29] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. 3
- [30] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022. 2, 4
- [31] Zhenguang Liu, Kedi Lyu, Shuang Wu, Haipeng Chen, Yanbin Hao, and Shouling Ji. Aggregated multi-gans for controlled 3D human motion prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2225–2232, 2021. 3
- [32] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. VIP: Towards universal visual reward and representation via value-implicit pre-training. *The Eleventh International Conference on Learning Representations*, 2023. 3, 5, 6
- [33] Takahiro Maeda and Norimichi Ukita. Motionaug: Augmentation with physical correction for human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6427–6436, 2022. 3
- [34] Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Tingfan Wu, Jay Vakil, et al. Where are we in the search for an artificial visual cortex for embodied intelligence? *Advances in Neural Information Processing Systems*, 36:655–677, 2023. 7
- [35] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Weakly-supervised action transition learning for stochastic human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8151–8160, 2022. 3
- [36] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017. 2
- [37] Russell Mendonca, Shikhar Bahl, and Deepak Pathak. Structured world models from human videos. *arXiv preprint arXiv:2308.10901*, 2023. 2
- [38] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3M: A universal visual representation for robot manipulation. *Conference on Robot Learning*, pages 892–909, 2023. 3, 5, 6
- [39] Wolfgang Prinz. Perception and action planning. *European Journal of Cognitive Psychology*, 9(2):129–154, 1997. 2
- [40] Lawrence Shapiro and Shannon Spaulding. Embodied Cognition. *The Stanford Encyclopedia of Philosophy*, 2024. 1
- [41] InternLM Team. InternLM: A multilingual language model with progressively enhanced capabilities, 2023. 4
- [42] Esther Thelen and Linda B Smith. *A dynamic systems approach to the development of cognition and action*. MIT press, 1994. 1
- [43] Esther Thelen, Gregor Schöner, Christian Scheier, and Linda B Smith. The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24(1):1–34, 2001. 1
- [44] Shashanka Venkataramanan, Mamshad Nayeem Rizve, João Carreira, Yuki M Asano, and Yannis Avrithis. Is ImageNet worth 1 video? Learning strong image encoders from 1 long unlabelled video. *The Twelfth International Conference on Learning Representations*, 2024. 2, 3, 5, 6, 7
- [45] Hongsong Wang, Jian Dong, Bin Cheng, and Jiashi Feng. PVRED: A position-velocity recurrent encoder-decoder for human motion prediction. *IEEE Transactions on Image Processing*, 30:6096–6106, 2021. 3
- [46] Yizhou Wang, Shixiang Tang, Feng Zhu, Lei Bai, Rui Zhao, Donglian Qi, and Wanli Ouyang. Revisiting the transferability of supervised pretraining: an mlp perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9183–9193, 2022. 2, 3
- [47] Yuqi Wang, Jiawei He, Lue Fan, Hongxin Li, Yuntao Chen, and Zhaoxiang Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024. 2
- [48] Manuel Wüthrich, Felix Widmaier, Felix Grimminger, Joel Akpo, Shruti Joshi, Vaibhav Agrawal, Bilal Hammoud, Majid Khadiv, Miroslav Bogdanovic, Vincent Berenz, et al. TriFinger: An open-source robot for learning dexterity. *Proceedings of Machine Learning Research*, 155:1871–1882, 2020. 2, 7
- [49] Jiannan Xiang, Guangyi Liu, Yi Gu, Qiyue Gao, Yuting Ning, Yuheng Zha, Zeyu Feng, Tianhua Tao, Shibo Hao, Yemin Shi, et al. Pandora: Towards general world model with natural language actions and video states. *arXiv preprint arXiv:2406.09455*, 2024. 2

- [50] Haitao Yan, Qiongjie Cui, Jiexin Xie, and Shijie Guo. Forecasting of 3D whole-body human poses with grasping objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1726–1736, 2024. [3](#)
- [51] Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter Abbeel. UniSim: Learning interactive real-world simulators. *The Twelfth International Conference on Learning Representations*, 2024. [2, 4](#)
- [52] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10082–10092, 2019. [3](#)
- [53] Yufei Zhang, Jeffrey O Kephart, and Qiang Ji. Incorporating physics principles for precise human motion prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6164–6174, 2024. [3](#)
- [54] Yang Zheng, Yanchao Yang, Kaichun Mo, Jiaman Li, Tao Yu, Yebin Liu, C Karen Liu, and Leonidas J Guibas. GIMO: Gaze-informed human motion prediction in context. In *European Conference on Computer Vision*, pages 676–694, 2022. [3](#)