

Intention-Guided Cognitive Reasoning for Egocentric Long-Term Action Anticipation

Qiaohui Chu^{1,2}, Haoyu Zhang^{1,2}, Meng Liu^{3*}, Yisen Feng¹, Haoxiang Shi^{1,2}, Liqiang Nie^{1*}

¹ Harbin Institute of Technology (Shenzhen)

² Pengcheng Laboratory

³ Shandong Jianzhu University

{qiaohuichu8599, zhang.hy.2019, mengliu.sdu, yisenfeng.hit, shihaoxiang1999, nieliqiang}@gmail.com

Abstract

Long-term action anticipation from egocentric video is critical for applications such as human-computer interaction and assistive technologies, where anticipating user intent enables proactive and context-aware AI assistance. However, existing approaches suffer from three key limitations: 1) underutilization of fine-grained visual cues from hand-object interactions, 2) neglect of semantic dependencies between verbs and nouns, and 3) lack of explicit cognitive reasoning, limiting generalization and long-term forecasting ability. To overcome these challenges, we propose INSIGHT, a unified two-stage framework for egocentric action anticipation. In the first stage, INSIGHT focuses on extracting semantically rich features from hand-object interaction regions and enhances action representations using a verb-noun co-occurrence matrix. In the second stage, it introduces a reinforcement learning-based module that simulates explicit cognitive reasoning through a structured process: *visual perception (think)* → *intention inference (reason)* → *action anticipation (answer)*. Extensive experiments on Ego4D, EPIC-Kitchens-55, and EGTEA Gaze+ benchmarks show that INSIGHT achieves state-of-the-art performance, demonstrating its effectiveness and strong generalization capability.

Code — <https://github.com/CorrineQiu/INSIGHT>

Extended version — <https://arxiv.org/abs/2508.01742>

Introduction

In real-world applications such as human-computer interaction (Azam and Desai 2024; Plizzari et al. 2024), augmented reality (Abreu et al. 2024; Xu et al. 2024), and assistive systems for visually impaired individuals (Lee et al. 2024; Xiao et al. 2025), AI agents must accurately interpret user intent and demonstrate effective long-term planning capabilities within egocentric vision scenarios. The ability to anticipate user actions well in advance allows AI systems to proactively adapt their behaviors, offer timely and contextually appropriate assistance, and avoid potential hazards or inefficiencies. Conversely, insufficient long-term anticipation can result in delayed or inaccurate responses, significantly impacting system reliability and user satisfaction.

*Corresponding authors.

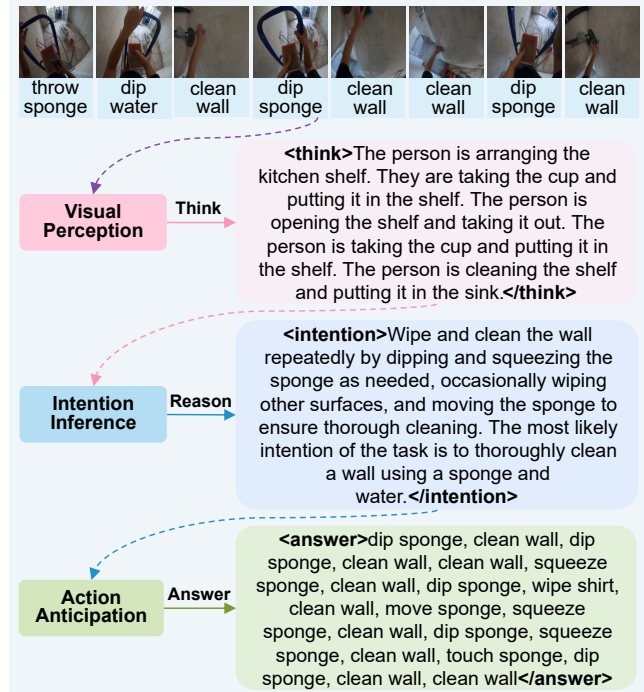


Figure 1: Illustration of the explicit cognitive reasoning process for long-term action anticipation.

To address these critical needs, the **Long-Term Action Anticipation (LTA)** task has been introduced. LTA focuses on predicting sequences of future actions likely to be executed by individuals wearing first-person cameras, facilitating more proactive and context-aware assistance.

Existing approaches to LTA can be broadly categorized into three categories: 1) *Lightweight Temporal Models* (Mascaro, Ahn, and Lee 2022; Ashutosh et al. 2023; Zhong et al. 2024). While computationally efficient, they suffer from limited model capacity and inadequate integration of prior knowledge, resulting in poor generalization. 2) *Large Language Models (LLMs)-based Methods* (Chen et al. 2023; Huang et al. 2023; Mittal et al. 2024). These methods extract visual content and generate textual descriptions processed by action recognition models or, optionally, Vision-

Language Models (VLMs). However, their purely text-based pipeline lacks direct visual grounding, limiting fine-grained scene comprehension and predictive accuracy, and is further constrained by the quality of large-scale pretraining data (Wang et al. 2025). 3) *Action-Embedded Visual Representation Methods* (Cao et al. 2025). These approaches incorporate action-related visual embeddings but neglect important egocentric cues and the co-occurrence of verb-noun pairs, both of which are critical for context-aware anticipation. Additionally, existing methods predominantly treat LTA as passive sequence prediction, lacking active cognitive reasoning. This limits their robustness and adaptability in dynamic environments.

To address limitations observed in prior research, we propose a unified two-stage framework named INSIGHT, which integrates **Hand-Object Semantic Action Recognition with Explicit Cognitive Reasoning for Anticipation**. In the first stage, we introduce a specialized feature extraction approach targeting Hand-Object Interaction (HOI) regions to capture crucial egocentric cues effectively. We further enhance contextual understanding by constructing a semantic co-occurrence transition matrix that explicitly models relationships between verbs and nouns. The resulting enriched visual features are then translated into textual prompts and integrated into the VLM to facilitate robust cross-modal representation. In the second stage, we present an innovative reinforcement learning-based intention inference module designed for explicit cognitive reasoning. This module simulates structured, human-like thought processes, i.e., **visual perception (think) → intention inference (reason) → action anticipation (answer)**, as illustrated in Figure 1. Unlike passive methods, INSIGHT proactively adapts its predictions by leveraging both observed contexts and inferred user intentions. Extensive evaluations on Ego4D, EPIC-Kitchens-55, and EGTEA Gaze+ demonstrate that INSIGHT significantly surpasses existing state-of-the-art methods.

The main contributions of our work are as follows: **1)** We propose a hand-object semantic action recognition module that captures fine-grained HOI cues and incorporates semantic co-occurrence correction, resulting in more robust visual representations. **2)** We introduce a reinforcement learning-based intention reasoning mechanism, enabling explicit, human-like cognitive reasoning for LTA. **3)** Extensive experiments across multiple benchmarks demonstrate that INSIGHT consistently outperforms state-of-the-art approaches, validating its effectiveness and generalizability.

Related Work

Long-Term Action Anticipation. Grauman et al. (Grauman et al. 2022) introduced a foundational end-to-end baseline for long-term temporal prediction, employing SlowFast as the visual encoder coupled with a Transformer for sequence modeling. This approach has inspired numerous subsequent studies, which generally fall into three categories based on modality handling and learning methodologies.

The first category emphasizes enhancing visual modeling through improved feature extraction, aggregation strategies, and advanced spatiotemporal architectures, with re-

lated ideas also explored in other domains (Zhang et al. 2023). For instance, lightweight models (Hussein, Gavves, and Smeulders 2019; Das and Ryoo 2022; Nawhal, Jyothi, and Mori 2022; Ishibashi et al. 2023) have been proposed to boost efficiency and representation quality. Additionally, Zhang et al. (Zhang et al. 2024a) introduced object-centric feature extraction using natural language queries, combined with dual-Transformer modules for encoding and temporal aggregation. Despite these innovations, such methods typically suffer from limited capacity and an absence of external priors, resulting in poor generalization in complex, dynamic, and cluttered egocentric scenarios.

The second category leverages LLMs for temporal reasoning based on textual inputs derived from action recognition models. For instance, Zhao et al. (Zhao et al. 2023), Pei et al. (Pei et al. 2024), and Chu et al. (Chu et al. 2025) fine-tuned LLMs such as LLaMA2-7B (Touvron et al. 2023) and Vicuna-7B (Zheng et al. 2023), using predicted action sequences as input. Kim et al. (Kim et al. 2024) further enhanced this setup by incorporating VLM-generated captions to enrich contextual understanding. Nonetheless, these approaches rely exclusively on text-based reasoning, neglect direct visual grounding (Liu et al. 2018a,b; Wu et al. 2025), and consequently compromising fine-grained scene comprehension, reducing accuracy in extended future predictions.

The third category aims to integrate visual semantics with LLM reasoning to enhance predictive accuracy. For example, Cao et al. (Cao et al. 2025) used VLMs to infer task intentions, subsequently employing these inferred intentions as attention queries within visual backbones for improved intention-aware feature aggregation. Although this hybrid method offers better semantic alignment, it overlooks critical egocentric details such as HOI regions and fails to capture the statistical co-occurrences between verbs and nouns, both essential elements for precise anticipation.

Large Language Model Reasoning. Existing LTA methods leveraging fine-tuned LLMs typically adopt supervised fine-tuning strategies (Lester, Al-Rfou, and Constant 2021; Liu et al. 2021; Zhang et al. 2025), relying on static priors—such as VLM-derived textual intentions or captions (Zhao et al. 2023; Kim et al. 2024). However, these approaches lack dynamic intention inference capabilities, rendering them fragile in complex, extended temporal contexts (Zhang et al. 2021). While Generalized Reinforcement Preference Optimization (GRPO) (Shao et al. 2024; Guo et al. 2025) offers a reward-driven alternative, its generic design fails to incorporate task-specific structure, reasoning constraints, and fine-grained success metrics, thereby limiting its effectiveness for LTA. To overcome this limitation, we propose adapting GRPO’s reward function to explicitly address the unique challenges and demands inherent to long-term action anticipation.

Method

An overview of INSIGHT is illustrated in Figure 2, comprising two key stages: hand-object semantic action recognition and explicit cognitive reasoning for anticipation.

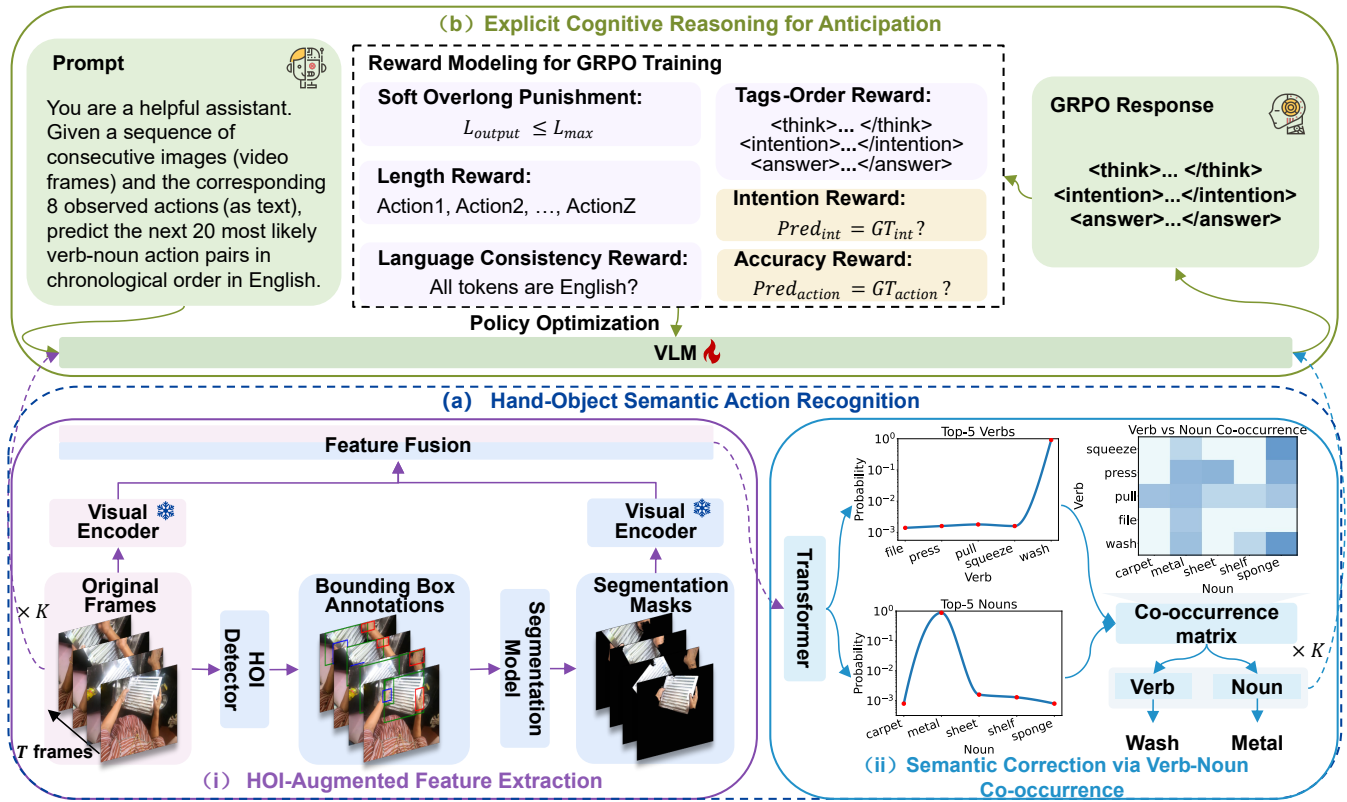


Figure 2: Overview of our two-stage framework, **INSIGHT**. **Stage 1: Hand-Object Semantic Action Recognition** leverages a HOI-augmented feature extraction module to focus on critical hand-object regions, alongside a semantic co-occurrence transition matrix that captures verb-noun relationships. The resulting enriched visual features are formatted as prompts for VLM. **Stage 2: Explicit Cognitive Reasoning for Anticipation** introduces a reinforcement learning-based intention inference mechanism that simulates a structured three-step cognitive process, i.e., visual perception (*think*), intention inference (*reason*), and action anticipation (*answer*), to enable dynamic reasoning and generate accurate long-term predictions.

Problem Definition

The goal of LTA is to predict future action sequences based on observed egocentric video segments. Given an observed sequence $\mathcal{V}_{obs} = \{S_k \mid k = 1, 2, \dots, K\}$, where S_k represents the k -th video segment, the task is to predict the subsequent Z actions, each composed of a verb-noun pair:

$$\{(\hat{v}_k, \hat{n}_k)\}_{k=K+1}^{K+Z} = f_\theta(\mathcal{V}_{obs}), \quad (1)$$

where (\hat{v}_k, \hat{n}_k) denotes the predicted verb-noun pair, and f_θ denotes our predictive model parameterized by θ .

Hand-Object Semantic Action Recognition

Action recognition (AR) plays a pivotal role in egocentric LTA, as it interprets ongoing actions within observed video sequences and provides critical prompts that guide future action predictions. However, most existing AR models in LTA adopt general-purpose visual encoders coupled with Transformer-based architectures, which often fail to capture the unique demands of egocentric perception. While some recent approaches incorporate co-occurrence matrixes to enforce semantic consistency (Zhong et al. 2024), they

still overlook a key aspect: the fine-grained visual information embedded in HOI regions. These regions are densely populated with action-relevant cues and are essential for distinguishing subtle, context-dependent behaviors, as evidenced by recent findings (Zhang, Gupta, and Zisserman 2023; Xu et al. 2023; Ohkawa et al. 2023; Zhang et al. 2024b). To overcome these limitations, we propose a novel action recognition framework tailored for egocentric LTA that explicitly focuses on capturing discriminative HOI features and incorporates verb-noun co-occurrence statistics as semantic priors. This dual-focus design enhances both the visual grounding and semantic coherence of action predictions, laying a stronger foundation for accurate and reliable long-term anticipation.

HOI-Augmented Feature Extraction. Traditional LTA approaches apply visual encoders (e.g., CLIP (Radford et al. 2021), EgoVLP (Lin et al. 2022), EgoVideo-V (Pei et al. 2024)) directly to entire frames. To improve the capture of relevant visual details in egocentric video, we introduce an HOI-focused feature extraction strategy. Specifically, given a video segment $S_k \in \mathcal{V}_{obs}$, we follow AntGPT (Zhao et al. 2023) and uniformly sample 4 frames per segment, denoted

as $F_{k,T}$ with $T \in \{1, \dots, 4\}$. Each frame $F_{k,T}$ undergoes HOI region detection using the pretrained 100DOH detector (Shan et al. 2020), refined with high-resolution masks from SAM2 (Ravi et al. 2024), producing precise masks $R_{k,T}$ representing active HOI regions.

We utilize a dual-stream EgoVideo-V architecture to encode both the full frame $F_{k,T}$ and the corresponding HOI region $R_{k,T}$, generating visual embeddings:

$$(\mathbf{X}_{k,T}^{ori}, \mathbf{X}_{k,T}^{mask}) = \text{EgoVideo-V}(F_{k,T}, R_{k,T}). \quad (2)$$

These embeddings are concatenated and fused through a linear Multi-Layer Perceptron (MLP), blending global scene context and detailed HOI features. The resulting integrated representation is processed by a Transformer-based recognition module, effectively capturing complex spatiotemporal relationships and significantly improving semantic accuracy for verb-noun predictions.

Semantic Correction via Verb-Noun Co-Occurrence.

To further enhance semantic consistency in action prediction, we implement a semantic correction module utilizing verb-noun co-occurrence statistics derived from training data. This module operates on the output of a Transformer-based model equipped with two parallel classifiers, one for verbs and one for nouns. However, independently predicted verb-noun pairs can result in improbable combinations (e.g., “drink + guitar”), diminishing prediction reliability. To mitigate this issue, we construct a verb-noun co-occurrence matrix $\mathbf{C} \in \mathbb{N}^{|\mathcal{V}| \times |\mathcal{N}|}$ from annotated labels:

$$\mathbf{C}_{v,n} = \sum_{k=1}^K \mathbf{1}_{\{v_k=v \wedge n_k=n\}}, \quad (3)$$

where v and n index the verb and noun categories in \mathcal{V} and \mathcal{N} , respectively. $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. The matrix \mathbf{C} is normalized row-wise and column-wise to obtain conditional probabilities:

$$\mathbf{P}_{v,n} = \begin{cases} \mathbf{P}_{v,n}^{(n|v)} = \frac{\mathbf{C}_{v,n}}{\sum_{n'} \mathbf{C}_{v,n'}} \in [0, 1]^{|\mathcal{V}| \times |\mathcal{N}|}, \\ \mathbf{P}_{v,n}^{(v|n)} = \frac{\mathbf{C}_{v,n}}{\sum_{v'} \mathbf{C}_{v',n}} \in [0, 1]^{|\mathcal{V}| \times |\mathcal{N}|}. \end{cases} \quad (4)$$

Given softmax probabilities $p(v_k)$ and $p(n_k)$ from classifiers, we calculate the corrected joint probability:

$$\tilde{p}(v_k, n_k) = p(v_k) \cdot p(n_k) \cdot \frac{1}{2}(\mathbf{P}_{v,n}^{(n|v)} + \mathbf{P}_{v,n}^{(v|n)}). \quad (5)$$

The final predicted action is determined through maximum a posteriori estimation:

$$(\hat{v}_k, \hat{n}_k) = \arg \max_{(v,n) \in \mathcal{V} \times \mathcal{N}} \tilde{p}(v_k, n_k). \quad (6)$$

This semantic correction significantly enhances the coherence and accuracy of predicted verb-noun pairs by aligning predictions with plausible semantic combinations.

Explicit Cognitive Reasoning for Anticipation

Most existing LTA methods adopt a pipeline wherein observed egocentric videos are first translated into textual descriptions using action recognition models, followed by fine-tuning LLMs to predict future actions leveraging semantic reasoning. Although effective, these methods employ

implicit reasoning and lack explicit modeling of decision-making processes, limiting generalization in diverse or extended scenarios.

To overcome these limitations, we extend GRPO, a reward-driven framework derived from PPO (Schulman et al. 2017), by introducing task-specific structured reasoning, namely “think \rightarrow reason \rightarrow answer”. This encourages active intention inference and adaptive decision-making for robust long-term anticipation. Our reward design integrates structural (format) and semantic (content) considerations to ensure precise, interpretable predictions.

Format Rewards. We define format-oriented reward functions to enforce structured, interpretable outputs:

1) Length Reward: A full reward is awarded if the generated verb-noun pair set \mathcal{L} matches or exceeds the target length Z :

$$S_{len} = \begin{cases} 1, & |\mathcal{L}| \geq Z, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

2) Tags-Order Reward: To enforce explicit reasoning stages, outputs must adhere to the structured format `<think>...</think>` for visual perception, `<intention>...</intention>` for intention inference, and `<answer>...</answer>` for final action prediction:

$$S_{fmt} = \begin{cases} 1, & \text{all tags are correct,} \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

3) Language Consistency Reward: Penalizes outputs containing non-English tokens to maintain linguistic consistency:

$$S_{lang} = \begin{cases} 1, & \text{all tokens are English,} \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

4) Soft Overlong Punishment: A soft penalty R_{Soft} discourages overly verbose outputs. If the generated length L falls within $[L_{max} - L_{cache}, L_{max}]$, the penalty decays linearly from 0 to -1 . If L exceeds L_{max} , the penalty is fixed at -1 . Here L_{max} denotes the maximum allowable length and L_{cache} is the buffer length.

Content Rewards. Content-oriented rewards are crafted to enhance action prediction accuracy and alignment of inferred intentions:

1) Accuracy Reward: We adopt edit distance (ED) as per Ego4D protocols (Grauman et al. 2022) to evaluate prediction accuracy. Given a predicted sequence s_{pred} and a ground truth sequence s_{true} , we normalize the prediction to the target length Z as follows: If $s_{pred} > Z$, truncate the sequence; if $s_{pred} < Z$, pad with a special token `<pad>`. The resulting normalized prediction sequence is defined as $\tilde{s}_{pred} = (\tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_Z)$. The edit distance between the normalized prediction and the ground truth sequence is computed as:

$$d_{ED}^Z = \text{ED}(\tilde{s}_{pred}, s_{true}), \quad (10)$$

where $\text{ED}(\cdot, \cdot)$ denotes the edit distance function. The final accuracy reward is normalized to the range $[0, 1]$ as follows:

$$S_{acc} = 1 - \frac{d_{ED}^Z}{|s_{true}|}, \quad S_{acc} \in [0, 1]. \quad (11)$$

Method	Venue	LLM	Visual Encoder	Ego4D-v2 (ED)		
				Verb ↓	Noun ↓	Action ↓
HAI-PUI	CVPR'24	–	–	0.7721	0.6733	0.9242
AntGPT	ICLR'23	LLaMA2-7B	EgoVLP	0.6728	0.6755	0.8931
PALM	ECCV'24	LLaMA2-7B	EgoVLP	0.7111	0.6465	0.8819
PaMsEgoAI	CVPR'23	–	SlowFast, CLIP	0.6702	0.6291	0.8753
EgoVideo	CVPR'24	Vicuna-7B	EgoVideo-V [†]	0.6576	0.6264	0.8619
ICVL	arXiv'25	LLaMA3-8B	ViT-L/14, BLIP2-OPT-2.7B	<u>0.6516</u>	<u>0.6194</u>	<u>0.8570</u>
INSIGHT	–	Qwen2.5-VL-7B	EgoVideo-V	0.6643	0.6092	0.8463

Table 1: Performance comparison between INSIGHT and prior state-of-the-art methods on the Ego4D-v2 validation set. † indicates models with visual encoders finetuned on the Ego4D dataset. Underlined values represent the current state-of-the-art, and bold values indicate the best (lowest) ED scores.

Following the EGO-TOPO (Nagarajan et al. 2020) setup on the EPIC-Kitchens-55 (EK-55) and EGTEA Gaze+ (EGTEA) datasets, we construct the accuracy reward using mean Average Precision (mAP).

2) Intention Reward: We hypothesize that explicitly articulating high-level task intentions improves long-term action anticipation. To encourage this behavior, we introduce an intention reward that aligns the model’s generated intention int_{gen} with a pseudo-ground-truth intention int_{gt} , automatically produced by GPT-4.1. This provides effective supervision without requiring manual annotation. We compute the cosine similarity between the Sentence-BERT (Reimers and Gurevych 2019) embeddings of the two intentions:

$$sim = \cos(Emb(int_{gen}), Emb(int_{gt})) \in [-1, 1], \quad (12)$$

The intention reward is normalized using a scaled sigmoid function:

$$S_{int} = \min\left(\frac{1}{1+\exp[-\gamma(sim-\beta)]} / \frac{1}{1+\exp[-\gamma(1-\beta)]}, 1\right), \quad (13)$$

where β and γ control the function’s center and sharpness, respectively.

Overall Reward. The overall reward function integrates both **format rewards** and **content rewards** by using a set of weight coefficients $\{\omega_i\}$ to jointly ensure structural validity and semantic precision in LTA, which is defined as:

$$R = \omega_1 S_{len} R_{task} + \omega_2 R_{Soft}, \quad (14)$$

where R_{task} is the weighted *task score*, which emphasizes action accuracy while incorporating auxiliary signals from intention alignment and output formatting:

$$R_{task} = \omega_3 S_{acc} + \omega_4 S_{int} + \omega_5 S_{lang} + \omega_6 S_{fmt}. \quad (15)$$

Experiment

Experimental Settings

Datasets. We evaluated our proposed method using three widely recognized egocentric video benchmarks, each encompassing distinct scenarios and action taxonomies:

1) Ego4D (Grauman et al. 2022) is among the most comprehensive egocentric datasets available, capturing diverse

daily-life scenarios. The *v1* release includes 3,670 hours of footage, with the *Forecasting* subset containing 1,723 annotated clips (116 hours) labeled with 115 verbs and 478 nouns. The updated *v2* dataset expands this subset to 3,472 clips (243 hours), annotated with 117 verbs and 521 nouns. We employed the official data splits from the *v2* version for our experiments.

2) EPIC-Kitchens-55 (Damen et al. 2018) comprises egocentric cooking activities recorded in realistic kitchen settings. This dataset includes 55 hours of densely annotated videos with 125 verbs and 352 nouns. Our experiments utilized standard data splits defined by EGO-TOPO (Nagarajan et al. 2020) to ensure comparability with previous studies.

3) EGTEA Gaze+ (Li, Liu, and Rehg 2018) consists of 86 cooking-focused egocentric videos totaling 26 hours. It provides annotations for 19 verbs and 51 nouns. Consistent with established evaluation protocols, we followed the official splits proposed in (Nagarajan et al. 2020).

Evaluation Metrics. We adopted evaluation protocols specific to each dataset, ensuring consistency and comparability with prior works. For the Ego4D dataset, we adhered to the official evaluation approach, which specifies an observed context of 8 segments and requires predicting 5 distinct future action sequences, each comprising 20 verb-noun pairs. We evaluated using edit distance (Damerau 1964) computed separately for verbs, nouns, and combined actions. The evaluation score is the minimal edit distance achieved between any of the predicted sequences and the ground-truth sequence.

For EK-55 and EGTEA, we followed the multi-label classification evaluation protocol proposed by Nagarajan et al. (Nagarajan et al. 2020). Specifically, we reported mAP across varying observation horizons. The model is given the first $P\%$ of a video as input and tasked with predicting verbs that occur in the remaining $(100 - P)\%$. We evaluated at $P \in \{25, 50, 75\}$, and reported mAP across three action categories: all actions (All), frequently occurring actions (Freq), and rare actions (Rare).

Implementation Details. We employed a frozen EgoVideo-V encoder to extract visual embeddings from full frames and corresponding HOI regions. These dual-stream features were linearly fused and processed by a Transformer

Method	Venue	EK-55 (mAP)			EGTEA (mAP)		
		ALL \uparrow	FREQ \uparrow	RARE \uparrow	ALL \uparrow	FREQ \uparrow	RARE \uparrow
Timeception	CVPR'19	35.6	55.9	26.1	74.1	79.7	59.7
VideoGraph	ICCV'19	22.5	49.4	14.0	67.7	77.1	47.2
EGO-TOPO	CVPR'20	38.0	56.9	29.2	73.5	80.7	54.7
Anticipatr	ECCV'22	39.1	58.1	29.1	76.8	83.3	55.1
AntGPT	ICLR'23	40.1	58.8	31.9	80.2	84.8	72.9
PALM	ECCV'24	40.4	59.3	30.3	80.7	85.0	73.5
ICVL	arXiv'25	<u>43.3</u>	<u>61.6</u>	<u>33.8</u>	<u>81.0</u>	<u>85.2</u>	<u>73.7</u>
INSIGHT	–	45.2	62.4	36.0	81.7	85.9	74.4

Table 2: Performance comparison between INSIGHT and prior state-of-the-art methods on the EK-55 validation set and the EGTEA test set. Underlined values indicate the current state-of-the-art, and bold values denote the highest mAP in each column.

architecture consisting of 4 layers and 8 attention heads per layer to capture spatiotemporal dependencies. The feature extraction stage was trained using a batch size of 8, a learning rate of $8e - 5$, and a plateau-based learning rate scheduler.

For cognitive reasoning, we used Qwen2.5-VL-Instruct-7B (Bai et al. 2025) as the backbone, incorporating structured prompting and trained end-to-end with Swift framework (Zhao et al. 2024) using our GRPO-based reinforcement learning objective. Experiments were conducted on six NVIDIA H20-SXM5-96GB GPUs.

Fine-tuning Hyperparameters included batch size of 24, learning rate of $3e - 6$, temperature of 0.9, KL divergence coefficient of 0.08, generation numbers of 5, buffer length of 256, and a maximum allowable length of 450. Reward function weights were set as follows: $\omega_1 = 0.90$, $\omega_2 = 0.10$, $\omega_3 = 0.85$, $\omega_4 = 0.05$, $\omega_5 = 0.05$, and $\omega_6 = 0.05$. The intention reward parameters were $\beta = 0.8$ and $\gamma = 40$. The fine-tuning lasted 500 steps, requiring about 90 GPU hours.

Performance Comparison

On Ego4D-v2. Table 1 compares INSIGHT against previous state-of-the-art methods on the Ego4D v2 validation set. INSIGHT outperforms the best existing method ICVL (Cao et al. 2025) in noun and combined action predictions by margins of 1.02% and 1.07%, respectively. The superior noun prediction was attributed to our HOI-focused feature extraction, capturing essential object manipulations. On “verb”, our score slightly above the best result. Notably, INSIGHT achieves this slight improvement in verb prediction using a frozen visual encoder, in contrast to EgoVideo (Pei et al. 2024)’s fine-tuned encoder and ICVL’s dual-model approach with ViT-L/14 and BLIP2-OPT-2.7B for enhanced visual learning. This suggests that the fine-tuned language model and the proposed cognitive reasoning module effectively compensate for visual ambiguities, leading to more accurate and context-aware long-term predictions.

On EK-55 and EGTEA. Table 2 further demonstrates INSIGHT’s superior performance on EK-55 and EGTEA benchmarks, improving on the baseline ICVL by 4.4%, 1.3%, and 6.5% (EK-55) and by 0.9%, 0.8%, and 1.0% (EGTEA) for All, Freq, and Rare actions, respectively. No-

Method	Ego4D-v2 (ED)		
	Verb \downarrow	Noun \downarrow	Action \downarrow
w/o HOI feature	0.6719	0.6158	0.8595
w/o Semantic correction	0.6716	0.6108	0.8587
w/o Cognitive reasoning	0.6750	0.6176	0.8612
w/o Intention	0.6685	0.6104	0.8571
INSIGHT (full model)	0.6643	0.6092	0.8463

Table 3: Ablation study analyzing the contribution of different modules within the INSIGHT framework.




Figure 3: Training visualization of the GRPO-based cognitive reasoning process.

tably, INSIGHT excels in rare classes, indicating that our cognitive reasoning and intention-alignment mechanisms substantially reduce tail-class confusion, ensuring robust generalization. Moreover, consistent improvements with a reduced head-tail gap showing better calibration, these gains persist across datasets with varying protocols, indicating the structured reasoning trace enhances long-term prediction performance.

Ablation Study

We conducted ablation studies to quantify the contribution of each component of INSIGHT. Comparison results are summarized in Table 3.

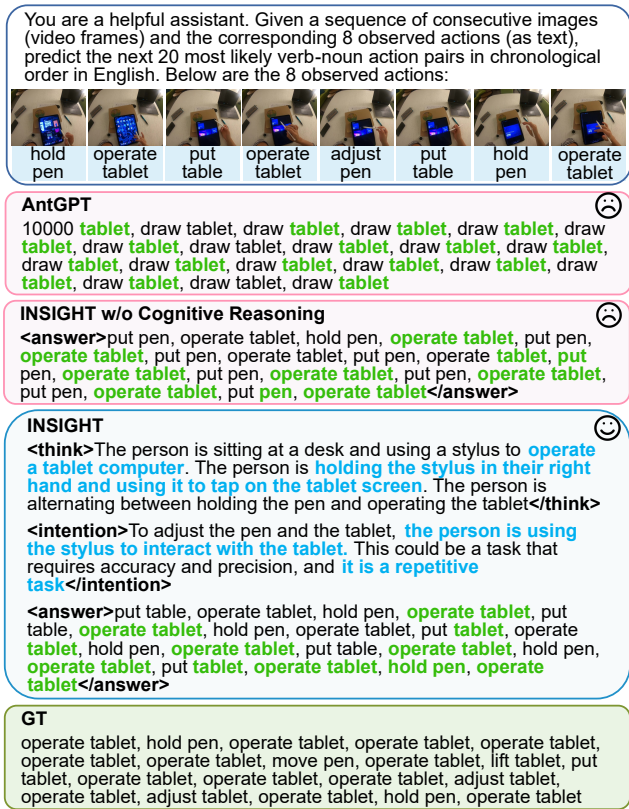


Figure 4: Case study on Ego4D for the LTA task. INSIGHT demonstrates intention-aware predictions with improved temporal coherence, fewer redundant actions, and more precise verb-noun pairings, leading to lower edit distance.

On Action Recognition. In the action recognition stage, removing HOI region features and using only raw frames (w/o HOI feature) leads to a clear degradation. This confirms the importance of explicitly modeling HOI cues, which significantly enhances recognition accuracy and supports downstream anticipation. Additionally, removing the verb-noun co-occurrence matrix (w/o Semantic correction) also causes a noticeable drop in performance. This indicates that structured semantic priors contribute to more discriminative and semantically consistent action representations.

On Anticipation. In the anticipation stage, discarding the structured reasoning trace (think → reason → answer) and directly predicting future actions (w/o Cognition reasoning) leads to a substantial drop in performance. This demonstrates that stepwise reasoning is crucial for enabling the model to internalize task intentions and temporal dependencies effectively. Furthermore, maintaining the reasoning pipeline but omitting supervision for the intention inference stage (w/o Intention) also results in performance degradation compared to the full model. This highlights the value of supervising intermediate reasoning steps, as it calibrates the cognitive process and improves overall prediction quality.

Overall, these results show that each module provides

complementary gains, and their integration delivers the highest accuracy. Notably, removing the structured reasoning trace (w/o Cognitive reasoning) has the most significant impact on performance, highlighting its importance. These findings align with prior observations that both high-quality action recognition and explicit cognitive reasoning are essential for robust and reliable long-term action anticipation.

Reward Analysis

Figure 3 presents learning dynamics under the adopted training configuration. (a) illustrates the progression of the overall reward, while (b) tracks the intention-specific reward; both panels show raw and smoothed curves to capture training trends. Training converges within 500 steps: rewards rise sharply during the initial 1 – 100 steps, followed by gradual improvement and stabilization between 100 – 300 steps. Importantly, the intention reward closely mirrors the overall reward trajectory, suggesting that intention supervision is well aligned with the task objective. Although raw curves exhibit fluctuations typical of on-policy learning and exploration, smoothed trends indicate consistent performance gains without signs of instability. These results demonstrate the robustness of GRPO under our selected hyperparameters and validate the utility of explicit intention guidance in improving policy quality and convergence.

Qualitative Comparison

We present a qualitative comparison on Ego4D-v2 by reproducing the AntGPT (Zhao et al. 2023) baseline, which, as shown in Figure 4, exhibits task-irrelevant tokens, repetitive predictions, and verb-noun mismatches. By contrast, INSIGHT without cognitive reasoning enhances lexical fidelity but suffers from temporal inconsistency, exemplified by repeated “put pen” actions without prior “hold pen”. The full INSIGHT model, equipped with cognitive reasoning, outperforms both by inferring task intention and aligning sub-goals, yielding precise verb-noun pairings, expanded verb diversity, and lower edit distances.

Conclusion

This paper introduces INSIGHT, a unified two-stage framework for LTA that combines Hand-Object Semantic Action Recognition with Explicit Cognitive Reasoning. INSIGHT first extracts discriminative visual features from HOI regions and reinforces semantic consistency using a verb-noun co-occurrence prior, enabling the construction of robust action representations. These representations are then processed through a reinforcement learning-based cognitive reasoning module that simulates a structured inference process “think → reason → answer” guided by intention supervision. Extensive experiments on Ego4D, EPIC-Kitchens-55, and EGTEA Gaze+ demonstrate that INSIGHT consistently outperforms prior state-of-the-art methods. As future work, we will capture egocentric dynamics by modeling hand motion trajectories and object state changes in HOI regions as priors, thereby strengthening visual grounding and enhancing long-term anticipation.

Acknowledgments

This work is supported by Shenzhen Science and Technology Program, No.:KQTD20240729102207002; Jiangsu Science and Technology Major Program, No.:BG2024041; the National Natural Science Foundation of China, No.:62376140, No.:625B2065 and No.:U23A20315; the Science and Technology Innovation Program for Distinguished Young Scholars of Shandong Province Higher Education Institutions, No.:2023KJ128; the Special Fund for Taishan Scholar Project of Shandong Province; the Major Key Project of Pengcheng Laboratory; the Fundamental Research Funds for the Central Universities, No.:HIT.DZJJ.2025048.

References

- Abreu, S.; Do, T. D.; Ahuja, K.; Gonzalez, E. J.; Payne, L.; McDuff, D.; and Gonzalez-Franco, M. 2024. Parse-ego4d: Personal action recommendation suggestions for egocentric videos. *arXiv preprint arXiv:2407.09503*.
- Ashutosh, K.; Girdhar, R.; Torresani, L.; and Grauman, K. 2023. Hiervl: Learning hierarchical video-language embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23066–23078.
- Azam, M. M.; and Desai, K. 2024. A Survey on 3D Egocentric Human Pose Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 1643–1654.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Cao, C.; Hu, L.; Yu, Y.; and Zhang, Y. 2025. Vision and Intention Boost Large Language Model in Long-Term Action Anticipation. *arXiv preprint arXiv:2505.01713*.
- Chen, G.; Zheng, Y.-D.; Wang, J.; Xu, J.; Huang, Y.; Pan, J.; Wang, Y.; Wang, Y.; Qiao, Y.; Lu, T.; et al. 2023. Videollm: Modeling video sequence with large language models. *arXiv preprint arXiv:2305.13292*.
- Chu, Q.; Zhang, H.; Feng, Y.; Liu, M.; Guan, W.; Wang, Y.; and Nie, L. 2025. Technical report for ego4d long-term action anticipation challenge 2025. *arXiv preprint arXiv:2506.02550*.
- Damen, D.; Doughty, H.; Farinella, G. M.; Fidler, S.; Furnari, A.; Kazakos, E.; Moltisanti, D.; Munro, J.; Perrett, T.; Price, W.; and Wray, M. 2018. Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In *European Conference on Computer Vision (ECCV)*.
- Damerau, F. J. 1964. A technique for computer detection and correction of spelling errors. *Communications of the ACM*, 7(3): 171–176.
- Das, S.; and Ryoo, M. S. 2022. Video+ clip baseline for ego4d long-term action anticipation. *arXiv preprint arXiv:2207.00579*.
- Grauman, K.; Westbury, A.; Byrne, E.; Chavis, Z.; Furnari, A.; Girdhar, R.; Hamburger, J.; Jiang, H.; Liu, M.; Liu, X.; et al. 2022. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18995–19012.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Huang, D.; Hilliges, O.; Van Gool, L.; and Wang, X. 2023. Palm: Predicting actions through language models@ ego4d long-term action anticipation challenge 2023. *arXiv preprint arXiv:2306.16545*.
- Hussein, N.; Gavves, E.; and Smeulders, A. W. 2019. Time-ception for complex action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 254–263.
- Ishibashi, T.; Ono, K.; Kugo, N.; and Sato, Y. 2023. Technical report for ego4d long term action anticipation challenge 2023. *arXiv preprint arXiv:2307.01467*.
- Kim, S.; Huang, D.; Xian, Y.; Hilliges, O.; Van Gool, L.; and Wang, X. 2024. Palm: Predicting actions through language models. In *European Conference on Computer Vision*, 140–158. Springer.
- Lee, J.; Tjahjadi, A. D.; Kim, J.; Yu, J.; Park, M.; Zhang, J.; Froehlich, J. E.; Tian, Y.; and Zhao, Y. 2024. CookAR: Affordance augmentations in wearable AR to support kitchen tool interactions for people with low vision. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 1–16.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, Y.; Liu, M.; and Rehg, J. M. 2018. 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)*, 619–635.
- Lin, K. Q.; Wang, J.; Soldan, M.; Wray, M.; Yan, R.; Xu, E. Z.; Gao, D.; Tu, R.-C.; Zhao, W.; Kong, W.; et al. 2022. Egocentric video-language pretraining. *Advances in Neural Information Processing Systems*, 35: 7575–7586.
- Liu, M.; Wang, X.; Nie, L.; He, X.; Chen, B.; and Chua, T.-S. 2018a. Attentive moment retrieval in videos. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, 15–24.
- Liu, M.; Wang, X.; Nie, L.; Tian, Q.; Chen, B.; and Chua, T.-S. 2018b. Cross-modal moment localization in videos. In *Proceedings of the 26th ACM international conference on Multimedia*, 843–851.
- Liu, X.; Ji, K.; Fu, Y.; Tam, W. L.; Du, Z.; Yang, Z.; and Tang, J. 2021. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*.
- Mascaro, E. V.; Ahn, H.; and Lee, D. 2022. Intention-Conditioned Long-Term Human Egocentric Action Forecasting. *arXiv preprint arXiv:2207.12080*.
- Mittal, H.; Agarwal, N.; Lo, S.-Y.; and Lee, K. 2024. Can't make an omelette without breaking some eggs: Plausible action anticipation using large video-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18580–18590.

- Nagarajan, T.; Li, Y.; Feichtenhofer, C.; and Grauman, K. 2020. Ego-topo: Environment affordances from egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 163–172.
- Nawhal, M.; Jyothi, A. A.; and Mori, G. 2022. Rethinking learning approaches for long-term action anticipation. In *European Conference on Computer Vision*, 558–576. Springer.
- Ohkawa, T.; He, K.; Sener, F.; Hodan, T.; Tran, L.; and Keskin, C. 2023. Assemblyhands: Towards egocentric activity understanding via 3d hand pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12999–13008.
- Pei, B.; Chen, G.; Xu, J.; He, Y.; Liu, Y.; Pan, K.; Huang, Y.; Wang, Y.; Lu, T.; Wang, L.; et al. 2024. Egovideo: Exploring egocentric foundation model and downstream adaptation. *arXiv preprint arXiv:2406.18070*.
- Plizzari, C.; Goletto, G.; Furnari, A.; Bansal, S.; Ragusa, F.; Farinella, G. M.; Damen, D.; and Tommasi, T. 2024. An outlook into the future of egocentric vision. *International Journal of Computer Vision*, 132(11): 4880–4936.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Reimers, N.; and Gurevych, I. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shan, D.; Geng, J.; Shu, M.; and Fouhey, D. F. 2020. Understanding human hands in contact at internet scale. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9869–9878.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, X.; Wu, J.; Lin, Z.; Zhang, F.; Zhang, D.; and Nie, L. 2025. Video DataFlywheel: Resolving the Impossible Data Trinity in Video-Language Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wu, J.; Liu, W.; Liu, Y.; Liu, M.; Nie, L.; Lin, Z.; and Chen, C. W. 2025. A survey on video temporal grounding with multimodal large language model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xiao, J.; Huang, N.; Qiu, H.; Tao, Z.; Yang, X.; Hong, R.; Wang, M.; and Yao, A. 2025. EgoBlind: Towards Egocentric Visual Assistance for the Blind People. *arXiv preprint arXiv:2503.08221*.
- Xu, F.; Zhou, T.; You, H.; and Du, J. 2024. Improving indoor wayfinding with AR-enabled egocentric cues: A comparative study. *Advanced Engineering Informatics*, 59: 102265.
- Xu, Y.; Li, Y.-L.; Huang, Z.; Liu, M. X.; Lu, C.; Tai, Y.-W.; and Tang, C.-K. 2023. Egozca: A new framework for egocentric hand-object interaction understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5273–5284.
- Zhang, C.; Fu, C.; Wang, S.; Agarwal, N.; Lee, K.; Choi, C.; and Sun, C. 2024a. Object-centric video representation for long-term action anticipation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 6751–6761.
- Zhang, C.; Gupta, A.; and Zisserman, A. 2023. Helping hands: An object-aware ego-centric video recognition model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13901–13912.
- Zhang, H.; Liu, M.; Gao, Z.; Lei, X.; Wang, Y.; and Nie, L. 2021. Multimodal dialog system: Relational graph-based context-aware question understanding. In *Proceedings of the 29th ACM international conference on multimedia*, 695–703.
- Zhang, H.; Liu, M.; Li, Y.; Yan, M.; Gao, Z.; Chang, X.; and Nie, L. 2023. Attribute-guided collaborative learning for partial person re-identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12): 14144–14160.
- Zhang, H.; Liu, M.; Li, Z.; Wen, H.; Guan, W.; Wang, Y.; and Nie, L. 2025. Spatial Understanding from Videos: Structured Prompts Meet Simulation Data. In *Advances in Neural Information Processing Systems*, 1–16.
- Zhang, H.; Liu, M.; Liu, Z.; Song, X.; Wang, Y.; and Nie, L. 2024b. Multi-factor adaptive vision selection for egocentric video question answering. In *Forty-first International Conference on Machine Learning*.
- Zhao, Q.; Wang, S.; Zhang, C.; Fu, C.; Do, M. Q.; Agarwal, N.; Lee, K.; and Sun, C. 2023. Antgpt: Can large language models help long-term action anticipation from videos? *arXiv preprint arXiv:2307.16368*.
- Zhao, Y.; Huang, J.; Hu, J.; Wang, X.; Mao, Y.; Zhang, D.; Jiang, Z.; Wu, Z.; Ai, B.; Wang, A.; Zhou, W.; and Chen, Y. 2024. SWIFT: A Scalable lightWeight Infrastructure for Fine-Tuning. *arXiv:2408.05517*.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36: 46595–46623.
- Zhong, Z.; Martin, M.; Diederichs, F.; and Beyerer, J. 2024. Querymamba: A mamba-based encoder-decoder architecture with a statistical verb-noun interaction module for video action forecasting@ ego4d long-term action anticipation challenge 2024. *arXiv preprint arXiv:2407.04184*.