



TeleEgo: Benchmarking Egocentric AI Assistants in the Wild

Jiaqi Yan*, Ruilong Ren*, Jingren Liu*, Shuning Xu, Ling Wang,
Yiheng Wang, Xinlin Zhong, Yun Wang, Long Zhang, Xiangyu Chen[†],
Changzhi Sun, Jixiang Luo, Dell Zhang[#], Hao Sun, Chi Zhang, Xuelong Li[#]

Institute of Artificial Intelligence (TeleAI), China Telecom

<https://github.com/TeleAI-UAGI/TeleEgo>

Abstract

*Egocentric AI assistants in real-world settings must process multi-modal inputs (video, audio, text), respond in real time, and retain evolving long-term memory. However, existing benchmarks typically evaluate these abilities in isolation, lack realistic streaming scenarios, or support only short-term tasks. We introduce **TeleEgo**, a long-duration, streaming, omni-modal benchmark for evaluating egocentric AI assistants in realistic daily contexts. The dataset features over 14 hours per participant of synchronized egocentric video, audio, and text across four domains: work & study, lifestyle & routines, social activities, and outings & culture. All data is aligned on a unified global timeline and includes high-quality visual narrations and speech transcripts, curated through human refinement. TeleEgo defines 12 diagnostic subtasks across three core capabilities: Memory (recalling past events), Understanding (interpreting the current moment), and Cross-Memory Reasoning (linking distant events). It contains 3,291 human-verified QA items spanning multiple question formats (single-choice, binary, multi-choice, and open-ended), evaluated strictly in a streaming setting. We propose Real-Time Accuracy (RTA) to jointly capture correctness and responsiveness under tight decision windows, and Memory Persistence Time (MPT) as a forward-looking metric for long-term retention in continuous streams. In this work, we report RTA results for current models and release TeleEgo, together with an MPT evaluation framework, as a realistic and extensible benchmark for future egocentric assistants with stronger streaming memory, enabling systematic study of both real-time behavior and long-horizon memory.*

1. Introduction

With the rapid advancement of artificial intelligence, egocentric AI assistants—those operating from a first-person perspective—are gradually transitioning from controlled experimental settings to real-world applications. To function effectively in such scenarios, these assistants must exhibit three tightly integrated capability: memory, streaming decision-making, and multimodal understanding. They must be able to retain and recall growing streams of past information; make timely judgments in the context of continuous audio-visual inputs; and interpret, in a unified manner, what the camera sees, what the microphone hears, and what the user expresses through language. Importantly, these capability are not exercised in isolation—they must work together in harmony. A model that remembers past events but acts at the wrong moment can still fail. Similarly, a system that processes the current frame but cannot identify the speaker or key objects is unlikely to succeed. Effective real-world AI assistants must therefore reason not only about what is happening, but also when and how to respond, based on long-term context and multimodal cues.

Despite these requirements, existing benchmarks evaluate these ability in isolation or in simplified settings (Table 1). First, some focus on offline long-term memory (e.g., X-LeBench [26]), while others test short-window streaming (e.g., StreamingBench [14], VStream-QA [24]), making it hard to assess the trade-off between memory and real-time performance. Second, true egocentric streaming evaluation is rare. Most datasets use third-person or static videos [10, 14, 16], avoiding challenges like self-motion and viewpoint shifts. Some exceptions (e.g., ODVBench [23]) use first-person footage, but with short sequences and limited multimodality. Third, few datasets of-

* Equal contribution. [†] Project lead. [#] Corresponding authors.

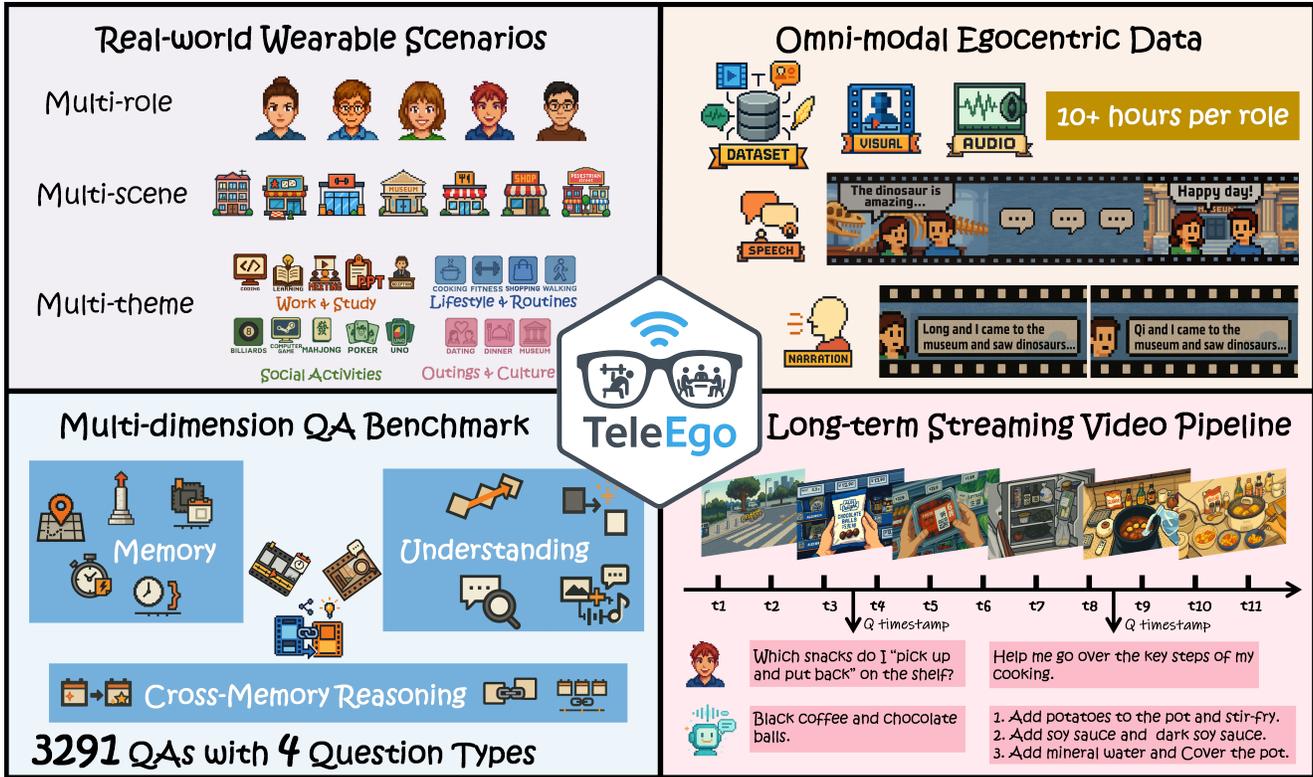


Figure 1. **An overview of the TeleEgo project.** Top-left: Scripted real-world wearable-camera scenarios covering multiple roles, scenes, and themes. Top-right: Omni-modal egocentric streaming data aligned to a shared timeline, comprising video, audio, and human-curated speech transcripts and visual narrations. Online multitask QA benchmark organized across three capability dimensions (Memory, Understanding, Cross-Memory Reasoning), containing 3,291 QA items across 4 question types. Bottom-right: Long-term streaming video pipeline—egocentric footage with query-time retrieval spanning seconds to days.

fer long, continuous, real-world recordings. Many are short clips or image sets [3, 23, 25]. While X-LeBench and Ego-Life [21] are longer, the former stitches clips, and the latter is recorded in closed and controlled environments.

To address these challenges, we present TeleEgo: a long-duration, streaming, and fully multimodal benchmark grounded in real-world scenarios, purpose-built to evaluate egocentric AI assistants (see Figure 1). TeleEgo consists of synchronized video, audio, and textual data collected from multiple participants, each contributing more than 14 hours of recordings. The dataset spans four major domains: work & study, lifestyle & routines, social activities, and outings & culture. All data streams are precisely aligned to a unified global timeline and enriched with manually curated speech transcripts and visual narrations to ensure high quality and semantic clarity. Building on this foundation, TeleEgo introduces 12 diagnostic tasks covering three core ability: *memory* (recalling past events), *understanding* (interpreting the present), and *cross-memory reasoning* (connecting distant moments). These tasks include a total of 3,291 human-verified QA items across various formats, including single-choice, binary, multiple-choice, and open-

ended questions. Each task is tied to a specific time point and a decision window, requiring models not only to answer correctly, but also to respond at the right time. All evaluations are conducted under streaming conditions, and a task is considered successful only if the model’s first correct answer falls within the allowed window. This setup prevents models from guessing too early or answering too late, offering a more accurate measure of real-time decision-making and responsiveness. We further introduce two evaluation metrics: Real-Time Accuracy, which measures whether the model produces a correct response within a given decision window, and Memory Persistence Time, a forward-looking metric that tracks how long past information remains correctly retrievable in a continuous stream without replaying evidences. In this work, we primarily report RTA on existing models and provide MPT as an evaluation protocol for future systems with stronger long-term streaming memory.

In general, our contributions are following:

- A unified, streaming, multimodal benchmark aligned to a global timeline with rich, real-world data across multiple participants and scenarios.
- A diagnostic task suite covering memory, understanding,

Benchmark	Focus Scene	Video Duration	Tasks	Omni-Modal	Streaming	Egocentric	Long-Mem QA
EgoExoLearn [9]	Ego-Exo Skill Assessment	avg. 13.4 min (ego) avg. 4.5 min (exo)	Cross-view Tasks	✗	✗	Partial	✗
EgoThink [3]	First-Person Thinking	–	Object, Activity, Localization, Reasoning, Forecasting and planning	✗	✗	✓	✗
OVBench [10]	Online Video Understanding	30 s-1 h	Spatiotemporal Understanding and Interpretation	✗	✓	✗	✗
StreamingBench [14]	Online Video Understanding	3 s-24 min	Real-Time Visual, Omni-Source and Contextual Understanding	✓	✓	✗	✗
EgoTextVQA [25]	Egocentric Scene-text	avg. 101.7 s	Identification and Reasoning	✗	✗	✓	✗
X-LeBench [26]	Extra Long Egocentric data	23 min–16.4 h	Temporal Localization, Summarization, Counting and Ordering	✗	✗	✓	✓
VStream-QA [24]	Online Video Stream	avg. 40 min	Event and Scene Understanding	✗	✓	Partial	✗
OVO-Bench [16]	Online Video Reasoning	avg. 263.42s	Backward Tracing, Real-Time Visual Perception, and Forward Active Responding	✗	✓	✗	✗
ODV-Bench [23]	Online Driving Video	5-90 s	Realtime Perception and Prediction	✗	✓	✓	✗
EgoLife [21]	Egocentric Assistant	avg. 44.3 h	EntityLog, EventRecall, HabitInsight, RelationMap and TaskMaster	✓	✗	✓	✓
M3-Bench [15]	Agent Memory and Reasoning	avg. 34 min	Agent-centric QA	✓	✓	Partial	✓
TeleEgo (Ours)	Real-world Settings	avg. 14.4 h	Memory, Understanding and Cross-Memory Reasoning	✓	✓	✓	✓

Table 1. **Related benchmarks for TeleEgo.** Comparison across duration, egocentricity, streaming protocol, modality coverage, and long-memory QA. TeleEgo (ours) uniquely satisfies all dimensions. ✓denotes supported, ✗ not applicable, and Partial partial support.

and cross-memory reasoning, with timestamped, multi-modal evidence for interpretability.

- An evaluation protocol for streaming inputs, using Real-Time Accuracy metric to assess contemporary models and a forward-looking Memory Persistence Time metric to study long-term retention.

2. Related Works

Egocentric Models and Benchmarks. Egocentric (first-person) vision research has grown significantly, evolving from early single-user recordings to diverse, large-scale datasets capturing daily life from a first-person perspective. Pioneering work such as EPIC-KITCHENS [5] and its extensions like VISOR [6] introduced large-scale, object-rich video datasets in home environments. Later, Ego4D and Ego-Exo4D [7, 8] broadened the scope to include tasks like episodic memory, future prediction, and skill learning, incorporating both egocentric and exocentric views. Recent datasets have begun to explore more complex cognitive tasks and assistant-like interactions. EgoLife [21] captures long-duration recordings from a single household, while EgoThink [3] and EgoExoLearn [9] focus on segment-level reasoning and teaching-following dynamics. Other work like MM-Ego [22] and EgoTextVQA [25] targets memory and text-based understanding. However, most of these evaluations remain offline and task-specific. TeleEgo advances egocentric benchmarking by combining multi-day, multi-role, and multi-theme recordings with dual text annotations aligned to a unified global timeline. Its online evaluation protocol measures not only real-time decision-making, but also the persistence of memory over time—enabling more realistic evaluations of assistant capabilities.

Streaming Video Understanding Benchmarks. To assess assistants in time-sensitive and dynamic environments, several benchmarks have emerged that focus on streaming video understanding. StreamingBench [14], OVBench, and OVO-Bench [10, 16] support online task formats, but typi-

cally span short episodes and lack sustained memory testing across events. ODV-Bench [23] emphasizes driving tasks and short-term prediction, prioritizing perception over memory. Some datasets, such as X-LeBench [26], extend video QA to longer contexts, but still operate in offline settings without real-time constraints. TeleEgo fills this gap by offering continuous, multi-day egocentric video streams with temporally grounded question-answer pairs. Its evaluation protocol emphasizes real-time responses and long-term memory recall, supporting deeper analyses of assistant performance in realistic, ever-evolving scenarios.

Omni-Modal Assistants. Recent omni-modal models seek to unify text, vision, audio, and speech for more flexible assistants. Closed-source systems such as GPT-4o [11] and Gemini 1.5/2.5 [4, 18] show strong empirical performance on speech and video, but their architectures, memory mechanisms, and streaming behavior are proprietary and largely unknown. Open-source omni-modal models have also progressed rapidly. LLaVA-OneVision [12] and InternVL-2.5 [2] support multi-image and video understanding, while Qwen2.5-Omni [20] and Baichuan-Omni [13] are optimized for audiovisual input and interactive responses. Yet they are mostly evaluated on offline, task-specific benchmarks or short clips rather than long, continuous egocentric streams. MiniCPM-o [19] moves closer to streaming by providing an online interface for token- and frame-level inputs, but relies on a growing KV cache without effective long-term memory control. Beyond single-model interfaces, M3-Agent [15] introduces a multimodal agent with explicit long-term memory, organizing episodic and semantic information in an entity-centric, multimodal form. However, it still assumes segment-wise video input instead of second-by-second egocentric streams. Despite these advances, few works test whether omni-modal assistants can both respond accurately under time-sensitive decision windows and retain information over long temporal spans in realistic, continuous settings. TeleEgo com-

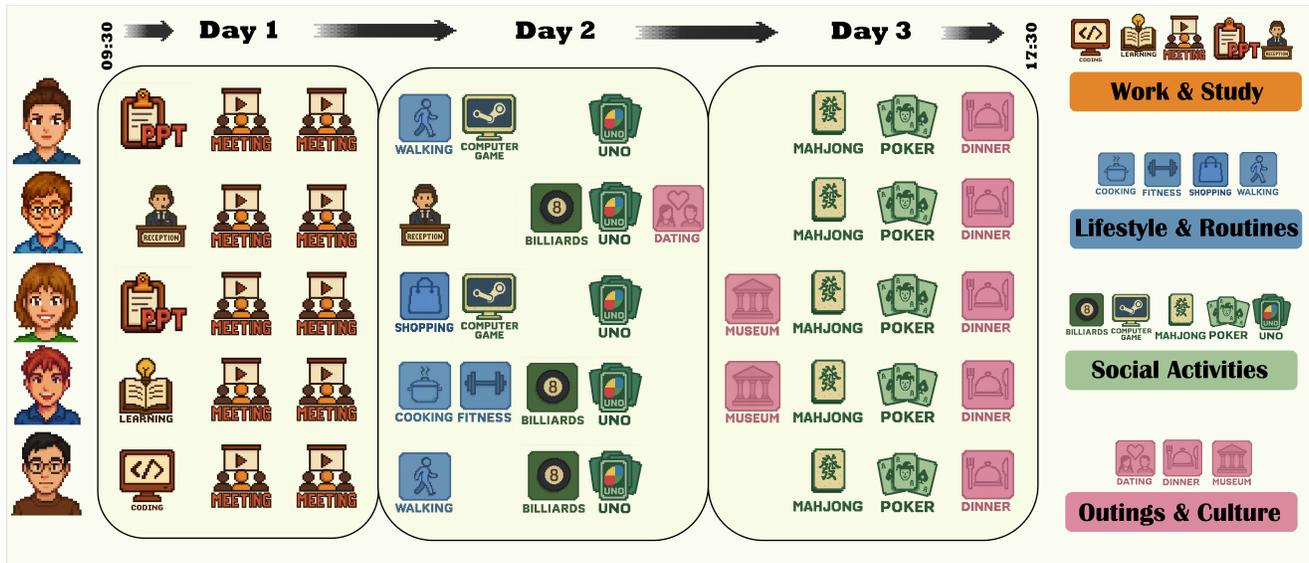


Figure 2. Scenario design and activity distribution in TeleEgo dataset. Each role engaged in diverse first-person activities across three recording days, systematically covering four themes, Work & Study, Lifestyle & Routines, Social Activities, and Outings & Culture. The design spans a wide spectrum of cognitive and social contexts, combining solo and multi-role interactions across indoor and outdoor environments. This structure ensures ecological diversity and supports analyses of long-term, cross-situational understanding.

binces egocentric recordings, streaming video contexts, and real-time memory-oriented evaluation to provide a unified testbed bridging egocentric datasets, streaming benchmarks, and multimodal models.

3. TeleEgo

3.1. Dataset Overview

To ensure broad diversity and real-world relevance, TeleEgo uses a carefully designed data collection protocol that spans multiple roles, themes, and tasks. We recruited five participants with balanced gender representation and a wide range of cultural, regional, and personality backgrounds, aiming to reflect a representative slice of the general population. Each participant wore a first-person camera over three consecutive days and recorded egocentric video following a set of predefined scenarios. The recordings include both solo activities and group interactions, taking place in diverse indoor and outdoor environments and across various social contexts. This approach goes beyond traditional single-household or uniform-group datasets, improving the generalizability and practical value of the benchmark. To capture the richness of everyday human experience, we structured the data around four common life themes (see Figure 2). These themes cover different cognitive demands, social situations, environmental settings, and physical activities.

Work & Study (e.g., giving a presentation, meetings, coding, learning, reception tasks): This category includes knowledge-based and goal-driven tasks that often involve tools, screens, and structured interactions. These activities

require focused attention, task switching, and formal turn-taking, making them ideal for evaluating cognitive workload and procedural behaviors.

Social Activities (e.g., playing UNO, Mahjong, Poker, video games, billiards): These scenarios feature multi-person interactions with competitive or turn-based structures. They are rich in gesture-speech coordination and quick context changes. This category is well-suited for analyzing gaze behavior, social cues, conversational grounding, and multimodal communication.

Lifestyle & Routines (e.g., shopping, exercising, walking, cooking): This theme involves semi-structured daily activities that combine object handling and movement in dynamic yet familiar environments. It supports research on long-term activity recognition, task progression, and inferring higher-level states such as fatigue or task completion.

Outings & Culture (dining out, dating, visiting museums): These scenarios occur in complex public environments with varying lighting, noise levels, and crowd density. They also involve subtle social norms and cultural practices. This category helps evaluate model robustness to occlusion, background noise, and unfamiliar contexts, while enabling understanding of social intent and etiquette.

Over the course of three days, each participant recorded a wide range of egocentric videos covering all four themes. The result is a rich, multi-role, multi-theme, multi-day dataset that supports research on long-term memory, context carryover, and generalization across different situations—key challenges in real-world perception.

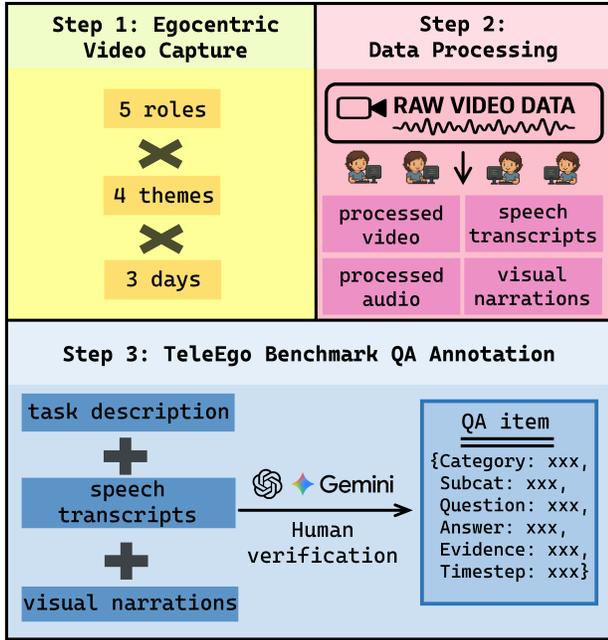


Figure 3. **TeleEgo construction pipeline.** Step 1: egocentric video capture across 5 roles, 4 themes and 3 days. Step 2: data processing into synchronized video, audio, speech, and narration captions. Step 3: AI tools generate candidate QA items from task descriptions and captions, followed by human verification.

3.2. Raw Data Processing

To ensure privacy and ethical use, all collected recordings go through a careful de-identification process. This includes blurring faces, removing speech from non-participants, and masking any sensitive visual or audio content. These steps preserve participant privacy while keeping the recordings natural and realistic, reflecting everyday first-person experiences. To support deeper multimodal understanding beyond raw audio and video, the TeleEgo dataset includes two types of time-aligned textual annotations.

Speech transcripts captures all verbal communication in multi-person settings. Spoken content is automatically transcribed, then manually verified and annotated with speaker identities, yielding temporally aligned conversational transcripts that preserve discourse structure and interaction dynamics. This produces complete dialogue corpora suitable for studying social cognition, turn-taking, and multimodal grounding in egocentric contexts.

Visual narrations consists of participants’ self-reported verbal descriptions of their ongoing activities and salient environmental details. When explicit actions are absent, narrations focus on attentional targets and key scene elements, providing semantic coverage of visual content such as object interactions, spatial relations, and contextual cues. Each narration is timestamped and aligned with the corresponding video segment, forming a natural-language layer

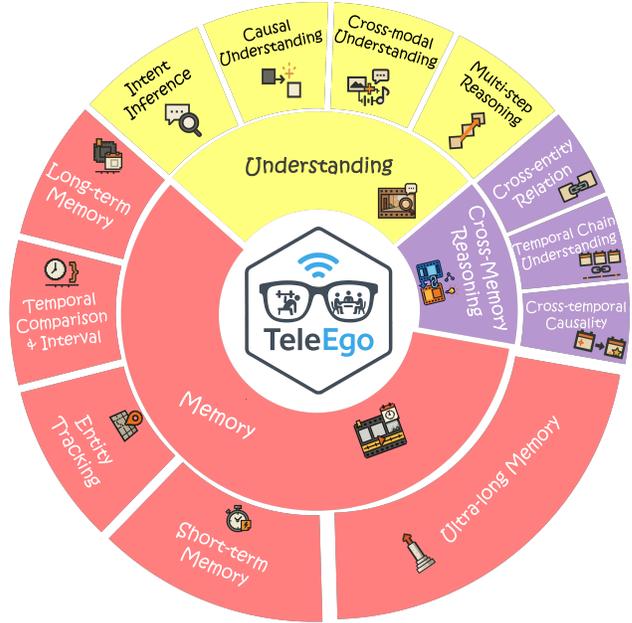


Figure 4. Hierarchical organization of the TeleEgo benchmark. The benchmark is organized around three cognitive dimensions: Memory, Understanding, and Cross-Memory Reasoning. Each dimension is further divided into fine-grained subcategories.

that parallels the perceptual stream.

Both streams are precisely synchronized with the video timeline, producing dual-layer annotations for visual and linguistic events. This structure enables rich cross-modal grounding. The processed multimodal data is then used by powerful AI tools to generate candidate QA items, which are further refined through human verification to build our benchmark system (see Figure 3). Collectively, TeleEgo offers real-world recordings that combine perception, language, and memory, furnishing comprehensive multimodal material for evaluating AI systems’ capacity to understand and retain complex first-person experiences.

3.3. Benchmark Task Design

To evaluate multidimensional cognitive abilities in egocentric video understanding, we introduce a benchmark that spans three key cognitive dimensions: memory, comprehension, and cross-memory reasoning (see Figure 4). These dimensions form a hierarchical structure that reflects different levels of cognition, ranging from momentary perception to long-term reasoning. This framework enables a systematic distinction between information retention, semantic understanding, and integrative reasoning across time and entities. To support this evaluation, we design 12 fine-grained question-answering subtasks, each corresponding to one of the three dimensions, allowing us to assess model performance across a wide range of cognitive scenarios. Examples of these subtasks are illustrated in Figure 5.

The **Memory** focuses on temporally grounded recall, assessing a model’s ability to retain, retrieve, and compare events over different time spans. Tasks range from short-term recall of transient object states and actions, to long-term and ultra-long memory over extended episodes, as well as continuous entity tracking and temporal interval reasoning. These tasks evaluate how well models maintain temporal coherence and represent evolving dynamics in first-person experiences.

The **Understanding** dimension measures a model’s capacity to grasp meaning and coherence within complex, context-rich scenarios. It goes beyond surface perception to assess how well a model understands causal structures and human intent. This includes recognizing cause-effect relationships, inferring latent motivations, and constructing a unified interpretation from temporally or spatially dispersed cues. It also requires integrating multimodal inputs into coherent semantic representations.

The **Cross-Memory Reasoning** dimension challenges models to combine information across disjoint time periods and entity contexts. Tasks require building global narrative structures, linking distant events into causal chains, inferring relational dynamics between interacting agents, and synthesizing long temporal sequences into structured, meaningful processes. This dimension represents the most complex aspect of egocentric cognition, requiring reasoning over long-range dependencies in continuous experiences.

3.4. Benchmark QA Annotation

We adopt four complementary QA formats: single-choice (mc_single), multiple-choice (mc_multi), binary, and open-ended. Each format serves a distinct purpose. The mc_single format allows for precise evaluation through carefully crafted distractors and unambiguous correct answers. The mc_multi format captures complex or uncertain scenarios by permitting multiple correct options. Binary questions offer high-precision evaluation at low annotation cost. Open-ended questions encourage free-form reasoning and compositional thinking, complementing the more structured formats. Together, these formats strike a balance between standardization and expressiveness, enabling scalable evaluation while supporting fine-grained behavioral probing. Our QA generation process begins with post-processed, time-aligned transcripts of speech and narration. We use state-of-the-art large language models (GPT-5-Thinking and Gemini-2.5-Pro) to draft initial QA candidates. For the Ultra-Long Memory subcategory, the models ingest full dual-stream transcripts and generate questions grounded in evidence spanning 10–60 minutes. For the remaining eleven subcategories, we segment each recording into 30-minute windows and prompt the models to generate QA pairs with evidence evenly distributed across each window. Human annotators then verify factual alignment with

the source video, correct timestamps, and remove ambiguous or low-quality items.

As shown in Figure 4, the TeleEgo benchmark is hierarchically structured along three key cognitive dimensions—memory, understanding, and reasoning—with twelve subcategories providing finer granularity. The final dataset contains 3,291 verified QA instances across all four formats (see Table 2). Subcategory distributions demonstrate balanced coverage across temporal, causal, and semantic reasoning challenges. Collectively, these design choices make TeleEgo a robust and discriminative benchmark for evaluating video understanding models.

4. Experiments

4.1. Evaluation Protocol

Online-only regime. All evaluations adhere to the streaming protocol. Video frames and audio are timestamped to a single global clock. For question-answer (QA) instance i , we define a decision window $[t_s^{(i)}, t_s^{(i)} + T]$, where T denotes the time margin. At test time, the model must respond within this window. We set $T = 5$ s in all experiments. We report two primary metrics:

1) Real-Time Accuracy (RTA, %): The percentage of QA items for which the model outputs a correct answer within the decision window.

2) Memory Persistence Time (MPT, minutes): For each item correctly answered at time t^* , we continue streaming without repeating the original evidence. At regular intervals, we re-query the same item. MPT is the time from t^* until the first failed recall. If the item is never answered correctly, $MPT = 0$. If it is never forgotten within the probing range, it is right-censored at the last probe. Figure 6 provides a schematic overview of the pipeline used to compute Memory Persistence Time (MPT).

Systems and settings. We evaluate six state-of-the-art vision-language or omni-modal models, covering both proprietary and open-source systems, as well as streaming-specialized designs: Gemini-2.5-Pro [4], GPT-4o [11], Qwen2.5-Omni [20], Videochat-Online [10], Qwen2.5-VL [1], and MiniCPM-o [19]. All models receive synchronized video, audio, and text inputs unless stated otherwise. For models lacking built-in speech recognition, we attach a Whisper-style ASR component [17] to transcribe audio.

4.2. Implementation Details

Setup and Streaming. All experiments are conducted on a single NVIDIA H200 GPU (140 GB). At inference time, for each role we order all videos by their start timestamps and concatenate them into one continuous stream, mirroring realistic personal-assistant usage where the assistant is

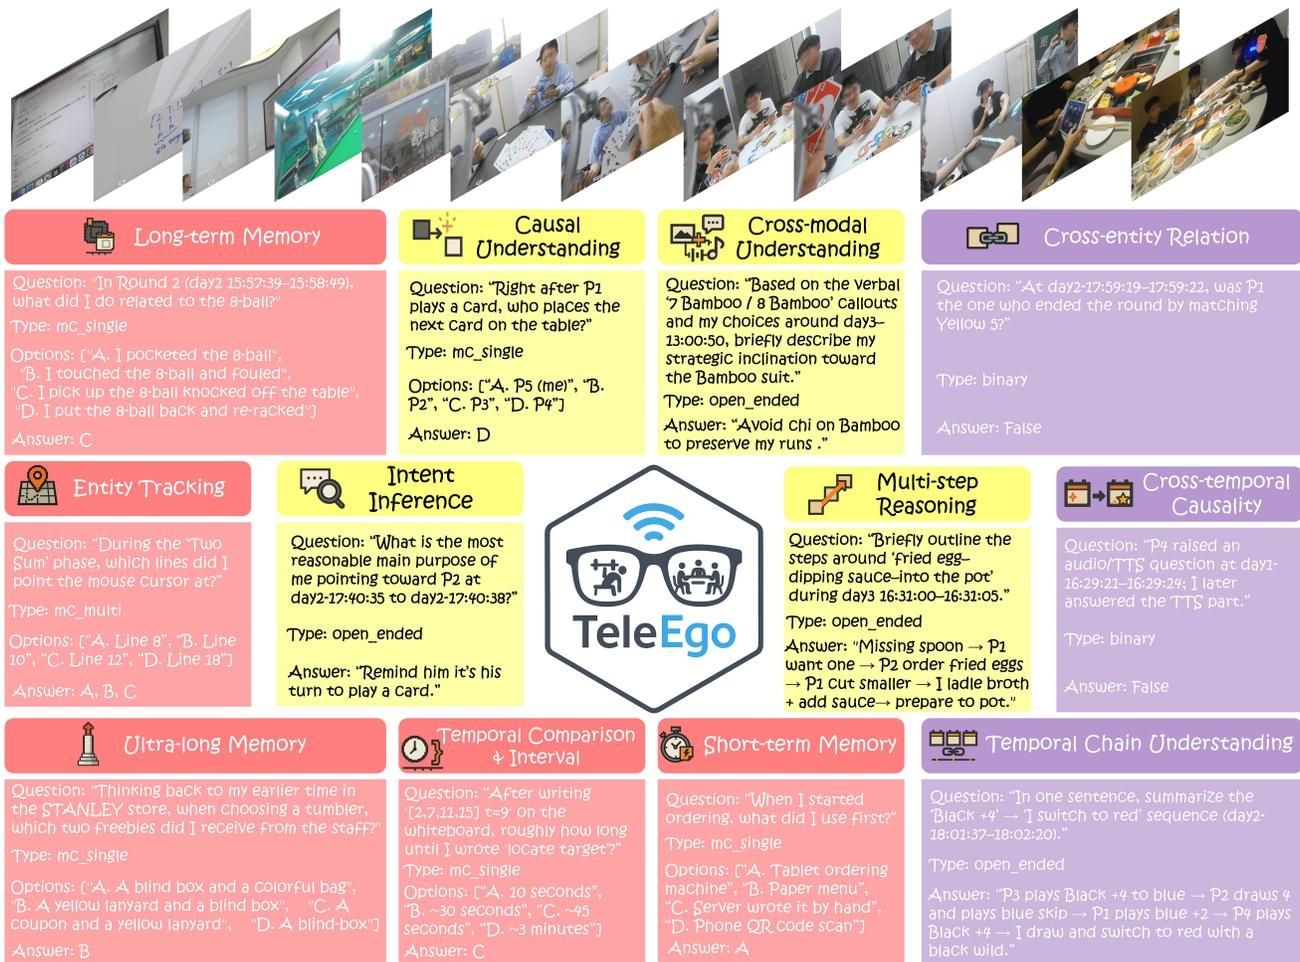


Figure 5. TeleEgo benchmark comprises twelve QA subcategories, illustrated here with one example per subcategory.

Overview Panels

QA Category			
Type	Count	Share (%)	
Memory	1937	58.8	
Understanding	897	27.3	
Cross-Memory Reasoning	457	13.9	

QA Type			
Type	Count	Share (%)	
mc_single	1743	53.0	
binary	614	18.6	
mc_multi	493	15.0	
open_ended	441	13.4	

Subcategory Distribution

Subcategory	Count	Share (%)	
Ultra-long Memory	722	21.9	
Short-term Memory	414	12.6	
Entity Tracking	289	8.8	
Temporal Comparison & Interval	259	7.9	
Long-term Memory	253	7.7	
Intent Inference	238	7.2	
Causal Understanding	225	6.9	
Cross-modal Understanding	219	6.8	
Multi-step Reasoning	215	6.5	
Cross-entity Relation	159	4.8	
Temporal Chain Understanding	152	4.6	
Cross-temporal Causality	146	4.4	

Table 2. Statistics of TeleEgo benchmark. Left: Overview of task categories and QA types, showing a balanced mixture across Memory, Understanding, and Cross-Memory Reasoning dimensions. Right: Subcategory-level distribution over twelve cognitive tasks, encompassing a total of 3,291 QA instances. The benchmark spans short- to ultra-long memory, causal and intent reasoning, and cross-temporal integration, providing a comprehensive foundation for evaluating multimodal and embodied intelligence models.

Method	Params	Omni	Streaming	Memory (%)						Understanding (%)					Cross-Memory Reasoning (%)				Overall
				UIM	StM	ET	TCI	LtM	All	II	CU	CmU	MsR	All	CeR	TCU	CtC	All	
GPT-4o [11]	-	✓	-	42.31	40.58	31.91	47.37	52.78	42.18	81.81	81.58	45.71	50.00	66.67	44.44	40.00	45.00	44.23	48.94
Gemini-2.5-Pro [4]	-	✓	-	49.04	45.59	34.04	47.37	44.44	45.05	63.64	55.26	37.14	28.57	48.03	40.74	40.00	45.00	42.31	45.55

Table 3. **RTA results of proprietary MLLMs on TeleEgo (single-video evaluation).** Columns are grouped into three capability blocks: Memory, Understanding, and Cross-Memory Reasoning, with an All column summarizing each block and an Overall column aggregating across blocks. “Omni” denotes integrated audio-video-text perception; “Streaming” denotes native support for streaming interaction. A ✓ indicates the capability is supported, ✗ indicates not supported, and “-” indicates unknown. Due to the extensive video duration (~14 hours) and API latency, both models were evaluated on the same participant’s video for fair comparison.

Method	Params	Omni	Streaming	Memory (%)						Understanding (%)					Cross-Memory Reasoning (%)				Overall
				UIM	StM	ET	TCI	LtM	All	II	CU	CmU	MsR	All	CeR	TCU	CtC	All	
VideoChat-Online [10]	4B	✗		30.74	26.29	19.35	29.84	22.78	26.69	57.35	44.13	35.44	27.89	42.34	18.79	32.00	41.60	29.43	31.28
Qwen2.5-VL [1]	8B	✗	✗	45.42	42.01	31.18	37.90	33.33	39.66	66.35	53.99	47.57	41.50	53.28	32.89	48.00	52.00	42.14	43.67
Qwen2.5-Omni [20]	7B	✓		44.21	42.75	35.48	41.13	37.97	41.20	72.51	61.50	50.00	48.98	59.07	37.58	56.00	61.60	49.16	46.96

Table 4. **RTA results of pseudo-streaming open-source MLLMs on TeleEgo.** These models lack cross-call memory mechanisms and process each input unit independently. They are categorized as “pseudo-streaming” because they lack native support for streaming interaction. To test the models’ intrinsic memory capabilities, they are evaluated under a strict streaming protocol where only the current frame and audio are provided at each second.

Method	Params	Omni	Streaming	Session Length	Memory (%)						Understanding (%)					Cross-Memory Reasoning (%)				Overall
					UIM	StM	ET	TCI	LtM	All	II	CU	CmU	MsR	All	CeR	TCU	CtC	All	
MiniCPM-o [19]	8B	✓	✓	1sec	51.99	51.60	35.13	47.58	44.30	47.54	75.83	61.50	51.94	44.22	59.59	28.19	48.00	64.80	45.15	50.60
				1min	50.60	55.53	37.28	53.63	51.05	50.11	80.57	69.48	56.80	51.02	65.64	33.56	36.00	66.40	47.49	54.10
				3min	48.19	52.58	37.28	49.19	45.99	47.31	73.93	63.38	53.88	40.14	59.33	32.89	32.00	66.40	46.82	50.57
				5min	43.00	44.23	33.69	47.58	37.55	41.71	63.51	56.34	44.17	43.54	52.64	23.49	32.00	56.80	38.13	44.34
				10min	28.50	28.01	23.66	30.65	26.16	27.60	41.71	40.38	34.95	29.93	37.32	15.44	36.00	37.60	26.42	30.15

Table 5. **RTA results of MiniCPM-o with varying session lengths on TeleEgo.** Session length determines how long the model retains memory of previously observed frames before reset. “1s” corresponds to no memory (each call is independent). Longer sessions allow more historical context but may cause attention dilution and generation quality degradation.

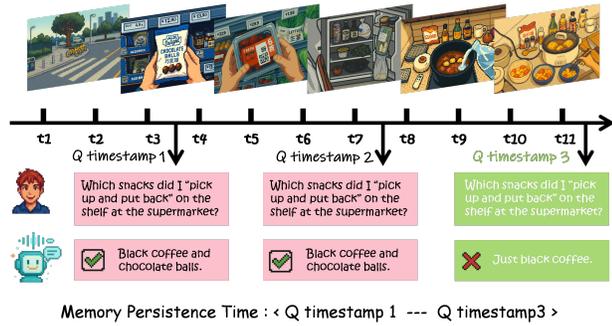


Figure 6. Illustration of MPT test pipeline. MPT is defined as the duration from the first correct response time to the first failure.

invoked intermittently around task- or context-specific segments rather than running continuously.

Scheduling and alignment. For each QA item, we take its evidence end time as the question timestamp (Qtimestamp) used for scheduling. Qtimestamps are rounded up to the nearest second. At test time, we pre-sort QA items by Qtimestamp and emit each item when the stream time first reaches its Qtimestamp; ties are resolved by source order,

and emissions are confined to the item’s decision window.

MPT implementation. For each item correctly answered at time t^* , we schedule up to ten recall evaluations at $t^* + r\Delta$ ($\Delta = 60$ s; $r = 1, \dots, 10$). At each evaluation, the original evidence is not replayed; only the ongoing stream is available. If an item fails an evaluation, it is removed from subsequent rounds; its horizon is the elapsed time from t^* to the first failed evaluation.

RTA Evaluation Metrics. Multiple choice questions (mc_single and mc_multi) are evaluated by exact match on option letters. Binary questions are evaluated by boolean equivalence. For open-ended questions we report an LLM-judge score (0–5) produced by GPT-4o.

4.3. Main Results

We group our RTA analysis into three model categories, each under a distinct evaluation setting.

Proprietary API Models. GPT-4o and Gemini-2.5-Pro are evaluated via API calls, with their internal implementations opaque. Due to the long video duration (~14 hours per participant) and API latency, we report RTA results

on a single participant (Table 3). GPT-4o shows strong Understanding performance (66.67%), especially on Intent Inference (81.81%) and Causal Understanding (81.58%). This strength aligns with the nature of Understanding tasks, which mainly require semantic reasoning and common-sense knowledge rather than precise temporal localization, allowing the model to exploit its general-purpose priors. In contrast, performance drops on Memory (42.18%) and Cross-Memory Reasoning (44.23%), where fine-grained temporal binding and evidence grounding are crucial.

Pseudo-Streaming Open-Source Models. VideoChat-Online, Qwen2.5-VL, and Qwen2.5-Omni lack genuine cross-call memory mechanisms (Table 4). VideoChat-Online’s Pyramid Memory Bank exists only within a single inference call and is reset at the end of each call. Qwen2.5-VL and Qwen2.5-Omni support multi-turn dialogue by naively accumulating raw conversation history, but provide no explicit streaming interface or memory management. In our evaluation, we deliberately avoid feeding uncompressed historical context into the models, since this merely shifts the memory burden to the context window rather than testing intrinsic memory capability. Instead, we adopt a strict streaming protocol: at each timestep, the model receives only the current frame and audio. Under this setting, Understanding tasks remain manageable, while Memory and Cross-Memory Reasoning become much harder. Among the three, Qwen2.5-Omni achieves the highest overall RTA (46.96%), primarily because it has stronger base model and avoids the ASR transcription errors.

Streaming Model: MiniCPM-o. MiniCPM-o is the only model in our study with an explicit streaming interface that accepts chunk-wise incremental input (Table 5). However, it lacks dynamic memory management: the KV cache grows without compression, eviction, or sliding window. This causes two issues in practice: GPU memory overflow as video length increases, and attention dilution that degrades generation quality in long sessions. We therefore reset the session at fixed intervals. RTA results indicate an optimal session length of about 1 minute (54.10% overall), with poorer performance at both shorter (50.60% at 1s) and longer intervals (30.15% at 10min). At longer intervals, the model often produces incoherent or off-topic responses, confirming that naive KV cache accumulation without memory management is insufficient for sustained streaming. This highlights an open challenge: efficiently managing long-term memory under bounded computational resources remains unresolved for streaming MLLMs.

4.4. MPT in Current Models

We proposed Memory Persistence Time (MPT) to measure how long a model retains information after initially answer-

ing correctly. However, the current evaluation does not include MPT results, as none of the tested models are suitable for this metric. For pseudo-streaming models (VideoChat-Online, Qwen2.5-VL, Qwen2.5-Omni), each inference call is independent with no cross-call memory. In this setting, MPT is ill-defined: if a memoryless model answers correctly in both the initial and recall phases, we cannot tell whether it “remembered” the information or simply “guessed correctly twice”. For MiniCPM-o, although the KV cache can persist across calls, the lack of dynamic memory management forces us to reset sessions every few minutes. When an initial test and its recall test span a session reset, the model’s memory is cleared in between. Any recall failure may then result from genuine forgetting or from the reset itself, making MPT measurements unreliable.

MPT is therefore a forward-looking metric that aligns with TeleEgo’s goal of evaluating realistic long-term egocentric assistants. At present, models with true long-horizon streaming memory are still rare: few can run continuously for hours without resets while maintaining stable generation quality. As stronger streaming MLLMs emerge, MPT will become central to measuring how long information is retained beyond the moment of first use. TeleEgo, together with our released MPT evaluation framework and code, provides a unified setting to assess both real-time correctness and long-term memory persistence, paving the way for future work on practical egocentric AI assistants.

5. Conclusion

We present TeleEgo, an online, omni-modal, first-person benchmark grounded in real-world use. It is built from continuous, multi-participant, multi-day recordings that align egocentric video, ambient speech/dialogue, and textual timelines under a unified global clock. A contract-based annotation scheme binds each query to its required modalities and to precise, time-stamped evidence spans, enabling auditable attribution. Centered on three capability axes—Memory, Understanding, and Cross-Memory Reasoning—TeleEgo employs a strict streaming-only evaluation protocol where models consume inputs once and responses must arrive within task-specific decision windows. We employ Real-Time Accuracy as the main metric in our experiments, capturing both correctness and timing under online decision constraints. In addition, we introduce Memory Persistence Time as a forward-looking protocol for probing long-horizon retention in continuous streams, and release an accompanying evaluation framework for future systems with stronger long-term streaming memory, along with failure breakdowns over retention, retrieval, alignment, and timing. We envision TeleEgo as an ecologically valid foundation for developing first-person assistants that must remember, align, and act in real time.

Acknowledgments. We thank the AI Glasses product team at China Telecom for their enthusiastic support and valuable assistance. We are grateful to Professor Haiwei Wu’s research group at the University of Electronic Science and Technology of China for their great help in data annotation, validation, and the construction of QA pairs.

References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 6, 8
- [2] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024. 3
- [3] Sijie Cheng, Zhicheng Guo, Jingwen Wu, Kechen Fang, Peng Li, Huaping Liu, and Yang Liu. Egothink: Evaluating first-person perspective thinking capability of vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14291–14302, 2024. 2, 3
- [4] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasapat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blisstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 3, 6, 8
- [5] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pages 720–736, 2018. 3
- [6] Ahmad Darkhalil, Dandan Shan, Bin Zhu, Jian Ma, Amlan Kar, Richard Higgins, Sanja Fidler, David Fouhey, and Dima Damen. Epic-kitchens visor benchmark: Video segmentations and object relations. *Advances in Neural Information Processing Systems*, 35:13745–13758, 2022. 3
- [7] 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
- [8] 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. 3
- [9] Yifei Huang, Guo Chen, Jilan Xu, Mingfang Zhang, Lijin Yang, Baoqi Pei, Hongjie Zhang, Lu Dong, Yali Wang, Limin Wang, et al. Egoexolearn: A dataset for bridging asynchronous ego-and exo-centric view of procedural activities in real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22072–22086, 2024. 3
- [10] Zhenpeng Huang, Xinhao Li, Jiaqi Li, Jing Wang, Xiangyu Zeng, Cheng Liang, Tao Wu, Xi Chen, Liang Li, and Limin Wang. Online video understanding: Ovbench and videochat-online. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3328–3338, 2025. 1, 3, 6, 8
- [11] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 3, 6, 8
- [12] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. 3
- [13] Yadong Li, Jun Liu, Tao Zhang, Song Chen, Tianpeng Li, Zehuan Li, Lijun Liu, Lingfeng Ming, Guosheng Dong, Da Pan, et al. Baichuan-omni-1.5 technical report. *arXiv preprint arXiv:2501.15368*, 2025. 3
- [14] Junming Lin, Zheng Fang, Chi Chen, Zihao Wan, Fuwen Luo, Peng Li, Yang Liu, and Maosong Sun. Streamingbench: Assessing the gap for mllms to achieve streaming video understanding. *arXiv preprint arXiv:2411.03628*, 2024. 1, 3
- [15] Lin Long, Yichen He, Wentao Ye, Yiyuan Pan, Yuan Lin, Hang Li, Junbo Zhao, and Wei Li. Seeing, listening, remembering, and reasoning: A multimodal agent with long-term memory. *arXiv preprint arXiv:2508.09736*, 2025. 3
- [16] Junbo Niu, Yifei Li, Ziyang Miao, Chunjiang Ge, Yuanhang Zhou, Qihao He, Xiaoyi Dong, Haodong Duan, Shuangrui Ding, Rui Qian, et al. Ovo-bench: How far is your video-llms from real-world online video understanding? In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 18902–18913, 2025. 1, 3
- [17] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR, 2023. 6
- [18] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024. 3
- [19] OpenBMB MiniCPM-o Team. Minicpm-o 2.6: A gpt-4o level mllm for vision, speech, and multimodal live streaming on your phone, 2025. 3, 6, 8
- [20] Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, et al. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*, 2025. 3, 6, 8
- [21] Jingfang Yang, Shuai Liu, Hongming Guo, Yuhao Dong, Xi-amengwei Zhang, Sicheng Zhang, Pengyun Wang, Zitang

- Zhou, Binzhu Xie, Ziyue Wang, et al. Egolife: Towards egocentric life assistant. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 28885–28900, 2025. [2](#), [3](#)
- [22] Hanrong Ye, Haotian Zhang, Erik Daxberger, Lin Chen, Zongyu Lin, Yanghao Li, Bowen Zhang, Haoxuan You, Dan Xu, Zhe Gan, et al. Mm-ego: Towards building egocentric multimodal llms for video qa. *arXiv preprint arXiv:2410.07177*, 2024. [3](#)
- [23] Xiangyu Zeng, Kefan Qiu, Qingyu Zhang, Xinhao Li, Jing Wang, Jiaxin Li, Ziang Yan, Kun Tian, Meng Tian, Xinhai Zhao, et al. Streamforest: Efficient online video understanding with persistent event memory. *arXiv preprint arXiv:2509.24871*, 2025. [1](#), [2](#), [3](#)
- [24] Haoji Zhang, Yiqin Wang, Yansong Tang, Yong Liu, Jiashi Feng, Jifeng Dai, and Xiaojie Jin. Flash-vstream: Memory-based real-time understanding for long video streams. *arXiv preprint arXiv:2406.08085*, 2024. [1](#), [3](#)
- [25] Sheng Zhou, Junbin Xiao, Qingyun Li, Yicong Li, Xun Yang, Dan Guo, Meng Wang, Tat-Seng Chua, and Angela Yao. Egotextvqa: Towards egocentric scene-text aware video question answering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3363–3373, 2025. [2](#), [3](#)
- [26] Wenqi Zhou, Kai Cao, Hao Zheng, Xinyi Zheng, Miao Liu, Per Ola Kristensson, Walterio Mayol-Cuevas, Fan Zhang, Weizhe Lin, and Junxiao Shen. X-lebench: A benchmark for extremely long egocentric video understanding. *arXiv preprint arXiv:2501.06835*, 2025. [1](#), [3](#)