

PARSE-Ego4D: Personal Action Recommendation Suggestions for Egocentric Videos

Steven Abreu^{1,2} Tiffany D. Do^{1,3} Karan Ahuja¹ Eric J. Gonzalez¹

Lee Payne¹ Daniel McDuff¹ Mar Gonzalez-Franco¹

¹Google ²University of Groningen ³University of Central Florida

s.abreu@rug.nl, tiffany.do@ucf.edu,

{karanahuja,ejgonz,leepayne,dmcduff,margon}@google.com




Figure 1: Examples of action suggestions for different videos in the PARSE-Ego4D dataset.

Abstract

Intelligent assistance involves not only understanding but also action. Existing ego-centric video datasets contain rich annotations of the videos, but not of actions that an intelligent assistant could perform in the moment. To address this gap, we release **PARSE-Ego4D**, a new set of personal action recommendation annotations for the Ego4D dataset. We take a multi-stage approach to generating and evaluating these annotations. First, we used a prompt-engineered large language model (LLM) to generate context-aware action suggestions and identified over 18,000 action suggestions. While these synthetic action suggestions are valuable, the inherent limitations of LLMs necessitate human evaluation. To ensure high-quality and user-centered recommendations, we conducted a large-scale human annotation study that provides grounding in human preferences for all of PARSE-Ego4D. We analyze the inter-rater agreement and evaluate subjective preferences of participants. Based on our synthetic dataset and complete human annotations, we propose several new tasks for action suggestions based on ego-centric videos. We encourage novel solutions that improve latency and energy requirements. The annotations in PARSE-Ego4D will support researchers and developers who are working on building action recommendation systems for augmented and virtual reality systems.

1 Introduction

Egocentric perception, the ability to capture and understanding of the world from a first-person perspective is gaining significant traction with the adoption of Augmented Reality (AR) and Head-Mounted Displays. Recent advancements in egocentric video understanding have opened new opportunities for research and application, including activity recognition [2, 30], object interaction analysis [11, 3, 52], and social interaction modeling [20]. However, a fundamental limitation of most existing systems is their reactive nature, driven by explicit user queries. We argue that the ability to take bespoke, proactive actions that anticipate a user’s needs is a core component of intelligent behavior without which these systems will be limited in their practical applications.

Public datasets have been highly consequential in the advancement of machine learning and artificial intelligence. However, older datasets, particularly in the field of computer vision, often included static, context agnostic, unimodal repositories of labeled data, *e.g.*, COCO [28] or Imagenet [41]. As ambitions in AI have become more complex and situated in the context of specific human-computer interaction scenarios, there has been a movement toward datasets that contain temporal, ecologically valid and multimodal data. This paradigm shift is exemplified in new datasets such as Ego4D [14] or Ego-Exo4D [15] which include thousands of hours of egocentric video streams. Several existing egocentric vision datasets provide rich annotations for tasks like activity recognition [7, 25, 6, 44, 10], object tracking [47], and for the analysis of interactions with other humans [40] and with the environment [5, 37]. These datasets play a crucial role in advancing research on egocentric perception. However, previous work focuses primarily on understanding and classifying video content. While valuable, such annotations don’t address how an intelligent system could suggest and take actions in the real or virtual world to assist the user. This ability to take appropriate action is a core component of intelligent behavior. Without this capability, systems can simply observe the world but have limited practical application as they rely on explicit user queries, as in existing work in visual question answering [12] and visual query localization [26]. The ability to generate bespoke or proactive actions, which could further our exploration of the environment, is currently missing.

To address this limitation and empower the development of proactive AI assistants, we release **PARSE-Ego4D**, a novel dataset designed to provide personal action recommendation annotations for egocentric videos. Herein, we consider personal suggestions that are context-dependent [13]. Our dataset is built upon the extensive Ego4D dataset [14], which contains 3,670 hours of first-person video recordings of a wide range of everyday activities. We leverage a two-stage annotation process, combining automated suggestions generated by a state-of-the-art large language model (Gemini Pro [48]) with meticulous human evaluation, to ensure the quality, relevance, and usefulness of the action recommendations. These annotations identify moments in the Ego4D video sequence when an assistant may be able to suggest a useful action (see more details in Section 3), creating a total of 18,360 possible action recommendations, which we call the *synthetic* dataset for it was created by an LLM and not yet grounded in human preferences. While the AI-assisted nature of these annotations allowed us to generate them at scale, the quality can be called into question. Consequently, we performed a large-scale human validation study that provides the necessary grounding in human preferences.

Using a 5-point Likert scale for human ratings, we found that 65% of all synthetically generated action suggestions were annotated with average scores above 3, and 42% were annotated with average scores above 4. Considering that our dataset aims at providing a footing to fine tune existing agents so they can provide better actions and personalized queries on-the-fly using real-time multi-modal data, the relatively high scoring validates our automatic captioning and annotation approach.

Our first study took 20 samples from our newly generated PARSE-Ego4D dataset and requested 20 human participants to evaluate our AI-generated queries and action suggestions with respect to five axes: (1) whether the query was *sensible* at all (to filter out hallucinations and mistakes from the LLM), (2) whether the suggestion would be helpful as an *implicit* suggestion if it was presented unsolicited to the user, (3) whether the action suggestion was *valuable* to the user (*e.g.*, by saving them time), (4) whether the suggested action was the *correct* action to take in response to the query, and (5) if the participant would personally be *likely* to take the presented action on their AR glasses (see Figure 4). In the large-scale annotation study, we requested 20% of the PARSE-Ego4D dataset to be annotated by 5 human raters, and the remaining 80% of the PARSE-Ego4D dataset to be annotated by 1 human rater. For the annotation study, we only evaluated the (1) sensibleness, (2) the helpfulness as an implicit action suggestion, and (3) the correctness of the action.

The current **PARSE-Ego4D** dataset aims at providing a basis for fine-tuning existing agents so they can provide better actions and queries on the fly using real-time multimodal data. Annotation, code and model responses can be found at: <https://parse-ego4d.github.io>.

2 Related Work

Human Computer Interaction Within the realm of Human-Computer Interaction (HCI), research on action recommendations has primarily focused on enhancing user experience and task efficiency [1]. Prior work has identified several key motivations for providing action suggestions in user interfaces (UIs): saving time by streamlining interactions [11, 52], improving discoverability of features and functionalities [46, 20], and enabling discrete interactions without explicit user input [49, 42] – an aspect that is particularly relevant for AR glasses.

Research on spatial UI transitions in AR has explored the balance between automation and user control in placing and manipulating UI elements [32], emphasizing the importance of user agency and control for a positive user experience. This underscores the need for easy error recovery mechanisms to mitigate the negative impact of incorrect predictions or actions. Explainability has emerged as a crucial aspect of action recommendations, particularly in the context of augmented reality (AR) systems. Xu et al. [55] introduced the XAIR framework, emphasizing the importance of providing clear and understandable explanations for AI-generated suggestions in AR environments. Their findings highlight that users prefer personalized explanations and that the timing, content, and modality of explanations should be carefully tailored to the user’s context and goals.

Machine Learning The increasing traction of egocentric devices through smart glasses, like Snap’s Spectacles [23] and Meta’s Ray-Ban Stories [36], and mixed reality head-mounted displays, like Apple’s Vision Pro [22] and Meta’s Quest [35], has spurred significant advancements in egocentric video [14] and user understanding [15, 45]. These devices provide a unique perspective on the user’s environment and activities, making them ideal platforms for personalized and context-aware AI assistants. The recent surge in multi-modal Large Language Models (M-LLMs) such as Gemini [48] and ChatGPT [39] has further propelled research in this area, particularly in the realm of visual perception and question answering. In the realm of egocentric video understanding, works like EgoOnly [51] have explored action detection without relying on exocentric (third-person) data, demonstrating the potential of understanding actions from a first-person perspective as a prerequisite for generating relevant action suggestions. Additionally, research in intent classification, such as IntentCapsNet [54], aims to discern user needs and preferences from egocentric videos, which can inform the generation of personalized suggestions.

Recent research has also focused on developing agents that can understand and execute instructions in interactive environments. In robotics, Instruct2Act [19] leverages LLMs to generate code that controls a robotic arm to manipulate objects based on multi-modal instructions. In UI interaction, approaches like CogAgent [17] have shown promising results in mapping natural language instructions to sequences of actions on mobile devices. Similarly, a plethora of LLM-based action agents are aiding in tasks such as knowledge discovery [38], web navigation [29], and shopping [56], among others.

Despite these advancements in understanding actions and executing instructions, there remains a gap in the development of proactive AI assistants for egocentric devices. Existing datasets like Ego4D [14] and EPIC-Kitchens [8] provide rich annotations for understanding activities and objects but do not offer a direct mapping to actionable recommendations.

The form factor and resource limitations of AR/VR devices, impose unique challenges on the machine learning models used in these systems. Energy efficiency, latency, and memory footprint are critical concerns for ensuring a positive user experience in these battery-powered and often mobile environments. Lightweight LLM models like Gemini XXS [48] are moving towards deployment on resource-constrained devices. Moreover, model compression techniques like quantization [21] have been applied to transformer architectures [50, 33] as well as pruning [34]. Furthermore, more efficient architectures are being developed that compete with transformers and offer better scaling with sequence length [4, 16, 9]. Model compression techniques and novel architectures for sequence modeling may provide a path towards efficient always-on foundation models running on resource-constrained AR/VR devices.

3 The PARSE-Ego4D Dataset

The PARSE-Ego4D dataset builds on top of the Ego4D dataset [14] and provides action suggestions that draw from the specification of available actions given in Section 3.2. After generating synthetic action suggestions using an LLM (Section 3.3), all action suggestions are rated through in a human annotation study (Section 3.4).

3.1 The Ego4D dataset

The Ego4D dataset is a massive ego-centric video dataset containing 3,670 hours of daily-life activity video from over 900 people across 74 locations and 9 countries. The data is split into $\approx 9,600$ videos with an average duration of 15-30 minutes and contains video streams from a head-mounted camera, as well as IMU and gaze data. The Ego4D dataset further contains rich annotations. All videos have dense written narrations in English for intervals of ≈ 10 seconds, as well as a summary for the whole video clip. Additionally, transcriptions, speech segmentation, user attention, speech target classification, speaker labeling, and episodic memory annotations are also provided for parts, or all, of the Ego4D dataset. We make use of the egocentric videos as well as the complete textual narrations from the Ego4D dataset.

Adding additional annotations and expanding the utility of such a dataset that already been collected is better than collecting a new dataset for two reasons. **(1)** It enables us to focus on the action suggestions without having to dedicate additional compute to labeling the narrations and captioning and labeling a whole new dataset. **(2)** Given the substantial investment made into this dataset, we can build on top of other projects that also have augmented the existing Ego4D [44, 47].

3.2 Available actions

To create a dataset with action suggestions, we first identify a set of possible actions that can be invoked from the AR/VR device, considering applications that future AR/VR devices are expected to support, such as:

- **Search:** an application that can take in the current camera input and a query (written or spoken) to run a multimodal search, and provide a written and/or spoken response.
- **Assistant search:** the AI assistant for the device, with access to system apps like “notes”, “timer”, “stopwatch”, “alarm”, “email”, “music”, “phone”, “contacts”, “messages”, “settings”, “calculator” and potentially more such as smart home access, notification access, and more.
- **Assistant local:** an application that can explicitly store memories and retrieve them later. Memories may be enrolled manually and explicitly, but they may also be enrolled passively and automatically as in the episodic memory tasks from the Ego4D dataset [14].
- **Language:** an application that can either transcribe what the user is hearing right now, translate what the user is reading or hearing, or determine what language is spoken.
- **Directions:** find relevant places nearby, plan routes, estimate distances and navigate to places.
- **Assistant guide:** an application that can give detailed and step-by-step instructions to the user.
- **Others:** For open-ended exploration, we also define the option to suggest actions that do not belong to the categories mentioned above. This may allow the LLM to come up with novel, creative use cases for AR glasses that are not covered by the available applications listed above. Actions that fall into this category are not included in the human annotation study.

3.3 Synthetic LLM annotation

In order to generate samples for action suggestions we used a prompt-engineered LLM, the Gemini Pro model [48]. We use prompt engineering for the LLM to use in-context learning to learn the annotation task. We pass textual narration sentences from the Ego4D annotations as input to the LLM, and request a JSON-formatted output in response. The process is illustrated in Figure 2. The system prompt to the LLM contains:

- **Task explanation:** the LLM is prompted to behave as a user experience researcher, helping to collect a dataset for useful interactions with AR glasses.

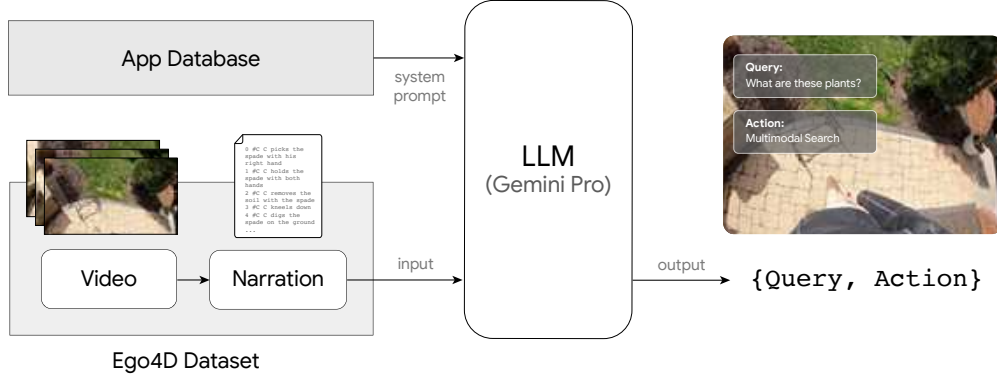


Figure 2: **PARSE-Ego4D** - We curated, annotated and open-source over 11,000 action suggestions for the Ego4D dataset. These annotations support researchers and developers who are working on building personalized action recommendation systems for augmented and virtual reality systems.

- **Input format:** the input format of the narrations is explained and an example is presented.
- **Available actions:** the set of available actions described in Section 3.2 is listed with example queries and the expected API format (this API format is not used for the annotation study).
- **Output format:** the expected JSON output format is described. The LLM is expected to return its thoughts to assess the situation and develop a rationale for the suggestion that it will return, the query that the user would ask along with the timestamp when this would be asked, and the corresponding action that the system should take in response to the query.

For every video clip in the Ego4D dataset, we split the entire video into batches of 200 narration sentences and pass these batches into the LLM. We drop 1897 short videos that have fewer than 50 sentences of narrations and do not generate any action suggestions for these. If the response of the LLM is not in valid JSON format, we ask the LLM to re-generate it to be valid. Once the LLM has generated a valid suggestion, we ask it to generate one more suggestion for the same input data. The complete system prompt is given in the Supplementary Materials.

The resulting dataset of synthetically generated action suggestions contains 32,155 action suggestions. After removing 7,491 duplicates (where the same batch of narrations gives the same query and action), we also remove 2,575 approximate duplicates. We classify a suggestion to be an approximate duplicate if it has an embedding distance $f(x_1, x_2) > 0.9$ using the normalized Gemini text embeddings from the Gemini API¹. This leaves 19,255 suggestions in our synthetic dataset, see Figure 3 (left).

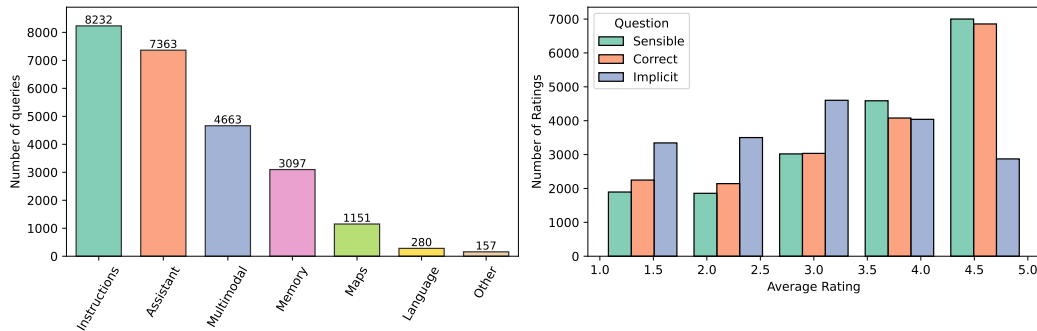


Figure 3: **Left:** Suggested actions by type. **Right:** Score distribution for different questions in the human annotation study, showing that there are more valid explicit suggestions than implicit suggestions.

¹ai.google.dev/gemini-api/docs/models/gemini#text-embedding



Query
e.g., "What am I cooking?"

Sensible
How appropriate/sensible is this user query?

Implicit
*How helpful would this suggestion be if the user **did not explicitly** ask the given query?*

Likely*
How likely would you personally be to ask this query to your AR glasses?

Action/App
e.g., Memory: can store memories and retrieve them later

Correct
Is the predicted action correct for this situation and user query?

Value*
How much value does the AI assistant on the AR/VR glasses add for the user in this scenario?

*included in personalization study only

Ego4D Video Clip
(presented with generated user query and suggested action for AR glasses)

Figure 4: Sketch of the survey that participants filled out in the human annotation study in order to verify the synthetically generated action suggestions in PARSE-Ego4D.

Every sample in the PARSE-Ego4D dataset contains a reference to the Ego4D video, a time range that corresponds to the narration sentence during which the action suggestion is invoked, the suggestion in the form of a (query, action) tuple, the name of the LLM that was used to generate the suggestion. Additionally, each sample also contains a parameter JSON that provides structured information that the suggested application may use. Furthermore, the dataset contains a rationale for each sample that was generated by the LLM as a form of chain-of-thought reasoning [53]. We do not include the action parameters or rationale in the human annotation study, but still provide them as part of the PARSE-Ego4D dataset.

3.4 Human annotation study

We annotate 20% of the synthetic action suggestion dataset gathered in Section 3.3 with 5 human raters which will be used as the test split. We annotate the remaining 80% of the dataset with 1 human rater each—of which 75% will be used as the train set and the other 5% as the validation set. In total, we received 36,171 annotations for 18,360 suggestions. The originally published benchmarks for the Ego4D dataset come with several different train/test/validation splits. However, these data splits are either based on subsets of the entire dataset, or based on specific scenarios, *e.g.*, hand-object interactions. As we are using the entirety of the Ego4D dataset, we chose a new random train/test/validation split.

The survey for participants of the annotation study is shown in Figure 4. In the large-scale annotation study, each sample is evaluated with three separate questions that each verify one dimension of the PARSE-Ego4D dataset. First, the sample is evaluated on being **sensible** to verify that the query makes sense in the given context. Second, query is being evaluated on being helpful as an **implicit** (or proactive) action suggestion. We expect that not all samples that score high on the **sensible** rating will also score highly on the **implicit** rating because we would expect users to have higher standards for implicit, proactive suggestions where false positives are disturbing or even annoying. Indeed, results from our annotation study confirm this, see Figure 3. Third, the action is evaluated for being **correct** given the query and context.

The release the PARSE-Ego4D dataset with all suggestions and their corresponding ratings from human annotators. For all downstream experiments, we filter the dataset to keep only suggestions that have (mean) ratings **sensible** ≥ 4 and **correct** ≥ 4 to use only verified, high-quality suggestions. If only the queries are used and actions are discarded, we suggest filtering for **sensible** ≥ 4 . For implicit, proactive suggestions we additionally filter for **implicit** ≥ 4 . Optionally, the cutoff for mean ratings can also be set at $\mu = 3$ instead of $\mu = 4$.

3.5 Subjective user study

In addition to providing annotations to verify and ground our synthetic action suggestions in human preferences, we ran two extended surveys for participants to assess their subjective preferences for different action suggestions. We ran one study with $N = 10$ participants and $M = 10$ samples, and

one study with $N = 20$ participants and $M = 20$ samples per participant. In these smaller subjective user studies, each participant is requested to answer all questions from the annotation survey shown in Figure 4. In addition to the questions outlined in the previous section, participants of the subjective user study were also asked to evaluate how likely they would personally be to ask the given query to their AR glasses, and how much value they think an AI assistant would add in the given scenario.

Table 1: Intraclass Correlation Coefficients (ICC) for the Annotation Questions.

Rating	ICC
Sensible	0.87
Helpful	0.73
Value	0.88
Likely	0.90
Correct	0.81

With these questions, we aim to better understand what kind of interactions different users value and to assess if there is a need for personalization in action recommendation systems based on our proposed action specification. Our results show that intraclass correlation coefficients (ICC) for the five annotation questions were above 0.7 for all questions and above 0.8 for all non-subjective questions from the study, thus showing high inter-rater agreement (see Table 1).

Although the ICC for the personal helpful question is lower than for other questions, the inter-rater agreement is still considerably high. We thus conclude that personalization may not be very important for building useful and valuable action recommendation systems of the sort that are described in this paper. However, we acknowledge that our user study was small and that the actions used in the annotations studies do not allow for the kind of personal data to be used that would be available to a real-world assistant on augmented and virtual reality

systems. We hypothesize that expanding the set of available actions and giving the AI assistant access to personal user data would strengthen the need for personalization in action suggestion systems.

3.6 Participants

Participants for both the subjective and annotation studies were recruited from Prolific, an online platform for crowdworkers, and were pre-screened for English fluency. For the larger subjective user study, we recruited 20 participants (10 male, 10 female) with an average age of 27.47 (SD=7.80). Participants were geographically diverse, residing in Poland (7), Portugal (6), Hungary (2), South Africa (2), Germany (1), Italy (1), Spain (1), and New Zealand (1).

The annotation study involved 1496 participants (749 male, 747 female), with an average age of 29.83 (SD=9.15). Figure 5 presents a demographic breakdown of our participants, including gender, race, age, and country of residence. Participants annotated up to 20 samples each and were compensated through Prolific with US\$0.13 per annotation for an average hourly wage of US\$8.79.

4 The PARSE-Ego4D Benchmark

We propose two tasks for action recommendation based on the PARSE-Ego4D dataset. Each task aims to build action recommendation systems either for (1) explicit user queries or (2) implicit user queries for proactive action suggestions, see Figure 1. Both tasks work towards building real-world action recommendation systems for augmented and virtual reality systems.

4.1 Task 1: Explicit Query-to-Action

Given the query from the PARSE-Ego4D dataset and the corresponding context from the Ego4D dataset, the task is to predict the action that the system should call on in order to handle the user query. The PARSE-Ego4D dataset provides human annotations for six kinds of actions, thus making this a classification task with $C = 6$ classes. Formally, the task is to approximate the function $f : (c, q) \mapsto a$ where $a \in \{1, \dots, C\}$ is the action, c is the context, and q is the text describing the user query. The task can be solved in language-only mode by using the textual narration from the Ego4D dataset, or in multimodal mode by using the raw videos from the Ego4D dataset. Thus, the context c can be either text (as the narrations), or a video stream, or a combination of multiple modalities. We report the baseline performance of this task with a prompt engineered Gemini Pro model used in zero-shot manner. The system prompt for this task is presented in the Supplementary Material.

Table 2: Baseline results for the PARSE-Ego4D benchmark tasks. All models use language-based input through the narrations provided by the Ego4D dataset annotations. The Gemini model is used in zero-shot manner. **Left:** shows the Explicit-Query-to-Action task where classification accuracy is reported. The constant model always predicts the most frequent action in the training dataset. **Right:** shows the Implicit Query-to-Action task where negative log likelihood is reported, see main text for explanation. The first random model randomly predicts one of the top-500 (query, action) pairs, the second random model chooses from all (query, action) pairs in the training dataset.

Model	Train	Val	Test
Gemini Pro	55.95%	54.43%	63.57%
Constant	42.75%	42.75%	42.75%

Model	Train	Val	Test
Gemini Pro	-43.43	-43.46	-42.50
Random (top)	-44.77	-45.07	-44.80
Random (all)	-53.68	-53.97	-53.39

4.2 Task 2: Implicit Query-to-Action

For increased autonomy of the AI assistant and easier interfacing for users of AR and VR systems, we further propose a new benchmark task to evaluate a system’s capability to make action suggestions without an explicit user query. Instead, the model only receives a signal of intent from the user, for example the press of an action button or the invocation of a hot word without an explicit query that specifies the user’s intent. The present dataset inherently contains such intent signals - which are the timesteps in the Ego4D data for which the PARSE-Ego4D dataset contains verified sensible suggestions.

The input is the context at a given point in time from the Ego4D dataset where the time is taken from the PARSE-Ego4D dataset, filtered as described in Section 3.4. As with the previous task, the context can be ingested either in language form, as narrations from the Ego4D dataset annotations, or in raw video form. We present baseline results for the language-based narrations as input. The output for this task can be an action suggestion, as shown in Figure 1. However, it is evident that all necessary information about such an action suggestion is also contained in the (query, action) pair that is provided in the PARSE-Ego4D dataset. As such, we propose to solve this task by learning the function $f : c \mapsto (q, a)$ where c is the context from the Ego4D dataset, and (q, a) is the (query, action) tuple. As this is an open-ended task with the final output being in natural language, we propose the use of the negative log-likelihood of the language model’s output on the (query, action) pair from the PARSE-Ego4D dataset, given the Ego4D context as input. We report the performance of a baseline LLM model on text-based narrations as context input, and provide two naive baseline methods for comparison, see Table 2. The system prompt for the LLM on this task is presented in the Supplementary Material.

5 Discussion and Limitations

Context only as textual narrations We generated the presented dataset based only on textual narrations from the Ego4D dataset that were provided by human annotators. Using a the few-shot learning ability of foundation models would, at the present time, be too computationally expensive on video data directly. However, it is conceivable to pass one, or a few, images from the video stream into the model, along with the textual narrations. It may also be advantageous to train a video-to-text model directly or fine-tune an existing model using our proposed dataset. Experiments using multimodal LLMs on our proposed benchmark tasks remain to be explored.

Efficient ML systems Our proposed experimental baselines use either naive methods or a state-of-the-art LLM that is too large to be deployed on AR/VR devices. We encourage future work to explore the tradeoffs between performance on the proposed tasks and the efficiency of the suggestion model. Novel efficient architectures for sequence modeling [4, 9, 16] may provide a path towards efficient AI assistants running on-device in resource-constrained environments such as those faced by AR/VR systems.

Moving beyond human annotations Despite in our approach we use LLMs to create the dataset through prompt engineering on the narration of videos, we still require a certain level of human

annotation to evaluate the quality of the dataset. This is inline with current recommendations that test the limits of how far can synthetic user experiences go [27]. It remains to be explored if new advances in self-training LLMs based on automated scalar feedback [43] or self-consistency [18] can be applied to our dataset to improve the performance of LLMs on our proposed tasks.

Multi-turn suggestions and bespoke UI The development of personalized action recommendation systems in egocentric video presents a unique challenge in the design of user interfaces (UI). Traditional Assistants rely on queries by user, often optimized for general use, may not be suitable for presenting contextually relevant suggestions unless users start doing multi-turn interactions. This necessitates the exploration of shortcuts and bespoke UI designs that can seamlessly integrate with the user’s context. In our research we propose implicit queries that can actually reduce the number of multi-turn queries or UI interactions needed.

Advanced LLM reasoning techniques. The creation of our PARSE-Ego4D dataset aligns with and could benefit from advancements in Large Language Model (LLM) reasoning techniques, specifically Chain-of-Thought (CoT) [53], Tree-of-Thought (ToT) [31, 57], and self-reflection [24]. These techniques hold the potential to enhance both the generation and evaluation of action suggestions, moving us closer to truly personalized AI assistants. CoT prompting encourages LLMs to generate intermediate reasoning steps before reaching a conclusion. This approach can be applied to action suggestions by prompting the LLM to explicitly consider the user’s context, goals, and preferences before recommending an action. For example, instead of directly suggesting “Turn on the lights”, the LLM might first reason about the time of day, the user’s location, and their recent activities. This could lead to more nuanced suggestions like, “It’s getting dark in the kitchen, would you like me to turn on the lights?”. ToT extends CoT by allowing the LLM to explore multiple reasoning paths in parallel. This could be beneficial for generating a wider range of action suggestions and evaluating their potential impact on the user. For instance, the LLM could consider different options for completing a task, weigh their pros and cons, and present the most suitable one to the user. Self-reflection enables LLMs to evaluate their own outputs and identify potential errors or biases. In the context of action suggestions, this could involve the LLM assessing the confidence of its recommendations and providing explanations to the user. This could increase user trust and allow them to understand the reasoning behind the suggestions. As LLM reasoning techniques advance, they will also open up new research avenues, such as developing LLM-based agents that can learn user preferences and adapt their suggestions over time. CoT and ToT prompting could be used in real-time to refine the LLM-generated action suggestions in PARSE-Ego4D, making them more contextually relevant, time bonded, and personalized.

6 Conclusion and Broader Impacts

In this work, we have introduced PARSE-Ego4D, a novel dataset that expands upon the existing Ego4D dataset by incorporating context-aware personal action recommendation annotations. By leveraging a two-stage annotation process combining automated suggestions from a large language model (Gemini Pro) and human evaluation, we have ensured the quality, relevance, and usefulness of these recommendations. Our comprehensive human evaluation not only validates the efficacy of the LLM-generated suggestions but also reveals insights into the nuances of user preferences in real-world scenarios, for example proposing a difference between implicit and explicit types of queries. Through this dataset, we aim to empower researchers and developers to build intelligent assistants capable of anticipating user needs and proactively offering personalized action suggestions, ultimately enhancing the user experience in egocentric video applications.

Our dataset also is free of personally identifiable information and given the very tailored prompt engineering eliminates the appearance of offensive content. Both aspects are also further enhanced by the reliance on the original Ego4D dataset. The annotations in PARSE-Ego4D will support future research on a variety of tasks, such as intent to action mapping, personalized suggestion learning, and user modeling. We believe that the release of this dataset will significantly advance the field of proactive AI assistance in egocentric video and contribute to the development of more intelligent and intuitive user experiences.

References

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. Guidelines for human-ai interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*, pages 1–13, 2019.
- [2] Dan Bohus, Sean Andrist, Ashley Feniello, Nick Saw, Mihai Jalobeanu, Patrick Sweeney, Anne Loomis Thompson, and Eric Horvitz. Platform for situated intelligence. *arXiv preprint arXiv:2103.15975*, 2021.
- [3] Dan Bohus, Sean Andrist, Nick Saw, Ann Paradiso, Ishani Chakraborty, and Mahdi Rad. Sigma: An open-source interactive system for mixed-reality task assistance research. *arXiv preprint arXiv:2405.13035*, 2024.
- [4] Aleksandar Botev, Soham De, Samuel L Smith, Anushan Fernando, George-Cristian Muraru, Ruba Haroun, Leonard Berrada, Razvan Pascanu, Pier Giuseppe Sessa, Robert Dadashi, et al. Recurrentgemma: Moving past transformers for efficient open language models. *arXiv preprint arXiv:2404.07839*, 2024.
- [5] Matthew Chang, Aditya Prakash, and Saurabh Gupta. Look ma, no hands! agent-environment factorization of egocentric videos. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 21466–21486. Curran Associates, Inc., 2023.
- [6] Dibyadip Chatterjee, Fadime Sener, Shugao Ma, and Angela Yao. Opening the vocabulary of egocentric actions. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 33174–33187. Curran Associates, Inc., 2023.
- [7] Jihoon Chung, Yu Wu, and Olga Russakovsky. Enabling detailed action recognition evaluation through video dataset augmentation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 39020–39033. Curran Associates, Inc., 2022.
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2020.
- [9] Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024.
- [10] Joseph DelPreto, Chao Liu, Yiyue Luo, Michael Foshey, Yunzhu Li, Antonio Torralba, Wojciech Matusik, and Daniela Rus. Actionsense: A multimodal dataset and recording framework for human activities using wearable sensors in a kitchen environment. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 13800–13813. Curran Associates, Inc., 2022.
- [11] Mustafa Doga Dogan, Eric J Gonzalez, Andrea Colaco, Karan Ahuja, Ruofei Du, Johnny Lee, Mar Gonzalez-Franco, and David Kim. Augmented object intelligence: Making the analog world interactable with xr-objects. *arXiv preprint arXiv:2404.13274*, 2024.
- [12] Chenyou Fan. Egovqa - an egocentric video question answering benchmark dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- [13] Giuseppe Ghiani, Marco Manca, Fabio Paternò, and Carmen Santoro. Personalization of context-dependent applications through trigger-action rules. *ACM Trans. Comput.-Hum. Interact.*, 24(2), apr 2017.
- [14] 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.

- [15] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. *arXiv preprint arXiv:2311.18259*, 2023.
- [16] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.
- [17] Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazhen Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. Cogagent: A visual language model for gui agents. *arXiv preprint arXiv:2312.08914*, 2023.
- [18] Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve, 2022.
- [19] Siyuan Huang, Zhengkai Jiang, Hao Dong, Yu Qiao, Peng Gao, and Hongsheng Li. Instruct2act: Mapping multi-modality instructions to robotic actions with large language model. *arXiv preprint arXiv:2305.11176*, 2023.
- [20] 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. *arXiv preprint arXiv:2403.16182*, 2024.
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *Journal of Machine Learning Research*, 18(187):1–30, 2018.
- [22] Apple Inc. Apple vision pro. <https://www.apple.com/apple-vision-pro/>, 2023. Accessed: 2024-06-03.
- [23] Snap Inc. Snap spectacles. <https://spectacles.com/>, 2024. Accessed: 2024-06-03.
- [24] Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. Towards mitigating llm hallucination via self reflection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1827–1843, 2023.
- [25] Baoxiong Jia, Ting Lei, Song-Chun Zhu, and Siyuan Huang. Egotaskqa: Understanding human tasks in egocentric videos. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 3343–3360. Curran Associates, Inc., 2022.
- [26] Hanwen Jiang, Santhosh Kumar Ramakrishnan, and Kristen Grauman. Single-stage visual query localization in egocentric videos. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 24143–24157. Curran Associates, Inc., 2023.
- [27] Jie Li. How far can we go with synthetic user experience research? *Interactions*, 31(3):26–29, may 2024.
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [29] Jiarun Liu, Wentao Hu, and Chunhong Zhang. Alltogether: Investigating the efficacy of spliced prompt for web navigation using large language models. *arXiv preprint arXiv:2310.18331*, 2023.
- [30] XingyuBruce Liu, JiahaoNick Li, David Kim, Xiang’Anthony’ Chen, and Ruofei Du. Human I/O: Towards a Unified Approach to Detecting Situational Impairments. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI. ACM, 2024.

- [31] Jieyi Long. Large language model guided tree-of-thought. *arXiv preprint arXiv:2305.08291*, 2023.
- [32] Feiyu Lu and Yan Xu. Exploring spatial ui transition mechanisms with head-worn augmented reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2022.
- [33] Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Li Dong, Ruiping Wang, Jilong Xue, and Furu Wei. The era of 1-bit llms: All large language models are in 1.58 bits, 2024.
- [34] Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 21702–21720. Curran Associates, Inc., 2023.
- [35] Inc. Meta Platforms. Meta quest virtual reality headset. <https://www.meta.com/quest/>, 2020. Accessed: 2024-06-03.
- [36] Inc. Meta Platforms and EssilorLuxottica. Ray-ban stories smart glasses. <https://www.meta.com/smart-glasses/>, 2023. Accessed: 2024-06-03.
- [37] Tushar Nagarajan, Santhosh Kumar Ramakrishnan, Ruta Desai, James Hillis, and Kristen Grauman. Egoenv: Human-centric environment representations from egocentric video. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 60130–60143. Curran Associates, Inc., 2023.
- [38] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- [39] OpenAI. Chatgpt. <https://chat.openai.com/>, 2022. Accessed: 2024-06-03.
- [40] Chirag Raman, Jose Vargas Quiros, Stephanie Tan, Ashraful Islam, Ekin Gedik, and Hayley Hung. Conflab: A data collection concept, dataset, and benchmark for machine analysis of free-standing social interactions in the wild. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 23701–23715. Curran Associates, Inc., 2022.
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [42] Albrecht Schmidt. Implicit human computer interaction through context. *Personal technologies*, 4:191–199, 2000.
- [43] Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, et al. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*, 2023.
- [44] Yale Song, Eugene Byrne, Tushar Nagarajan, Huiyu Wang, Miguel Martin, and Lorenzo Torresani. Ego4d goal-step: Toward hierarchical understanding of procedural activities. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 38863–38886. Curran Associates, Inc., 2023.
- [45] Yale Song, Eugene Byrne, Tushar Nagarajan, Huiyu Wang, Miguel Martin, and Lorenzo Torresani. Ego4d goal-step: Toward hierarchical understanding of procedural activities. *Advances in Neural Information Processing Systems*, 36, 2024.
- [46] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. Sensecape: Enabling multilevel exploration and sensemaking with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18, 2023.

- [47] Hao Tang, Kevin J Liang, Kristen Grauman, Matt Feiszli, and Weiyao Wang. Egotracks: A long-term egocentric visual object tracking dataset. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 75716–75739. Curran Associates, Inc., 2023.
- [48] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [49] Daniel Vogel and Ravin Balakrishnan. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*, pages 137–146, 2004.
- [50] Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Huaijie Wang, Lingxiao Ma, Fan Yang, Ruiping Wang, Yi Wu, and Furu Wei. Bitnet: Scaling 1-bit transformers for large language models, 2023.
- [51] Huiyu Wang, Mitesh Kumar Singh, and Lorenzo Torresani. Ego-only: Egocentric action detection without exocentric transferring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5250–5261, 2023.
- [52] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Frujeri, et al. Holoassist: an egocentric human interaction dataset for interactive ai assistants in the real world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20270–20281, 2023.
- [53] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [54] Congying Xia, Chenwei Zhang, Xiaohui Yan, Yi Chang, and Philip S Yu. Zero-shot user intent detection via capsule neural networks. *arXiv preprint arXiv:1809.00385*, 2018.
- [55] Xuhai Xu, Anna Yu, Tanya R Jonker, Kashyap Todi, Feiyu Lu, Xun Qian, João Marcelo Evangelista Belo, Tianyi Wang, Michelle Li, Aran Mun, et al. Xair: A framework of explainable ai in augmented reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–30, 2023.
- [56] Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757, 2022.
- [57] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36, 2024.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) Limitations are discussed in Section 5.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) The broader impacts of our work are discussed in Section 7.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[N/A\]](#) No theoretical results are included.
 - (b) Did you include complete proofs of all theoretical results? [\[N/A\]](#) No theoretical results are included.
3. If you ran experiments (e.g. for benchmarks)...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#) Code, data and instructions for reproducing the experimental results can be found on our project page.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) Our work builds on and contributes to the Ego4D dataset which we have cited.
 - (b) Did you mention the license of the assets? [\[Yes\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[Yes\]](#) We have provided detailed information about the instructions and information given to annotators recruited for this work.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[Yes\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[Yes\]](#) Yes, we have included details of the instructions and screenshots of the annotation interface.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#) Risks associated with the annotation task were deemed very minimal.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[Yes\]](#) Details of the tasks and compensation are included in Section 3.6 "Participants".

A Appendix

A.1 Dataset Availability

The dataset is available on the PARSE-Ego4D GitHub repository on: <https://github.com/parse-ego4d/parse-ego4d.github.io/tree/main/dataset/>.

A.2 Human Annotation Demographics

A visualization of the demographics from our human annotation study is presented in Figure 5.

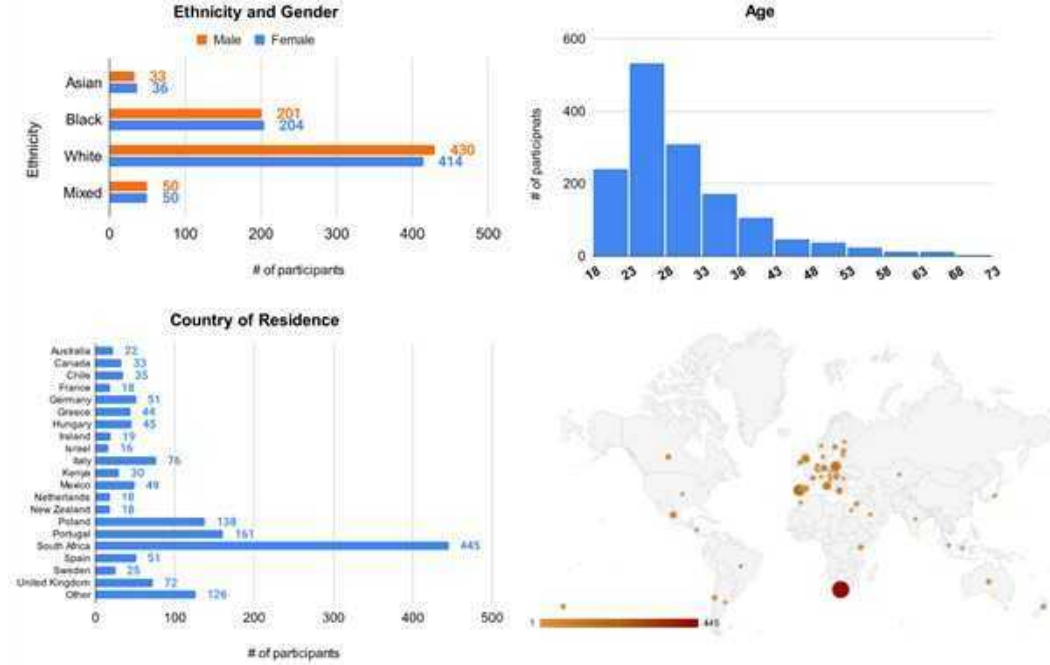


Figure 5: A demographic breakdown of our participants in the annotation study, including ethnicity, gender, and age. Countries with fewer than 15 participants are listed in "Other".

A.3 Annotation Metrics

The human annotations are used to filter the suggestions in PARSE-Ego4D so that samples above a certain mean rating for each question are accepted. Table 3 shows an overview of how many samples are accepted at different mean ratings.

A.4 Annotation Interface Screenshots

The human annotation study was run using Prolific, with participants filling out the survey on Qualtrics. The survey design is illustrated in Figure 4 and Figure 6 shows screenshots of the survey that human participants filled out.

Filter	Percentage	Number of Suggestions
All samples	100.00%	18,360
$\text{sensible} \geq 3$	78.10%	14,340
$\text{sensible} \geq 3.5$	63.10%	11,586
$\text{sensible} \geq 4$	58.31%	10,705
$\text{correct} \geq 3$	74.56%	13,689
$\text{correct} \geq 3.5$	59.54%	10,932
$\text{correct} \geq 4$	54.80%	10,061
$\text{implicit} \geq 3$	59.38%	10,903
$\text{implicit} \geq 3.5$	37.61%	6,905
$\text{implicit} \geq 4$	33.26%	6,107
$\{\text{sensible}, \text{correct}\} \geq 3$	65.00%	11,934
$\{\text{sensible}, \text{correct}\} \geq 3.5$	47.17%	8,660
$\{\text{sensible}, \text{correct}\} \geq 4$	42.32%	7,770
$\{\text{sensible}, \text{correct}, \text{implicit}\} \geq 3$	48.22%	8,854
$\{\text{sensible}, \text{correct}, \text{implicit}\} \geq 3.5$	27.65%	5,076
$\{\text{sensible}, \text{correct}, \text{implicit}\} \geq 4$	24.02%	4,410

Table 3: Number of suggestions in PARSE-Ego4D above a mean rating for different metrics. The filter $\{\text{sensible}, \text{correct}\}$ is applied for Task 1, whereas the $\{\text{sensible}, \text{correct}, \text{implicit}\}$ filter is applied for Task 2.



(a) Introductory instructions.



(b) Example video.



(c) Task video.



(d) Task questions.

Figure 6: Screenshots of the human annotation task.