
Learning to Plan with Personalized Preferences

Anonymous Authors¹

Abstract

Effective integration of Artificial Intelligence (AI) agents into daily life requires them to understand and adapt to individual human preferences, particularly in collaborative roles. Although recent studies on embodied intelligence have advanced significantly, they typically adopt generalized approaches that overlook personal *preferences in planning*. We address this limitation by developing agents that not only learn preferences from few demonstrations but also learn to adapt their planning strategies based on these preferences. Our research leverages the observation that preferences, though implicitly expressed through minimal demonstrations, can generalize across diverse planning scenarios. To systematically evaluate this hypothesis, we introduce Preference-based Planning (PBP) benchmark, an embodied benchmark featuring hundreds of diverse preferences spanning from atomic actions to complex sequences. We demonstrate that incorporating learned preferences as intermediate representations in planning significantly improves the agent’s ability to construct personalized plans. These findings establish preferences as a valuable abstraction layer for adaptive planning, opening new directions for research in preference-guided plan generation and execution.

1. Introduction

The field of embodied Artificial Intelligence (AI) is rapidly advancing, driven by significant progress in foundation models for vision and language (Bommasani et al., 2021; Peng et al., 2023; Achiam et al., 2023; Bai et al., 2023). These advances enable AI systems to autonomously collaborate with or assist humans in daily tasks, particularly in domestic settings (Driess et al., 2023; Leal et al., 2023; Zitkovich et al.,

2023; Ahn et al., 2024). However, recent approaches utilizing natural language instructions (Mu et al., 2023; Zitkovich et al., 2023; Singh et al., 2023) face fundamental limitations in capturing human preferences (Zhu et al., 2016). While natural language is our primary means of communication, its inherent ambiguity creates a gap between instructions and intended executions (Yuan et al., 2022; Jiang et al., 2022; 2021; Yuan et al., 2020). For instance, when a user requests help in preparing an apple, the agent needs to understand specific preferences about apple selection, washing requirements, cutting style, and container choice—details that vary significantly across individuals; see also Figure 1 for a graphical illustration.

Preference, central to personalization (Slovic, 1995), remains inadequately addressed in embodied Artificial Intelligence (AI). Integrating personalized preferences is crucial for tailoring agent actions to individual users, thereby enhancing the effectiveness and satisfaction of embodied assistants (Lee et al., 2012; Leyzberg et al., 2014). Moreover, preferences guide human-like decision-making and intelligent behavior. Psychological research emphasizes that understanding preferences is vital for interpreting human behaviors (Fawcett & Markson, 2010) and facilitating social interactions (Gerson et al., 2017; Liberman et al., 2021), suggesting that preference understanding could enable more grounded planning in embodied assistants.

Learning human preferences in real-world settings presents unique challenges (Peng et al., 2024). Humans typically communicate their needs succinctly (Levinson, 1983), without exhaustive preference details (Lichtenstein & Slovic, 2006), and many preferences include unconscious or instinctive elements difficult to articulate (Epstein, 1994; Simonson, 2008). A more practical approach is to infer preferences from observed human choices and decision-making patterns, as illustrated in Figure 1, where a robot assistant can learn users’ preferences and behavior habits from previous observations.

In this paper, we focus on developing agents capable of learning preferences from human behavior and subsequently planning actions guided by these learned preferences. While previous studies like NeatNet (Kapelyukh & Johns, 2022) and SAND (Yuan et al., 2023) have explored preference-based learning, they are limited to specific tasks (e.g., rear-

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

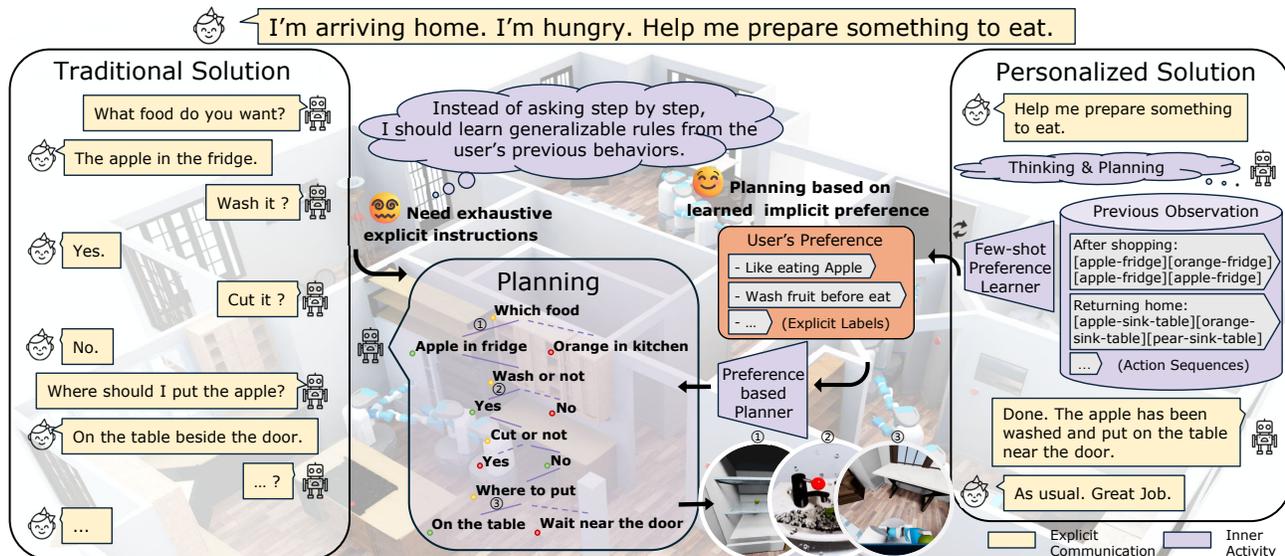


Figure 1: An example of preference-based planning in a food preparation scenario. When the assistant receives a natural language instruction for a food preparation task, it can follow one of two approaches: (Left, traditional methods) The assistant verifies details with the user at each step through exhaustive communication; or (Right, our personalized approach) it first learns from previous user action sequences to infer explicit preference labels and then generates a personalized plan based on the learned preferences. The planning tree (middle) illustrates how preferences guide the whole decision-making process across multiple dimensions. By learning preferences as a key intermediate representation from minimal human demonstrations, our approach enables AI agents to deliver personalized and adaptable assistance without explicit step-by-step instructions.

rangement) and fail to generalize across different situations. To address this limitation, we introduce Preference-based Planning (PBP), a comprehensive embodied benchmark built upon NVIDIA Omniverse and OmniGibson (Li et al., 2023). PBP provides realistic simulation and real-time rendering for thousands of daily activities across 50 scenes, featuring a parameterized vocabulary of 290 diverse preferences. These preferences span multiple levels, from specific action-level preferences (e.g., preferred glass type, water temperature) to task sequence-level preferences (e.g., task ordering, subtask prioritization).

Given the expensive nature of data collection (Akgun et al., 2012) and the few-shot nature of preference acquisition, we frame preference learning as a few-shot learning from demonstration task. In this framework, agents must respond to ambiguous instructions by formulating plans aligned with preferences demonstrated in limited example sequences. Specifically, an agent needs to analyze behavioral data, identify consistent patterns, and extrapolate these patterns to higher-level preference abstractions that can generalize across various tasks (Chao et al., 2011). Furthermore, when confronted with new tasks, the agent should leverage these learned preferences to generate adaptive action sequences that align with user preferences while maintaining task efficiency.

With the PBP benchmark developed, we challenge existing learning agents on their ability to learn human prefer-

ence and subsequently conduct preference-based planning. Our systematic evaluation of State-of-the-Art (SOTA) algorithms on PBP reveals that preferences serve as valuable abstractions of human behaviors, and their incorporation as intermediate planning steps significantly enhances agent adaptability. Through extensive experimentation, we demonstrate that symbol-based approaches show promise in scalability, yet significant challenges remain in both preference learning and planning. These challenges stem from the complexity of planning intricate activities and the nuanced nature of learning preferences through perception. Our analysis particularly highlights the difficulties in few-shot preference learning and preference-guided planning, establishing preferences as a crucial abstraction layer between high-level goals and low-level actions. We present this work as a foundation for addressing these challenges in preference-based embodied AI.

2. Related Work

2.1. Theoretical Foundations of Human Preferences

Preference theory originates from psychological research, where it describes predictable patterns in human behavior that can be modeled mathematically (Kahneman, 1982). These preferences reflect individual attitudes towards available choices in decision-making (Lichtenstein & Slovic, 2006) and operate both consciously and unconsciously to shape behavior (Coppin et al., 2010). A fundamental prin-

principle is that underlying preferences can be inferred from consistent behavioral patterns (Sen, 1973), enabling systematic analysis of decision-making processes. This framework has extended beyond psychology into economics, where Rational Choice Theory (Scott et al., 2000) models decision-making based on rational self-interest (Zey, 1998). Building on this, Utility Theory provides a mathematical foundation for modeling how preferences relate to attitudes toward rewards and risks (Mongin, 1997; Aleskerov et al., 2007). These theoretical foundations establish preferences as fundamental elements in shaping both individual behavior and broader societal dynamics. In recent years, these preference models have found new applications in artificial intelligence and robotics, particularly in developing human-centric AI assistants capable of understanding and adapting to individual user preferences.

2.2. Preference in Embodied Task Planning

The application of preferences in embodied task planning encompasses two distinct approaches. The first focuses on **general** preference-based planning, where robots leverage commonsense knowledge to execute universally accepted behavioral norms. Object rearrangement exemplifies this approach, with systems organizing items based on common occurrence patterns and spatial relationships (Taniguchi et al., 2021; Sarch et al., 2022). The second approach emphasizes **personalized preferences**, where embodied agents align their actions with individual user habits. This includes personalized object placement strategies (Abdo et al., 2015; Kapelyukh & Johns, 2022; Wu et al., 2023), preference-aware table setting (Puig et al., 2021b), and multi-agent coordination where agents maintain individual preferences while achieving optimal coordination (Shu & Tian, 2019).

Our work extends these approaches by considering preferences across diverse situations and scenes. Beyond spatial arrangements, we address temporal action sequences, state transitions during interactions, and few-shot preference learning. This comprehensive framework enables robust preference modeling and adaptation in real-world scenarios.

2.3. Embodied Assistants

The development of intelligent embodied assistants has evolved from basic Vision-and-Language Navigation (VLN) tasks (Anderson et al., 2018; Chen et al., 2019; Thomason et al., 2020) to complex interactive scenarios. ALFRED (Shridhar et al., 2020) introduced object manipulation, state tracking, and temporal dependencies between instructions, while platforms like Habitat (Savva et al., 2019; Puig et al., 2023b) and AI2-THOR (Kolve et al., 2017) emphasize active perception, long-term planning, and interactive learning in realistic environments. Recent research has shifted toward implicit-instruction scenarios, particularly in house-

keeping tasks (Kapelyukh & Johns, 2022; Kant et al., 2022; Sarch et al., 2022; Wu et al., 2023), where robots must reason about object arrangements without explicit directives. Works on proactive assistance (Patel & Chernova, 2023; Patel et al., 2023; Puig et al., 2023a) further explore anticipating temporal patterns in humans’ daily routines.

Methodologically, recent advances utilize Large Language Models (LLMs) as few-shot planners to generate language-based action sequences from limited demonstrations (Song et al., 2023; Driess et al., 2023; Ding et al., 2023; Zhang et al., 2024). Foundation Vision-Language Models (VLMs) have enhanced robotic systems’ perception and reasoning capabilities (Ahn et al., 2024; Leal et al., 2023; Gu et al., 2023; Brohan et al., 2022; Zitkovich et al., 2023; Xu et al., 2024a), enabling understanding of complex visual and linguistic inputs in everyday tasks. However, while these foundation models excel at reasoning from text or image information, their ability to learn individual preferences from limited demonstrations and plan adaptively remains an open challenge, particularly in multi-step tasks requiring personalized execution strategies.

3. Formulating Preference-based Planning

Tasks in PBP mirror real-world watch-and-help scenarios (Puig et al., 2021a), where an agent observes a few demonstrations of a user performing tasks that reveal preferences. The agent must then complete similar tasks in different setups while adhering to the demonstrated preferences.

Preference-based planning comprises two key components: few-shot **preference learning** of user preferences and subsequent **planning** guided by these learned preferences. Since humans, even infants, can naturally detect others’ preferences from limited decisions (Choi & Luo, 2023), and collecting extensive personal demonstrations is impractical in daily life, we formulate this as few-shot learning from demonstration.

Given a user with preference \mathbf{p} , the agent observes the user performing tasks from a first-person perspective, denoted as \mathbf{O} . These observations span multiple demonstrations. Formally, \mathbf{O} contains both state and action observations: $\mathbf{O} = \{(\mathcal{S}_i, \mathcal{A}_i, \mathcal{M})_N\}$, where \mathcal{S}_i denotes the egocentric observation sequence in the i -th demonstration, \mathcal{A}_i represents the action sequence, and \mathcal{M} optionally provides a bird’s-eye view of the entire scene map.

In the first stage, the objective is to learn the preference representation demonstrated through user actions:

$$\mathbf{p} = f(\mathbf{O}; \theta_f), \quad (1)$$

where \mathbf{p} denotes the learned preference representation here. It can either be a hidden representation or an explicit textual label, depending on the task settings.

The learned preference \mathbf{p} should then guide planning when the agent faces different setups with varying objects, room layouts, or entire scenes. Specifically, the agent optimizes:

$$\mathcal{L} = \sum_{i=1} \ell(g(s_i, f(\mathbf{O}; \theta_f); \theta_g), a_i), \quad (2)$$

where $g(\cdot)$ represents a potentially parameterized planning function that maps the current state and preference representation to the next action, and a_i denotes the ground-truth action demonstrating the user’s preference at the current stage.

4. The Preference-based Planning (PBP) Benchmark

Built on NVIDIA’s Omniverse and OmniGibson simulation environment (Li et al., 2023), our PBP benchmark enables realistic simulation of thousands of daily activities. It spans 50 distinct scenes and encodes 290 unique preferences, with a comprehensive test set of 5000 instances. Below, we detail the preference structure and test set construction.

4.1. Definition of Preferences

We organize preferences in a three-tiered hierarchical structure that captures varying degrees of specificity across tasks. Figure 2 provides an overview of all preferences and their distribution, while Figure 3 illustrates concrete examples of preferences and corresponding agent actions. The 290 preferences are distributed across three levels: 80 sequence-level, 135 option-level, and 75 action-level preferences.

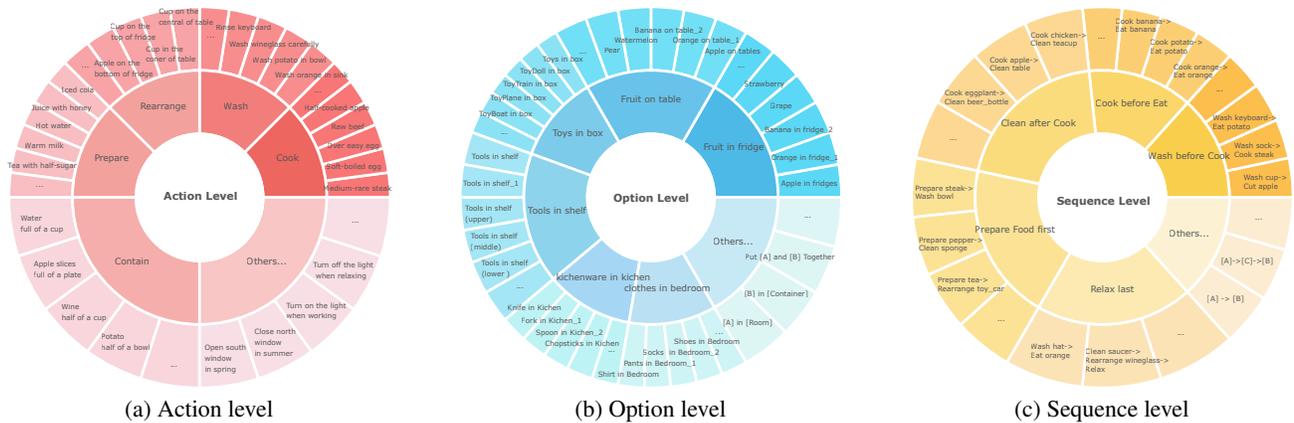


Figure 2: **Hierarchical organization of user preferences.** Our framework organizes preferences in a three-tiered structure, visualized through sunburst diagrams: (a) Action level captures fine-grained execution details within specific tasks, from quantity preferences in “Contain” (e.g., “half a cup” vs. “full cup”) to environmental controls (e.g., lighting and window operations). (b) Option level represents spatial preferences for object categories, encoding both storage decisions (e.g., table vs. fridge for fruits) and organizational choices (e.g., shelf levels and boxes for tools/toys). (c) Sequence level defines temporal relationships between tasks, encompassing both basic preparation sequences (e.g., “Prepare Food first”) and conditional orderings (e.g., “Clean after Cook,” “[A]->[B]”). Each diagram’s hierarchical structure branches from general categories to specific instances, revealing detailed preference patterns upon closer inspection. (Vector graphics; zoom in for details.)

Action Level These bottom-level preferences govern fine-grained execution details within specific sub-tasks, such as water quantity preferences when filling cups or shelf placement choices for books.

Option Level Middle-level preferences encode alternative approaches to sub-tasks. For instance, in “storing-nonperishable-food,” users may prefer cabinet storage versus table placement. These preferences can bind to different objects and may compose multiple action-level preferences.

Sequence Level Top-level preferences define task ordering and prioritization. They capture temporal dependencies between sub-tasks, such as cleaning furniture before rearranging kitchen utensils, followed by dinner preparation upon returning home.

4.2. Constructing PBP Test Set

Our PBP benchmark includes a default test set for systematic model evaluation. Following the formulation in Section 3, we structure PBP tasks as few-shot learning-from-demonstration problems. Each test point comprises several (typically three) unique demonstrations with egocentric observations of action sequences and their corresponding preference labels. As illustrated in Figure 4, a demonstration includes an egocentric video of agent activity, a bird’s-eye-view map tracking agent position, and frame-level action annotations. We also provide third-person view recordings for enhanced visualization. We prioritize the egocentric perspective for two reasons: 1) it offers a clear view with minimal occlusions, and 2) it aligns with human perception,

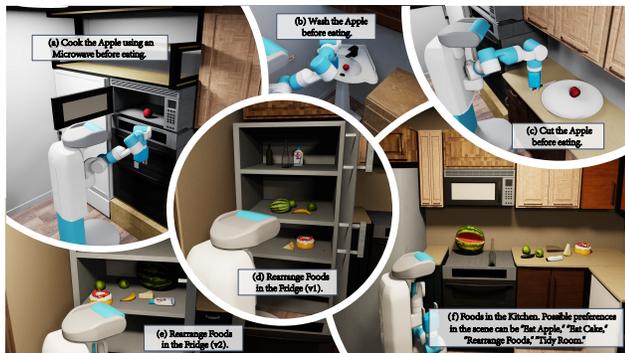


Figure 3: **Example of preferences and their corresponding actions in PBP.** At the primitive action level, we demonstrate preferences through basic tasks: (a) cooking using microwave, (b) washing in the sink, and (c) cutting into halves. At the option level, we showcase different approaches to object rearrangement, where users can prefer either (d) grouping objects by their categories (v1) or (e) placing them on the same layer of the fridge (v2). At the sequence level, we illustrate how preferences guide task ordering: (f) shows a user’s preference to have fruits first, followed by specific cleaning tasks.

facilitating transfer to real-world data from head-mounted devices.

The test set construction follows a two-stage process. First, we build a reusable and extensible demonstration pool. To generate each demonstration, we randomly assign a preference from our defined primitives to one of 50 OmniGibson scenes, then sample relevant objects within the chosen scene. We generate multi-perspective observations using rule-based planners for high-level planning and predefined scripts for low-level execution (*e.g.*, Inverse Kinematics (IK) for grasping, A* for movement).

Second, we construct test points by sampling preferences and retrieving relevant demonstrations from the pool. To reflect real-world few-shot scenarios, each preference pairs with demonstrations that share the same high-level preference but vary in scene settings or object selections. We also include unrelated demonstrations to prevent sampling bias.

The default test set contains 5,000 test points, drawing from a pool of 15,000 unique recordings. Unless specified otherwise, all experiments use this default set. The benchmark also supports custom test point generation through flexible demonstration sampling, preference definition, and third-person view video creation.

4.3. Models

Our evaluation focuses primarily on multimodal models that incorporate LLMs and demonstrate strong few-shot learning capabilities. The LLM component serves as a knowledge base that can enhance preference learning through commonsense reasoning. We also include symbol-based LLM

models for ablation studies to analyze how different modalities impact PBP performance. Most models evaluated can function in both end-to-end and two-stage pipeline configurations. See Appendix D for detailed implementations.

5. Experiments

5.1. Experimental Setup

We evaluate preference learning capabilities across two distinct settings: end-to-end and two-stage approaches. In the end-to-end setting, models directly map raw state inputs to action outputs. Leveraging models’ in-context learning abilities, we provide demonstrations alongside current state information as input and evaluate the generated action sequences against ground truth.

The two-stage setting introduces an intermediate step where models first learn to predict explicit preference labels during training. These predicted labels then serve as preference representations for subsequent planning stages. For black-box models, we employ carefully designed prompts rather than fine-tuning approaches.

All demonstration videos maintain consistent technical specifications across models and agents: egocentric perspective, 512×512 resolution, and 8 fps frame rate. Video duration matches the corresponding action sequence length. For LLM inference, we use conservative decoding parameters: temperature of 0.05, top-k of 1, and top-p of 0.05. All experiments run on a single machine with 8 NVIDIA A100 GPUs.

5.2. End-to-end Action Preference Learning

We first evaluate model performance in the end-to-end setting, where models generate actions directly from previous demonstrations and current state information. To quantify performance, we use Levenshtein distance to measure discrepancies between generated and ground truth action sequences, treating each individual action as a token.

As shown in Table 1 (the End-to-end row), video-based models produce Levenshtein distances approaching the average ground truth sequence lengths (15.80 at option level, 35.87 at sequence level). These high distances indicate that the models generate predominantly inconsistent action sequences, suggesting a failure to grasp preferences embedded in demonstration videos. While symbol-based models show modest improvements, their performance gains remain limited.

These findings expose a fundamental limitation in current models: they struggle to extract underlying relationships from perceptual inputs without explicit intermediate guidance. The models appear to learn individual, isolated actions rather than cohesive action patterns that reflect implicit pref-

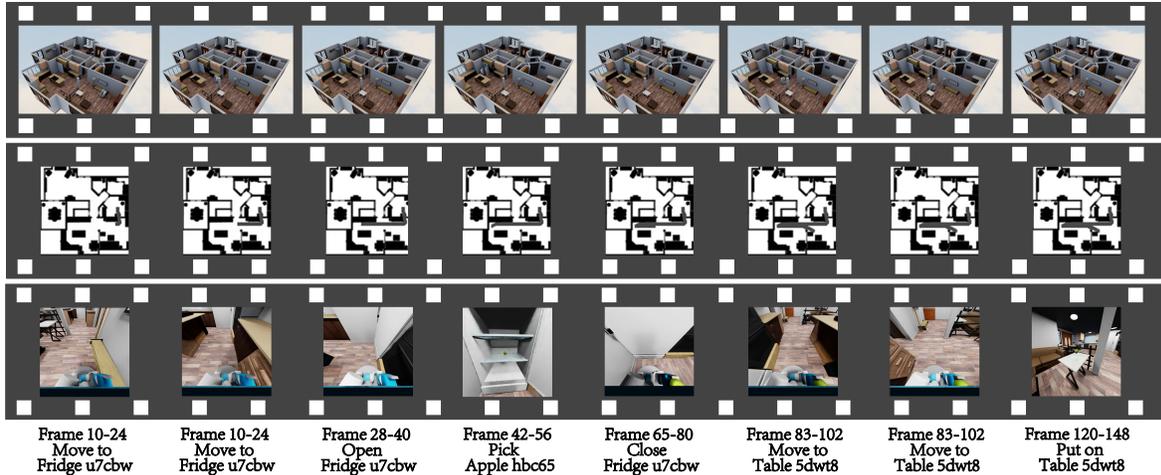


Figure 4: **Example of a demonstration in PBP.** The robot in the demonstration is executing the task “Pick Apple from Fridge and place on Table”. **Top:** A third-person view video provides an overhead perspective of the entire scene. **Middle:** The bird’s-eye-view map displays the robot’s relative position within the scene. **Bottom:** The egocentric video captures the robot’s first-person observations during task execution. **Text:** The per-frame action annotations contain Omniverse object IDs, which ensure each object reference is unique and enable the model to identify specific objects precisely.

Table 1: **Levenshtein distance between generated and ground truth action sequences.** **End-to-end** represents models directly generating action sequences from demonstration-preference pair examples. **Second-stage** indicates generation using both demonstrations and previously inferred preference labels based on demonstrations. **Second-stage (gt)** uses demonstrations alongside ground truth preference labels for sequence generation.

		VIDEO-BASED INPUT				SYMBOL-BASED INPUT	
		ViViT	LLaVA-Next	EILEV	GPT-4V	Llama3-8B	GPT-4
Option Level	End-to-end	15.49±1.29	15.94±3.41	12.88±2.20	15.63±2.31	14.74±3.21	12.23±2.96
	Second-stage	-	12.46±3.23	12.89±3.74	8.37±2.19	9.67±5.16	2.23±2.79
	Second-stage (gt)	-	3.28±5.29	11.18±4.20	1.26±2.55	8.22±5.58	0.12±3.12
Sequence Level	End-to-end	34.04±11.84	34.76±11.25	33.10±12.21	33.75±11.15	31.79±7.32	27.85±6.57
	Second-stage	-	30.02±13.54	33.03±13.61	27.52±9.48	25.46±5.93	16.45±4.00
	Second-stage (gt)	-	18.92±14.18	26.57±12.21	11.36±8.05	19.02±7.10	12.29±3.12
Overall	End-to-end	24.76	25.35	22.99	24.69	23.26	20.04
	Second-stage	-	21.24	22.96	17.94	17.56	9.34
	Second-stage (gt)	-	11.10	18.88	6.31	13.62	6.21

ferences. This significant gap underscores the inherent challenge of performing end-to-end preference learning solely from demonstrations.

5.3. Two-Stage Learning-Planning

Given the limitations of end-to-end learning, we implement a two-stage approach to decompose the preference learning problem. The first stage focuses on preference prediction, where we provide models with auxiliary preference token labels and train them to predict hidden preferences explicitly. These preference tokens, as discussed in Section 4.1, maintain sufficient semantic content for translation into primitive actions.

Results from the first stage (Table 2) reveal significant performance variations across models. At the option level, GPT-4V achieves superior performance with 48.48% accuracy, demonstrating strong capability in interpreting demon-

strated preferences. Among symbol-based models, the stark contrast between DAG-Opt’s limited performance and the improved results from Llama3-8B and GPT-4 highlights the advantage of next-token prediction over dependency learning for preference inference. Models with language components consistently show improved preference understanding compared to end-to-end learning.

The second stage involves generating action sequences based on both demonstrations and predicted preference labels from the first stage, introducing potential error propagation. Results in Table 1 (Second-stage row) and Figure 5 show significant improvements when models receive explicit preferences. For comprehensive evaluation, we include planning results using ground truth preference labels (Second-stage (gt) row). GPT-4V and GPT-4 achieve near-zero Levenshtein distances, indicating almost perfect alignment with ground truth action sequences.

Analysis of both stages reveals distinct challenges across

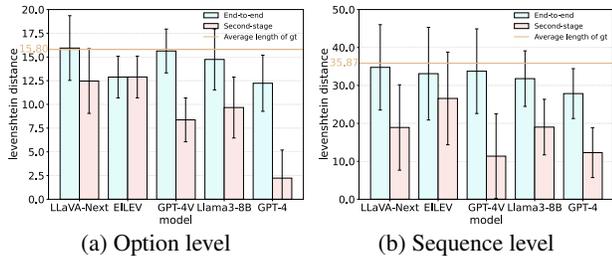


Figure 5: **Levenshtein distance between generated and ground truth action sequences.** Results shown for both (a) option level and (b) sequence level under two conditions: **End-to-end** bars represent direct sequence generation from previous observations, while **Second-stage** bars show performance when models receive predicted preference labels. The **—** line indicates average ground truth sequence length (option level: 15.80, sequence level: 35.87). Lower distances indicate better performance. Results demonstrate significantly improved performance under the two-stage approach compared to end-to-end generation.

model types. Vision-based models like LLaVA-Next and GPT-4V struggle with preference inference but excel in action planning given preference labels, suggesting difficulty in abstracting preferences from visual input. Symbol-based models perform well in both preference inference and preference-guided planning, yet underperform in end-to-end settings. This indicates that models may lack innate preference-based reasoning capabilities but can effectively plan when preferences are explicitly provided.

To isolate the impact of prior knowledge versus in-context learning, we conduct ablation studies by removing demonstrations and testing preference prediction on isolated test sequences. Results in Table 2 (bottom) show significant performance degradation compared to few-shot learning (Table 2 (top)), particularly at the sequence level. This suggests that while models may encode basic task-specific preferences, they rely heavily on demonstrations to recognize complex preference patterns in varied sequences.

5.4. Generalization

While human actions may vary across different objects and scenes, underlying preferences often remain consistent.

Table 3: **Models’ generalization ability.** *direct* denotes experiments conducted *without* generalization. *orig* denotes the original experiments conducted *with* generalization cases. Also the accuracy of preference prediction.

	LLaVA-Next	EILEV	GPT-4V	GPT-4
Option Level <i>direct</i>	33.25	46.93	53.24	86.32
Option Level <i>orig</i>	36.87	38.33	48.48	86.27
Sequence Level <i>direct</i>	33.12	37.53	39.42	70.27
Sequence Level <i>orig</i>	24.85	32.69	37.50	68.42

We evaluate the models’ ability to generalize preference learning across varying visual contexts. The original test set inherently tests generalization by randomly sampling scenes and objects when rendering video demonstrations for each preference. To gain additional insights, we conduct complementary experiments with controlled conditions where demonstration and test videos are rendered in identical rooms with the same objects. This controlled setting enables direct performance comparisons under consistent conditions. We evaluate Efficient In-context Learning on Egocentric Videos (EILEV), Large Language and Vision Assistant (LLaVA), and GPT-4 series models on this variant of PBP, as these models previously demonstrated strong few-shot reasoning capabilities. Results are summarized in Table 3.

Symbol-based reasoning (GPT-4) demonstrates consistent performance regardless of scene or object variations, while vision-based models show greater sensitivity to scene changes. This distinction stems from the nature of our predefined preferences, which are sufficiently abstract and general to apply across diverse scenes and objects. Vision-based models, however, tend to anchor their few-shot learned preferences to specific visual features of scenes or objects. When these visual elements change, preference recognition accuracy may deteriorate. This contextual dependence remains a persistent challenge for vision-based models, which often overfit to scene-specific features from training videos.

Analysis of test points across *direct* and *gen* conditions (Figure 6) reveals two key findings: (i) Preference learning performance correlates with scene characteristics, with certain scenes proving consistently challenging across both

Table 2: **Preference prediction accuracy in few-shot and ablative settings.**

		VIDEO-BASED INPUT				SYMBOL-BASED INPUT		
		ViViT	LLaVA-Next	EILEV	GPT-4V	DAG-Opt	Llama3-8B	GPT-4
Few-shot	Option Level	9.38	36.87	38.33	48.48	10.15	72.98	86.27
	Sequence Level	4.24	24.85	32.69	37.50	13.49	67.18	68.42
	Sequence Level	6.81	30.86	35.51	42.99	11.82	70.08	77.34
Ablative	Option Level	9.16	15.47	4.77	29.42	3.84	39.50	73.87
	Sequence Level	4.38	8.13	0.00	0.00	1.28	6.25	9.42
	Sequence Level	6.77	11.8	2.38	14.71	2.56	22.88	41.64

385 conditions. (ii) While *direct* cases show better performance
 386 overall, failure patterns differ between conditions, particu-
 387 larly for vision-based models. This suggests models rely
 388 heavily on visual context consistency—including object ar-
 389 rangement and scene layout—for accurate predictions, indi-
 390 cating potential superficial learning rather than true pref-
 391 erence understanding. Symbol-based reasoning maintains
 392 robust performance across varied scenes due to the gen-
 393 eral nature of predefined preferences, whereas vision-based
 394 models’ strong dependence on specific visual contexts limits
 395 their generalization capability.

5.5. Ablations on Demonstration Numbers

397 We examine the effect of demonstration quantity on model
 398 performance through an ablation study (Figure 7). Re-
 399 sults show that increasing demonstration numbers generally
 400 improves preference learning and planning effectiveness.
 401 This improvement is most evident in second-stage planning,
 402 where models achieve lower sequence distances by more
 403 accurately replicating human actions. Models like GPT-4,
 404 Llama3, and EILEV show consistent performance gains
 405 with additional demonstrations. However, we observe that
 406 excessive demonstrations (*e.g.*, 5-demo cases for GPT-4 and
 407 EILVE) can sometimes impair first-stage prediction accu-
 408 racy. Despite these occasional exceptions, the overall trend
 409 confirms our intuition: more demonstrations enhance learn-
 410 ing and planning performance. These findings highlight
 411 the importance of demonstration quantity in developing ef-
 412 fective personalized planning systems that align with user
 413 preferences.
 414
 415

6. Conclusion

416 We investigate methods for embodied agents to learn and
 417 implement human preferences through behavioral obser-
 418 vation and user interaction. We present Preference-based
 419 Planning (PBP), a comprehensive embodied benchmark de-
 420
 421
 422
 423
 424

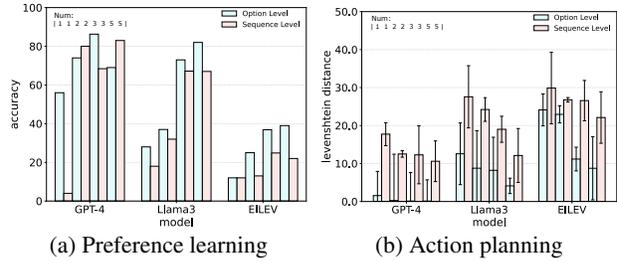


Figure 7: **Ablation study on the number of demonstrations.** Models are evaluated across both of the two stages within PBP task: (a) first-stage preference learning and (b) second-stage action planning. We evaluate both Option Level and Sequence Level tasks. The number of few-shot demonstrations varies from [1, 2, 3, 5], presented left to right. For (a), higher accuracy indicates better performance. For (b), lower distance indicates better performance. Increased demonstration quantity generally improves both preference learning capability and planning effectiveness.

425 signed to capture the complexity of real-world human pref-
 426 erences. We also develop an evaluation framework to assess
 427 models’ preference learning and implementation capabili-
 428 ties. Our findings demonstrate that preferences effectively
 429 abstract human behaviors and guide planning processes.
 430 While current models still face challenges in preference in-
 431 ference and adaptive planning from limited observations,
 432 incorporating preference-based reasoning improves both ef-
 433 fectiveness and generalization. We aim to stimulate further
 434 research in this crucial yet understudied domain of develop-
 435 ing preference-aware embodied agents.

436 **Limitations and Future Work** Our work’s primary lim-
 437 itation stems from its reliance on synthetic data. Despite
 438 Omniverse’s high-quality scene rendering, the simulator
 439 cannot fully replicate real-world complexity and variabil-
 440 ity. Furthermore, human-defined preference labels may not
 441 completely capture preference subtleties and diversity. We
 442 are addressing these limitations by collecting real-world
 443 preference demonstrations using head-worn devices.

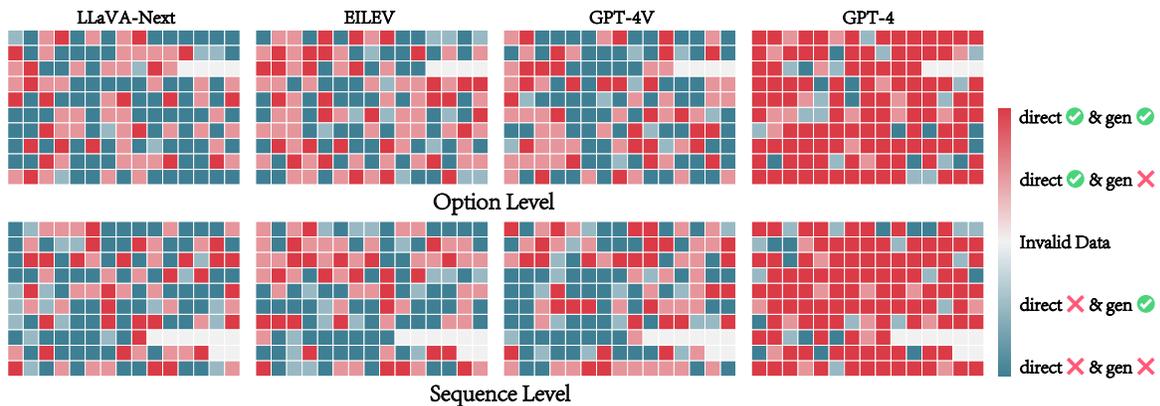


Figure 6: **Analysis of test samples in *direct* and *generalization* settings.** Lines represent distinct scenes, with grid colors indicating different sample statuses.

Impact Statement

The development of personalized, preference-based embodied AI agents has the potential to significantly enhance human-AI collaboration, improving efficiency and user satisfaction in a variety of daily tasks. As AI agents increasingly adapt to individual behaviors, ensuring that these systems respect user autonomy and data privacy is crucial. Given our focus on private scenarios, we anticipate minimal negative societal impact from this research.

References

- Abdo, N., Stachniss, C., Spinello, L., and Burgard, W. Robot, organize my shelves! tidying up objects by predicting user preferences. In *International Conference on Robotics and Automation (ICRA)*, 2015. [3](#)
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. [1](#), [A9](#)
- Ahn, M., Dwibedi, D., Finn, C., Arenas, M. G., Gopalakrishnan, K., Hausman, K., Ichter, B., Irfan, A., Joshi, N., Julian, R., et al. Autort: Embodied foundation models for large scale orchestration of robotic agents. *arXiv preprint arXiv:2401.12963*, 2024. [1](#), [3](#)
- Akgun, B., Cakmak, M., Jiang, K., and Thomaz, A. L. Keyframe-based learning from demonstration: Method and evaluation. *International Journal of Social Robotics*, 4:343–355, 2012. [2](#)
- Aleskerov, F., Bouyssou, D., and Monjardet, B. *Utility maximization, choice and preference*, volume 16. Springer Science & Business Media, 2007. [3](#)
- Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., and Van Den Hengel, A. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [3](#)
- Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., and Schmid, C. Vivit: A video vision transformer. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2021. [A6](#)
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. [1](#)
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. [1](#)
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. [3](#)
- Chao, C., Cakmak, M., and Thomaz, A. L. Towards grounding concepts for transfer in goal learning from demonstration. In *2011 IEEE International Conference on Development and Learning (ICDL)*, volume 2, pp. 1–6. IEEE, 2011. [2](#)
- Chen, H., Suhr, A., Misra, D., Snavely, N., and Artzi, Y. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [3](#)
- Choi, Y. and Luo, Y. Understanding preferences in infancy. *Wiley Interdisciplinary Reviews: Cognitive Science*, 14(4):e1643, 2023. [3](#)
- Coppin, G., Delplanque, S., Cayeux, I., Porcherot, C., and Sander, D. I’m no longer torn after choice: How explicit choices implicitly shape preferences of odors. *Psychological science*, 21(4):489–493, 2010. [2](#)
- Ding, Y., Zhang, X., Amiri, S., Cao, N., Yang, H., Kaminski, A., Esselink, C., and Zhang, S. Integrating action knowledge and llms for task planning and situation handling in open worlds. *Autonomous Robots*, 47(8):981–997, 2023. [3](#)
- Driess, D., Xia, F., Sajjadi, M. S., Lynch, C., Chowdhery, A., Ichter, B., Wahid, A., Tompson, J., Vuong, Q., Yu, T., et al. Palm-e: An embodied multimodal language model. In *Proceedings of International Conference on Machine Learning (ICML)*, 2023. [1](#), [3](#)
- Epstein, S. Integration of the cognitive and the psychodynamic unconscious. *American psychologist*, 49(8):709, 1994. [1](#)
- Fawcett, C. A. and Markson, L. Children reason about shared preferences. *Developmental psychology*, 46(2): 299, 2010. [1](#)
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., and Crawford, K. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021. [A1](#)
- Gerson, S. A., Bekkering, H., and Hunnius, S. Do you do as i do?: Young toddlers prefer and copy toy choices of similarly acting others. *Infancy*, 22(1):5–22, 2017. [1](#)

- 495 Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari,
496 A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X.,
497 et al. Ego4d: Around the world in 3,000 hours of egocen-
498 tric video. In *Proceedings of Conference on Computer*
499 *Vision and Pattern Recognition (CVPR)*, 2022. [A5](#), [A7](#)
- 500 Gu, J., Kirmani, S., Wohlhart, P., Lu, Y., Arenas, M. G.,
501 Rao, K., Yu, W., Fu, C., Gopalakrishnan, K., Xu, Z., et al.
502 Rt-trajectory: Robotic task generalization via hindsight
503 trajectory sketches. *arXiv preprint arXiv:2311.01977*,
504 2023. [3](#)
- 505 Jiang, K., Stacy, S., Wei, C., Chan, A., Rossano, F., Zhu,
506 Y., and Gao, T. Individual vs. joint perception: a prag-
507 matic model of pointing as communicative smithian help-
508 ing. In *Annual Meeting of the Cognitive Science Society*
509 *(CogSci)*, 2021. [1](#)
- 510 Jiang, K., Dahmani, A., Stacy, S., Jiang, B., Rossano, F.,
511 Zhu, Y., and Gao, T. What is the point? a theory of mind
512 model of relevance. In *Annual Meeting of the Cognitive*
513 *Science Society (CogSci)*, 2022. [1](#)
- 514 Kahneman, D. The psychology of preferences. *Scientific*
515 *American*, 1982. [2](#)
- 516 Kant, Y., Ramachandran, A., Yenamandra, S., Gilitschenski,
517 I., Batra, D., Szot, A., and Agrawal, H. Housekeep: Tidy-
518 ing virtual households using commonsense reasoning. In
519 *Proceedings of European Conference on Computer Vision*
520 *(ECCV)*, 2022. [3](#)
- 521 Kapelyukh, I. and Johns, E. My house, my rules: Learn-
522 ing tidying preferences with graph neural networks. In
523 *Conference on Robot Learning (CoRL)*, 2022. [1](#), [3](#)
- 524 Kolve, E., Mottaghi, R., Han, W., VanderBilt, E., Weihs, L.,
525 Herrasti, A., Deitke, M., Ehsani, K., Gordon, D., Zhu, Y.,
526 et al. Ai2-thor: An interactive 3d environment for visual
527 ai. *arXiv preprint arXiv:1712.05474*, 2017. [3](#)
- 528 Leal, I., Choromanski, K., Jain, D., Dubey, A., Varley,
529 J., Ryoo, M., Lu, Y., Liu, F., Sindhvani, V., Vuong,
530 Q., et al. Sara-rt: Scaling up robotics transformers with
531 self-adaptive robust attention. *arXiv preprint*
532 *arXiv:2312.01990*, 2023. [1](#), [3](#)
- 533 Lee, M. K., Forlizzi, J., Kiesler, S., Rybski, P., Antanitis, J.,
534 and Savetsila, S. Personalization in hri: A longitudinal
535 field experiment. In *ACM/IEEE International Conference*
536 *on Human-Robot Interaction (HRI)*, 2012. [1](#)
- 537 Levinson, S. C. *Pragmatics*. Cambridge UP, 1983. [1](#)
- 538 Leyzberg, D., Spaulding, S., and Scassellati, B. Personal-
539 izing robot tutors to individuals' learning differences. In
540 *ACM/IEEE International Conference on Human-Robot*
541 *Interaction (HRI)*, 2014. [1](#)
- 542 Li, C., Zhang, R., Wong, J., Gokmen, C., Srivastava, S.,
543 Martín-Martín, R., Wang, C., Levine, G., Lingelbach, M.,
544 Sun, J., et al. Behavior-1k: A benchmark for embodied
545 ai with 1,000 everyday activities and realistic simulation.
546 In *Conference on Robot Learning (CoRL)*, 2023. [2](#), [4](#), [A2](#)
- 547 Liberman, Z., Kinzler, K. D., and Woodward, A. L. Ori-
548 gins of homophily: Infants expect people with shared
549 preferences to affiliate. *Cognition*, 212:104695, 2021. [1](#)
- Lichtenstein, S. and Slovic, P. The construction of prefer-
ence: An overview. *The construction of preference*, 1:
1–40, 2006. [1](#), [2](#)
- Liu, H., Li, C., Li, Y., and Lee, Y. J. Improved baselines with
visual instruction tuning. In *Proceedings of Conference*
on Computer Vision and Pattern Recognition (CVPR),
2024. [A7](#)
- Mongin, P. Expected utility theory. *Handbook of Economic*
Methodology, pp. 342–350, 01 1997. [3](#)
- Mu, Y., Zhang, Q., Hu, M., Wang, W., Ding, M., Jin,
J., Wang, B., Dai, J., Qiao, Y., and Luo, P. Embod-
iedgpt: Vision-language pre-training via embodied chain
of thought. In *Proceedings of Advances in Neural Infor-*
mation Processing Systems (NeurIPS), 2023. [1](#)
- Patel, M. and Chernova, S. Proactive robot assistance via
spatio-temporal object modeling. In *Conference on Robot*
Learning (CoRL), 2023. [3](#)
- Patel, M., Prakash, A. G., and Chernova, S. Predicting
routine object usage for proactive robot assistance. In
Conference on Robot Learning (CoRL), 2023. [3](#)
- Peng, Y., Han, J., Zhang, Z., Fan, L., Liu, T., Qi, S., Feng,
X., Ma, Y., Wang, Y., and Zhu, S.-C. The tong test:
Evaluating artificial general intelligence through dynamic
embodied physical and social interactions. *Engineering*,
34:12–22, 2024. [1](#)
- Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma,
S., and Wei, F. Kosmos-2: Grounding multimodal
large language models to the world. *arXiv preprint*
arXiv:2306.14824, 2023. [1](#)
- Puig, X., Shu, T., Li, S., Wang, Z., Liao, Y.-H., Tenenbaum,
J. B., Fidler, S., and Torralba, A. Watch-and-help: A chal-
lenge for social perception and human-ai collaboration.
In *Proceedings of International Conference on Learning*
Representations (ICLR), 2021a. [3](#)
- Puig, X., Shu, T., Li, S., Wang, Z., Liao, Y.-H., Tenenbaum,
J. B., Fidler, S., and Torralba, A. Watch-and-help: A chal-
lenge for social perception and human-ai collaboration.
In *Proceedings of International Conference on Learning*
Representations (ICLR), 2021b. [3](#)

- 550 Puig, X., Shu, T., Tenenbaum, J. B., and Torralba, A. Nopa:
551 Neurally-guided online probabilistic assistance for build-
552 ing socially intelligent home assistants. In *International*
553 *Conference on Robotics and Automation (ICRA)*, 2023a.
554 3
- 555 Puig, X., Undersander, E., Szot, A., Cote, M. D., Yang,
556 T.-Y., Partsey, R., Desai, R., Clegg, A., Hlavac, M., Min,
557 S. Y., et al. Habitat 3.0: A co-habitat for humans, avatars,
558 and robots. In *Proceedings of International Conference*
559 *on Learning Representations (ICLR)*, 2023b. 3
- 560 Sarch, G., Fang, Z., Harley, A. W., Schydlow, P., Tarr, M. J.,
561 Gupta, S., and Fragkiadaki, K. Tidee: Tidying up novel
562 rooms using visuo-semantic commonsense priors. In
563 *Proceedings of European Conference on Computer Vision*
564 *(ECCV)*, 2022. 3
- 565 Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans,
566 E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J.,
567 et al. Habitat: A platform for embodied ai research. In
568 *Proceedings of International Conference on Computer*
569 *Vision (ICCV)*, 2019. 3
- 570 Scott, J. et al. Rational choice theory. *Understanding con-*
571 *temporary society: Theories of the present*, 129:126–138,
572 2000. 3
- 573 Sen, A. Behaviour and the concept of preference. *Econom-*
574 *ica*, 40(159):241–259, 1973. 3
- 575 Shridhar, M., Thomason, J., Gordon, D., Bisk, Y., Han, W.,
576 Mottaghi, R., Zettlemoyer, L., and Fox, D. Alfred: A
577 benchmark for interpreting grounded instructions for ev-
578 eryday tasks. In *Proceedings of Conference on Computer*
579 *Vision and Pattern Recognition (CVPR)*, 2020. 3
- 580 Shu, T. and Tian, Y. M³rl: Mind-aware multi-agent man-
581 agement reinforcement learning. In *International Confer-*
582 *ence on Learning Representations*, 2019. 3
- 583 Simonson, I. Will i like a “medium” pillow? another look
584 at constructed and inherent preferences. *Journal of Con-*
585 *sumer Psychology*, 18(3):155–169, 2008. 1
- 586 Singh, I., Blukis, V., Mousavian, A., Goyal, A., Xu, D.,
587 Tremblay, J., Fox, D., Thomason, J., and Garg, A.
588 Progprompt: Generating situated robot task plans using
589 large language models. In *International Conference on*
590 *Robotics and Automation (ICRA)*, 2023. 1
- 591 Slovic, P. The construction of preference. *American Psy-*
592 *chologist*, 50(5):364, 1995. 1
- 593 Song, C. H., Wu, J., Washington, C., Sadler, B. M., Chao,
594 W.-L., and Su, Y. Llm-planner: Few-shot grounded plan-
595 ning for embodied agents with large language models. In
596 *Proceedings of International Conference on Computer*
597 *Vision (ICCV)*, 2023. 3
- 598 Sucas, I. A., Moll, M., and Kavraki, L. E. The open motion
599 planning library. *IEEE Robotics & Automation Magazine*,
600 19(4):72–82, 2012. A4
- 601 Taniguchi, A., Isobe, S., El Hafi, L., Hagiwara, Y., and
602 Taniguchi, T. Autonomous planning based on spatial con-
603 cepts to tidy up home environments with service robots.
604 *Advanced Robotics*, 35(8):471–489, 2021. 3
- 605 Thomason, J., Murray, M., Cakmak, M., and Zettlemoyer,
606 L. Vision-and-dialog navigation. In *Conference on Robot*
607 *Learning (CoRL)*, 2020. 3
- 608 Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux,
609 M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E.,
610 Azhar, F., et al. Llama: Open and efficient foundation
611 language models. *arXiv preprint arXiv:2302.13971*, 2023.
612 A9
- 613 Wu, J., Antonova, R., Kan, A., Lepert, M., Zeng, A., Song,
614 S., Bohg, J., Rusinkiewicz, S., and Funkhouser, T. Tidy-
615 bot: Personalized robot assistance with large language
616 models. *Autonomous Robots*, 47(8):1087–1102, 2023. 3
- 617 Xu, M., Jiang, G., Liang, W., Zhang, C., and Zhu, Y. Active
618 reasoning in an open-world environment. *Advances in*
619 *Neural Information Processing Systems*, 36, 2024a. 3
- 620 Xu, M., Jiang, G., Liang, W., Zhang, C., and Zhu, Y. Inter-
621 active visual reasoning under uncertainty. *Advances in*
622 *Neural Information Processing Systems*, 36, 2024b. A8
- 623 Yu, K. P., Zhang, Z., Hu, F., and Chai, J. Efficient in-context
624 learning in vision-language models for egocentric videos.
625 *arXiv preprint arXiv:2311.17041*, 2023. A7
- 626 Yuan, L., Gao, X., Zheng, Z., Edmonds, M., Wu, Y. N.,
627 Rossano, F., Lu, H., Zhu, Y., and Zhu, S.-C. In situ bidi-
628 rectional human-robot value alignment. *Science Robotics*,
629 7(68), 2022. 1
- 630 Yuan, T., Liu, H., Fan, L., Zheng, Z., Gao, T., Zhu, Y., and
631 Zhu, S.-C. Joint inference of states, robot knowledge,
632 and human (false-)beliefs. In *International Conference*
633 *on Robotics and Automation (ICRA)*, 2020. 1
- 634 Yuan, Y., Wang, H., Ding, J., Jin, D., and Li, Y. Learning
635 to simulate daily activities via modeling dynamic human
636 needs. In *Proceedings of the ACM Web Conference*, 2023.
637 1
- 638 Zey, M. *Rational choice theory and organizational theory:*
639 *A critique*. Sage, 1998. 3
- 640 Zhang, C., Jia, B., Edmonds, M., Zhu, S.-C., and Zhu,
641 Y. Acre: Abstract causal reasoning beyond covariation.
642 In *Proceedings of Conference on Computer Vision and*
643 *Pattern Recognition (CVPR)*, 2021. A8

605 Zhang, H., Du, W., Shan, J., Zhou, Q., Du, Y., Tenenbaum,
606 J. B., Shu, T., and Gan, C. zhang2023building. In *Pro-*
607 *ceedings of International Conference on Learning Repre-*
608 *sentations (ICLR)*, 2024. 3

609 Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M.,
610 Chen, S., Dewan, C., Diab, M., Li, X., Lin, X. V.,
611 et al. Opt: Open pre-trained transformer language models.
612 *arXiv preprint arXiv:2205.01068*, 2022. A7

614 Zheng, X., Aragam, B., Ravikumar, P. K., and Xing, E. P.
615 Dags with no tears: Continuous optimization for structure
616 learning. In *Proceedings of Advances in Neural Informa-*
617 *tion Processing Systems (NeurIPS)*, 2018. A8

619 Zhu, Y., Jiang, C., Zhao, Y., Terzopoulos, D., and Zhu, S.-C.
620 Inferring forces and learning human utilities from videos.
621 In *Proceedings of Conference on Computer Vision and*
622 *Pattern Recognition (CVPR)*, 2016. 1

623 Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F.,
624 Wu, J., Wohlhart, P., Welker, S., Wahid, A., et al. Rt-2:
625 Vision-language-action models transfer web knowledge
626 to robotic control. In *Conference on Robot Learning*
627 *(CoRL)*, 2023. 1, 3

628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659

A. Dataset Card

We follow the datasheet proposed in [Gebru et al. \(2021\)](#) for documenting our proposed PBP:

1. Motivation

(a) **For what purpose was the dataset created?**

The benchmark was created to evaluate existing learning agents on their ability to understand and adapt to various human preferences. Specifically, it aims to test the agents' proficiency in few-shot learning from demonstrations, where they must respond to ambiguous task instructions and formulate adaptive task plans based on limited examples of user preferences. The benchmark is designed to highlight the challenges and gaps in current AI systems' capabilities in planning activities and abstracting human preferences, ultimately driving advancements towards developing more intelligent and personalized embodied agents.

(b) **Who created the dataset and on behalf of which entity?**

N/A.

(c) **Who funded the creation of the dataset?**

N/A.

(d) **Any other Comments?**

None.

2. Composition

(a) **What do the instances that comprise the dataset represent?**

Each instance contains an egocentric video of an agent's activity, its bird's-eye-view map of the position of the agent, and a frame-level textual annotation of the current action, as shown in [Figure 4](#). Additionally, we provide a rendered third-person view of the entire process.

(b) **How many instances are there in total?**

15000.

(c) **Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?**

No. The dataset contains a set of demonstrations rendered within the simulator, The users can render more diverse instances if they want. We have provided the rendering instructions.

(d) **What data does each instance consist of?**

The instances that comprise the benchmark represent various types of human preferences applied to different tasks within a realistic embodied scene. Each instance is designed to challenge the learning agents to understand and adapt to these preferences based on a few demonstration examples, reflecting the diverse and hierarchical nature of user preferences in real-world scenarios. See above for data details.

(e) **Is there a label or target associated with each instance?**

Yes.

(f) **Is any information missing from individual instances?**

No.

(g) **Are relationships between individual instances made explicit?**

Yes.

(h) **Are there recommended data splits?**

No.

(i) **Are there any errors, sources of noise, or redundancies in the dataset?**

No.

(j) **Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?**

Self-contained.

(k) **Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)?**

No.

- 715 (l) **Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might**
 716 **otherwise cause anxiety?**
 717 No.
- 718 (m) **Does the dataset relate to people?**
 719 No.
- 720 (n) **Does the dataset identify any subpopulations (e.g., by age, gender)?**
 721 No.
- 722 (o) **Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in**
 723 **combination with other data) from the dataset?**
 724 No.
- 725 (p) **Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or**
 726 **ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations;**
 727 **financial or health data; biometric or genetic data; forms of government identification, such as social**
 728 **security numbers; criminal history)?**
 729 No.
- 730 (q) **Any other comments?**
 731 None.

3. Collection Process

- 735 (a) **How was the data associated with each instance acquired?**
 736 We render PBP using NVIDIA’s Omniverse and OmniGibson simulation environment (Li et al., 2023).
- 737 (b) **What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual**
 738 **human curation, software program, software API)?**
 739 The data for each instance in the benchmark was acquired by sampling preferences from a predefined set and
 740 constructing tasks paired with a few demonstrations that shared high-level preferences but differed in specific
 741 objects and scenes. Each sampled preference was randomly assigned to one of the 50 scenes provided by
 742 OmniGibson, with relevant objects sampled within the scene. Egocentric observation and action sequences of
 743 an embodied agent were generated as the agent performed tasks guided by a rule-based planner using planning
 744 primitives like inverse kinematics for grasping and the A* algorithm for movement.
- 745 (c) **If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic**
 746 **with specific sampling probabilities)?**
 747 N/A.
- 748 (d) **Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were**
 749 **they compensated (e.g., how much were crowdworkers paid)?**
 750 N/A.
- 751 (e) **Over what timeframe was the data collected?**
 752 N/A.
- 753 (f) **Were any ethical review processes conducted (e.g., by an institutional review board)?**
 754 The dataset raises no ethical concerns.
- 755 (g) **Does the dataset relate to people?**
 756 No.
- 757 (h) **Did you collect the data from the individuals in question directly, or obtain it via third parties or other**
 758 **sources (e.g., websites)?**
 759 N/A.
- 760 (i) **Were the individuals in question notified about the data collection?**
 761 N/A.
- 762 (j) **Did the individuals in question consent to the collection and use of their data?**
 763 N/A.
- 764 (k) **If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent**
 765 **in the future or for certain uses?**
 766 N/A.

- (l) **Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted?**

Yes.

- (m) **Any other comments?**

None.

4. Preprocessing, Cleaning and Labeling

- (a) **Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?**

N/A.

- (b) **Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?**

N/A.

- (c) **Is the software used to preprocess/clean/label the instances available?**

N/A.

- (d) **Any other comments?**

None.

5. Uses

- (a) **Has the dataset been used for any tasks already?**

No, the dataset is newly proposed by us.

- (b) **Is there a repository that links to any or all papers or systems that use the dataset?**

No, the dataset is new.

- (c) **What (other) tasks could the dataset be used for?**

This dataset could be used for research topics like embodied AI and human-computer interaction.

- (d) **Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?**

N/A.

- (e) **Are there tasks for which the dataset should not be used?**

N/A.

- (f) **Any other comments?**

None.

6. Distribution

- (a) **Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?**

No before it is made public.

- (b) **How will the dataset be distributed (e.g., tarball on website, API, GitHub)?**

On our project website upon acceptance.

- (c) **When will the dataset be distributed?**

Upon acceptance.

- (d) **Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?**

Under CC BY-NC ¹ license.

- (e) **Have any third parties imposed IP-based or other restrictions on the data associated with the instances?**

No.

- (f) **Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?**

No.

- (g) **Any other comments?**

None.

¹<https://creativecommons.org/licenses/by-nc/4.0/>

7. Maintenance

- (a) **Who is supporting/hosting/maintaining the dataset?**
The authors.
- (b) **How can the owner/curator/manager of the dataset be contacted (e.g., email address)?**
N/A.
- (c) **Is there an erratum?**
Future erratum will be released through the website.
- (d) **Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?**
Yes.
- (e) **If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)?**
N/A. The dataset does not relate to people.
- (f) **Will older versions of the dataset continue to be supported/hosted/maintained?**
Yes.
- (g) **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?**
Yes. We will release the source code as well as a licence on our project website after acceptance.
- (h) **Any other comments?**
None.

B. Dataset Statistics

The length of the simulations in the dataset ranges from 1 to 5 minutes, depending on the tasks recorded. And the videos are recorded at 30 fps.

B.1. Preferenes

See [Table A1](#) for the preference statistics in PBP.

Table A1: Dataset Statistics in PBP.

	Action_Level	Option_Level	Sequence_Level
Preference Num	75	135	80
Video Num	5000	5000	5000
Sub-task Num	1	2-3	2-3

B.2. Actions

See [Table A2](#) for the action statistics in PBP. We implement 17 action primitives in PBP to assist with model planning and dataset rendering. These action primitives have parameters that simplify tasks and are considered the lowest-level actions. Each sub-task contains 8 to 20 such lowest-level actions. Generally, most of these actions consist of two parts: the robot movement part and the arm (gripper) execution part. For robot movement, we use the A* algorithm to find paths and avoid collisions. We build a connection map during scene initialization for navigation, taking the robot’s width into consideration. For the arm (gripper) execution, we primarily use the IK algorithm to compute arm movements. However, since IK cannot handle complex tasks, such as picking objects from the fridge, we also leverage the Open Motion Planning Library (OMPL) planner ([Sucan et al., 2012](#)) with forward planning to assist in planning the arm positions.

B.3. More Dataset Details and Discussion

Dataset production The process of producing data is mainly explained in [Section 4.2](#). In summary, we follow the order of “sample preference - sample scene - sample objects to be manipulated - generate actions guided by a rule-based planner”.

Length and FPS of the simulations The length of the simulations ranges from 1 to 5 minutes, depending on the tasks recorded. The videos are recorded at 30 fps.

Actions contained in each simulation The number of actions in simulations varies among different preference levels. There is 1 subtask for action-level, 2-3 subtasks for option-level, and 2-3 subtasks for sequence-level preferences. Each subtask contains 8-20 actions.

Scenes and rooms Each scene contains various types of rooms. The main differences between scenes are the type, number, and layout of both rooms and furniture. Additionally, each room may contain different objects and have unique layouts. Details of the scenes and rooms can be found in Omnigibson’s official documentation (<https://behavior.stanford.edu/omnigibson/>), as we directly adopt these scenes from the open-sourced project.

290 preference types Considering that preferences in household activities are not only multi-dimensional but also hierarchical, we first define a hierarchy of preferences from the perspective of how things happen in a life scenario, that is, from each specific action to a sub-task consisting of several actions, and then to the sequence combining these sub-tasks. The next step is to expand each level with typical tasks and actions. The detailed definition of the 290 preferences can be found in [Section 4.1](#).

The egocentric view Collecting both egocentric observations and third-person views is feasible in PbP or similar environments built on simulators like iGibson. However, in real-world scenarios, it is generally easier to gather egocentric observations of human daily activities, as these can be efficiently captured through wearable devices. Additionally, there are numerous egocentric-view datasets available, such as Ego4D ([Grauman et al., 2022](#)), which further facilitate this approach. While third-person views can provide a different perspective, they often encounter issues such as occlusion. Although research based on third-person views is essential for applications involving real robots, focusing on egocentric views in the current work allows for a more straightforward exploration of preference learning and planning. Nevertheless, third-person view data can be obtained by integrating additional cameras, as outlined in our provided code.

Action ground truth In experiments involving vision input, we do not explicitly provide the action sequence of the user. In the symbolic-based experiment, we provide the action sequence to reduce the perception cost to concentrate more effectively on the inference and planning aspects of the study.

Table A2: Action Primitives in PbP.

Action List	Explanation
Move_to_[]	Move to a specified location, or a specified room, or a specified object
Rotate_to_[]	Rotate to a specified orientation or a specified object
Pick_[]	Pick up an object using the gripper, <i>e.g.</i> , “Pick_apple”
Place_[]	Place an object at a location, <i>e.g.</i> , “Place_apple_on_table”
Fill_[]_with_[]	Fill a container with a substance, <i>e.g.</i> , “Fill_glass_with_water”
Pour_[]	Pour a substance from a container, <i>e.g.</i> , “Pour_milk”
Open_[]	Open an object, <i>e.g.</i> , “Open_door”
Close_[]	Close an object, <i>e.g.</i> , “Close_fridge”
Cut_[]	Cut an object, <i>e.g.</i> , “Cut_carrot”
Cook_[]	Cook an item, <i>e.g.</i> , “Cook_pasta”
Wash_[]	Wash an object, <i>e.g.</i> , “Wash_dishes”
Clean_[]	Clean a surface or object, <i>e.g.</i> , “Clean_counter”
Cover_[]	Cover an object, <i>e.g.</i> , “Cover_bowl”
Uncover_[]	Uncover an object, <i>e.g.</i> , “Uncover_bowl”
Toggle_on_[]	Turn on a device, <i>e.g.</i> , “Toggle_on_light”
Toggle_off_[]	Turn off a device, <i>e.g.</i> , “Toggle_off_stove”
Wait_[]	Wait some time

C. Experiment Details

C.1. Case Study

We also provide a case with preference *Put fruit on the bed* in the following table. We present a simplified version of the demonstrations, where all video outputs have been translated into symbol-based action sequences for ease of understanding. Video-based models such as LLaVA-Next and GPT-4V struggle with comprehending preferences and tend to replicate certain action patterns from the video demonstration, such as “move to” and “pick up”. Llama3 demonstrates a partial understanding and execution of the preference. It correctly moves to each fruit (grape, banana), picks them up, and places them on the bed. However, it also interacts with the pencil and places it on the bed, which is not required by the preference. Ideally, the pencil should be placed on the table, similar to the pen. On the other hand, GPT-4(Symbol) accurately interacts with the grape and banana by moving to each fruit, picking it up, and placing it on the bed. This demonstrates a better understanding and execution of the preference compared to the other models.

Case	ID	Scene	Actions
Context	Demo 1	Beechwood	move to orange, pick up orange, move to bed with orange, place orange on bed, move to apple, pick up apple, move to bed with apple, place apple on bed, move to dvd, pick up dvd, move to table with dvd, place dvd on table.
	Demo 2	Rs	move to banana, pick up banana, move to bed with banana, place banana on bed, move to club sandwich, pick up club sandwich, move to countertop with club sandwich, place club sandwich on countertop, move to orange, pick up orange, move to countertop with orange, move to bed with orange, place orange on bed.
	Demo 3	Merom	move to pen, pick up pen, move to table with pen, place pen, move to strawberry, pick up strawberry, move with to bed with strawberry, place strawberry on bed, move to textbook, pick up textbook, move with textbook, move to bookshelf with textbook, place textbook in bookshelf.
Test Scene		Beechwood	grape, pencil, banana
Answer	LLaVA-Next	Beechwood	move to grape, move to bed with grape, move to pencil, pick up pencil, move to banana, pick up banana
	GPT-4V	Beechwood	move to grape, pick up grape, move to bed with grape, place grape on bed
	Llama3	Beechwood	move to grape, pick up grape, move to bed with grape, place grape on bed, move to pencil, pick up pencil, move to bed with pencil, place pencil on bed, move to banana, pick up banana, move to bed with banana, place banana on bed
	GPT-4(Symbol)	Beechwood	move to grape, pick up grape, move to bed with grape, place grape on bed, move to banana, pick up banana, move to bed with banana, place banana on bed

Table A3: Case Study with preference *Put fruit on the bed*.

D. Baseline Details

D.1. ViViT

As a baseline, we employ the pure-Transformer-based Video Vision Transformer (ViViT) (Arnab et al., 2021), an end-to-end trainable model with proven capabilities in extracting spatial and temporal information from video inputs. Since it lacks a LLM component, ViViT likely serves as a lower bound for commonsense understanding in PBP tasks.

Inspired by Vision Transformer, ViViT extracts spatio-temporal tokens from the input video and outputs video classification labels for classification. We adopt the ViViT implementation from the official GitHub repo <https://github.com/google-research/scenic>.

Specifically, we utilize a ViViT with an image size of 224 and a patch size of 16. We extract 2 frames per second from the

990 input video and pad them with the last frame. The Transformer architecture with 3 attention heads operates on features of
 991 hidden size of 192 and depth of 4. Each attention head operates on a dimension of 64. We train our model for 30 epochs
 992 with a learning rate $3e-5$. For the few-shot setting, we concatenate the demo videos temporally.

994 D.2. LLaVA

995 Building on more sophisticated architectures, LLaVA (Liu et al., 2024) represents an end-to-end trainable large multimodal
 996 model that integrates vision and text for comprehensive visual-language understanding. We specifically evaluate LLaVA-
 997 NeXT, which has been finetuned to excel at zero-shot video understanding tasks.

999 Following the official implementation of LLaVA from <https://github.com/LLaVA-VL/LLaVA-NeXT>, we test the LLaVA-
 1000 NeXT-Video-7B-DPO model which is designed for video understanding. Specifically, we run the model following the
 1001 default inference settings, with vicuna_v1 as the prompt mode, a sample frame number of 32, and a spatial pooling stride of
 1002 2. The textual prompts are as follows²:

1005 “Stage One / Preference Prediction”

1006 You are a robot assistant that can help summarize the host's preference.

1007 All possible preferences are: {ALL POSSIBLE PREFERENCES}

1008 Now there are some previous video demos:

1009 [VIDEO_DEMO_1] The preference is [PREFERENCE_1]

1010 [VIDEO_DEMO_2] The preference is [PREFERENCE_2]

1011 [VIDEO_DEMO_3] The preference is [PREFERENCE_3]

1012 Now, please summarize the preference from the last video: [TEST_CASE]

1013 Question: What's the user's preference? Choose from the preference listed before:

1015 “Stage Two / Planning”

1016 You are a robot assistant. Please view the demos and help generate action sequence.

1017 All possible preferences are: {ALL POSSIBLE ACTIONS}

1018 Now there are some previous video demos:

1019 [VIDEO_DEMO_1]

1020 [VIDEO_DEMO_2]

1021 [VIDEO_DEMO_3]

1022 Now you are in the scene with [SCENE DESCRIPTIONS]. Your action sequence is:

1025 D.3. EILEV

1026 For specialized egocentric video processing, we incorporate EILEV (Yu et al., 2023), which achieves in-context learning
 1027 through architectural modifications to a pretrained VLM. Our implementation uses OPT-2.7B (Zhang et al., 2022) as
 1028 the language backbone. The model’s pretraining on Ego4D (Grauman et al., 2022) aligns well with PBP’s egocentric
 1029 perspective.

1031 Following the official implementation from <https://github.com/yukw777/EILEV.git>, we test the EILEV model in PBP.
 1032 There are two reasons we chose EILEV among other VLMs as one of our baselines: 1) EILEV elicits in-context learning
 1033 through a series of architectural modifications and a unique training process, 2) EILEV is trained using ego-centric data,
 1034 which is compatible with PBP’s input. The textual prompts are as follows. Since EILEV requires the input of the videos and
 1035 texts to follow a certain pattern for better in-context learning, there are some small modifications to the prompt:

1037 “Stage One / Preference Prediction”

1038 You are a robot assistant that can help summarize the host's preference.

1039 All possible preferences are: {ALL POSSIBLE PREFERENCES}

1041 ²For the textual prompts, we aim to maintain consistency across all LLMs, although some baselines may have additional requirements
 1042 for the input format. The prompt design is mainly motivated by OpenAI Cookbook [git@github.com:openai/openai-cookbook.git](https://github.com/openai/openai-cookbook.git).
 1043 We omitted the prompt tuning process, as we found that minor changes in the prompt were unlikely to significantly impact the results.
 1044 Conversely, selecting the proper demonstrations in the few-shot examples has a much greater influence on the results.

1045 Quesiton: What's the user's preference? Choose from the preference listed before:
1046 Now there are some prevous video demos:
1047 [VIDEO_DEMO_1] The preference is [PREFERENCE_1]
1048 [VIDEO_DEMO_2] The preference is [PREFERENCE_2]
1049 [VIDEO_DEMO_3] The preference is [PREFERENCE_3]
1050 [TEST_CASE]
1051
1052 “Stage Two / Planning”
1053 You are a robot assistant. Please view the demos and help generate action sequence.
1054 All possible preferences are: {ALL POSSIBLE ACTIONS}
1055 Now there are some prevous video demos:
1056 [VIDEO_DEMO_1]
1057 [VIDEO_DEMO_2]
1058 [VIDEO_DEMO_3]
1059 Now you are in the scene with [SCENE DESCRIPTIONS]. Your action sequence is:

1060
1061
1062 **D.4. GPT-4V**

1063 To benchmark against state-of-the-art visual-language models, we evaluate GPT-4V using the Azure OpenAI API (version
1064 “gpt-4-turbo-2024-04-09”). Due to image token limitations, we implement video input subsampling while maintaining
1065 temporal coherence.
1066 We run our GPT-4 model through the AzureOpenAI API using the GPT version “gpt-4-turbo-2024-04-09”. The API has a
1067 limit of 10 images per request. Consequently, for the zero-shot setting, we resample each input video to 8 frames of size
1068 224. For the few-shot setting, where we need to input 3 extra video demonstrations, we concatenate 4 images into a frame,
1069 thereby obtaining 4 videos in 8 frames, maintaining the same frame number as the previous setting. We test the model with
1070 a temperature of 0.05. The textual prompts are as follows:

1071
1072 “Stage One / Preference Prediction”
1073 You are a robot assistant that can help summarize the host's preference.
1074 All possible preferences are: {ALL POSSIBLE PREFERENCES}
1075 Now there are some prevous video demos:
1076 [VIDEO_DEMO_1] The preference is [PREFERENCE_1]
1077 [VIDEO_DEMO_2] The preference is [PREFERENCE_2]
1078 [VIDEO_DEMO_3] The preference is [PREFERENCE_3]
1079 [VIDEO_DEMO_3] The preference is [PREFERENCE_3]
1080 Now, please summarize the preference from the last video: [TEST_CASE]
1081 Quesiton: What's the user's preference? Choose from the preference listed before:

1082
1083 “Stage Two / Planning”
1084 You are a robot assistant. Please view the demos and help generate action sequence.
1085 All possible preferences are: {ALL POSSIBLE ACTIONS}
1086 Now there are some prevous video demos:
1087 [VIDEO_DEMO_1]
1088 [VIDEO_DEMO_2]
1089 [VIDEO_DEMO_3]
1090 Now you are in the scene with [SCENE DESCRIPTIONS]. Your action sequence is:

1091
1092 Beyond multimodal approaches, we also evaluate single-modal models that process only action sequences:

1093
1094 **D.5. DAG-Opt**

1095 We approach symbolic reasoning by framing the problem as a DAG-Optimization task that uncovers dependency relations
1096 between actions and preferences (Zheng et al., 2018). Our implementation uses a score-based NOTEARS model to learn a
1097 generalized Structural Equation Model (SEM), following previous few-shot reasoning frameworks (Zhang et al., 2021; Xu
1098 et al., 2024b) based on causal dependency structures.
1099

1100 We implement the DAG-Opt baseline following <https://github.com/xunzheng/notears.git>. Specifically, we imple-
 1101 ment a nonlinear NOTEARS using MLP in evaluation.

1102 To assess pure language understanding, we evaluate advanced LLMs including Llama3 (Touvron et al., 2023) and GPT-4
 1103 (Achiam et al., 2023) using only action sequence inputs. This approach treats actions as high-level abstractions of egocentric
 1104 videos, reducing visual complexity while maintaining task semantics. We benchmark Llama3-8B as our baseline against
 1105 GPT-4-Turbo as the current state-of-the-art, employing prompt designs informed by the OpenAI Cookbook for optimal
 1106 few-shot performance.
 1107

1108 **D.6. Llama3-8B**

1110 We test the Llama3 series model with the official scripts from <https://github.com/meta-llama/llama3>. Specifically,
 1111 we test the 8B instruction-tuned variant “Meta-Llama-3-8B-Instruct” on PBP. We test the model with a temperature of 0.05.
 1112 The textual prompts are as follows:
 1113

1114 “Stage One / Preference Prediction”

1115 You are a robot assistant that can help summarize the host's preference.

1116 Please read the following text file and summarize the user's preference.

1117 All possible preferences are: {ALL POSSIBLE PREFERENCES}

1118 [TEXT_ANNOTATION_1] The preference is [PREFERENCE_1]

1119 [TEXT_ANNOTATION_2] The preference is [PREFERENCE_2]

1120 [TEXT_ANNOTATION_3] The preference is [PREFERENCE_3]

1121 Now, please summarize the preference from the last tet file: [TEST_CASE]

1122 Quesiton: What's the user's preference? Choose from the preference listed before:
 1123

1124 “Stage Two / Planning”

1125 You are a robot assistant. Please read the following text files and help generate action sequence.

1126 All possible preferences are: {ALL POSSIBLE ACTIONS}

1127 Now there are some previous video demos:

1128 [TEXT_ANNOTATION_1] (action sequence)

1129 [TEXT_ANNOTATION_2] (action sequence)

1130 [TEXT_ANNOTATION_3] (action sequence)

1131 Now you are in the scene with [SCENE DESCRIPTIONS]. Your action sequence is:
 1132

1133 **D.7. GPT-4**

1134 We use “gpt-4-turbo-2024-04-09” with a temperature of 0.05. The textual prompts are as follows:
 1135

1136 “Stage One / Preference Prediction”

1137 You are a robot assistant that can help summarize the host's preference.

1138 Please read the following text file and summarize the user's preference.

1139 All possible preferences are: {ALL POSSIBLE PREFERENCES}

1140 [TEXT_ANNOTATION_1] The preference is [PREFERENCE_1]

1141 [TEXT_ANNOTATION_2] The preference is [PREFERENCE_2]

1142 [TEXT_ANNOTATION_3] The preference is [PREFERENCE_3]

1143 Now, please summarize the preference from the last tet file: [TEST_CASE]

1144 Quesiton: What's the user's preference? Choose from the preference listed before:
 1145

1146 “Stage Two / Planning”

1147 You are a robot assistant. Please read the following text files and help generate action sequence.

1148 All possible preferences are: {ALL POSSIBLE ACTIONS}

1149 Now there are some previous video demos:

1150 [TEXT_ANNOTATION_1] (action sequence)

1151 [TEXT_ANNOTATION_2] (action sequence)

1152 [TEXT_ANNOTATION_3] (action sequence)

1153 Now you are in the scene with [SCENE DESCRIPTIONS]. Your action sequence is:
 1154