

Embodied Multi-Modal Agent trained by an LLM from a Parallel TextWorld

Yijun Yang^{1,5,4} Tianyi Zhou² Kanxue Li³ Dapeng Tao³ Lusong Li⁴
Li Shen^{4,*} Xiaodong He⁴ Jing Jiang⁵ Yuhui Shi^{1,*}

¹Southern University of Science and Technology ²University of Maryland, College Park
³Yunnan University ⁴JD Explore Academy ⁵University of Technology Sydney

<https://github.com/stevenyangyj/Emma-Alfworld>

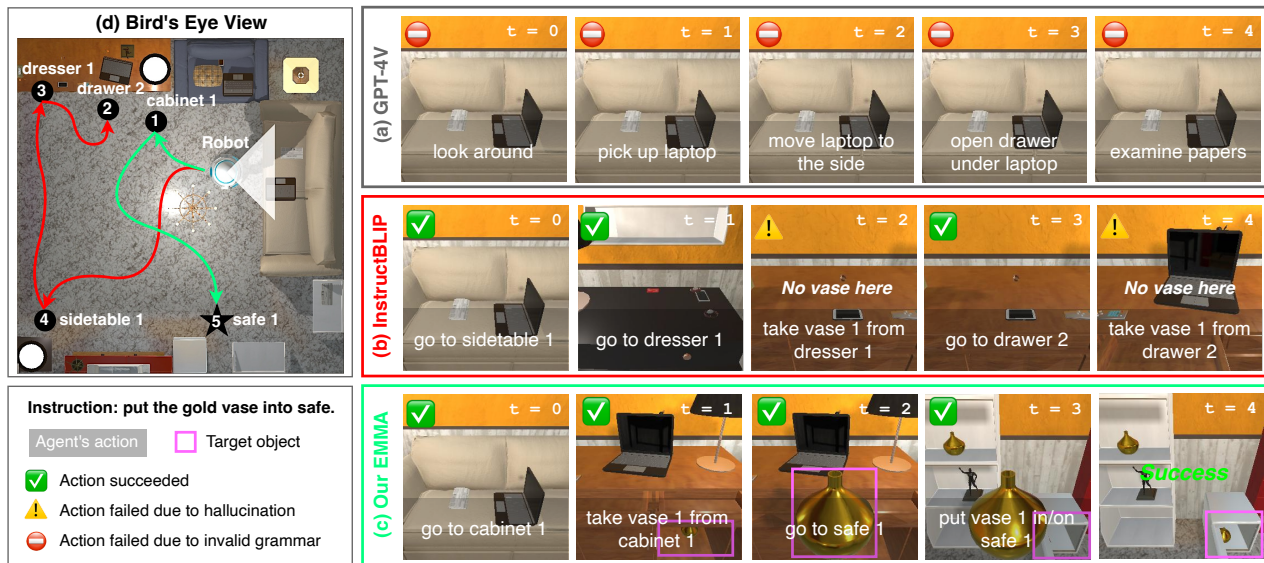


Figure 1: **Comparison of three VLM-based agents in the visual environment from ALFWorld.** Given the task instruction and the current-step observation as inputs, a VLM agent is expected to predict an action, e.g., “go to cabinet 1”, towards completing the task. (a) GPT-4V(ision). (b) InstructBLIP finetuned by behavior cloning on a static dataset of demonstrations by a rule-based expert in the visual simulator. (c) Our Embodied Multi-Modal Agent (EMMA) trained via cross-modality interactive imitation learning from an LLM expert.

Abstract

While large language models (LLMs) excel in a simulated world of texts, they struggle to interact with the more realistic world without perceptions of other modalities such as visual or audio signals. Although vision-language models (VLMs) integrate LLM modules (1) aligned with static image features, and (2) may possess prior knowledge of world dynamics (as demonstrated in the text world), they have not been trained in an embodied visual world and thus cannot align with its dynamics. On the other hand, training an embodied agent in a noisy visual world without expert guidance is often challenging and inefficient. In this paper, we train a VLM agent living in a visual world using an LLM agent excelling in a parallel text world. Specifically, we distill LLM’s reflection outcomes (improved actions by analyzing mistakes) in

a text world’s tasks to finetune the VLM on the same tasks of the visual world, resulting in an Embodied Multi-Modal Agent (EMMA) quickly adapting to the visual world dynamics. Such cross-modality imitation learning between the two parallel worlds is achieved by a novel Dagger-DPO algorithm, enabling EMMA to generalize to a broad scope of new tasks without any further guidance from the LLM expert. Extensive evaluations on the ALFWorld benchmark’s diverse tasks highlight EMMA’s superior performance to SOTA VLM-based agents, e.g., 20%-70% improvement in the success rate.

1. Introduction

Embodied multi-modal agents have been acknowledged as a pivotal stride towards achieving Artificial General Intelligence (AGI), as they encompass the potential for a broad scope of intelligent activities [27, 63]. The rising of founda-

*Corresponding author

tion models brings a glimmer of hope for constructing such agents [11, 32, 34, 35, 38, 71, 76], and notable efforts from the community have tried to harness them in many decision-making scenarios, e.g., autonomous driving [15, 66], daily household robots [6, 13, 77], and complex manipulation tasks [25, 29, 42, 57].

LLMs can perform as reflex agents interacting with a text world via verbalized descriptions of the world states and textual actions. In addition, they exhibit great potential in planning, reflection, and reward shaping. This is mainly due to their prior knowledge and semantic abstraction of the world. However, LLM-based agents cannot be directly applied in a visual world. Although current vision-language models (VLMs) try to align LLMs with the visual modality, their pretraining only focuses on static alignment between image-text pairs so the resulting agents have not been aligned well with the dynamics of the visual world. As shown in Fig. 1 (a), even the privileged VLM, e.g., GPT-4V [43], fails to accomplish tasks in an embodied ALFWorld environment [56]. In such a zero-shot setting, GPT-4V tends to mainly rely on its language prior of these objects detected in the current step, rather than the alignment between the visual input and the environment dynamics conditioned on the task instruction.

In this paper, we study how to train a VLM towards an embodied agent in a visual world by aligning it with the visual world dynamics and distilling the skills of an LLM agent in a parallel text world. Our overarching goal is to build such an **Embodied Multi-Modal Agent (EMMA)** that can take a textual task instruction (e.g., from human users) and pixel observations of the state per step to produce a sequence of actions leading to the efficient completion of the task. This is a challenging problem due to (1) the sparsity of task reward [2–4, 14], (2) noisy visual representations, (3) the hallucination of VLMs [60], and (4) the misalignment of VLM’s static representations to the visual world dynamics. While offline distillation and imitation from a powerful LLM agent in a parallel text world can potentially overcome the former two challenges [75], effectively mitigating the remaining challenges necessitates online finetuning of the VLM agent within an interactive and visual world [18, 31, 70].

To this end, we finetune a VLM agent by imitation learning from an LLM expert (e.g., built on ChatGPT) launched in a parallel text world on the same tasks. Specifically, in each step of EMMA interacting with the visual world, we convert its visual observation into an equivalent textual description sent to the LLM agent, which produces an action for EMMA to imitate. Such cross-modality interactive imitation learning is based on DAgger [51], which overcomes the cumulative errors and distribution shifts caused by behavior cloning (BC). As depicted in Fig. 1 (b), an InstructBLIP [11] agent finetuned by BC on 170K expert demonstrations produced by a rule-based expert in the

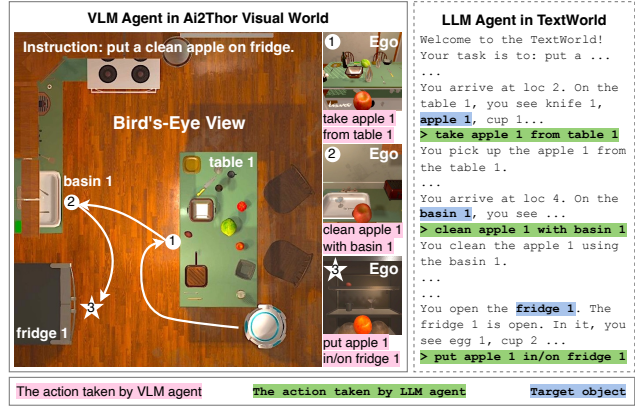


Figure 2: **An example of tasks generated for the two parallel worlds.** A VLM agent in the visual world and an LLM agent in the text world as household robots instructed to clean an apple and then put it into the fridge. Zoom in for more details.

visual world still fails to take correct actions based on visual observations. We further improve the DAgger’s objective to be the direct preference optimization (DPO) [48], which maximizes the preference of LLM-expert’s action (positive) over VLM-student’s action (negative) in each interaction step. To collect better teaching signals retrospective to the VLM student’s actions, the LLM expert is composed of an LLM actor prompted to output expert actions, and an LLM critic prompted for reflection feedback on the VLM agent’s historical trajectories. We maintain a long-term memory storing the feedback, which is then used to induce the LLM actor to improve actions for imitation in future episodes.

We evaluate EMMA and compare it with vision-only agents, LLM agents, and VLM agents deployed to the ALFWorld benchmark [56], which includes numerous tasks in both visual and textual environments. Extensive evaluations highlight that EMMA substantially outperforms state-of-the-art (SOTA) VLM agents in visual-only environments by 20%-70% in terms of success rate. In addition, EMMA is the only VLM agent that can generalize to open-vocabulary and free-form test tasks, shedding novel insights on using LLM feedback to train more versatile and generalizable embodied agents in multi-modality environments.

2. Embodied Multi-Modal Agent

Fig. 3 illustrates the main idea of our “Embodied Multi-Modal Agent (EMMA)”, whose detailed training procedures are given in Alg. 1. The agent is built upon a modularized VLM, which can follow instructions and interact with the environment through pixel observations and textual actions. To overcome these challenges associated with training EMMA, such as sparse reward or distribution shift, we explore the construction of an LLM expert from a parallel text world (Sec. 2.2) for providing EMMA with step-by-step guidance.

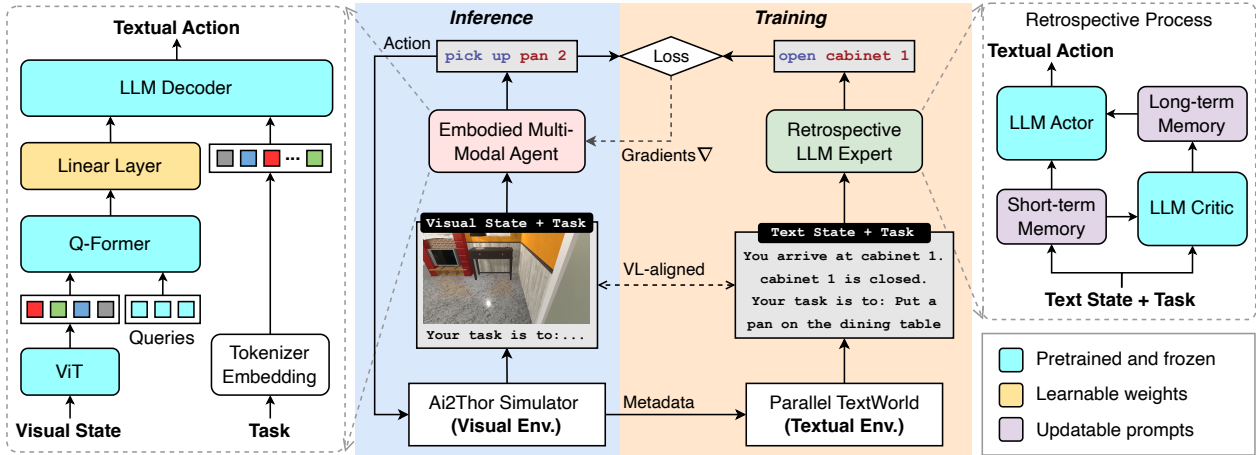


Figure 3: **Embodied Multi-Modal Agent (EMMA) trained by an LLM expert via cross-modality imitation learning.** EMMA takes a textual task instruction and pixel observations as its input state per step to generate a sequence of actions using a VLM. Then, we convert each pixel observation into a textual equivalent as the context of an LLM expert to produce improved actions for EMMA to imitate.

Lastly, Sec. 2.3 further discusses how to harness the LLM expert to train EMMA via cross-modality imitation learning.

2.1. EMMA in the Visual World

In visual environments, EMMA π_θ is designed to process a task instruction x_{task} (e.g., provided by human users) and the pixel observation s_v^t at each interaction step t . It is expected to generate a sequence of high-level textual actions $\{x_a^t \sim \pi_\theta(\cdot | x_{\text{task}}, s_v^t)\}_{t=0}^T$ ¹ towards efficient completion of the task. To achieve this, we draw inspiration from recent advances of large pretrained VLMs [11, 32, 34, 38, 76], and modularize EMMA’s architecture into four components: (1) a ViT to encode s_v into visual embeddings, (2) a querying transformer (Q-Former) tailored to extract the most relevant visual features via the cross-attention between the visual embeddings and query tokens, (3) a linear projection layer to align visual features to text embeddings, (4) an LLM decoder taking the concatenation of the instruction tokens and the output of the linear projection layer to autoregressively generate the action x_a . In order to reduce computational overhead and prevent catastrophic forgetting [37], we adopt the pretrained ViT, Q-Former, and LLM from InstructBLIP [11] and keep them frozen at the finetuning stage. We only update the linear projection layer, as illustrated in Fig. 3. Such a modularized architecture enables EMMA to integrate any existing pretrained vision models and LLMs in a flexible and computationally-efficient way.

However, deploying EMMA into a complex visual world remains challenging. One of the main obstacles is that the direct use of any pretrained VLM is suboptimal because existing pretraining only focuses on static alignment between

image-text pairs [11, 32, 34], so that the resulting agent may struggle to reason about the dynamics of the world. As discussed in Sec. 1, even the SOTA VLM, i.e., GPT-4V, fails to accomplish tasks in an embodied ALFWorld environment [56]. In such a zero-shot setting, GPT-4V tends to rely on the linguistic prior of the currently detected objects, rather than the given task instruction and the underlying dynamics of the environment, to guide the interactions. Moreover, finetuning a pretrained VLM on a pre-collected demonstration dataset is also suboptimal due to the diversity of environments and tasks [77], the lack of large-scale expert annotations [44] as well as the challenges posed by the distribution shift issue [31]. A seemingly natural solution to the above challenges is reinforcement learning from environmental feedback (RLEF) [68], in which reward signals rely on decomposing a task into a sequence of reasonable sub-goals and checking their completion. However, in real-world scenarios, most sub-goals cannot be defined or described precisely, so the reward is sparse; hence, we do not expect RLEF to be effective.

To this end, we propose to leverage interactive imitation learning (IL) to align EMMA with the dynamics of any environment, which however results in two critical algorithmic challenges: (1) How to obtain a high-quality, accessible, and scalable expert that EMMA can query during IL (Sec. 2.2)? (2) Designing an effective strategy to train EMMA using this expert in complex, diverse, and potentially open-ended environments (Sec. 2.3).

2.2. LLM Expert from a Parallel TextWorld

Thanks to a series of prompting techniques in the realm of in-context learning, such as chain-of-thought [62], tree-of-thought [72], ReAct [73], and Reflexion [54], pretrained

¹For brevity, we omit x_{task} in the rest of this paper.

Algorithm 1 DAgger-DPO with a single task instruction

- 1: **initialize:** $i = 0$, $\mathcal{D} \leftarrow \emptyset$, EMMA π_θ , LLM actor M_a with a FIFO memory pool $\mathcal{P} \leftarrow \emptyset$, LLM critic M_c
 - 2: **input:** max trials I , training epochs I_e , visual and textual envs E_v, E_l , instruction x_{task} , reference agent π_{ref}
 - 3: Initialize π_θ to π_{ref} \triangleright **Behavior Cloning Initialization**
 - 4: **while** task not completed or $i < I$ **do**
 - 5: Get $\tau_v^i = [x_{\text{task}}, s_v^0, x_a^0, \dots, s_v^T, x_a^T]$ via E_v with π_θ
 - 6: Get $\tau_l^i = [x_{\text{task}}, s_l^0, x_a^0, \dots, s_l^T, x_a^T]$ via E_l with τ_v^i
 - 7: Generate retrospective feedback $\mathcal{P}_i = M_c(\tau_l^i)$
 - 8: Update \mathcal{P} with \mathcal{P}_i (i.e., $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{P}_i$)
 - 9: **for** $t = 0$ to T **do** \triangleright **Dataset Aggregation**
 - 10: $x_a^* = M_a(\mathcal{P}, x_{\text{task}}, \dots, x_a^{t-1}, s_l^t)$
 - 11: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(x_{\text{task}}, s_v^t, x_a^t, x_a^*)\}$
 - 12: **for** $j = 0$ to $I_e - 1$ **do** \triangleright **Gradient Descent on θ**
 - 13: Sample a mini-batch τ from \mathcal{D}
 - 14: Update θ by minimizing Eq. (2) with π_{ref} and τ
 - 15: $i \leftarrow i + 1$
 - 16: **output:** π_{θ^*}
-

LLMs have demonstrated impressive zero-shot performance across many decision-making scenarios [6, 13, 15, 25, 29, 42, 57, 66, 77]. Despite the great potential in serving as high-quality and scalable experts, they are only able to interact with the environment via textual descriptions of the states, rather than using raw pixel observations like EMMA. To bridge this gap, we convert each pixel observation s_v into a textual equivalent by extracting its metadata from the simulator [30], which is composed of attributes such as Observed Objects, Observed Relations, Inventory, and Locations. We then employ the Planning Domain Definition Language (PDDL) [1] to describe this metadata and create an equivalent textual description/state s_l using the TextWorld engine [10]. Additional details are available in Appendix 7 and an example of this process is illustrated in Fig. 2. This methodology enables the utilization of any pretrained LLM agent that generates a sequence of actions facilitating the training of EMMA through cross-modality IL between the two agents.

2.3. Training EMMA via Cross-Modality Imitation

Given an LLM expert from the parallel TextWorld, we aim to train a VLM agent π_θ in the visual world to imitate its behaviors closely. This is equal to minimizing the following objective under the distribution of states induced by π_θ .

$$\theta^* = \arg \min_{\theta \in \Theta} \mathbb{E}_{\pi_\theta} [\mathcal{L}_{\text{imit}}(\pi_\theta(x_a | s_v), x_a^*)], \quad (1)$$

in which the choice of loss function $\mathcal{L}_{\text{imit}}$ is dependent on specific scenarios. For instance, it may be the expected cross-entropy loss for the discrete action space, or the expected

MSE loss for the continuous action space. In our case, we select DPO [48] loss due to its proven superior performance to the cross-entropy on aligning models with expert preferences within the discrete language space. Hence, Eq. (1) can be extended to the formulation below.

$$\theta^* = \arg \min_{\theta \in \Theta} -\mathbb{E}_{\pi_\theta} [\mathcal{L}_{\text{imit}}(\pi_\theta, \pi_{\text{ref}}, s_v, x_a, x_a^*)], \quad (2)$$
$$\mathcal{L}_{\text{imit}}(\cdot) \triangleq \log \sigma \left(\beta \log \frac{\pi_\theta(x_a^* | s_v)}{\pi_{\text{ref}}(x_a^* | s_v)} - \beta \log \frac{\pi_\theta(x_a | s_v)}{\pi_{\text{ref}}(x_a | s_v)} \right)$$

where x_a^* is the action given by an expert while x_a is the action given by π_θ , σ is the logistic function, and β is a hyperparameter controlling the deviation from π_{ref} , i.e. the reference agent obtained by behavior cloning on a demonstration dataset produced by a rule-based expert in the visual world (see Appendix 9 for complete details). This regularization is essential, as it prevents the agent from deviating too far from the distribution on which the expert is accurate, as well as maintaining the generation diversity and avoiding premature convergence to some easy tasks [48]. In practice, the VLM agent π_θ is also initialized to π_{ref} for stabilizing the training process. Since environmental dynamics is both unknown and complex, we cannot compute the distribution of states visited by π_θ and can only sample it by rolling out the agent. Hence, Eq. (2) is a non-i.i.d. supervised learning problem due to the dependence of the state distribution on π_θ itself, in which naïve behavior cloning faces issues like cumulative error and distribution shift [26]. To address this, we employ an interactive IL algorithm, DAgger [51], which provably converges to the optimal agent π_{θ^*} .

As discussed in Sec. 2.2 and Sec. 2.3, we can harness an LLM expert to generate a sequence of actions, serving as x_a^* in Eq. (2). However, these actions may be suboptimal due to sparse environmental feedback [54, 68] or defective in-context instructions [9]. To collect better teaching signals, we introduce a “retrospective LLM expert” that is composed of two specialized models: an actor (M_a), built upon an API LLM and prompted to generate actions based on the task instruction and state observations; and a critic (M_c), also based on the same LLM, but designed to analyze EMMA’s historical interactions and provide reflective feedback. A long-term memory \mathcal{P} is maintained to store the feedback generated by M_c , which is then used to prompt M_a for improved actions. The complete procedure is detailed in Line 7-10 of Alg. 1, and all prompts are provided in Sec. 6 of the Appendix.

3. Experiments

3.1. Experimental Setup

Environments. We base our experiments on the ALFWorld benchmark [56], a cross-modality simulation platform that encompasses a wide range of embodied household tasks. For




Figure 4: **Comparison of success rate between EMMA and the retrospective LLM expert.** As the number of trials increases, the LLM expert progressively improves its performance through retrospective processes, and also the gap between the two agents decreases, indicating the effectiveness of cross-modality imitation. We separately plot the curve per task type in Fig. 9 of Appendix.

each task, ALFWorld integrates a visual environment, rendered by the Ai2Thor simulator [30], with a corresponding textual environment. Textual environments employ the Planning Domain Definition Language (PDDL) [1] to translate each pixel observation from the simulator into an equivalent text-based observation, and then construct an interactive world using the TextWorld engine [10]. Fig. 2 provides an illustrative example of the tasks featured in ALFWorld. The tasks within the benchmark are categorized into 6 types: Pick & Place, Clean & Place, Heat & Place, Cool & Place, Look in Light, and Pick Two Objects & Place. Each task requires an agent to execute a series of text-based actions, such as “go to safe 1”, “open safe 1”, or “heat egg 1 with microwave 1”, following a predefined instruction. These actions involve navigating and interacting with the environment. To provide a comprehensive understanding, we have visualized an example of each task type in Fig. 8 of the Appendix. A task in this benchmark may involve interactions with over 10 objects and require more than 30 steps for a human expert to solve, thus challenging an agent’s capabilities in long-horizon planning, instruction following, and the utilization of commonsense knowledge. For a fair comparison, we follow the same setting as prior work [41, 54–56, 73] and evaluate all baselines using 134 out-of-distribution (OOD) tasks.

Baselines. To verify the effectiveness of cross-modality imitation learning, we compare our EMMA with several baselines and SOTA agents using the ALFWorld benchmark with both visual and textual environments. The compared agents can be divided into three categories: vision models, language models, and vision-language models. Concretely, vision models, including ResNet-18 [20] and MCNN-FPN [21], utilize pretrained vision encoders to extract salient features

Agent	Visual Env.	Textual Env.	Success Rate with Template Task Instruction							
			Avg.	Pick	Clean	Heat	Cool	Look	Pick2	
VMs	ResNet-18* [56]	✓	✗	0.06(-)	-	-	-	-	-	-
	MCNN-FPN* [56]	✓	✗	0.05(-)	-	-	-	-	-	-
LMs	BUTLER* [56]	✗	✓	0.26(-)	0.31(-)	0.41(-)	0.60(-)	0.27(-)	0.12(-)	0.29(-)
	GPT-BUTLER [41]	✗	✓	0.69(18.8)	0.62(18.4)	0.81(18.4)	0.85 (12.7)	0.78(15.6)	0.50(24.4)	0.47(26.6)
	ReAct [73]	✗	✓	0.54(20.6)	0.71(18.1)	0.65(18.8)	0.62(18.2)	0.44(23.2)	0.28(23.7)	0.35(25.5)
	Reflexion [54]	✗	✓	0.91 (18.7)	0.96 (17.4)	1.00 (17.0)	0.81(19.4)	0.83(21.6)	0.94(16.9)	0.88 (21.6)
	DEPS* [61]	✗	✓	0.76(-)	0.93(-)	0.50(-)	0.80(-)	1.00 (-)	1.00 (-)	0.00(-)
	AutoGen* [64]	✗	✓	0.77(-)	0.92(-)	0.74(-)	0.78(-)	0.86(-)	0.83(-)	0.41(-)
VLMs	MiniGPT-4 [76]	✓	✗	0.16(26.9)	0.04(29.0)	0.00(30.0)	0.19(26.3)	0.17(26.7)	0.67(17.7)	0.06(28.9)
	BLIP-2 [34]	✓	✗	0.04(29.5)	0.00(30.0)	0.06(29.3)	0.04(29.9)	0.11(28.2)	0.06(29.2)	0.00(30.0)
	LLaMA-Adapter [17]	✓	✗	0.13(27.5)	0.17(26.4)	0.10(28.6)	0.27(24.2)	0.22(26.7)	0.00(30.0)	0.00(30.0)
	InstructBLIP [11]	✓	✗	0.22(26.2)	0.50(21.5)	0.26(25.0)	0.23(27.2)	0.06(28.9)	0.17(26.8)	0.00(30.0)
	EMMA (Ours)	✓	✗	0.82(19.5)	0.71 (19.3)	0.94 (17.5)	0.85 (19.6)	0.83 (19.9)	0.88 (19.6)	0.67 (22.4)
Human Performance* [55]			✓	✗	0.91 (-)	-	-	-	-	-

Table 1: **Comparison with the state of the arts.** *-reported in previous work. VMs = vision models, LMs = language models, VLMs = vision-language models. “Visual Env.” and “Textual Env.” refer to the visual and the parallel textual environments from ALFWorld [56], respectively. ✓✗ denotes the corresponding environment used/not used to deploy the agent. The highest scores for each task in the same type of environment are highlighted in **bold**. The average interaction steps are given in the (parentheses). **EMMA substantially outperforms other SOTA VLM agents in the visual environments, and its success also directs a promising way to achieve human-level performance in ALFWorld.**

from each pixel observation. The extracted features then serve as input for a Multi-Layer Perceptron (MLP) policy, which is trained by behavior cloning on a pre-collected demonstration dataset. Unlike vision models performing in the visual environment, language models complete the same tasks but in a parallel, text-based environment. BUTLER [56] employs a transformer seq2seq model enhanced with a pointer softmax mechanism [19]. This architecture aggregates previous observations as input to generate text-based actions token-by-token. GPT-BUTLER [41], a variant of the GPT-2 model [45], is initially pretrained on a static demonstration dataset and further finetuned using data collected online. ReAct [73] takes a novel approach by utilizing LLMs to generate reasoning traces and task-specific actions in an interleaved manner. This method aids the agent in developing, tracking, and updating its action plans interactively. Reflexion [54] similarly employs an LLM, but it focuses on reflecting upon environmental feedback. It maintains this reflective text in an episodic memory buffer, enhancing the agent’s ability to improve actions in subsequent trials. Similar to Reflexion, a concurrent work, DEPS [61], also corrects errors in previous LLM-generated actions by integrating descriptions of the action execution process and providing self-explanations for the feedback. Moreover, beyond the single-agent framework, AutoGen [64] exhibits the potential to accomplish a broad spectrum of tasks through the cooperation of multiple LLM agents. Finally, we consider a range of vision-language models, such as MiniGPT-4 [76], BLIP-2 [34], LLaMA-Adaptor [17], and InstructBLIP [11], as agents to interact with the visual environment. Unlike pure vision or language models, VLMs are designed to process and integrate both visual and textual data, offering a

more holistic understanding of the environment. To align these agents with the specific requirements of the ALFWorld benchmark, we finetune them on a pre-collected demonstration dataset. This finetuning process is crucial as it enables the agent to comprehend and adhere to ALFWorld’s unique grammar and to develop a basic “gamesense”.

3.2. Training Details

The architectural design of EMMA is depicted in Fig. 3. At its core, EMMA employs a Query Transformer (Q-Former) to process visual data. This Q-Former extracts features using a frozen ViT encoder. Its output consists of 32 visual tokens, which are then passed through a linear projection layer before being fed to a frozen LLM decoder. Similar to other VLM agents, EMMA is also finetuned on a pre-collected demonstration dataset, aligning its basic ability with the ALFWorld benchmark. In order to align EMMA with the dynamics of ALFWorld, we train it by imitating an LLM expert (see Alg. 1). We choose *text-davinci-003* developed by OpenAI as our LLM expert because of its established capabilities in reasoning and planning [54, 61, 64, 73]. In this setup, *text-davinci-003* serves dual roles: it serves as an actor, providing EMMA with expert actions, and as a critic, analyzing EMMA’s historical trajectories. This analysis generates retrospective feedback, which is then incorporated into the actor’s long-term memory, leading to improved actions in future trials. Further details about the hyperparameters and prompts used in our training procedure are available in Table 2 of the Appendix.

3.3. Comparison with State of the Art

EMMA sets new SOTA performance in visual ALFWorld. In this section, we compare EMMA with 12 other representative agents on the ALFWorld benchmark, spanning both visual and textual environments (refer to Table 1). We assess two key metrics: the success rate, which is the percentage of trials completed successfully, and the average number of interaction steps required for task completion, with a lower number indicating higher efficiency. EMMA demonstrates superior performance in both metrics, significantly outperforming all VLM agents in visual environments. This achievement underscores the effectiveness of our cross-modality imitation learning approach, as depicted in the learning curve shown in Fig. 4. Furthermore, EMMA’s performance markedly exceeds that of VM agents, highlighting the crucial role of the prior knowledge embedded in VLMs. Intriguingly, EMMA’s performance is comparable with LLM agents that operate using perfectly semantic descriptions of visual observations. This is largely attributed to EMMA’s training strategy of imitating an expert LLM agent, proving to be more efficient than learning from scratch in a purely visual setting. As a result, EMMA stands out as the only VLM agent that substantially surpasses SOTA LLM agents, such




Figure 5: **Comparison of robustness between VLM and LLM agents.** Given a specific noise rate, we crop a random portion of the pixel observation and resize it to a given size as the new observation. Similarly, we randomly replace some tokens in the textual observation with arbitrary ones.

as AutoGen [64] and DEPS [61], in these environments. And its success also directs a potential way to achieve human-level performance in the visual environments of ALFWorld. **EMMA is more robust to noisy observations than LLM agents.** While LLM agents exhibit a higher success rate with fewer interaction steps in textual environments, as indicated in Table 1, we hypothesize that this superior performance largely relies on their precise semantic abstraction of the environment. However, such an abstraction might not be feasible in real-world applications. To verify this assumption, we set up a more practical scenario where observations are deliberately perturbed at a specific noise rate. We then compare the robustness of EMMA and a SOTA LLM agent, Reflexion, under these noisy observations. To generate noisy observations, a random portion of the visual observation is cropped, and then used to replace the original observation. Similarly, in the textual observation, random tokens are substituted with arbitrary ones. As illustrated in Fig. 5, with the noise rate increasing, EMMA’s performance remains significantly more robust than Reflexion. This could be attributed to the vision encoder in the VLM, which is adept at filtering out visual noises. On the other hand, textual noises are directly processed by the LLM, which can substantially impair the performance of LLM-based agents. This finding highlights the potential advantages of VLM agents like EMMA in visual scenarios, in which data is often imperfect and noisy.

3.4. Ablation Study

Retrospection improves EMMA over time. To assess the importance of the retrospective LLM expert, we present the average success rate of EMMA after each trial and compare it




Figure 6: **Ablation study.** *left:* The ability of LLM expert to reflect on and learn from past actions appears to be a key factor in achieving the impressive results we observed. *middle:* The “w/o BC initialization” variant uses a pretrained VLM instead of the reference agent π_{ref} to initialize EMMA. *right:* The ‘w/ CE Loss’ variant replaces the DPO loss (described in Eq. 2) with a token-level cross-entropy (CE) loss. Please refer to Sec. 3.4 for a thorough discussion.

with a key variation: “EMMA w/o Retrospection”, as shown in Fig. 6 (left). This variant of EMMA is trained using the same procedure as the original EMMA but removes the retrospective process. Instead, it relies solely on a plain LLM actor to provide relabeled actions. The results show that EMMA with the retrospective mechanism significantly outperforms its counterpart. This finding is crucial as it indicates that the retrospective process is not just a supplementary feature but a fundamental component of EMMA’s architecture that contributes substantially to its enhanced performance.

EMMA benefits from BC initialization. Through an ablation study, we evaluate the impact of behavior cloning (BC) initialization, a process described in line 3 of Alg. 1. Fig. 6 (middle) demonstrates that EMMA, when deprived of BC initialization, experiences a slight reduction in the average success rate across 134 unseen tasks compared to its original setup. Despite this decrease, EMMA without BC initialization still outperforms other VLM agents, as clearly shown when compared with the results in Table 1. Furthermore, Fig. 10 in the Appendix breaks down the result by task type. It reveals a consistent but slight drop in performance across 5 out of the 6 task types. These results reflect that while BC initialization contributes positively to EMMA’s overall performance, it is not critical for achieving the notable results we have reported.

DPO enables more effective imitation learning. To evaluate the effectiveness of using DPO loss in Eq. (1), we conducted an ablation study with an alternative version of EMMA, referred to “EMMA w/ Cross Entropy (CE) Loss”. In this variant, EMMA is optimized using the token-level CE loss, a common objective for finetuning VLMs. The results, as depicted in Fig. 6 (right), reveal that EMMA w/ CE Loss does not achieve the same high success rate as the original EMMA with DPO loss, suggesting that DPO loss contributes to enhancing the upper performance bound of EMMA. In addition, we noted that EMMA w/ CE Loss exhibits faster convergence in the initial training stages than the original

EMMA. This premature convergence leads to the agent paying attention to the expert actions from easier tasks, usually addressed in the first few training epochs. This can suppress EMMA’s exploration and learning on more complex tasks, resulting in worse performance.

3.5. Generalization to Free-form Task Instructions

The ability of AI agents to precisely follow task instructions given in open-vocabulary and free-form text is crucial for real-world problems. To assess this, we conducted an additional experiment focusing on the generalization capabilities of EMMA, an agent trained with template task instructions. In this experiment, we re-evaluated the performance of EMMA and other baseline agents on 134 unseen tasks, using human-annotated instructions instead of template ones. These human-annotated instructions include a large amount of OOD verbs and objects, presenting a more realistic and challenging scenario for the agents to address. To underscore this challenge, we compared the vocabulary distribution between the template and human-annotated instructions, as shown in Fig. 7 (right). Moreover, we provide a comprehensive analysis of the vocabulary used across both instruction types in Fig. 11 of the Appendix. In Fig. 7 (left), EMMA demonstrates a slight performance decline, while a significant degradation is observed in other baselines. We also note that Reflexion, a SOTA LLM agent, exhibits exceptional generalization to those OOD instructions. According to these empirical results in Fig. 7, we have the following conclusions: (1) EMMA obtains and benefits from the generalization capabilities inherent in the SOTA LLM agent through cross-modality imitation learning; (2) Our work sheds novel insights on using LLM feedback to train more versatile and generalizable embodied agents.

4. Related Work

Agents based on Foundation Models. Recent research has increasingly focused on harnessing the capabilities of large

Agent	Visual Env.	Textual Env.	Success Rate with Free-form Task Instruction							
			Avg.	Pick	Clean	Heat	Cool	Look	Pick2	
LMs	BUTLER* [56]	✗	✓	0.03(-)	0.10(-)	0.22(-)	0.05(-)	0.17(-)	0.00(-)	0.00(-)
	GPT-BUTLER [41]	✗	✓	0.31(24.7)	0.25(25.4)	0.42(23.5)	0.46(20.9)	0.44(21.8)	0.00(30.0)	0.12(28.9)
	ReAct [73]	✗	✓	0.37(23.6)	0.46(22.0)	0.39(23.8)	0.50(20.5)	0.28(24.6)	0.22(26.4)	0.29(25.9)
	Reflexion [54]	✗	✓	0.78 (17.0)	0.84 (15.8)	0.74 (18.4)	0.77 (17.2)	0.61 (19.4)	0.94 (13.8)	0.82 (16.7)
VLMs	MiniGPT-4 [76]	✓	✗	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)
	BLIP-2 [34]	✓	✗	0.01(29.7)	0.04(29.1)	0.03(29.5)	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)
	LLaMA-Adapter [17]	✓	✗	0.02(29.6)	0.04(29.3)	0.03(29.4)	0.04(29.3)	0.00(30.0)	0.00(30.0)	0.00(30.0)
	InstructBLIP [11]	✓	✗	0.01(29.8)	0.00(30.0)	0.03(29.3)	0.00(30.0)	0.00(30.0)	0.00(30.0)	0.00(30.0)
	EMMA (Ours)	✓	✗	0.68 (22.0)	0.72 (21.7)	0.65 (22.7)	0.72 (22.0)	0.56 (23.9)	0.67 (23.3)	0.76 (18.1)




Figure 7: **Generalization to open-vocabulary, free-form task instructions.** *Left:* These instructions are provided by different human annotators using Amazon Mechanical Turk [55]. *Right:* Frequency distribution of verbs for human-annotated and templated task instructions.

pre-trained foundation models to build AI agents [65, 69]. These models (e.g., LLMs), benefiting from their common-sense knowledge inherited from Internet-scale pretraining, are able to reason actions according to descriptions of the external environments. For example, given a set of task instructions, LLMs can be elaborately prompted to perform as agents generating high-level step-by-step plans [7, 24, 54, 64, 73], and each step can be parsed into a sequence of robotic actions that are executed via pretrained policies or available APIs [6, 36, 59]. Furthermore, by using VLMs [32], plans can also be conditioned on visual inputs that are transformed into language descriptions [7, 16, 25, 57] or token embeddings aligned with LLMs [13, 42, 68, 77]. However, existing foundation models are usually pretrained on static text or text-image datasets and thus may struggle to align with the dynamics of the world. To bridge this gap, we study how to finetune a VLM to be an embodied agent dynamically aligned with the world by distilling the cross-modality knowledge from an LLM expert. The work most closely related to ours is EUREKA [40], which also explores using source information provided by the simulator as the context of an LLM to aid agent training. Instead of directly mimicking the output of the LLM as we did, EUREKA harnesses the coding LLM to generate a desired reward function for a given task and optimizes a policy against the reward function using RL, leading to a more complex and expensive training procedure [47, 52, 53].

Imitation Learning. Imitation learning is the study of algorithms that improve performance by mimicking an expert’s decisions and behaviors. We summarize three main categories of existing methods in the following: (1) behavior cloning (BC), (2) inverse reinforcement learning (IRL), and (3) the combination of imitation and reinforcement learning. The naïve BC [5] ignores the change in distribution and simply trains a policy that only performs well under the distribution of states visited by the expert. Following works, such as dataset aggregation [51] or policy aggregation [28, 49], have been proposed to address the limitations of BC. Another line of work, IRL, is a more complicated algorithm framework that learns the reward function from

expert demonstrations and then improves the policy using RL with the learned reward. A representative method in this category is generative adversarial imitation learning (GAIL) [23], in which a policy and a discriminator compete with each other in order to maximize the likelihood of the policy’s behavior matching the expert. The third category of methods usually leverages an IL policy to initialize RL and continues to boost its performance via online collected data from RL [8, 50, 58]. This simple combination significantly improves RL’s sample efficiency and IL’s upper performance bound constrained by the expert. Nevertheless, all of the above methods assume that the expert and the imitator understand the world in the same modality, and thus overlook the fact that the complementary knowledge from other modalities often boosts the model’s accuracy and generalization dramatically [46, 67].

5. Conclusion

We create **EMMA**, an **Embodied Multi-Modal Agent**, by finetuning a VLM in an embodied visual world with interactive imitation learning from an LLM expert in a parallel text world, who produces better actions and retrospective feedback to VLM’s trajectories. Such imitation learning exhibits substantial advantages over vision or VLM policies directly finetuned in the visual world or finetuned by behavior cloning of a rule-based expert, and SOTA API VLMs such as GPT-4V(ision). As a result, EMMA achieves a comparable success rate and much better robustness in the noisy visual world than its LLM teacher in the easier text world. Furthermore, EMMA shows powerful generalization to open-vocabulary and free-form task instructions, highlighting its potential in real-world scenarios.

Acknowledgment

This work is partially supported by the National Key R&D Program of China under the Grant No. (2023YFE0106300) and (2017YFC0804002), Shenzhen Fundamental Research Program under the Grant No. (JCYJ20200109141235597), and Program for Guangdong Introducing Innovative and Entrepreneurial Teams under Grant No. (2017ZT07X386).

References

- [1] Constructions Aeronautiques, Adele Howe, Craig Knoblock, ISI Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, David Wilkins SRI, Anthony Barrett, Dave Christian-son, et al. Pddl: the planning domain definition language. *Technical Report, Tech. Rep.*, 1998. 4, 5, 3
- [2] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. 2017. 2
- [3] Shuang Ao, Tianyi Zhou, Guodong Long, Qinghua Lu, Liming Zhu, and Jing Jiang. Co-pilot: Collaborative planning and reinforcement learning on sub-task curriculum. 2021.
- [4] Shuang Ao, Tianyi Zhou, Jing Jiang, Guodong Long, Xuan Song, and Chengqi Zhang. Eat-c: environment-adversarial sub-task curriculum for efficient reinforcement learning. In *ICML*, 2022. 2
- [5] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence*, 1995. 8
- [6] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 2, 4, 8
- [7] Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *CoRL*, 2023. 8
- [8] Kai-Wei Chang, Akshay Krishnamurthy, Alekh Agarwal, Hal Daumé III, and John Langford. Learning to search better than your teacher. In *ICML*, 2015. 8
- [9] Lichang Chen, Jiu-hai Chen, Tom Goldstein, Heng Huang, and Tianyi Zhou. Instructzero: Efficient instruction optimization for black-box large language models. *arXiv preprint arXiv:2306.03082*, 2023. 4
- [10] Marc-Alexandre Côté, Akos Kádár, Xingdi Yuan, Ben Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew Hausknecht, Layla El Asri, Mahmoud Adada, et al. Textworld: A learning environment for text-based games. In *Computer Games Workshop, IJCAI*, 2019. 4, 5
- [11] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *NeurIPS*, 2023. 2, 3, 5, 8
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 3
- [13] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 2, 4, 8
- [14] Meng Fang, Tianyi Zhou, Yali Du, Lei Han, and Zhengyou Zhang. Curriculum-guided hindsight experience replay. *NeurIPS*, 2019. 2
- [15] Daocheng Fu, Xin Li, Licheng Wen, Min Dou, Pinlong Cai, Botian Shi, and Yu Qiao. Drive like a human: Rethinking autonomous driving with large language models. *arXiv preprint arXiv:2307.07162*, 2023. 2, 4
- [16] Jensen Gao, Bidipta Sarkar, Fei Xia, Ted Xiao, Jiajun Wu, Brian Ichter, Anirudha Majumdar, and Dorsa Sadigh. Physically grounded vision-language models for robotic manipulation. *arXiv preprint arXiv:2309.02561*, 2023. 8
- [17] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. 5, 8, 3
- [18] Jiayi Guan, Guang Chen, Jiaming Ji, Long Yang, Zhijun Li, et al. Voce: Variational optimization with conservative estimation for offline safe reinforcement learning. *NeurIPS*, 36, 2023. 2
- [19] Çağlar Gülçehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. Pointing the unknown words. In *ACL*, 2016. 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, 2017. 5
- [22] Malte Helmert. The fast downward planning system. *Journal of Artificial Intelligence Research*, 2006. 3
- [23] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *NeurIPS*, 2016. 8
- [24] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *ICML*, 2022. 8
- [25] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. *arXiv preprint arXiv:2307.05973*, 2023. 2, 4, 8
- [26] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017. 4
- [27] Marcus Hutter. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer Science & Business Media, 2004. 1
- [28] Hal Daumé III, John Langford, and Daniel Marcu. Search-based structured prediction. *Mach. Learn.*, 2009. 8
- [29] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. VIMA: robot manipulation with multimodal prompts. In *ICML*, 2023. 2, 4
- [30] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. 4, 5, 1

- [31] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020. 2, 3
- [32] Chunyuan Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, and Jianfeng Gao. Multimodal foundation models: From specialists to general-purpose assistants. *arXiv preprint arXiv:2309.10020*, 2023. 2, 3, 8
- [33] Dongxu Li, Junnan Li, Hung Le, Guangsen Wang, Silvio Savarese, and Steven C.H. Hoi. LAVIS: A one-stop library for language-vision intelligence. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, 2023. 3
- [34] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 2, 3, 5, 8
- [35] Chen Liang, Jiahui Yu, Ming-Hsuan Yang, Matthew Brown, Yin Cui, Tuo Zhao, Boqing Gong, and Tianyi Zhou. Module-wise adaptive distillation for multimodality foundation models. 2024. 2
- [36] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *ICRA*, 2023. 8
- [37] Yong Lin, Lu Tan, Hangyu Lin, Zeming Zheng, Renjie Pi, Jipeng Zhang, Shizhe Diao, Haoxiang Wang, Han Zhao, Yuan Yao, et al. Speciality vs generality: An empirical study on catastrophic forgetting in fine-tuning foundation models. *arXiv preprint arXiv:2309.06256*, 2023. 3
- [38] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 2, 3
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 3
- [40] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931*, 2023. 8
- [41] Vincent Micheli and François Fleuret. Language models are few-shot butlers. In *EMNLP*, 2021. 5, 8
- [42] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *arXiv preprint arXiv:2305.15021*, 2023. 2, 4, 8
- [43] OpenAI. GPT-4V(ision) system card. https://cdn.openai.com/papers/GPTV_System_Card.pdf, 2023. 2
- [44] Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, et al. Open x-embodiment: Robotic learning datasets and rt-x models, 2023. 3
- [45] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 5
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 8, 3
- [47] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023. 8
- [48] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *NeurIPS*, 2024. 2, 4
- [49] Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *AISTATS*, 2010. 8
- [50] Stéphane Ross and J Andrew Bagnell. Reinforcement and imitation learning via interactive no-regret learning. *arXiv preprint arXiv:1406.5979*, 2014. 8
- [51] Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011. 2, 4, 8
- [52] John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In *ICML*, 2015. 8
- [53] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 8
- [54] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *NeurIPS*, 2023. 3, 4, 5, 6, 8
- [55] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. ALFRED: A benchmark for interpreting grounded instructions for everyday tasks. In *CVPR*, 2020. 5, 8, 3, 6
- [56] Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew J. Hausknecht. Alfworld: Aligning text and embodied environments for interactive learning. In *ICLR*, 2021. 2, 3, 4, 5, 8, 1
- [57] Theodore Sumers, Kenneth Marino, Arun Ahuja, Rob Fergus, and Ishita Dasgupta. Distilling internet-scale vision-language models into embodied agents. In *ICML*, 2023. 2, 4, 8
- [58] Wen Sun, J. Andrew Bagnell, and Byron Boots. Truncated horizon policy search: Combining reinforcement learning & imitation learning. In *ICLR*, 2018. 8
- [59] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023. 8
- [60] Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Yuancheng Xu, Feihong He, Jaehong Yoon, Taixi Lu, Gedas Bertasius, Mohit Bansal, et al. Mementos: A comprehensive benchmark for multimodal large language model reasoning over image sequences. *arXiv preprint arXiv:2401.10529*, 2024. 2

- [61] Zihao Wang, Shaofei Cai, Anji Liu, Xiaojian Ma, and Yitao Liang. Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents. *arXiv preprint arXiv:2302.01560*, 2023. 5, 6
- [62] 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 *NeurIPS*, 2022. 3
- [63] Michael Wooldridge and Nicholas R Jennings. Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2):115–152, 1995. 1
- [64] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*, 2023. 5, 6, 8
- [65] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*, 2023. 8
- [66] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kenneth KY Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*, 2023. 2, 4
- [67] Zihui Xue, Zhengqi Gao, Sucheng Ren, and Hang Zhao. The modality focusing hypothesis: Towards understanding cross-modal knowledge distillation. In *ICLR*, 2023. 8
- [68] Jingkan Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Chencheng Jiang, Haoran Tan, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, et al. Octopus: Embodied vision-language programmer from environmental feedback. *arXiv preprint arXiv:2310.08588*, 2023. 3, 4, 8
- [69] Sherry Yang, Ofir Nachum, Yilun Du, Jason Wei, Pieter Abbeel, and Dale Schuurmans. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023. 8
- [70] Yijun Yang, Jing Jiang, Tianyi Zhou, Jie Ma, and Yuhui Shi. Pareto policy pool for model-based offline reinforcement learning. In *ICLR*, 2022. 2
- [71] Yijun Yang, Tianyi Zhou, Jing Jiang, Guodong Long, and Yuhui Shi. Continual task allocation in meta-policy network via sparse prompting. In *ICML*, 2023. 2
- [72] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023. 3
- [73] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *ICLR*, 2023. 3, 5, 6, 8, 2
- [74] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023. 3
- [75] Victor Zhong, Dipendra Misra, Xingdi Yuan, and Marc-Alexandre Côté. Policy improvement using language feedback models. *arXiv preprint arXiv:2402.07876*, 2024. 2
- [76] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 2, 3, 5, 8
- [77] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *CoRL*, 2023. 2, 3, 4, 8