

ON THE STRENGTHS AND WEAKNESSES OF DATA FOR OPEN-SET EMBODIED ASSISTANCE

Pradyumna Tambwekar*

Distyl AI
pradyumna.tambwekar@distyl.ai

Andrew Silva

Toyota Research Institute
andrew.silva@tri.global

Deepak Gopinath

Toyota Research Institute
deepak.gopinath@tri.global

Jonathan DeCastro

Toyota Research Institute
jonathan.decastro@tri.global

Xiongyi Cui

Toyota Research Institute
xiongyi.cui@tri.global

Guy Rosman

Toyota Research Institute
guy.rosman@tri.global

ABSTRACT

Embodied foundation models are increasingly performant in real-world domains such as robotics or autonomous driving. These models are often deployed in interactive or assistive settings, where it is important that these assistive models generalize to new users and new tasks. Diverse interactive data generation offers a promising avenue for providing data-efficient generalization capabilities for interactive embodied foundation models. In this paper, we investigate the generalization capabilities of a multimodal foundation model fine-tuned on diverse interactive assistance data in a synthetic domain. We explore generalization along two axes: a) assistance with unseen categories of user behavior and b) providing guidance in new configurations not encountered during training. We study a broad capability called **Open-Set Corrective Assistance**, in which the model needs to inspect lengthy user behavior and provide assistance through either corrective actions or language-based feedback. This task remains unsolved in prior work, which typically assumes closed corrective categories or relies on external planners, making it a challenging testbed for evaluating the limits of assistive data. To support this task, we generate synthetic assistive datasets in Overcooked and fine-tune a LLaMA-based model to evaluate generalization to novel tasks and user behaviors. Our approach provides key insights into the nature of assistive datasets required to enable open-set assistive intelligence. In particular, we show that performant models benefit from datasets that cover different aspects of assistance, including multimodal grounding, defect inference, and exposure to diverse scenarios.

1 INTRODUCTION

Embodied foundation models, once trained, are shown to generalize to new tasks or modalities with minimal additional task-specific data (Zheng et al., 2025; Liu et al., 2025; Gao et al., 2024). However, training schemes for these models often involve vast amounts of complex multimodal data, which are expensive and laborious to collect (de Melo et al., 2022). Therefore, many recent approaches leverage intelligent synthetic data-generation/augmentation paradigms to increase the diversity of the tasks and data that the model gets exposed to during training. (Ahn et al., 2024). In recent work, foundational robotics models trained on large-scale synthetic datasets have demonstrated competitive generalization performance on tasks such as robotic manipulation (Deng et al., 2025). To transform embodied foundation models into fielded artifacts that interact with the real world, understanding and developing novel methods and modalities of synthetic data is a critical next step (Ren et al., 2025).

*Work completed while at Toyota Research Institute




Figure 1: We simulate synthetic users in Overcooked to generate multimodal (image + text) gameplay trajectories, which used to distill complementary synthetic datasets. These datasets are designed to (1) ground actions to environmental outcomes and (2) support behavior understanding and correction from trajectories. By training an embodied model on these data, we evaluate whether an embodied foundation model can generalize to unseen defective behaviors and novel task configurations.

This paper focuses specifically on embodied models which are aimed at assistive collaboration, i.e. *interactive* foundation models. Interactive foundation models must be able to inspect and understand temporally-extended user-behavior via a multimodal history of the user’s interactions with the assistive agent and environment (Sinha et al., 2024; Liu et al., 2024). We focus on an unsolved task which we denote as Open-Set Corrective Assistance, wherein an assistive model needs to inspect user-behavior and provide assistance to correct any behavior that is counterproductive to the task-goals, without a prespecified finite set of tasks nor a finite set of counterproductive defects in behavior. Large-scale, real-world data collection in these settings remains challenging, as long-form interaction data in embodied environments is inherently noisy and often requires substantial monitoring and scaffolding (Sumner et al., 2025). Furthermore, models trained in these settings need to generalize along multiple dimensions, i.e. differing behavior-types as well as differing tasks or domains (Sui et al., 2025). Collectively, these constraints motivate diverse assistive data as an important tool for evaluating the scaling of interactive foundation models in embodied environments.

In this paper, we focus on testing the capabilities of a foundation model, trained on assistive data, so as to probe the influence of diverse embodied data on open-set assistive models. Recent work on assistive modeling has provided methods to generate corrective actions which a user can employ to remediate their suboptimal behavior (Verghese et al., 2025; Gopinath et al., 2024), or provide language-based coaching to enable the user to understand and correct any deficiencies (Liu et al., 2023b; Guan et al., 2024). However, these methods are limited by the requirement of a closed-set of possible corrections to select from, or are dependent on a domain-specific planner to execute high-level corrective strategies. By exploring this problem in an open-set scenario (without the assumption of a fixed set of corrections), we directly probe whether foundation models trained solely on a limited task data can generalize their domain knowledge towards previously unseen behaviors and task configurations.

In this paper, we instruction-tune (Ouyang et al., 2022) an assistive model capable of generating *either* corrective actions or open-ended feedback to correct a deficiency in a user’s behavior, entirely in natural language. Our model-architecture leverages Llama-3 suite models as the base, equipped with a ViT encoder to embed the image-based state representations across a trajectory into the embedding-space of the multimodal language model (MLM) (Liu et al., 2023a). Our model is trained entirely on synthetic trajectories in Overcooked, to facilitate diverse behavior exploration. We collected a set of gameplay trajectories by deploying synthetic users in procedurally generated overcooked configurations. We then utilized these trajectories to curate two types of datasets and (1) Grounding datasets focused on enabling the model to accurately interpret the observation space and (2) Task-specific datasets focused on analyzing and correcting defective user-behavior. These data span a wide range of behavior types and task configurations, enabling clear stratification and a direct evaluation of the assistive model’s generalization capabilities.

Our evaluations seek to analyze the generalization capabilities of an assistive foundation model along multiple axes. First, we evaluate on withheld defective categories and find that, in both zero-shot and few-shot settings, our trained model outperforms a GPT-4o-based behavior critic at predicting corrective feedback and actions. Next, we study task-level generalization, where the model must provide open-ended assistance on entirely new recipes. This task is extremely difficult, as it requires the model to (1) compose learned individual concepts in completely new ways, such as learning to fill

a soup pot with steaks instead of tomatoes, and (2) coach a user on how to perform this completely new task more efficiently. In this demanding setting, our approach outperforms the GPT-4o baseline when providing corrective actions in the zero-shot setting and achieves superior performance on both feedback and corrective action prediction tasks when exposed to few-shot training examples. Together, our results show that carefully constructed assistive data can enable robust generalization, while also revealing clear limits of data-driven approaches for open-set interactive assistance. Our contributions can be summarized as follows:

1. We develop a synthetic data generation framework that simulates long-horizon user trajectories in *Overcooked*, producing complementary datasets that endow an embodied foundation model with grounding, actuation, and assistive capabilities.
2. We train a multimodal model to perform Open-Set Assistance, and show that the model, leveraging diverse assistive data along with an LLM backbone can generalize to unseen task configurations and defective categories.
3. We derive insights into effective dataset design for embodied assistance, highlighting the role of multimodal compositionality, spatial reasoning, and task decomposition.

2 RELATED WORK

2.1 DATA GENERATION FOR EMBODIED FOUNDATION MODELS

Synthetic data are used primarily to fill in gaps in underrepresented regions of the data distribution, such as missing domains (Chen et al., 2025a) or safety-critical long-tail settings (Ding et al., 2023). The proficiency of SOTA language and visual models has allowed researchers to adapt insights from these models to scale up the grounding, multimodal understanding, and actuation capabilities of “in-the-wild” embodied foundation models (Ahn et al., 2024; Tang et al., 2025; Mandlekar et al., 2023; Singh et al., 2024). Another approach is to directly optimize the data generation process by indexing key performance metrics in a downstream task (Huang et al., 2025; Wang et al., 2024b). In addition to modeling based approaches, recent research has proposed numerous novel simulation mechanisms to allow accelerated data generation for tasks that were once considered infeasible (de Melo et al., 2022; Li et al., 2024; Chen et al., 2025b; Yang et al., 2025). These methods highlight the significant utility of targeted data generation in systematically addressing deficiencies in model capabilities and improving generalization across tasks.

2.2 SIMULATING USER-BEHAVIOR IN SEQUENTIAL DECISION MAKING SETTINGS

Modeling user-behavior is a critical component of an assistive model. The first component of user-modeling is simulating user-behavior in an embodied environment. A popular method of building synthetic agents has been through co-play or behavior cloning (Strouse et al., 2021; Yan et al., 2023; Carroll et al., 2019; Wang et al., 2024a). Other methods seek to directly model a distribution of user-actions which take into account style and preferences, given an input-state (Pearce et al., 2023; Laidlaw & Dragan, 2022). Finally, a more recent suite of methods utilize LLMs to function as user-models which condense user-preferences through demonstrations or added-context (Yun et al., 2025; Zhang & Soh, 2023; Zhang et al., 2024). These methods provide critical tools for simulating interactions with complex synthetic user’s to build models capable of understanding and adapting to human-behavior (Anthis et al., 2025).

2.3 ASSISTIVE MODELING

Addressing user-deficiencies in embodied settings requires inferring the user’s high-level strategy, and intent Hoffman et al. (2024), identifying any deficiencies, and generating appropriate feedback. Relevant prior work includes methods to discover user-traits from their underlying policies (Ankile et al., 2023), or condition the agent’s policy on an estimated distribution over the partner’s intention (DeCastro et al., 2024; Wang et al., 2024a; Zhao et al., 2022). Recent work has also developed bespoke Theory-of-Mind modules which provide cognitive scaffolds to ground predictions of human-behavior (Liu et al., 2024; Cross et al., 2024). Upon identifying a user’s intention, a model can then move on to correcting the deficiencies in their behavior; either through language-based critique or

corrective guidance. Behavior critique methods process behavior trajectories and generate actionable feedback or reflections (Liu et al., 2023b; Guan et al., 2024; Gopinath et al., 2024). Behavior corrections provide guidance actions by via a behavior model, or select a corrective strategy, to physically demonstrate correct behavior (Verghese et al., 2025; Yang et al., 2024; Sinha et al., 2024).

Despite substantial progress in synthetic data generation, user simulation, and assistive modeling, existing work often studies these components in isolation and relies on closed sets of corrective strategies, external planners, or real-world data. Furthermore, testing an assistive model’s ability to generalize corrective guidance across diverse unseen tasks and failure modes remains largely unexplored (Zapata et al., 2025). In this work, we introduce *open-set assistance* as a challenging testbed for evaluating such generalization and for informing the design of future embodied assistive models.

3 OPEN-SET CORRECTIVE ASSISTANCE

In this work, we seek to build a model which can provide Open-Set Corrective Assistance, i.e. our assistive model needs to predict either coaching-based assistance or correction-based guidance based on the instruction provided, and the model is never given a preset list of possible actions or feedback classes. In the rest of this section, we define the task of Open-Set Corrective Assistance (§ 3.1), and provide details regarding our model architecture (§ 3.2) to facilitate such assistance.

3.1 PROBLEM DEFINITION

In this paper, we seek to build an assistive model capable of inspecting multimodal traces of user-behavior, and generatively providing assistance, either through language or physical actuation. We train our model on a dataset of multimodal trajectories in Overcooked, along with the assistance in the form of language or a physical action, i.e. $\mathcal{D}_{train} = \{\tau_i, [c|a]\}_{i=1}^{|\mathcal{D}_{train}|}$. Each trajectory, τ_i , is comprised of an interleaved sequence of state-action pairs ($\langle s_j^i, a_j^i \rangle$), where the states are RGB images of the overcooked game-state, and actions are represented by language, i.e. “move left,” “no action,” “interact,” etc. $[c|a]$ represent the modalities of assistance the model can provide which is either language-based coaching advice (c), or a remedial action (a). Note that coaching advice (c) is delivered in a terminal setting, whereas remedial actions (a) are generated only after the full input trajectory has been observed, rather than in an intermittent manner.

Each trajectory in \mathcal{D}_{train} is a user’s gameplay history with an underlying defective category, $\delta_k \in \Delta$, where Δ corresponds to the set of all defective categories. Defective categories, or *defects*, are defined as an impairment in or characteristic of a user’s behavior which may adversely impact cognitive processes such as attention, memory, decision-making, planning, etc. In this paper, we focus on defects that fall under the category of “cognitive impairment”, in that they are short-term conditions that impair decision-making, rather than a condition that is progressive and permanent (Dhakal & Bobrin, 2025). Specifically, we focus on defects that manifest as improper planning/sequencing of operations, or impairments in visuo-spatial skills leading to difficulty in identifying spatial relationships and interacting with the environment. Examples of some defects that we design are “Player does not know when the tomatoes in the pot have started cooking,” (Visuo-spatial) or “Player thinks it is best to only serve a single dish” (Planning). Our training dataset has a total of seventeen defects, $|\Delta| = 17$, where one of the possible defective categories is a null setting, i.e., “No Defect.” Although we do not work with an infinitely large set of defective categories, we selected a wide range of defects that encompass the variety of possible behaviors exhibited by an overcooked player. The full set of defects is included in the Appendix (§ C).

We aim to train a unified model comprised of a base LLM, π_θ , and a visual encoder, ϕ , to provide corrective guidance, i.e. $c_i \sim \pi_\theta(\cdot | \phi(\tau_i), p_t)$. p_t represents the associated task prompt for either action-prediction or feedback-prediction. For the action-prediction task, p_t also includes a reference trajectory of the user’s gameplay on a simple map, τ_r , to clarify the user’s general strategy (described in further detail in § 4.1) and allow the model to produce guidance aligned with the user’s overarching strategy. The full model-instructions are provided in the Appendix, § E).

3.2 MODEL ARCHITECTURE

We employ a projection-based multimodal model-structure akin to Llava (Liu et al., 2023a). We utilize Llama-3 (Grattafiori et al., 2024) as our base-LLM, and a ViT-base model (Dosovitskiy et al., 2020) as our image-encoder. The image-encoder generates visual features for the entire batch of images in the trajectory, $Z_i = \phi(s_i)$. We then utilize a language-projection layer, W_p , to project the image embeddings into the space of the language model, $X_i^s = W_p \cdot Z_i$. These image-tokens, X^s are interleaved with the language embeddings from the LLM, X^a to make up the entire set of trajectory tokens which are input into the Llama model, $X^\tau = X^s \otimes X^a$. For decoding, we employ a text-only decoding framework where both the coaching and corrective-actions are decoded as language tokens.

4 DATA DEFINITIONS

We synthetically generate a dataset of trajectories to train our model to perform Open-Set corrective assistance. Our goal is to derive clearly stratified synthetic data to accurately test our model’s ability to adapt to new tasks and defective behaviors. We begin by describing the strategies employed to generate each component of our dataset, i.e. trajectories, corrections, etc. (§ 4.1). We then conclude with a comprehensive description of each of our training sets used during model-training (§ 4.2)

4.1 DATA GENERATION

Synthetic Users. Our goal was to generate a diverse dataset which covered a range of varying strategies for playing Overcooked. We first designed an “Overcooked API” which provided skills to our synthetic agents e.g. `pickup(dish, counter)`, `move_to(sink)`, etc. Using this API, we created a set of 5 rule-based heuristics ($\mathbb{H} = \{H_1, H_2, H_3, H_4, H_5\}$), where each heuristic represents a different set of priorities, or strategies for employing a task, e.g. “*Always prioritize steak preparation over soup preparation*”, or “*Always make sure dishes are cleaned and ready before starting any meal-prep*”, etc. Note that each of these heuristics were designed to represent diverse preferences for playing Overcooked, they were not designed to be optimal.

Defective Rollouts. We instrumented our agent APIs with defective wrappers to enable altering of a specific aspect of their behavior, such as “the user no longer being able to identify when a cooking pot is full,” or “the user ignoring trip hazards during their path planning.” We roll out trajectories, \mathcal{T} , with these defective wrappers and heuristics, injecting additional stochasticity so the synthetic user takes a random action with 20% probability. To ensure diversity of rollouts, we apply noisy rollouts, wherein there is a small probability of a random action being selected in the trajectory at every step. This noise pushes the player into atypical states, yielding more varied behaviors across trajectories. Furthermore, to prevent overfitting to repetitive patterns, each user–defect pair is executed only once in each map configuration. Our training dataset uses a set of 450 procedurally generated maps during the rollouts (The prompt utilized to generate the map-configurations is provided in the Appendix, § B.1).

Ground Truth Correction Generation. We seek to assign a ground truth corrective-action or coaching snippet to each trajectory in the dataset. The corrective action is computed by predicting the next action of the heuristic, given the last state, without the defective wrapper. The ground truth coaching snippets are generated synthetically using an independent GPT-4o model, M . First, we create a set of boilerplate coaching snippets, $C^{raw} = \{M(\delta) \forall \delta \in \Delta\}^N$, by prompting the language model, M , to generate a coaching snippet which identifies and corrects a given defect, for each defect in the dataset. This process is repeated N times per defect, which in our case is set to 5. Next, for each seed coaching snippet, we use M to tailor the response to a specific “persona” ($p \in \mathcal{P}$), to expose the model to a more diverse set of coaching snippets. Each persona adapts the raw coaching snippet to exhibit a specific style such as “urgency,” “frustration,” “encouragement,” “brevity,” etc., $C_\delta^{final} = \{M(p, c^{raw}) \forall p \in \mathcal{P} \text{ and } \forall c^{raw} \in C_\delta^{raw}\}^N \forall \delta \in \Delta$. In between each of these steps, we adopt a self-evaluation step wherein the model validates whether the generated coaching snippet adequately addresses the original defect. This validation is done in an ensemble fashion, where the same language model is asked to vote five times, a data-point is only selected if at least four out of five instances passed the validation test. Finally, each trajectory in our corpus is randomly assigned one of the synthetically generated coaching snippets corresponding to the defect in the given trajectory.

Reasoning Trace Generation Finally, we augment the trajectories in our dataset with synthetic reasoning traces. The inclusion of reasoning traces in our dataset enables us to measure whether teaching our model to reason on the behavior exhibited in a trajectory can improve its ability to generalize to new defective categories or assist on previously unseen tasks. We collect reasoning traces by prompting M to analyze a trajectory by parsing an abstract list of events that occurred in the trajectory. These abstracted events, $E(\tau_i)$, include any important outcomes or changes in the state, based on the player’s actions and are computed during the initial trajectory rollout. Each reasoning trace includes three components: (1) A *summary* of the trajectory describing what the user tried to do, (2) A description of the player’s *successes* over the course of the trajectory, and (3) The set of *challenges* faced by the player while performing their task. The language model performs this reasoning procedure, $r_i^{sum}, r_i^{suc}, r_i^{ch} \sim M_{reason}(E(\tau_i))$, three times for each example, and chooses the best output for each category based on a self-assessed measure of confidence with regards to the output’s relevance to the original set of events.

4.2 TRAINING DATASETS

We curate two sets of datasets for model-training: (1) Grounding Datasets - Datasets which equip the model with the ability to ingest a multi-modal trajectory and ground the events which occur in the trajectory, and (2) Task-Specific Datasets - Datasets which teach the model to analyze trajectories, identify defects and provide assistance.

4.2.1 GROUNDING DATASETS

We generate visual question-answering (VQA) datasets, by sampling images or sequences of images from trajectories generated by our synthetic agents. We designed questions to enable the model to visually analyze the observation space to make assessments about a given state or identify important events which occurred across a sequence of states. We design three VQA datasets for this task,

- **Image-QA:** In the Image-QA dataset, $\mathcal{D}^I = \{I_i, \{q_j, a_j\}_{j=1}^6\}_{i=1}^N$, each datapoint is comprised of an image, randomly sampled from a trajectory of a synthetic heuristic, associated with six question-answer pairs associated with the given state, which are randomly sampled from a larger set of questions. These questions probe the model’s grounding of spatial information in the environment, which is a task MLLMs often struggle with (Rajabi & Kosecka, 2024; Zhang et al., 2025). The entire dataset is comprised of 55,000 examples.
- **Trajectory-QA:** The Trajectory-QA dataset is a dataset comprised of trajectories, with six associated question-answer pairs, $\mathcal{D}^T = \{\hat{\tau}_i, \{q_j, a_j\}_{j=1}^6\}_{i=1}^N$. The trajectory in each datapoint, $\hat{\tau}_i$, is selected by extracting a sub-trajectory from a full-rollout, of length, $K \sim \mathcal{N}(15, 2)$. These questions seek to improve the model’s temporal grounding capabilities, i.e., identify how the actions taken in the trajectory affect the environment. This dataset has a total of 54,000 datapoints.
- **Video-QA:** Similar to the previous dataset, the Video-QA dataset, \mathcal{D}^V is comprised of trajectories augmented with question-answer pairs. However, in this dataset, the trajectories are just a sequence of images in the form of a video, $\hat{\tau}_i = \{s_k\}_{k=1}^K$. We utilize the same template for the question-answer pairs as that of the Trajectory-QA dataset. The purpose of the dataset is to ensure that the model adequately relies on its visual encoding capabilities rather than overfitting to the sequence of actions provided in the trajectory. This dataset is comprised of 55,000 datapoints.

The entire set of questions for each of these grounding datasets can be found in Appendix F. We label the set of all these grounding datasets as \mathcal{D}_{ground} .

4.2.2 TASK-SPECIFIC DATASETS

By utilizing our synthetic data generation procedures described in §4.1, we assemble three datasets which directly pertain to corrective assistance (see Appendix Figure 2).

- **Coaching:** We formulate the coaching task using the dataset, $\mathcal{D}^{coach} = \{\tau_i^{1:t}, c_{\delta_i}, [r_i]\}$ where each datapoint consists of a trajectory $\tau_i^{1:T} \in \mathcal{T}$, a coaching signal c_{δ_i} , which is randomly selected from the set of available coaching snippets for the given defect ($C_{\delta_i}^{final}$),

and an optional reasoning trace $[r_i]$ associated with the trajectory. Note that the coaching snippet implicitly aligns with the behavior exhibited in the trajectory, since the trajectory is generated by applying a particular defect, but this alignment is not explicitly enforced while assigning the coaching snippet to the trajectory. This set contains 26,000 training examples.

- **Corrections:** We formulate the correction task as follows, $\mathcal{D}^{\text{correct}} = \{\tau_{ref}, \tau_i^{t-10:t}, \hat{a}_{t+1}, [r_i]\}$. The model needs to predict the next *corrected* action the user should take to steer them away from their defective gameplay. For the corrections task, the model is also provided a reference trajectory, τ_{ref} , comprised of a rollout of the specific user acting in a small map (different to the any of the maps utilized to simulate trajectories) to highlight the user’s high-level strategy. Finally, we restrict the input trajectory to the 10 preceding steps, which allows the reference trajectory to fit within the model’s available context window. This dataset contains 27,000 training examples.
- **Defect-Delineation:** Our third task task, Defect Delineation, seeks to improve the model’s ability to discern how different defects manifest in user behavior. This task was inspired by the Next-Sentence Prediction pretraining task (Devlin et al., 2019), leveraged to build a greater semantic understanding of sentence-similarity. Similarly, in Defect-Delineation the model inspects two trajectories and predicts whether the two trajectories have the same defect, and also the specific defect in each trajectory, $\mathcal{D}^{DD} = \{\tau^{(1)}, \tau^{(2)}, \delta^{(1)}, \delta^{(2)}, \mathbb{I}\{\delta^{(1)} = \delta^{(2)}\}\}_{i=1}^{|\mathcal{D}^{DD}|}$. Each trajectory is randomly sampled from our dataset of trajectories, \mathcal{T} . The output is formatted as a text-based answer, “Yes. The defect in both trajectories is <defect>” or “No. The defect in trajectory 1 is <defect1>. The defect in trajectory 2 is <defect2>.” Our dataset includes five negative samples for every positive sample, as the task of distinguishing and correctly predicting the defects in two different trajectories is more challenging. This dataset contains 20,000 training examples.

Our default training dataset, \mathcal{D}_{train} , includes all three of these datasets with equal weighting.

5 EXPERIMENTS

5.1 BASELINES

A suitable baseline for Open-Set Corrective Assistance must be able to process video trajectories, perform visual grounding to interpret task instructions, and provide the corresponding assistance. We adopt a methodology for behavior correction, developed in prior work, which empowers pretrained VLMs to provide language-based coaching based on videos of a robot’s policy (Guan et al., 2024). In our work, we extended this baseline to generate corrections as well as coaching.

1. Behavior Critic (GPT-4o) - We prompted GPT-4o to perform assistance based on a downsampled trajectory. In addition to the trajectory, the model is provided the rules of overcooked for additional grounding. Furthermore, unlike our model, which needs to implicitly identify new defects without any prior knowledge, the GPT-4o behavior-critic is also provided the full list of potential defective behaviors that could be found in the behavior.
2. Behavior Critic (GPT-4o) + language summary - We extend the original behavior critic by adding reasoning traces from our dataset to the prompt as a part of the input to help improve the model’s ability to identify the defect and provide assistance.

The prompts for both of these baselines are included in the appendix (§ B.5).

5.2 EVALUATION TASKS

For the novel task-formulation of open-set assistance, our evaluations seek to comprehensively examine how well our assistive model can perform across multiple axes of generalization. In all our experiments (unless otherwise stated), we perform few-shot training to fine-tune our model on a small set of examples prior to testing on held-out data (10 examples per new defect).

- **Novel Defective Modalities** - This evaluation tests how well our model can utilize its understanding of the task to provide actionable assistance in previously unseen failure modes

Table 1: **Generalization across defects and tasks.** We compare assistive performance when generalizing to (left) unseen defective behaviors and (right) novel task configurations (recipes). Models trained entirely on diverse assistive data outperform behavior-critic baselines when adapting to new defects with only a few examples, with performance saturating at the 1B scale. In contrast, generalization to new tasks benefits substantially from model scaling, likely due to the stronger multi-modal grounding capacity offered by the larger model.

Baseline	Held-Out Defects		Task Generalization	
	Coaching	Corrections	Coaching	Corrections
Behavior Critic	21.00	20.40	34.21	9.17
Behavior Critic + Summaries	55.70	19.80	71.05	15.83
Ours 1B	76.60	55.70	50.88	50.83
Ours 8B	77.80	54.60	85.96	56.67

Table 2: The 8B variant of our model outperforms the behavior critic baseline in all settings except while coaching in the task generalization setting. Coaching users to perform new tasks proves to be a challenging setting for the model, however, the inclusion of reasoning traces provides an 8% boost in performance.

Baseline	Held-out Defects		Task Generalization	
	Coaching	Corrections	Coaching	Corrections
Behavior Critic	21.00	20.40	34.21	9.17
Ours 8B	34.30	39.70	0.00	55.00
Ours 8B + reasoning	18.10	32.60	8.77	43.30

(defects). For instance, during training, the model may encounter a defective modality such as “Player believes it is best to serve only soup,” while a held-out test defect might be “Player believes it is best to serve only steak.” If the model has learned the underlying concept of a player omitting key components of a recipe, it should be able to generalize this understanding to novel contexts involving different missing components.

- **New Tasks (Overcooked Recipes)** - During training the agent is always following a single recipe. In this evaluation, we measure how well the model is able to assist users performing tasks, i.e. completing two new recipes, within Overcooked. This evaluation probes the model’s ability to adapt its domain knowledge to a related task within the same domain through compositionality, grounding unseen recipes to new object configurations and layouts to identify previously unseen failure modes. All recipes are provided in Appendix D. .

Finally, we perform two ablations pertaining to the training set configuration. First, we compare multi-dataset training on the full \mathcal{D}_{train} against single-task training on \mathcal{D}^{coach} or $\mathcal{D}^{correct}$ to assess whether jointly learning related assistive tasks leads to more robust and adaptable assistive behavior. Second, we assess whether co-training with our grounding datasets (\mathcal{D}_{ground}) improves the model’s ability to localize actions and events required for novel recipes and translate them into assistance.

5.3 METRICS

We utilize two metrics to measure the accuracy of the model for each modality of assistance.

- **LLM-as-judge (Coaching)** - We leverage an LLM to identify whether or not a generated coaching snippet adequately corrects the underlying, ground-truth defect. The full prompt used for enabling GPT-4o to evaluate our model-outputs is included in the appendix.
- **Accuracy (Corrections)** - To evaluate the correction-generation capabilities of our model, we simply measure the accuracy when compared to the ground truth action.

5.4 RESULTS

Our model can learn to assist users with unseen defective behaviors with just 10 training examples per defect – Table 1 reports the performance of our model on the dataset comprised of trajectories with novel defective behaviors. First, we find that our model, when trained on the three task-specific datasets (\mathcal{D}^{coach} , $\mathcal{D}^{correct}$, \mathcal{D}^{DD}) outperforms an off-the-shelf GPT-4o model for both tasks. Incorporating the visual-reasoning traces as an input to GPT-4o’s input, boosts the model’s ability to accurately coach by 34%, however, it is still inferior to our model. Upgrading the base LLaMA model from one billion (1B) to eight billion (8B) parameters results in only negligible performance differences (+1 for coaching and -1 for corrections). This suggests that the model has reached a saturation point on the dataset for the held-out defect evaluation.

Our model can generalize to new in-domain tasks – We report findings on our evaluation of the model’s ability to assist users performing new recipes in Table 1. First, our 1B model surpasses the base behavior critic (+15/+41%) but falls short of the summary-augmented version (-20/+35%). Unlike in the case of the held-out defect evaluation, performance of the model increases, for both corrections and coaching, when scaling up from the 1B LLaMa model to the 8B LLaMa model. This indicates that a greater degree of multimodal compositionality is required for this task, which larger VLMs are better equipped to handle.

Table 3: **Effect of dataset ablations on assistive generalization.** Training jointly across coaching, correction and defect delineation improves downstream assistive performance compared to individually training on \mathcal{D}^{coach} or $\mathcal{D}^{correct}$.

Training Data	Unseen Defects		Novel Tasks	
	Coaching	Corrections	Coaching	Corrections
1B (\mathcal{D}_{train})	76.60	55.70	50.88	50.83
1B (\mathcal{D}^{coach})	38.10	–	54.39	–
1B ($\mathcal{D}^{correct}$)	–	29.40	–	33.34

Inducing reasoning priors for assistance improves zero-shot coaching performance while generalizing to new tasks – When our model is evaluated in a zero-shot setting, we see a reduction in performance on the coaching task for both sets of evaluations (see Table 2). Our zero-shot model still outperforms the behavior critic on both assistive tasks on the held-out defect evaluation. However, when reasoning is added, model performance worsens. This is likely due to mode-collapse; since the model has not been trained to reason with reinforcement learning, the model may struggle to condition its assistance on out-of-distribution reasoning traces, even if the reasoning traces correctly ground the input trajectory. For the task-generalization evaluation, the model outperforms both versions of GPT-4o while predicting corrections, however, completely fails to coaching while assisting on novel-tasks (0.00%). Coaching the user to perform new tasks is much more challenging for any assistive model, as the model needs to infer mistakes from the new recipes and map failures identified in these trajectories to these inferences. Without any training on held-out data, the model is unable to make these inferences. Notably, incorporating reasoning traces mitigates this issue, improving coaching performance by 8%, however, the model is still unable to match the behavior critic’s performance on task-generalization.

Training on multiple datasets simultaneously improves the model’s ability to assist – Table 3 presents a dataset-analysis to understand the effect of each dataset on the model’s ability to perform open-set assistance. These experiments are performed with our 1B-scale model. First, we measure the impact of multi-task training with all three datasets in \mathcal{D}_{train} . Having the model simultaneously train on both tasks, as well as the auxiliary defect delineation task, wherein the model is prompted to identify and delineate defects, generally improves downstream performance on both tasks.

Co-training with grounding datasets can improve visual compositionality, thereby benefiting task generalization - We hypothesized that task generalization requires stronger visual compositionality to translate unseen behaviors into effective assistance. To test this, we co-train with the grounding datasets, introduced in Section 4.2, and find that such co-training, both jointly and individually, improves downstream assistive performance on novel tasks (Table 4). Improved performance may arise from the model’s increased ability to recognize key in-game events, such as the

Table 4: **Inclusion of additional grounding datasets improves model’s ability to assist in new task configurations.** This figure depicts the performance of the assistive model, on the task-generalization experiment, after co-training with varying configurations of the grounding datasets, \mathcal{D}_{ground} . Among the grounding datasets, the Trajectory-QA dataset, \mathcal{D}^T , appears to be the most effective, whereas the Video-QA dataset, \mathcal{D}^V , appears to reduce performance.

Baseline	Coaching	Corrections
1B	50.88	50.83
1B ($\mathcal{D}_{ground} + \mathcal{D}_{train}$)	58.77	53.33
1B ($\mathcal{D}^I + \mathcal{D}_{train}$)	69.30	46.67
1B ($\mathcal{D}^V + \mathcal{D}_{train}$)	46.49	42.50
1B ($\mathcal{D}^T + \mathcal{D}_{train}$)	78.95	52.50

completion of a delivery or proximity to the delivery location, which in turn supports the inference of new failure modes from trajectories of players performing unfamiliar tasks. Among the grounding datasets, \mathcal{D}^T appears to be the most effective, whereas co-training with only \mathcal{D}^V seems to reduce task performance. This finding implies that access to action sequences may be critical for analyzing new failure modes in novel recipes. \mathcal{D}^I also appears to provide benefits; unlike the Yes/No questions in the Trajectory and Video datasets, the Image-QA dataset requires more complex analyses of the game state, including path planning and object disambiguation, which may improve the model’s ability to process image-based states.

6 DISCUSSION

This work examines the capabilities of diverse assistive data in training foundation models to perform embodied open-set assistive tasks. We generate a synthetic dataset and train a model to perform a representative task, Open-Set Corrective Assistance, in which the model must provide corrective guidance by observing embodied trajectories of user behavior, without access to a closed set of corrective categories. We evaluate models, trained entirely on synthetic data, in two complementary settings: (1) generalization to trajectories containing defective behaviors not seen during training, and (2) generalization to novel task configurations (recipes) in Overcooked.

Firstly, our dataset ablations yield several insights into the composition of training data required to support generalization. In particular, task-level generalization places greater demands on multimodal compositionality, requiring data that captures higher-level relationships between actions, goals, and environmental outcomes. We find that multi-task co-training across coaching, correction, and defect delineation consistently improves downstream assistive performance, while grounding datasets that contain re-usable information pertaining to action outcomes are especially beneficial for generalization to connect outcomes to deficiencies in new task configurations. Additionally, inducing reasoning priors through behavior summaries improves assistive performance, particularly for coaching and zero-shot evaluation. However, performance degrades when reasoning traces fall outside the training distribution, indicating that current forms of reasoning supervision are brittle. Together, these results suggest that effective training datasets for assistive modeling should emphasize decompositional structure, explicitly separating perception, interpretation, and actuation, rather than relying solely on end-to-end demonstrations.

7 LIMITATIONS

Finally, there are two important limitations to cover pertaining to pending avenues of exploration and model-capabilities. First, it is important to highlight that the capabilities of a diverse trained model to provide assistance on real-human trajectories remains unexplored. Having explored the data configurations that is helpful in promoting generalization to novel tasks and defective behaviors, a critical next step is to collect human trajectories to understand the quality, quantity and design of synthetic data required to promote sim2real generalization on complex, embodied assistive tasks requiring foundation models. Secondly, our model training procedure leverages a standard instruction-tuning setup, absent of any feedback-based post-training. Applying alignment-based approaches

could produce assistance that better captures the intrinsic reward model of the user that is being assisted. Recently, a suite of methods have been developed to apply reinforcement-learning-based optimization methods to learning in complex multi-turn or embodied settings (Silva et al., 2025; Ding et al., 2025; Boyle et al., 2025). Applying these alignment methods would prove a natural fit to producing assistance that more effectively reasons about user-behavior and thereby exponentiate the grounding capabilities imbued through large-scale training on synthetic data.

8 CONCLUSION

In this paper, we seek to quantify the capabilities of diverse assistance for training assistive model for complex embodied tasks, requiring a multimodal foundation model. We study this question through Open-Set Corrective Assistance, a representative and largely unsolved setting in which a model must identify and correct user behavior through language feedback or corrective actions without access to a closed set of defect categories. Using Overcooked as our testbed, we collect a set of behavior trajectories by simulating gameplay with a set of synthetic users. We then processed these trajectories to produce multi-faceted training data, to imbue the model with grounding, actuation and assistive capabilities. Next, we trained our model on these datasets, so as to probe it on two clear stratifications for generalization; (1) Unseen failure modes (defects) and (2) Novel task configurations. Through our experiments, we show that our trained model can assist in both these novel environment configurations better than a behavior-critic from prior work. Our dataset ablations highlight the importance of designing datasets which individually teach the compounding skills required to perform a task, which in our case were grounding and multi-turn behavior analysis. Overall, our work establishes a foundation for future research on open-set assistive models, providing findings beneficial towards extending this work to assistance on real-world interactive data.

REFERENCES

- Michael Ahn, Debidatta Dwibedi, Chelsea Finn, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Karol Hausman, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, et al. AutoRT: Embodied foundation models for large scale orchestration of robotic agents. *arXiv preprint arXiv:2401.12963*, 2024.
- Lars L Ankile, Brian S Ham, Kevin Mao, Eura Shin, Siddharth Swaroop, Finale Doshi-Velez, and Weiwei Pan. Discovering user types: Mapping user traits by task-specific behaviors in reinforcement learning. *arXiv preprint arXiv:2307.08169*, 2023.
- Jacy Reese Anthis, Ryan Liu, Sean M Richardson, Austin C Kozlowski, Bernard Koch, James Evans, Erik Brynjolfsson, and Michael Bernstein. LLM social simulations are a promising research method. *arXiv preprint arXiv:2504.02234*, 2025.
- Liam Boyle, Nicolas Baumann, Paviththiren Sivasothilingam, Michele Magno, and Luca Benini. Robotxr1: Enabling embodied robotic intelligence on large language models through closed-loop reinforcement learning. *arXiv preprint arXiv:2505.03238*, 2025.
- Micah Carroll, Rohin Shah, Mark K Ho, Tom Griffiths, Sanjit Seshia, Pieter Abbeel, and Anca Dragan. On the utility of learning about humans for human-ai coordination. *Advances in neural information processing systems*, 32, 2019.
- Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation, 2025a. URL <https://arxiv.org/abs/2506.18088>.
- Yurui Chen, Junge Zhang, Ziyang Xie, Wenye Li, Feihu Zhang, Jiachen Lu, and Li Zhang. S-nerf++: Autonomous driving simulation via neural reconstruction and generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(6):4358–4376, 2025b. doi: 10.1109/TPAMI.2025.3543072.

- Logan Cross, Violet Xiang, Agam Bhatia, Daniel LK Yamins, and Nick Haber. Hypothetical minds: Scaffolding theory of mind for multi-agent tasks with large language models. *arXiv preprint arXiv:2407.07086*, 2024.
- Celso M. de Melo, Antonio Torralba, Leonidas Guibas, James DiCarlo, Rama Chellappa, and Jessica Hodgins. Next-generation deep learning based on simulators and synthetic data. *Trends in Cognitive Sciences*, 26(2):174–187, 2022. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2021.11.008>. URL <https://www.sciencedirect.com/science/article/pii/S136466132100293X>.
- Jonathan DeCastro, Andrew Silva, Deepak Gopinath, Emily Sumner, Thomas M Balch, Laporsha Dees, and Guy Rosman. Dreaming to assist: Learning to align with human objectives for shared control in high-speed racing. *arXiv preprint arXiv:2410.10062*, 2024.
- Shengliang Deng, Mi Yan, Songlin Wei, Haixin Ma, Yuxin Yang, Jiayi Chen, Zhiqi Zhang, Taoyu Yang, Xuheng Zhang, Wenhao Zhang, et al. Graspvla: a grasping foundation model pre-trained on billion-scale synthetic action data. *arXiv preprint arXiv:2505.03233*, 2025.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.
- Aayush Dhakal and Bradford D. Bobrin. *Cognitive Deficits*. StatPearls Publishing, Treasure Island (FL), 2025. URL <http://europepmc.org/books/NBK559052>.
- Wenhao Ding, Chejian Xu, Mansur Arief, Haohong Lin, Bo Li, and Ding Zhao. A survey on safety-critical driving scenario generation—a methodological perspective. *IEEE Transactions on Intelligent Transportation Systems*, 24(7):6971–6988, 2023. doi: 10.1109/TITS.2023.3259322.
- Yifeng Ding, Hung Le, Songyang Han, Kangrui Ruan, Zhenghui Jin, Varun Kumar, Zijian Wang, and Anoop Deoras. Empowering multi-turn tool-integrated reasoning with group turn policy optimization. *arXiv preprint arXiv:2511.14846*, 2025.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Haoxiang Gao, Zhongruo Wang, Yaqian Li, Kaiwen Long, Ming Yang, and Yiqing Shen. A survey for foundation models in autonomous driving. *arXiv preprint arXiv:2402.01105*, 2024.
- Deepak Gopinath, Xiongyi Cui, Jonathan DeCastro, Emily Sumner, Jean Costa, Hiroshi Yasuda, Allison Morgan, Laporsha Dees, Sheryl Chau, John Leonard, et al. Computational teaching for driving via multi-task imitation learning. *arXiv preprint arXiv:2410.01608*, 2024.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Lin Guan, Yifan Zhou, Denis Liu, Yantian Zha, Heni Ben Amor, and Subbarao Kambhampati. Task success is not enough: Investigating the use of video-language models as behavior critics for catching undesirable agent behaviors. *arXiv preprint arXiv:2402.04210*, 2024.
- Guy Hoffman, Tapomayukh Bhattacharjee, and Stefanos Nikolaidis. Inferring human intent and predicting human action in human–robot collaboration. *Annual Review of Control, Robotics, and Autonomous Systems*, 7, 2024.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. LoRA: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

- Binyuan Huang, Yuqing Wen, Yucheng Zhao, Yaosi Hu, Yingfei Liu, Fan Jia, Weixin Mao, Tiancai Wang, Chi Zhang, Chang Wen Chen, et al. Subjectdrive: Scaling generative data in autonomous driving via subject control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 3617–3625, 2025.
- Cassidy Laidlaw and Anca Dragan. The boltzmann policy distribution: Accounting for systematic suboptimality in human models. *arXiv preprint arXiv:2204.10759*, 2022.
- Gen Li, Kaifeng Zhao, Siwei Zhang, Xiaozhong Lyu, Mihai Dusmanu, Yan Zhang, Marc Pollefeys, and Siyu Tang. Egogen: An egocentric synthetic data generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14497–14509, 2024.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023a.
- Jijia Liu, Chao Yu, Jiaxuan Gao, Yuqing Xie, Qingmin Liao, Yi Wu, and Yu Wang. Llm-powered hierarchical language agent for real-time human-ai coordination. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS '24*, pp. 1219–1228, Richland, SC, 2024. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864.
- Yang Liu, Weixing Chen, Yongjie Bai, Xiaodan Liang, Guanbin Li, Wen Gao, and Liang Lin. Aligning cyber space with physical world: A comprehensive survey on embodied AI. *IEEE/ASME Transactions on Mechatronics*, 30(6):7253–7274, 2025. doi: 10.1109/TMECH.2025.3574943.
- Zeyi Liu, Arpit Bahety, and Shuran Song. Reflect: Summarizing robot experiences for failure explanation and correction. *arXiv preprint arXiv:2306.15724*, 2023b.
- Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. *arXiv preprint arXiv:2310.17596*, 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Tim Pearce, Tabish Rashid, Anssi Kanervisto, Dave Bignell, Mingfei Sun, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Ida Momennejad, Katja Hofmann, et al. Imitating human behaviour with diffusion models. *arXiv preprint arXiv:2301.10677*, 2023.
- Navid Rajabi and Jana Kosecka. GSR-BENCH: A benchmark for grounded spatial reasoning evaluation via multimodal llms. *arXiv preprint arXiv:2406.13246*, 2024.
- Xuanchi Ren, Yifan Lu, Tianshi Cao, Ruiyuan Gao, Shengyu Huang, Amirmojtaba Sabour, Tianchang Shen, Tobias Pfaff, Jay Zhangjie Wu, Runjian Chen, et al. Cosmos-drive-dreams: Scalable synthetic driving data generation with world foundation models. *arXiv preprint arXiv:2506.09042*, 2025.
- Andrew Silva, Pradyumna Tambwekar, Deepak Edakkattil Gopinath, Jonathan DeCastro, Guy Rosman, and Avinash Balachandran. Stability of preference alignment for multi-turn control with LLM policies. In *First Workshop on Multi-Turn Interactions in Large Language Models*, 2025.
- Ritvik Singh, Jingzhou Liu, Karl Van Wyk, Yu-Wei Chao, Jean-Francois Lafleche, Florian Shkurti, Nathan Ratliff, and Ankur Handa. Synthetica: Large scale synthetic data for robot perception. *arXiv preprint arXiv:2410.21153*, 2024.
- Rohan Sinha, Amine Elhafsi, Christopher Agia, Matthew Foutter, Edward Schmerling, and Marco Pavone. Real-time anomaly detection and reactive planning with large language models. *arXiv preprint arXiv:2407.08735*, 2024.
- DJ Strouse, Kevin McKee, Matt Botvinick, Edward Hughes, and Richard Everett. Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, 34: 14502–14515, 2021.

- Xiuchao Sui, Daiying Tian, Qi Sun, Ruirui Chen, Dongkyu Choi, Kenneth Kwok, and Soujanya Poria. From grounding to manipulation: Case studies of foundation model integration in embodied robotic systems. *arXiv preprint arXiv:2505.15685*, 2025.
- Emily Sumner, Deepak E Gopinath, Laporsha Dees, Patricio Reyes Gomez, Xiongyi Cui, Andrew Silva, Jean Costa, Allison Morgan, Mariah Schrum, Tiffany L Chen, et al. Simcoachcorpus: A naturalistic dataset with language and trajectories for embodied teaching. *arXiv preprint arXiv:2509.14548*, 2025.
- Grace Tang, Swetha Rajkumar, Yifei Zhou, Homer Rich Walke, Sergey Levine, and Kuan Fang. Kalie: Fine-tuning vision-language models for open-world manipulation without robot data. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9507–9515, 2025. doi: 10.1109/ICRA55743.2025.11128156.
- Mrinal Verghese, Brian Chen, Hamid Eghbalzadeh, Tushar Nagarajan, and Ruta Desai. User-in-the-loop evaluation of multimodal llms for activity assistance. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1144–1154. IEEE, 2025.
- Chenxu Wang, Zilong Chen, and Huaping Liu. On the utility of external agent intention predictor for human-ai coordination. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS '24*, pp. 2546–2548, Richland, SC, 2024a. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864.
- Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. Robogen: towards unleashing infinite data for automated robot learning via generative simulation. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org, 2024b.
- Xue Yan, Jiaxian Guo, Xingzhou Lou, Jun Wang, Haifeng Zhang, and Yali Du. An efficient end-to-end training approach for zero-shot human-ai coordination. *Advances in Neural Information Processing Systems*, 36:2636–2658, 2023.
- Xuemeng Yang, Licheng Wen, Tiantian Wei, Yukai Ma, Jianbiao Mei, Xin Li, Wenjie Lei, Daocheng Fu, Pinlong Cai, Min Dou, et al. Drivearena: A closed-loop generative simulation platform for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 26933–26943, 2025.
- Zhaojing Yang, Miru Jun, Jeremy Tien, Stuart J Russell, Anca Dragan, and Erdem Bıyık. Trajectory improvement and reward learning from comparative language feedback. *arXiv preprint arXiv:2410.06401*, 2024.
- Taedong Yun, Eric Yang, Mustafa Safdari, Jong Ha Lee, Vaishnavi Vinod Kumar, S Sara Mahdavi, Jonathan Amar, Derek Peyton, Reut Aharoni, Andreas Michaelides, et al. Sleepless nights, sugary days: Creating synthetic users with health conditions for realistic coaching agent interactions. *arXiv preprint arXiv:2502.13135*, 2025.
- Mireya Zapata, Diego Guevara, Jennifer Obregón, Hugo Arias-Flores, and Carlos Ramos-Galarza. Advancements in assistive robotics: A systematic review of inclusive technologies for people with disabilities. *IEEE Access*, 13:171871–171888, 2025.
- Bowen Zhang and Harold Soh. Large language models as zero-shot human models for human-robot interaction. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7961–7968, 2023. doi: 10.1109/IROS55552.2023.10341488.
- Ceyao Zhang, Kaijie Yang, Siyi Hu, Zihao Wang, Guanghe Li, Yihang Sun, Cheng Zhang, Zhaowei Zhang, Anji Liu, Song-Chun Zhu, et al. Proagent: building proactive cooperative agents with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17591–17599, 2024.
- Wanyue Zhang, Yibin Huang, Yangbin Xu, JingJing Huang, Helu Zhi, Shuo Ren, Wang Xu, and Jiajun Zhang. Why do MLLMs struggle with spatial understanding? a systematic analysis from data to architecture. *arXiv preprint arXiv:2509.02359*, 2025.

Michelle Zhao, Reid Simmons, and Henny Admoni. Coordination with humans via strategy matching. In *International Conference on Intelligent Robots and Systems (IROS)*, pp. 9116–9123. IEEE, 2022.

Jinliang Zheng, Jianxiong Li, Dongxiu Liu, Yinan Zheng, Zhihao Wang, Zhonghong Ou, Yu Liu, Jingjing Liu, Ya-Qin Zhang, and Xianyuan Zhan. Universal actions for enhanced embodied foundation models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 22508–22519, 2025.

A EXPERIMENTAL SETUP

The models which used the 1 billion parameter Llama3 models were trained on either a g6e.12xlarge instance on AWS or a local cluster with four Nvidia RTX 6000 (49 GB memory). The models which used the 8 billion parameter Llama3 models were trained on the p5.48xlarge AWS instance, as these models utilized over 80 GB of GPU-memory. Both the visual encoder, and language model are trained using Low-Rank Adaptation (LoRA) (Hu et al., 2022). The multimodal projection layer was trained from scratch. The specific values for our experimental parameters are provide in Table 5.

Table 5: Experimental parameters per baseline. All used image + llm LoRA with r, alpha = 16. Image -> LLM projection layer trained from scratch. 500 warmup steps.

Model-name	Training Datasets	Data Aug	Learning rate	Training Steps	Per-device batch size	Weight Decay
Ours (1B)	\mathcal{D}_{train}	No	$5e-5$	15000	1	0.01
1B + grounding datasets	$\mathcal{D}_{train} + \mathcal{D}_{ground}$	Yes	$5e-5$	16500	1	0.01
1B Coaching Only	\mathcal{D}_{coach}	Yes	$5e-5$	5000	1	0.01
1B Correction Only	$\mathcal{D}_{correct}$	No	$5e-5$	5000	1	0.01
Ours (1B) + I-QA	$\mathcal{D}_{train} + \mathcal{D}^I$	Yes	$5e-5$	16500	1	0.01
Ours (1B) + T-QA	$\mathcal{D}_{train} + \mathcal{D}^T$	Yes	$5e-5$	16500	1	0.01
Ours (1B) + V-QA	$\mathcal{D}_{train} + \mathcal{D}^V$	Yes	$5e-5$	16500	1	0.01




Figure 2: This figure provides an abstract depiction of each task synthesized in this paper to train our assistive model. Top to Bottom - Trajectory-QA, Video-QA, Image-QA, Corrections, Coaching, Defect Delineation

B PROMPTS

B.1 MAP GENERATION

We synthetically generated maps to rollout our defective heuristics and obtain diverse behavior trajectories (see Figure 5). Each map was checked to ensure satisfaction with common-sense and validity checks to ensure that they yield reasonable trajectories.

```
System: Map Generation
I need you to create new overcooked maps. I will provide you with a sample map. You need to generate new maps, with the same number of each object.
```

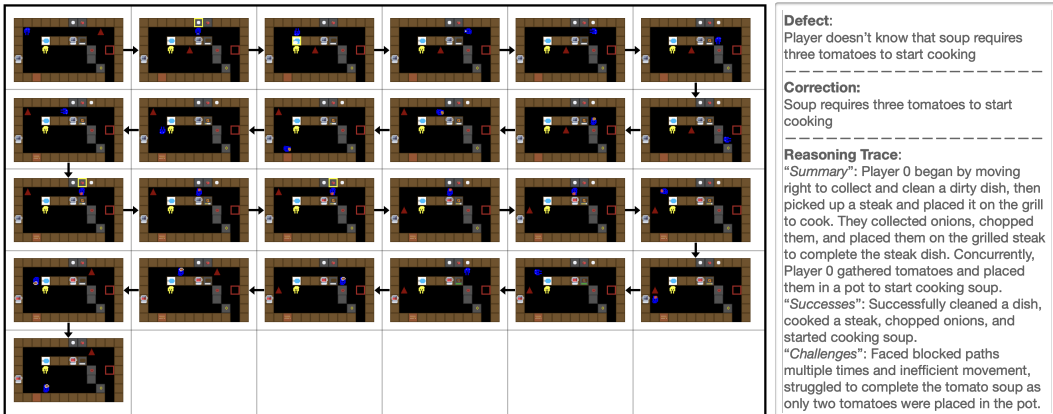


Figure 3: Dataset Example

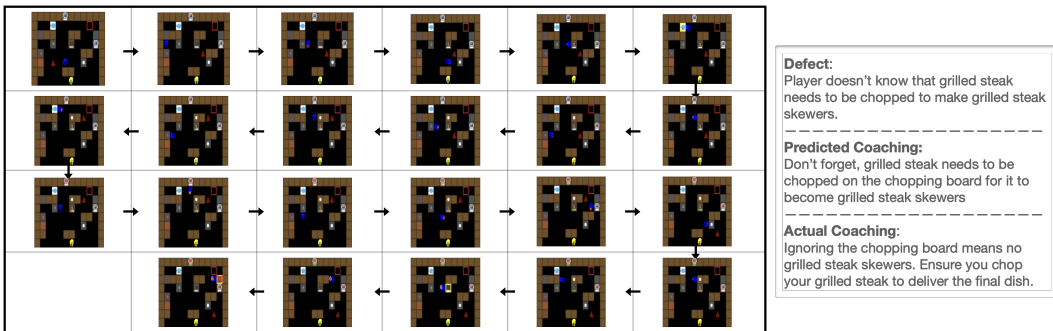


Figure 4: Model Output

```
Here is a dictionary which identifies each symbol in the map: \n
{map_key}
```

```
Generate a map while keeping the following things in mind:
(1) Make sure there is at least one of each of the following
symbols, ["O", "D", "T", "C", "N", "S", "H", "G", "I"]
(2) Make sure that Player 0 can access all target locations, ["O",
"D", "T", "C", "N", "S", "H", "G", "I"].
(3) You are allowed to change the size of the outer grid, as long as
the number of the "special" locations remain the same.
(4) Limit the map size to a 12 x 12 grid.
(5) Make sure you have at least one "X" and one "Y" in the map.
(6) Including spaces, each row should have the same number of
characters.
(7) Generate only one map.
```

```
Your response should just include the map.
```

User: Map Generation

```
Here are a few examples of a few possible overcooked map-layouts:
```

```
Map 6:
```

Figure 5: Prompt utilized to enable GPT-4o to generate overcooked Map configurations.

B.2 COACHING

We leveraged GPT-4o to synthetically generate language-based corrective feedback for each trajectory, to teach our model to coach a synthetic user with defective behavior. The process of generating coaching snippets was two-fold: (1) We first generated a set of initial coaching snippets given a description of the defect (see Figure 6), (2) Next, we augmented these coaching snippets using a set of directives such as “Make the advice more concise” or “Make the advice more humorous.” (see Figure 7)

System Prompt

```
I need to generate corrective feedback to enable a user to fix a
specific defect in their strategy for the game overcooked. I will
provide you a description of the defect and you need to provide some
options for how you would assist a user in correcting this defect .
Please craft the exact statement you would say to a user.
```

```
Each correction should be brief (maximum of two sentences). You
should address only the specific defect mentioned, avoiding general
advice or including irrelevant information. For example, discuss
grills only if the defect pertains to grills, if you are asked to
correct a defect pertaining to general interactions with dispensers
do not specifically mention any individual dispenser (onion, dish,
tomato, steak, etc.).
```

```
Note: Make sure to keep your feedback grounded in overcooked rather
than an actual kitchen. Make sure your corrections are consistent
with the rules of overcooked: {rules}
```

User Prompt

Please give me five different potential corrections for the following defect. Your corrections should be in the following format:

```
Correction1: ....
Correction2: ....
Correction3: ....
Correction4: ....
Correction5: ....

Defect: {defect}
```

Figure 6: Prompts used for generating initial coaching snippet from the defect

System Prompt

I have a set of corrections to help a user fix a defect in their gameplay strategy. Your job is to rephrase this set of corrections according to a directive I will provide to you. The instructions will be things like "make the correction more polite," "make the correction more direct," "make the correction more urgent," etc. However, you must ensure that the meaning of the correction should still remain the same, i.e., the correction should still be helpful in correcting the given defect in the user's strategy. Make sure that your augmented corrections are still grounded in overcooked and not in an actual kitchen setting. Try not to add unnecessary filler words if you can avoid it. The input will be provided to you in the following format:

```
## Player Defect
Defect: <description of the defect>

## Initial Corrections
Correction1: <correction1>
Correction2: <correction2>
Correction3: <correction3>
Correction4: <correction4>
Correction5: <correction5>

## Directive - <directive>

## Augmented Corrections (TODO)
Augmented Correction1: <updated correction1>
Augmented Correction2: <updated correction2>
Augmented Correction3: <updated correction3>
Augmented Correction4: <updated correction4>
Augmented Correction5: <updated correction5>
```

Make sure your corrections are consistent with the rules of overcooked: {rules}

User Prompt

Please update the given set of corrections for the following defect using the directive provided. Note that you need to augment each correction individually according to the directive, i.e., you are generating one augmented correction for every correction you have been provided. --

```
## Player Defect
Defect: {defect}

## Initial Corrections
Correction1: {corrections[0]}
Correction2: {corrections[1]}
Correction3: {corrections[2]}
Correction4: {corrections[3]}
Correction5: {corrections[4]}

## Directive - {directive}
```

Figure 7: Prompts used for augmenting initial coaching snippets based on a set of directives.

B.3 REASONING GENERATION

We prompt GPT-4o to generate reasoning traces by using the list of events which occurred in a trajectory (see Figure 8). Each reasoning trace is comprised of three components; (1) Summary, (2) Successes, (3) Challenges. This information is included in the “format instructions” portion of the prompt. The model is also provided with the list of rules in *Overcooked* before reasoning over the trajectory. The specific set of rules included in the prompt vary based on the recipe the user is attempting to make (see Figures 9, 10, 11, 12).

Reasoning Trace Generation Prompt

You are an expert game analyst. Ground your analysis in the following rules of *Overcooked*: {overcooked_rules}"

Generate a structured summary of an *Overcooked* gameplay trajectory with the following format requirements: {format_instructions} After summarizing, for each field perform self-verification: indicate if it's supported by the events and rules ('valid') and a confidence score between 0-1, under 'verification'

Respond strictly with the JSON object, no extra text.
Events (chronological list): {events}

Figure 8: Reasoning Generation Prompt

Overcooked Rules (original recipe)

Overcooked Rules:

1. Players must pick up ingredients from dispensers and place them in the cooking ranges to cook them.
2. Tomatoes are placed in the cooking pots to make soup.
3. Players must interact with dispensers to pick up ingredients.
4. Onions need to be cut at the cutting board.
5. Steak needs to be cooked on the grill.
6. There is one visible trip hazard and one invisible trip hazard on the map.
7. If the player walks over a hazard while holding an object, they may drop the object.
8. Players must deliver cooked food to the delivery location.
9. A player can only hold one object at a time.
10. Dishes are cleaned at the sink.

11. A player can perform an action at a location by interacting with the location, i.e. cleaning dishes at the sink, or cutting onions at the cutting board.
12. Soup starts cooking when there are three tomatoes in the pot.
13. Steak starts cooking when there is one steak on the grill.
14. Once food is cooking, you can no longer place ingredients on it until the food is cooked and picked up.
15. If you want to drop an ingredient, you have to drop it on a counter.
16. You are only controlling one player in overcooked.
17. Interacting with cooking ranges before the food is cooked will not do anything.
18. Cut onions can only be picked up once you are holding grilled steak.
19. The only ingredients in the game are steak, tomatoes, and onions.
20. The game has no audio, i.e. the player only receives visual cues while playing the game.

Figure 9: Overcooked rules

Overcooked Rules (alternate1 recipe)

Overcooked Rules:

1. Players must pick up ingredients from dispensers and place them in the cooking ranges to cook them.
2. Tomatoes are placed in the cooking pots to make soup.
3. Players must interact with dispensers to pick up ingredients.
4. Onions need to be cut at the cutting board.
5. Steak needs to be cooked on the grill.
6. There is one visible trip hazard and one invisible trip hazard on the map.
7. If the player walks over a hazard while holding an object, they may drop the object.
8. Players must deliver cooked food to the delivery location.
9. A player can only hold one object at a time.
10. Dishes are cleaned at the sink.
11. A player can perform an action at a location by interacting with the location, i.e. cleaning dishes at the sink, or cutting onions at the cutting board.
12. Soup starts cooking when there are three tomatoes in the pot.
13. Steak starts cooking when there is one steak on the grill.
14. The two recipes you need to make are soup with chopped onions and a grilled steak.
15. If you want to drop an ingredient, you have to drop it on a counter.
16. You are only controlling one player in overcooked.
17. Interacting with cooking ranges before the food is cooked will not do anything.
18. Chopped onions can only be picked up once you are holding soup.
19. The only ingredients in the game are steak, tomatoes, and onions.
20. The game has no audio, i.e. the player only receives visual cues while playing the game.
21. Chopped onions must be picked up from the chopping board while holding soup.

Figure 10: Overcooked rules (Alternate 1)

Overcooked Rules (alternate2 recipe)

Overcooked Rules:

1. Players must pick up ingredients from dispensers and place them in the cooking ranges to cook them.

2. Tomatoes are placed in the cooking pots to make soup.
3. Players must interact with dispensers to pick up ingredients.
4. One onion needs to be placed in the cooking pot to make tomato and onion soup.
5. Steak needs to be cooked on the grill.
6. There is one visible trip hazard and one invisible trip hazard on the map.
7. If the player walks over a hazard while holding an object, they may drop the object.
8. Players must deliver cooked food to the delivery location.
9. A player can only hold one object at a time.
10. Dishes are cleaned at the sink.
11. A player can perform an action at a location by interacting with the location, i.e. cleaning dishes at the sink, or cutting onions at the cutting board.
12. Soup starts cooking when there are two tomatoes and one onion in the cooking pot.
13. Steak starts cooking when there is one steak on the grill.
14. The only recipe you need to make is tomato, onion and steak stew.
15. If you want to drop an ingredient, you have to drop it on a counter.
16. You are only controlling one player in overcooked.
17. Interacting with cooking ranges before the food is cooked will not do anything.
18. You can only place an onion in the cooking pot, once there are two tomatoes in the pot.
19. The only ingredients in the game are steak, tomatoes, and onions.
20. The game has no audio, i.e. the player only receives visual cues while playing the game.
21. You need to chop grilled steak on the chopping board.
22. To complete your stew, you first need to pick up the cooked tomato and onion soup from the pot, and then pick up chopped steak from the chopping board while holding the soup.

Figure 11: Overcooked rules (Alternate 2)

Overcooked Rules (alternate2 recipe)

Overcooked Rules:

1. Players must pick up ingredients from dispensers and place them in the cooking ranges to cook them.
2. Tomatoes are placed in the cooking pots to make soup.
3. Players must interact with dispensers to pick up ingredients.
4. One onion needs to be placed in the cooking pot to make tomato and onion soup.
5. Steak needs to be cooked on the grill.
6. There is one visible trip hazard and one invisible trip hazard on the map.
7. If the player walks over a hazard while holding an object, they may drop the object.
8. Players must deliver cooked food to the delivery location.
9. A player can only hold one object at a time.
10. Dishes are cleaned at the sink.
11. A player can perform an action at a location by interacting with the location, i.e. cleaning dishes at the sink, or cutting onions at the cutting board.
12. Soup starts cooking when there are two tomatoes and one onion in the cooking pot.
13. Steak starts cooking when there is one steak on the grill.
14. The two recipes you need to make are tomato and onion soup and grilled steak skewers.

```
15. If you want to drop an ingredient, you have to drop it on a
counter.
16. You are only controlling one player in overcooked.
17. Interacting with cooking ranges before the food is cooked will
not do anything.
18. You can only place an onion in the cooking pot, once there are
two tomatoes in the pot to make tomato and onion soup.
19. The only ingredients in the game are steak, tomatoes, and
onions.
20. The game has no audio, i.e. the player only receives visual ques
while playing the game.
21. You need to chop grilled steak on the chopping board to make
grilled steak skewers.
22. Once you have chopped the grilled steak, you can pick it up and
serve the grilled steak skewers.
23. Ingredients cannot be removed once they are placed in the
cooking pot or on the grill.
```

Figure 12: Overcooked rules (Alternate 3)

B.4 LLM-AS-JUDGE

We prompt GPT-4 to serve as an evaluator to judge the accuracy of the coaching trace generated by our model in comparison to the actual underlying defect in the trajectory (see Figure 13).

System: LLM-Judge

```
You are an evaluator reviewing whether coaching advice generated by
an AI assistant correctly addresses a specific gameplay defect in
Overcooked.
### Task
You are given the following:

- Generated Advice: The assistant's response to a gameplay
issue.
- Actual Defect: The player's observed gameplay problem.
- Reference Coaching Advice: A human-written or high-quality
model coaching response meant to guide users with this type of
defect.

Your job is to decide whether the Generated Advice specifically
and adequately addresses the Actual Defect.

The Reference Coaching Advice is provided as helpful context, to
better understand the defect being resolved. But the generated
advice does not need to match it exactly.
---
### Instructions
- Output Label: [Yes] if the generated advice directly and
sufficiently addresses the defect.
- Output Label: [No] if it fails to address, only partially
addresses, or is too vague, off-topic, or incorrect in relation
to the defect.
- Include a short Think step to reason through your decision.
---
### Format
### User
Generated Advice: <generated_answer>
Actual Defect: <actual_defect>
Reference Coaching Advice: <reference_advice>

### Assistant
Think: <your brief analysis of the advice and defect relationship>
Label: [Yes/No]
```

```
---
### Examples

#### Example 1
### User
Generated Advice: Make sure you place the ingredients in the cooking
range to start
Actual Defect: The player does not know what to do with the
ingredients.
Reference Coaching Advice: Place ingredients in the pot to cook
them. Don't leave them on the counters.

### Assistant
Think: The advice, though shorter, tells the player what to do with
the ingredients, aligning with the defect.
Label: [Yes]
---
#### Example 2

### User
Generated Advice: Avoid inefficient movement.
Actual Defect: The player makes unnecessary movements/turns while
navigating the map.
Reference Coaching Advice: Stick to efficient paths and avoid
zig-zagging or turning back.

### Assistant
Think: The generated advice is high-level but directly relevant to
the defect.
Label: [Yes]
---
#### Example 3

### User
Generated Advice: Remember to place three tomatoes in the pot to
cook the soup.
Actual Defect: The player thinks it is best to only serve soup.
Reference Coaching Advice: It would be helpful to serve both soup
and steak, rather than soup alone.

### Assistant
Think: The generated advice focuses on preparing soup, which does
not address the misconception about ignoring other dishes. The
reference advice suggests serving both soup and steak, which is not
mentioned in the generated advice.
Label: [No]
---
#### Example 4

### User
Generated Advice: Place three tomatoes in the pot to start the
cooking process.
Actual Defect: The player tries to deliver cooked food by placing it
on a counter.
Reference Coaching Advice: Deliver completed dishes by placing them
at the delivery location, not counters.

### Assistant
Think: The advice is unrelated to the problem of delivering food to
the right location.
Label: [No]
---
#### Example 5

### User
Generated Advice: Remember to slice onions on the board.
Actual Defect: The player tries to place onions in the sink.
```

```

Reference Coaching Advice: Use the cutting board for prepping onions
before cooking.

### Assistant
Think: The advice provides the correct instruction that addresses
the mistake with the handling of onions.
Label: [Yes]

User: LLM-Judge

Generated Advice: <|generated_answer|>
Actual Defect: <|answer|>
Reference Coaching Advice: <|reference_answer|>

```

Figure 13: LLM-as-judge to evaluate coaching outputs

B.5 BEHAVIOR-CRITIC BASELINE

Our baseline behavior-critic leverages GPT-4o to perform coaching or correction based on a down-sampled video of the defective trajectory. The prompts for both tasks can be found in Figure 14. We also enhanced our behavior critic by including the synthetically generated reasoning traces, as per Section B.3. The prompts for this enhanced behavior critic can be found in Figure 15.

```

Behavior Critic (Coaching)

**Task:**
You will be provided a short video clip (sequence of images) of a
human playing Overcooked.
The player may or may not have a defect in their gameplay strategy
which is hampering their performance. Your job is to analyze this
trajectory and provide corrective feedback to enable the user to
understand and fix the defect in their overcooked gameplay.

Please share your analysis of what is occurring in the video with
regards to the potentially defective gameplay trajectory, prior to
providing your feedback. Note that you need to analyze the actions
of the blue player.

**Example Responses (you should follow the same JSON output
format):**
{examples}

**Note**
The images are key frames extracted from a video and are presented
in the same order as the original footage.
Each player will have at most one defect. The potential defects that
the player may have will be one of the following categories:
{defects}
To help you perform the task, here are some important rules to know
about playing Overcooked:
**{overcooked_rules}
**Now, please look at the following sequence of images carefully.**
{extra}

```

Corrections

```
**Task:**
You are an AI assistant tasked with improving a user's performance
in Overcooked.
To help you perform the task, here are some important rules to know
about playing Overcooked:
**{overcooked_rules}

Given a defective seed trajectory, provide corrective actions, that
the user should perform next, that align with the user's strategy.
Use the accompanying reference trajectory, a non-defective
demonstration on a reference map, to infer their strategy.

**Example Responses (you should follow the same JSON output
format):**
{examples}

**Note**
You need to predict the immediate next action the user should take
to begin to fix the defect in their behavior. You must choose from
one of the following 8 actions: [move forward, move backward, move
left, move right, turn left, turn right, interact, no action]
**Now, please look at the following sequence of images carefully.**
{extra}
```

Figure 14: Prompts for GPT-4o baselines without reasoning.

Coaching Prompt

```
**Task:**
You will be provided a short video clip (sequence of images) of a
human playing Overcooked.
The player may or may not have a defect in their gameplay strategy
which is hampering their performance. Your job is to provide
corrective feedback to enable the user to understand and fix the
defect in their overcooked gameplay.
In addition to the sequence of images, you will also be provided
with an analysis of the player's trajectory, which will provide a
summary, along with the challenges, successes faced by the player
during the trajectory.

**Example Responses (you should follow the same JSON output
format):**
{examples}
**Note**
The images are key frames extracted from a video and are presented
in the same order as the original footage.
Each player will have at most one defect. The potential defects that
the player may have will be one of the following categories:
**{defects}
To help you perform the task, here are some important rules to know
about playing Overcooked:
**{overcooked_rules}

**Now, please look at the following sequence of images carefully.**
{extra}
```

Corrections Prompt

```

**Task:**
You are an AI assistant tasked with improving a user's performance
in Overcooked.
To help you perform the task, here are some important rules to know
about playing Overcooked:
**{overcooked_rules}
Given a defective seed trajectory, provide corrective actions, that
the user should perform next, that align with the user's strategy.
In addition to the sequence of images, you will also be provided
with an analysis of the player's trajectory, which will provide a
summary, along with the challenges, successes faced by the player
during the trajectory.
Use the accompanying reference trajectory, a non-defective
demonstration on a reference map, to infer their strategy.

**Example Responses (you should follow the same JSON output
format):**
{examples}

**Note**
You need to predict the immediate next action the user should take
to begin to fix the defect in their behavior. You must choose from
one of the following 8 actions: [move forward, move backward, move
left, move right, turn left, turn right, interact, no action]
**Now, please look at the following sequence of images carefully.**

{extra}

```

Figure 15: Prompts for GPT-4o baselines with reasoning.

C DEFECTS

This section provides the full list of defects included in the training-set, held-out-defect dataset and the task-generalization dataset (see Figure 16).

Defects

```

## Training Defects
1: "Player does not know that the cooking pot needs to be full for
the soup to start cooking.",
2: "Player does not know that you need to interact with dispensers
to pick up ingredients.",
3: "Player believes that leaving cooked food on an empty counter
means they have delivered the food.",
4: "Player does not know that they need to face the object to
successfully interact with any object in the game.",
5: "The player thinks it is best to only serve soup.",
6: "The player thinks tomatoes can be placed on the grill.",
7: "The player thinks onions should be placed directly on the
grill.",
8: "Player does not know that cut onions need to be placed on the
cooked steak before serving.",
9: "Player mistakenly thinks that the teammate will always walk
around them.",
10: "Player does not know that dishes need to be cleaned before
being used to pick up food.",
11: "Player does not walk around the visible trip hazards when they
are holding an object.",
12: "Player does not remember or keep track of where the invisible
trip hazard is." ,
13: "Player attempts to pick up food before it is finished
cooking.",

```

```

14: "Player is overly cautious about trip hazards when they are
holding something, and doesn't go near spaces adjacent to the
hazard.",
15: "Player tries to directly hand ingredients they are holding to
their teammate.",
16: "Player tries to wash ingredients in the sink.",
"No Defect": "No Defect"

## Heldout defects
1: "Player does not know that ingredients need to be placed in the
cooking ranges to cook.",
2: "The player never serves cooked soup.",
3: "The player thinks it is best to only serve steak.",
4: "The player thinks onions can be cut in the cooking pot.",
5: "The player places ingredients on the counter instead of in the
cooking range.",

```

Defects for Task Generalization Eval

```

## Alternate Task 1
1: "Player forgets to place chopped onions in the cooked soup before
serving the soup.",
2: "Player tries to place chopped onions on grilled steak.",

## Alternate Task 2
1: "Player tries to place onions on the chopping board instead of in
the cooking pot.",
2: "Player tries to serve cooked tomato and onion soup before
picking up chopped steak from the chopping board.",

## Alternate Task 3
1: "Player tries to place three tomatoes into the cooking pot,
instead of two tomatoes and an onion.",
2: "Player doesn't know that grilled steak needs to be chopped to
make grilled steak skewers.",

```

Figure 16: List of defects used in each dataset

D RECIPES

Our training dataset uses the original overcooked recipe, i.e. the player can cook either grilled steak and onions or tomato soup. For the task-generalization experiment, the user is prompted to make three alternate recipes:

1. Grilled Steak + Soup and chopped onions
2. Tomato, Onion and Steak Stew
3. Tomato and Onion Soup + Grilled Steak Skewers

The descriptions for each recipe, as utilized in our prompts, are provided in Figure 17)

Recipes

Original Recipe: Grilled Steak and Onions: The player needs to cook steak on the grill, chop onions on the chopping board, and combine them with the cooked steak. \nTomato Soup: The player needs to place three tomatoes in the cooking pot.

Task-Gen Recipe 1: Grilled Steak: The player needs to cook steak on the grill and serve it. \n Soup and Chopped Onions: The player needs to make soup by placing three tomatoes in the cooking pot, and add chopped onions after picking up the soup.

Task-Gen Recipe 2: Tomato, Onion and Steak Stew: The player needs to cook steak on the grill and chop it on the chopping board. The player also needs to cook tomato and onion soup, by cooking two tomatoes and one onions in the cooking pot. Finally, the chopped steak should be added to the soup prior to serving.

Task-Gen Recipe 3: Tomato and Onion Soup: The player needs to cook two tomatoes and one onion in the cooking pot. Grilled Steak Skewers: The player needs to cook steak on the grill, then chop it prior to serving.

Figure 17: Recipes

E MODEL INSTRUCTIONS

The exact instructions included in the prompt while tuning our model to perform coaching or corrections are provided in Figure 18.

System

You are an AI assistant tasked with providing feedback to a user to improve their performance on the game Overcooked.
<recipe prompt>

You will be given the user's gameplay trajectory, which includes a specific defect affecting their performance. Your job is to analyze this trajectory and provide corrective feedback to enable the user to understand and fix the defect in their overcooked gameplay.

Defective trajectory: <defective_trajectory>

Corrective Feedback:

User:

You are an AI assistant tasked with improving a user's performance in Overcooked.
<recipe prompt>

Given a defective seed trajectory, provide corrective actions, that the user should perform next, that align with the user's strategy. Use the accompanying reference trajectory, a non-defective demonstration on a reference map, to infer their strategy.

Reference Trajectory: <reference_trajectory>

Defective Trajectory: <defective_trajectory>

Corrective Actions:

Figure 18: Model prompts for our model

F VISUAL-QA QUESTIONS

As explained in Section 4.2, we utilize visual QA datasets to improve our model's ability to ground and perform task-generalization. This section provides the entire list of questions in the Image/Video/Trajectory-QA datasets (see Figure 19). The Video-QA and Trajectory-QA datasets use the same set of predominantly yes/no questions, whereas the Image-QA dataset features more complex questions requiring a wider range of answers..

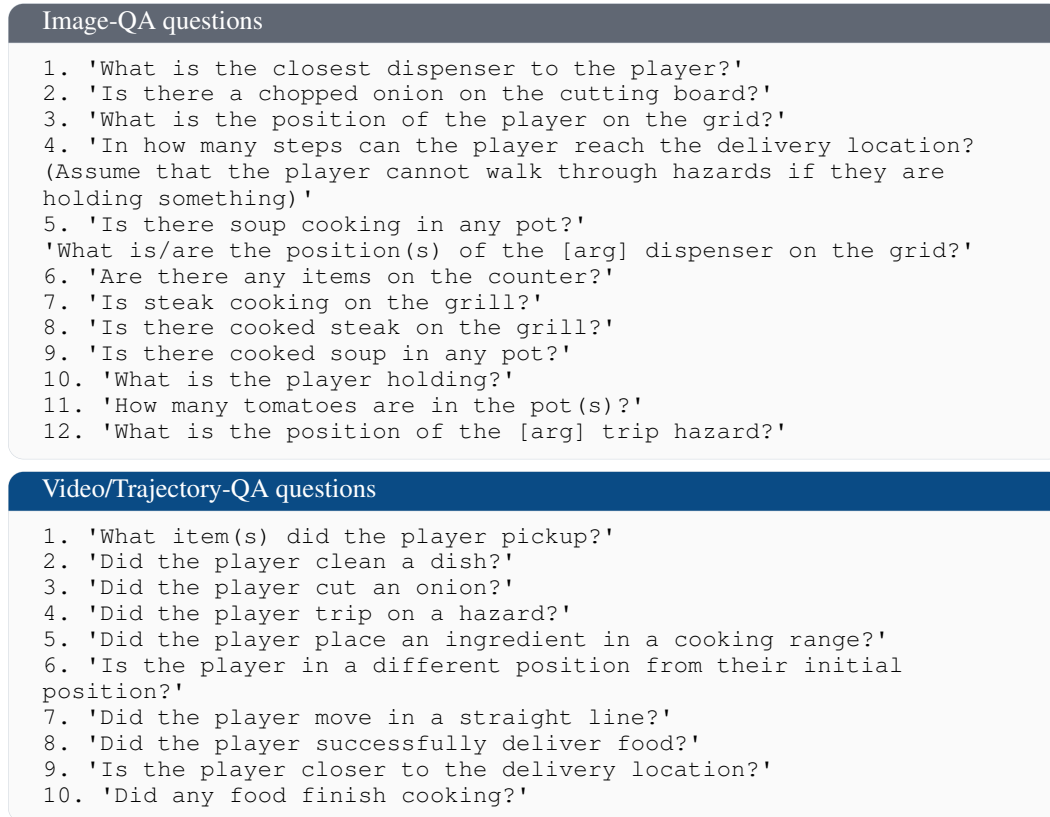


Figure 19: Questions utilized in each of the visual grounding datasets