

# Do What? Teaching Vision-Language-Action Models to Reject the Impossible

Wen-Han Hsieh\* Elvis Hsieh\* Dantong Niu  
Trevor Darrell Roi Herzig David M. Chan  
University of California, Berkeley

{hense1219, htelvis92, niudantong.88, trevordarrell, roeiherz, davidchan}@berkeley.edu

## Abstract

Recently, Vision-Language-Action (VLA) models have demonstrated strong performance on a range of robotic tasks. These models rely on multimodal inputs, with language instructions playing a crucial role—not only in predicting actions, but also in robustly interpreting user intent, even when the requests are impossible to fulfill. In this work, we investigate how VLAs can recognize, interpret, and respond to false-premise instructions—natural language commands that reference objects or conditions absent from the environment. We propose—Instruct-Verify-and-Act (IVA)—a unified framework that (i) detects when an instruction cannot be executed due to a false premise, (ii) engages in language-based clarification or correction, and (iii) grounds plausible alternatives in perception and action. Towards this end, we construct a large-scale instruction tuning setup with structured language prompts and train a VLA model capable of handling both accurate and erroneous requests. Our approach leverages a contextually augmented, semi-synthetic dataset containing paired positive and false-premise instructions, enabling robust detection and natural language correction. Our experiments show that IVA can improve false premise detection accuracy by 97.56% over baselines, while increasing successful responses in false-premise scenarios by 50.78%.

## 1 Introduction

Vision-Language-Action (VLA) models (Niu et al., 2024; Kim et al., 2024; Li et al., 2024; Niu et al., 2025) represent a significant advancement in robotics, enabling agents to perform tasks using multimodal inputs by integrating visual perception, natural language understanding, and action generation. However, as these models are increasingly deployed in open-ended, real-world environments, they must handle diverse and often ambiguous instructions issued by users. Consider a household

\* Equal contribution

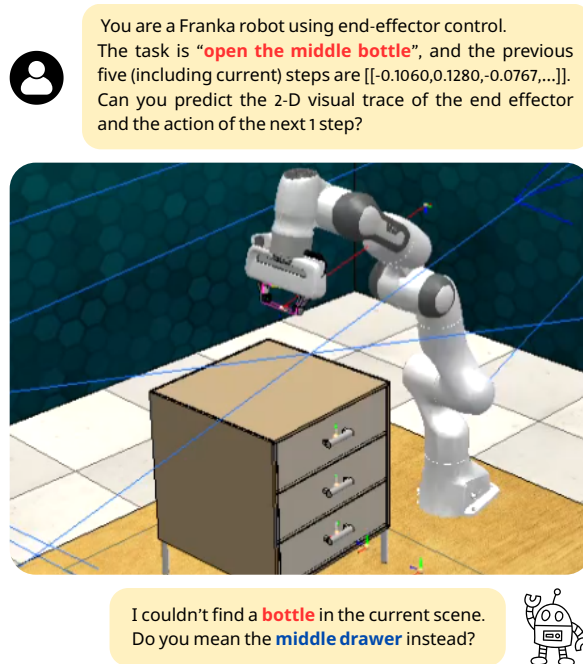


Figure 1: **Our Instruct-Verify-and-Act (IVA) framework is designed to handle false-premise instructions.** When the robot receives a command referencing a non-existent object (a *bottle*), it detects the false premise, and generates a clarifying response that corrects the instruction, suggesting a valid alternative (a *drawer*).

robot presented with the command, “Bring me the red mug on the kitchen table,” when no such mug exists. In this scenario, the robot’s ability to recognize the impossibility of the request, explain the issue, and suggest plausible alternatives is essential for safe and effective human-robot interaction. Yet, most existing VLAs lack mechanisms to detect or respond appropriately to instructions grounded in false premises - commands that reference objects, attributes, or relations not present in the environment.

While large multimodal models have made significant strides in visual grounding and instruction following, prior research in both natural language processing and robotics has typically assumed

that user instructions are feasible and contextually grounded. In domains such as visual question answering and dialogue, models have been studied for their robustness to unanswerable or contradictory queries have been developed which can evaluate and correct for false premises in the case of question answering (Johnson et al., 2017; Suhr et al., 2017; Liu et al., 2019; Hudson and Manning, 2019; Gurari et al., 2018; Mahendru et al., 2017; Whitehead et al., 2022; Mashrur et al., 2023, 2022; Li et al., 2020; Prabhakar et al., 2018; Mahendru et al., 2017; Karnik et al., 2024; Ren et al., 2023). However, in the context of embodied agents and robotics, the issue of handling false-premise instructions - recognizing unfulfillable commands and producing helpful language-based corrections - remains largely unexplored. Existing robotics benchmarks and VLAs focus on execution success given correct instructions, without testing how these systems respond when user intent is unattainable or ambiguous.

To address this gap, we develop a unified VLA model capable of interpreting and correcting false-premise instructions in robotic settings. Our framework combines large-scale instruction tuning with contextually-augmented datasets, enabling the model to detect unfulfillable requests, provide natural language feedback, and suggest alternative courses of action. See Figure 1. We evaluate our approach across eight robotic tasks, measuring false premise detection accuracy and the rate of successful responses in false-premise scenarios. Our method achieves an improvement of 97.56% in false premise detection over established baselines, and increases success in false-premise handling by 50.78%. These results demonstrate that language-aware robots can move beyond simple execution - reasoning about user intent, clarifying ambiguity, and interacting naturally even when confronted with the impossible.

## 2 Related Work

**Vision-Language-Action Models (VLAs).** VLAs integrate visual perception, language understanding, and action generation to produce robot control sequences from visual observations and language instructions. Models such as LLARVA (Niu et al., 2024),  $\pi$  (Physical Intelligence et al., 2025), Gemini Robotics (Gemini Robotics et al., 2025), OpenVLA (Kim et al., 2024), RT-2 (Brohan et al., 2023b), PaLM-E embodied multimodal model (Driess et al., 2023), and LLaRA (Li et al.,

2024) fine-tune pretrained large multimodal models (LMMs) to predict actions, often using special tokens to represent the action space. These models share the goal of adapting LMMs for robotic control as they differ in the choice of LMM and action encoding methods. While existing VLAs typically leverage language decoders pretrained for high-level tasks (e.g., image captioning (Kim et al., 2024) and VQA (Brohan et al., 2023a)), when deployed in diverse, real-world environments, they exhibit a critical limitation: the inability to handle instructions grounded in false premises. Our work addresses this gap by explicitly training VLA models to recognize when referenced objects or conditions are absent and to respond appropriately.

**False Premises.** The problem of detecting false premises has long been studied in NLP, particularly in QA benchmarks such as SQUAD 2.0 (Rajpurkar et al., 2018) and False QA (Hu et al., 2023), with research primarily focused on model uncertainty for unanswerable questions (Raina and Gales, 2022; Sulem et al., 2022). This concept of false premises has recently spread from NLP to other areas of interest, including visual question answering (Ray et al., 2016), image/text matching (Feng and Lapata, 2012; Xu et al., 2015; Ordonez et al., 2011; Karpathy and Fei-Fei, 2015; Fang et al., 2015), image-grounded conversation (Mostafazadeh et al., 2017), tool usage (Toor et al., 2019) and hallucination detection (Rohrbach et al., 2018). Research in human-robot interaction has explored handling imperfect instructions and clarifying ambiguous instructions (Deits et al., 2013; Park et al., 2024; Prabhakar et al., 2018; Shi et al., 2024; Tellex et al., 2014), however this short paper is the first to explore false premise tasks in the context of vision-language-action models (VLAs).

## 3 A VLA that Instructs, Verifies, and Acts

Our method builds upon the LLARVA model (Niu et al., 2024), a VLA architecture designed for robotic instruction following. LLARVA leverages instruction-tuned large multimodal models (LMMs) to jointly interpret visual observations, natural language instructions, and robot proprioception, outputting robot actions along with intermediate visual representations called *visual traces*.

**Input.** LLARVA takes a visual observation  $o_t$ , an RGB image at timestep  $t$ , and a structured natural-language instruction  $l_t$  as input. The language instruction encodes the robot type (e.g.,

Franka Panda), control mode (e.g., end-effector or joint control), task description (e.g., “close the drawer”), proprioceptive states (internally-sensed joint-angle vectors that indicate the robot’s current position before it plans the next move) from the previous  $h$  timesteps, and the number of future actions ( $n$ ) to predict, the instruction tuning template is shown as follows:

“You are a [Robot] robot using [Control Mode] control. The task is [Task Description], and the previous [h] steps are [Proprioceptive States]. Can you predict the trajectory of the end-effector and the action of the next [n] steps?”

**Model.** LLARVA integrates three main components as follows:

- **Vision Encoder:** A frozen pretrained visual encoder (CLIP ViT-L/14) encodes image observations  $o_t$  into visual tokens.
- **Language Encoder:** Tokenizes and embeds the language instruction  $l_t$ , forming language tokens.
- **Multimodal Decoder:** An autoregressive transformer decoder combines visual and language tokens, generating predictions for robotic actions  $A_{t:t+n-1}$  and future visual traces  $P_{t:N}$ , formally:

$$\pi(o_t, l_t) \rightarrow A_{t:t+n-1}, P_{t:N}$$

where  $A_{t:t+n-1}$  represents predicted robot actions for the next  $n$  steps, and  $P_{t:N}$  indicates predicted 2-D visual trajectories of the robot’s end-effector from timestep  $t$  to episode end  $N$ .

LLARVA is initially pretrained on large-scale vision-action instruction data from Open X-Embodiment (OXE) and subsequently fine-tuned for specific robotic tasks, enabling strong generalization across diverse tasks and environments.

**False Premise Instruction Dataset.** Following Niu et al. (2024), we utilize image-action pairs from the OXE dataset (Vuong et al., 2023), providing rich visual, language, and action representations. However, prior work did not explicitly address reasoning about the feasibility of user instructions. To address this gap, we constructed a dedicated dataset explicitly curated to handle false premise scenarios and their corresponding corrections.

Our false-premise instruction dataset is generated from robotic trajectories available in RLBench

(James et al., 2019). Each task in our dataset includes two categories of false premise instructions:

**In-Domain False Premise.** These instructions involve geometrically similar and contextually plausible objects derived from related tasks, making the intended correction relatively intuitive. For instance, in the task of closing a jar, if the user prompt is “The task is *close the blue safe*”, the model is expected to respond, “I don’t see a safe in the current scene. Do you mean jar?”.

**Out-of-Domain False Premise.** These instructions contain clearly infeasible or nonsensical requests involving objects or scenarios impossible within the given context. For example, during the open drawer task, if the user prompts “The task is *open the top elephant*,” the model should identify the absurdity and respond appropriately, such as “I couldn’t find a elephant in the current scene,” subsequently terminating the interaction since the request is fundamentally invalid.

For training purposes, our dataset composition strategically includes approximately 20% of episodes containing Out-of-Domain false premises, and 65% of episodes containing In-Domain false premises injected into 10% of their respective steps. This deliberate distribution ensures comprehensive exposure and training of the model in recognizing and handling various types of erroneous instructions.

**IVA Training.** We adopt an end-to-end instruction-tuning approach, closely following LLARVA’s training methodology, to train our model using the newly curated false-premise dataset. While keeping both the vision and language encoders frozen, we fine-tune the auto-regressive transformer decoder using standard LoRA adapters.

Specifically, our training procedure utilizes 800 episodes per task, with each episode containing image observations  $o_t$ , language instructions  $l_t$ , ground-truth robotic actions  $\hat{A}_{t:t+n-1}$ , and visual traces  $\hat{P}_{t:N}$ . The training data includes a mixture of true-premise and false-premise instructions, with roughly 20% of episodes containing Out-of-Domain false premises and approximately 65% containing In-Domain false premises, introduced at 10% of steps within each episode. Given  $o_t$  and  $l_t$ , the model predicts actions and visual traces auto-regressively:

$$p(\hat{A}_{t:t+n-1}, \hat{P}_{t:N} | o_t, l_t) = \prod_{i=1}^{|R|} p_{\theta}(x_i | o_t, l_t)$$

where  $\theta$  denotes trainable parameters,  $x_i$  is the

Table 1: Comparison of IVA and LLaRVA across RLBench tasks: overall success rate, FP (false premise) detection rate (In-Domain/Out-of-Domain), and TP (true premise) success rate. The overall success rate is computed as the average of true and false premises success rates.

Task	IVA			LLaRVA		
	Overall Success	FP Detection (In-Domain/Out-of-Domain)	TP Success	Overall Success	FP Detection (In-Domain/Out-of-Domain)	TP Success
meat off grill	58%	100% / 100%	16%	2%	0% / 0%	4%
open drawer	61%	100% / 80%	32%	20%	0% / 0%	40%
push buttons	68%	100% / 100%	36%	16%	0% / 0%	32%
put money in safe	64%	100% / 100%	28%	20%	0% / 0%	40%
reach and drag	80%	100% / 100%	60%	22%	0% / 0%	44%
slide block	96%	100% / 100%	92%	44%	0% / 0%	88%
sweep to dustpan	94%	100% / 100%	88%	30%	0% / 0%	60%
turn tap	61%	100% / 80%	32%	20%	0% / 0%	40%
close jar	50%	100% / 100%	0%	0%	0% / 0%	0%

predicted token at timestep  $i$ , and  $R$  represents the full response sequence. We compute the training loss as the standard cross-entropy between predicted tokens and ground-truth annotations.

Unlike LLaRVA’s two-stage procedure (pre-training and fine-tuning), we train our model end-to-end on our unified dataset, ensuring simultaneous learning of accurate robotic action prediction, robust false-premise detection, and appropriate language-based correction responses.

## 4 Results

We evaluated IVA on two aspects: its ability to detect and correct false-premise (FP) instructions, and its performance on standard, true-premise (TP) tasks.

**Experiment Setup.** We conducted experiments on 9 RLBench tasks, generating 25 episodes per task with randomly varied object positions. Each episode was paired with two type of user prompts: one standard and one containing a false premise. The model was fine-tuned for 8 hours on 8 A100 GPUs. During fine-tuning, the model received the front camera view and the previous 5 joint positions as input, and predicted both the visual trace and the next action step, represented as an 8-dimensional vector (7 joint velocities plus a binary gripper state). We report the success/failure rates for one fixed seed per task on the validation set, comparing with LLaRVA as a baseline. The results are shown in Table 1.

**Evaluation Procedure.** We evaluate IVA in a single-pass, end-to-end setting on all 225 episodes (25 episodes  $\times$  9 RLBench tasks) with randomized object poses. For each episode, we generate a true- or false-premise instruction and score IVA’s full

response as follows:

(1) **Detection Stage.** IVA first outputs a textual response, which we parse to classify the instruction as either *Accept* (True-Premise) or *Clarify/Refuse* (False-Premise).

- **Accept (True-Premise):** Scored 1 if the subsequent executed trajectory exactly matches the ground-truth; otherwise 0.
- **Clarify/Refuse (False-Premise):** Scored 1 for an explicit out-of-domain refusal (e.g., “I’m sorry, that object isn’t here...”). For in-domain false premises, scored 1 if the object is correctly re-targeted, and 0 if any impossible action is attempted.

For multi-step episodes, we average the per-step Detection scores to yield a single FP score per episode.

(2) **Execution Stage.** Whenever IVA “Accepts,” we execute the predicted 8D joint-velocity sequence. Each RLBench task’s built-in success detector then labels the outcome as success (1) or failure (0).

(3) **Overall Accuracy.** We average the 225 episode scores (Detection + Execution) to yield a single accuracy metric that jointly rewards correct task execution and safe refusal/clarification behavior.

**False Premise Detection and Correction.** A major challenge in real-world robotics is handling instructions that reference unavailable objects or impossible actions. To address this, we fine-tuned IVA on a curated false-premise dataset and evaluated it alongside the true-premise-trained baseline. Test prompts included both “In-Domain” false premises (plausible but absent objects, e.g., “open the middle

block” when only a drawer is present) and “Out-of-Domain” false premises (impossible requests, e.g., “open the top chicken” for drawer task).

IVA achieved perfect detection (100%) on In-Domain false-premise instructions, consistently identifying and correcting these cases. For Out-of-Domain false premises, IVA reached a detection accuracy of 97.78%. In both scenarios, IVA generated contextually appropriate clarifications—such as “I don’t see a tree in the current scene. Do you mean jar?”—for In-Domain false premise and terminate the interaction when face the Out-of-Domain false premise.

**Performance on True-Premise Tasks.** To ensure that enhanced false-premise reasoning does not compromise standard performance, we tested IVA on tasks with only true-premise instructions. IVA maintained a success rate of  $42.67\% \pm 8.34\%$  compared to the baseline’s  $38.67\% \pm 8.55\%$ . The observed difference lies well within the variance range, suggesting that the slight drop in performance is statistically insignificant. Action prediction accuracy and overall task completion remained fairly stable, confirming that robust false-premise handling does not significantly degrade general competence.

## 5 Conclusion

In this work, we introduced the Instruct-Verify-and-Act framework, enabling VLA models to robustly detect, clarify, and correct false-premise instructions in robotic settings. Our approach achieves strong performance in both false-premise detection and standard task execution, demonstrating the value of explicit false-premise reasoning in VLA models. We hope these results demonstrate the potential of language-aware robots to engage more naturally and safely with users, even when faced with ambiguous or impossible commands. Future work will focus on extending IVA’s capabilities to more complex environments and real-world deployments. Finally, we do not anticipate a specific negative impact, but, as with any machine learning method, we recommend exercising caution.

## Limitations

While our IVA framework demonstrates strong performance in detecting and correcting false-premise instructions within vision-language-action (VLA) models, several limitations remain that should be addressed in future work:

**Dataset Scope and Realism.** Our false-premise instruction dataset is primarily generated from the RL Bench environment, which, despite its diversity, remains a simulated domain with a constrained set of objects, scenes, and tasks. Consequently, the distribution and complexity of false-premise instructions may not fully capture the ambiguity and diversity encountered in real-world human-robot interaction. Moreover, the proportion of In-Domain and Out-of-Domain false premises is artificially balanced to facilitate training and evaluation, and may not reflect the true incidence of such instructions in open-world settings.

**Generalization to Real-World Deployment.** Although IVA achieves high detection rates in simulation, its robustness in real-world deployments is not yet validated. Domain shift—such as differences in visual appearance, sensor noise, or language usage—could degrade performance. Additionally, our framework assumes that visual observations and proprioceptive data are accurately and reliably captured, which may not always hold in practical robotics scenarios.

**Correction and Clarification Strategies.** The natural language responses generated by IVA are limited to clarifications and suggestions based on the specific types of false premises represented in the training data. The model’s ability to propose truly creative or contextually appropriate alternatives remains limited, especially for Out-of-Domain or out-of-distribution instructions. In more complex environments, nuanced reasoning about task feasibility, user intent, and multi-turn clarification dialogues may be necessary.

**Instruction and Environment Complexity.** The instructions used for evaluation are relatively short and structured, and the environments contain a modest number of distractor objects. Real human instructions can be longer, more ambiguous, and embedded in broader conversational contexts. Our current framework does not explicitly handle multi-turn dialogues, implicit user intent, or ambiguous references beyond the immediate instruction.

**Scalability and Efficiency.** Our approach leverages instruction-tuned large multimodal models with frozen vision and language encoders, which, while effective, may impose computational and memory overhead unsuitable for some real-time or resource-constrained robotic applications.

## Acknowledgments

We would like to thank Zhang-Wei Hong and Alexander Pan for their helpful feedback and discussions. Authors, as part of their affiliation with UC Berkeley, were supported in part by the National Science Foundation, US Department of Defense, and/or the Berkeley Artificial Intelligence Research (BAIR) industrial alliance program. This research was also developed with funding from the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-25-3-0133. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of any sponsor, the Department of Defense, or the U.S. Government.

## References

- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, and 35 others. 2023a. [RT-2: Vision-language-action models transfer web knowledge to robotic control](#). *Preprint*, arxiv:2307.15818 [cs].
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, and 1 others. 2023b. [RT-2: Vision-language-action models transfer web knowledge to robotic control](#). *arXiv preprint arXiv:2307.15818*.
- Robin Deits, Stefanie Tellex, Pratiksha Thaker, Dimitar Simeonov, Thomas Kollar, and Nicholas Roy. 2013. Clarifying commands with information-theoretic human-robot dialog. *Journal of Human-Robot Interaction*, 2(2):58–79.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, and 1 others. 2023. [Palm-e: An embodied multimodal language model](#). *arXiv preprint arXiv:2303.03378*.
- Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, and 1 others. 2015. From captions to visual concepts and back. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1473–1482.
- Yansong Feng and Mirella Lapata. 2012. Automatic caption generation for news images. *IEEE transactions on pattern analysis and machine intelligence*, 35(4):797–812.
- Team Gemini Robotics, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, and 1 others. 2025. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3608–3617.
- Shengding Hu, Yifan Luo, Huadong Wang, Xingyi Cheng, Zhiyuan Liu, and Maosong Sun. 2023. Won’t get fooled again: Answering questions with false premises. In *Proceedings of the 61st Annual Meeting of the ACL (Volume 1: Long Papers)*, pages 5626–5643.
- Drew A Hudson and Christopher D Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.
- Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. 2019. [RLbench: The robot learning benchmark and learning environment](#). *Preprint*, arXiv:1909.12271.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910.
- Sathwik Karnik, Zhang-Wei Hong, Nishant Abhangi, Yen-Chen Lin, Tsun-Hsuan Wang, and Pulkit Agrawal. 2024. [Embodied red teaming for auditing robotic foundation models](#). *ArXiv*, abs/2411.18676.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137.
- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. 2024. [OpenVLA: An open-source vision-language-action model](#). *Preprint*, arxiv:2406.09246 [cs].
- Mengdi Li, Cornelius Weber, and Stefan Wermter. 2020. Neural networks for detecting irrelevant questions during visual question answering. In *Artificial Neural Networks and Machine Learning—ICANN 2020: 29th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 15–18, 2020, Proceedings, Part II 29*, pages 786–797. Springer.

- Xiang Li, Cristina Mata, Jongwoo Park, Kumara Kahatapitiya, Yoo Sung Jang, Jinghuan Shang, Kanchana Ranasinghe, Ryan Burgert, Mu Cai, Yong Jae Lee, and Michael S. Ryoo. 2024. *LLaRA: Supercharging robot learning data for vision-language policy*. Preprint, arxiv:2406.20095 [cs].
- Runtao Liu, Chenxi Liu, Yutong Bai, and Alan L Yuille. 2019. Clevr-ref+: Diagnosing visual reasoning with referring expressions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4185–4194.
- Aroma Mahendru, Viraj Prabhu, Akrit Mohapatra, Dhruv Batra, and Stefan Lee. 2017. The promise of premise: Harnessing question premises in visual question answering. *arXiv preprint arXiv:1705.00601*.
- Akib Mashrur, Wei Luo, Nayyar A Zaidi, and Antonio Robles-Kelly. 2022. Semantic multi-modal reprojection for robust visual question answering. In *2022 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–6. IEEE.
- Akib Mashrur, Wei Luo, Nayyar A Zaidi, and Antonio Robles-Kelly. 2023. Robust visual question answering via semantic cross modal augmentation. *Computer Vision and Image Understanding*, page 103862.
- Nasrin Mostafazadeh, Chris Brockett, Bill Dolan, Michel Galley, Jianfeng Gao, Georgios P Spithourakis, and Lucy Vanderwende. 2017. Image-grounded conversations: Multimodal context for natural question and response generation. *arXiv preprint arXiv:1701.08251*.
- Dantong Niu, Yuvan Sharma, Giscard Biamby, Jerome Quenum, Yutong Bai, Baifeng Shi, Trevor Darrell, and Roei Herzig. 2024. Llarva: Vision-action instruction tuning enhances robot learning. In *Conference on Robot Learning*.
- Dantong Niu, Yuvan Sharma, Haoru Xue, Giscard Biamby, Junyi Zhang, Ziteng Ji, Trevor Darrell, and Roei Herzig. 2025. Pre-training auto-regressive robotic models with 4d representations. *arXiv preprint arXiv:2502.13142*.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. Im2text: Describing images using 1 million captioned photographs. *Advances in neural information processing systems*, 24.
- Jeongeun Park, Seungwon Lim, Joonhyung Lee, Sangbeom Park, Minsuk Chang, Youngjae Yu, and Sungjoon Choi. 2024. Clara: Classifying and disambiguating user commands for reliable interactive robotic agents. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. To appear.
- Team Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, and 1 others. 2025. pi-0.5: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*.
- Prakruthi Prabhakar, Nitish Kulkarni, and Linghao Zhang. 2018. Question relevance in visual question answering. *arXiv preprint arXiv:1807.08435*.
- Vatsal Raina and Mark J. F. Gales. 2022. Answer uncertainty and unanswerability in multiple-choice machine reading comprehension. In *Findings of the Association for Computational Linguistics: EMNLP 2022*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 784–789.
- Arijit Ray, Gordon Christie, Mohit Bansal, Dhruv Batra, and Devi Parikh. 2016. Question relevance in vqa: Identifying non-visual and false-premise questions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 919–924.
- Allen Z. Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, Zhenjia Xu, Dorsa Sadigh, Andy Zeng, and Anirudha Majumdar. 2023. Robots that ask for help: Uncertainty alignment for large language model planners. *Proceedings of Machine Learning Research*, 229. Publisher Copyright: © 2023 Proceedings of Machine Learning Research. All Rights Reserved.; 7th Conference on Robot Learning, CoRL 2023 ; Conference date: 06-11-2023 Through 09-11-2023.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*.
- Lucy X. Shi, Zheyuan Hu, Tony Z. Zhao, Archit Sharma, Karl Pertsch, Jianlan Luo, Sergey Levine, and Chelsea Finn. 2024. Yell at your robot: Improving on-the-fly from language corrections. *arXiv preprint arXiv:2403.12910*.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 217–223.
- Elior Sulem, Jamaal Hay, and Dan Roth. 2022. Yes, no or idk: The challenge of unanswerable yes/no questions. In *Proceedings of the 2022 Conference of the North American Chapter of ACL: Human Language Technologies (NAACL)*, pages 1075–1085.
- Stefanie Tellex, Ross A. Knepper, Adrian Li, Thomas M. Howard, Daniela Rus, and Nicholas Roy. 2014. Asking for help using inverse semantics. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 63–70.
- Andeep S Toor, Harry Wechsler, and Michele Nappi. 2019. Question action relevance and editing for visual question answering. *Multimedia Tools and Applications*, 78:2921–2935.

Quan Vuong, Sergey Levine, Homer Rich Walke, Karl Pertsch, Anikait Singh, Ria Doshi, Charles Xu, Jianlan Luo, Liam Tan, Dhruv Shah, and 1 others. 2023. Open x-embodiment: Robotic learning datasets and rt-x models. In *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@CoRL2023*.

Spencer Whitehead, Suzanne Petryk, Vedaad Shakib, Joseph Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. 2022. Reliable visual question answering: Abstain rather than answer incorrectly. In *European Conference on Computer Vision*, pages 148–166. Springer.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR.

## Appendix

The appendix consists of the following further discussion:

- [Appendix A](#) provides information on the code release, including links to the code bases and datasets used in the project.
- [Appendix B](#) details the use of AI in the creation of this manuscript.
- [Appendix C](#) visualizes examples of tasks used for training and evaluation.
- [Appendix D](#) shows the example training dataset of In-Domain false premise, Out-of-Domain false premise, and true premise.
- [Appendix E](#) provides the generation process of infeasible instructions.

## A Code Release

We make the code and data for our analysis available in our [project page](#). We release both under the MIT license.

## B Disclosure of AI Usage

The authors acknowledge the use of artificial intelligence (AI) tools in the preparation of this manuscript. Specifically, Microsoft Copilot, OpenAI ChatGPT, and Google Gemini Pro were utilized for general editing and code generation / completion purposes. All generated code and text was verified for correctness by one or more of the authors.

## C Qualitative examples

We visualize qualitative examples of 9 tasks from RL Bench in [Figure C.2](#).

## D Example Data

### • In-Domain false premise:

**human:** “<image>\You are a Franka robot using joint control. The task is “take the drawer off the grill”, and the previous five (including current) steps are [[0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0.0098, 0.1741, -0.0053, -0.8438, -0.0026, 1.2311, 0.7985]]. Can you predict action of the next 1 step?”

**gpt:** “I don’t see drawer in the current scene. Do you mean chicken?”

**human:** “Yes, You are a Franka robot using the joint control. The task is “take the chicken off the grill”, and the previous five (including current) steps are [[0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0.0098, 0.1741, -0.0053, -0.8438, -0.0026, 1.2311, 0.7985]]. Can you predict action of the next 1 step?”

**gpt:** “2D visual trace: [[61, 51], [61, 51], [61, 51], [61, 51], [62, 51], [60, 11], [59, 13], [57, 16], [57, 21], [55, 28], [55, 35], [56, 37], [57, 39], [59, 46], [60, 51], [58, 46], [56, 40], [57, 37], [57, 35], [61, 37], [65, 37], [72, 35], [78, 34], [80, 34], [78, 37], [76, 45], [79, 51]]. The next action step: [-0.0004, -0.0204, -0.0007, -0.0588, -0.0004, 0.0213, 0.0058, 1.0]”

### • Out-of-Domain false premise:

**human:** “<image>\nYou are a Franka robot using joint control. The task is “take the elephant off the grill”, and the previous five (including current) steps are [[0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [-0.0152, 0.1717, -0.0032, -0.8589, 0.0182, 1.2171, 0.7744]]. Can you predict action of the next 1 step?”

**gpt:** “I couldn’t find an elephant in the current scene.”

### • True premise:

**human:** “<image>\nYou are a Franka robot using joint control. The task is “take the chicken off the grill”, and the previous five (including current) steps are [[0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0], [0.0098, 0.1741, -0.0053, -0.8438, -0.0026, 1.2311, 0.7985]]. Can you predict action of the next 1 step?”



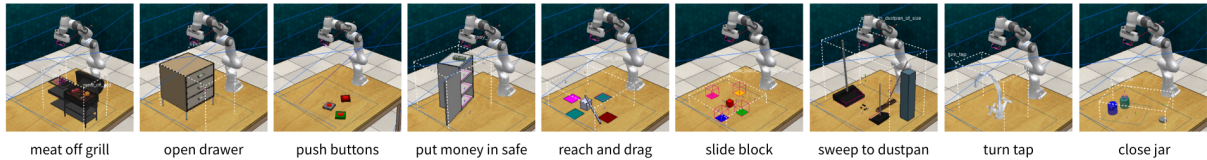


Figure C.2: Qualitative examples of 9 tasks.

0, 0, 0, 0, 0], [0, 0, 0, 0, 0, 0, 0],  
 [0.0115, 0.1585, -0.0003, -0.8588,  
 0.0045, 1.2363, 0.8086]]. Can you  
 predict action of the next 1 step?"

**gpt:** "2D visual trace: [[73, 10],  
 [78, 9], [83, 11], [90, 13], [96,  
 16], [100, 21], [103, 24], [106,  
 29], [107, 32], [109, 35], [110,  
 38], [109, 40], [108, 49], [108,  
 51], [108, 47], [110, 39], [109,  
 38], [100, 37], [87, 34], [71, 32],  
 [59, 32], [57, 32], [58, 42], [57,  
 48]]. The next action step: [0.0173,  
 0.0007, -0.0033, -0.0291, -0.0006,  
 0.0108, -0.056, 1.0]"

## E Infeasible Instruction Generation

To generate "infeasible" instructions, we first curate two distractor pools of nouns:

1. **In-Domain:** objects that appear in our RL Bench scenes (e.g., "blue safe", "drawer", "mug").
2. **Out-of-Domain:** objects never seen in the scenes (e.g., "sofa", "durian", "elephant"), drawn from a list generated by GPT.

We then rewrite 85% of the original true-premise prompts using one of two LLM-based rewriters:

1. *In-Domain FP* (65% overall): replace the target OBJECT by sampling from the In-Domain pool, ensuring the new noun is absent from that episode's scene.
2. *Out-of-Domain FP* (20% overall): replace the target OBJECT with a noun sampled from the Out-of-Domain pool.

Finally, we manually reviewed 200 randomly selected rewritten prompts to verify grammaticality and correctness.