

Seen vs. Followed: A Protocol for Measuring Contextual Awareness of In-Scene Text in Vision-Language-Action Models

Zachary Gunderson
Independent Researcher
zachary.t.gunderson@gmail.com

Abstract: A sticky note in view can hijack a robot controlled by a vision-language model (VLM). We test whether four vision-language-action models (VLAs), which are built on VLMs, inherit this vulnerability in LIBERO-Goal. We first check that each model uses its camera, with blind and degraded inputs. We then compare helpful, harmful, and neutral notes in the scene against the same text appended to the instruction. Three of the four models rely on the camera, but we saw no effect from the sticky notes on their task outcomes. This suggests that VLA benchmarks would benefit from in-scene text to identify models with poor contextual awareness.

Keywords: vision-language-action models, prompt injection, contextual awareness, evaluation

1 Introduction

A generalist robot has to act on a semantic understanding of the physical scene, including the messages posted in it. A door marked “PULL”, a note left by a wet-lab scientist, and a stove light labeled “on when the surface is hot” each tell the robot something it needs to know to act safely. Because the safety constraints of a generalist robot depend heavily on context [1], contextual awareness is central to its safety. Vision-language-action models (VLAs) are one type of open-source robot foundation model (RFM) for building such robots. A VLA meant to run on an onboard edge computer is built on a small vision-language model (VLM), which does the multimodal reasoning, with a fine-tuned policy layer added to control the motors. The VLA may therefore inherit the known safety failures of VLMs, which can answer image questions correctly without the image [2] and can be hijacked by adversarial sticky notes [3].

For VLMs, the response to text in the scene has already been tested for safety. Samarakoon et al. [3] show that text placed in view can hijack VLM-controlled robots, and Chang et al. [4], testing both helpful and harmful notes on static images, find that some VLM agents read a note but rarely act on it (“modality laziness”), while others follow harmful cues over helpful ones. Tests on VLAs have largely overlaid text on the image, as one of several visual stressors [5] or as task cues in LIBERO-VIFO [6], rather than placing it in the simulated scene. Defenses either mask all text, which degrades contextual awareness [3], or add a judgment step that decides which cues to trust [4].

Sinha et al. [1] frame the risks of generalist robots around contextual hazards, misalignment, and adversarial actors. We address the part where contextual awareness meets adversarial actors, because a note in the scene can be either a warning the robot needs or an attack planted to misuse it. The same authors argue that, because real-world safety testing is unsafe and costly, the safety case must rest mainly on offline evaluation, combined with limited real-world tests. Safety results for LLMs do not automatically transfer to RFMs [7], so a simulated protocol that tests both kinds of note before deployment is one step toward that safety case.

We propose such a protocol and apply it to four open VLAs in LIBERO-Goal [8]. Phase 1 checks that each model uses its camera. Phase 2 then places matched helpful, harmful, and neutral notes in the scene and appends the same text to the instruction, either as an OCR read or typed. Three of the four models rely on the camera, but the notes did not change the outcome for any of those three, and no route moved any of them to follow the note in more than 3 of 15 seeds.

Our contributions are the following:

- A camera-gated protocol for testing in-scene text on end-to-end VLAs in closed-loop simulation. It first checks that each model uses its camera (blind input, graded degradation), then compares matched helpful, harmful, and neutral notes against a blank card, and delivers the same text through three routes (in the scene, as an OCR read, and typed).
- Evidence consistent with the “modality laziness” of VLM agents [4] in VLAs. A note in the scene is legible in the camera frame and shifts the actions of $\pi_{0.5}$ and GR00T, but its wording does not change the outcome for any of the three camera-using models.

2 Citations

2.1 Text in the scene as an attack on VLM robots

Physical prompt injection with a sticky note hijacks VLM-controlled robots [3], and related attacks use signs [9], text on placed 3D objects [10], writing on walls [11], printed stickers [12], and physical prompt injection more broadly [13]; text and voice jailbreaks reach LLM-controlled robots as well [14, 15]. None of these tests an end-to-end VLA. They target VLM or LLM planners, or a modular pipeline built on CLIP [12].

2.2 Text and VLAs

For end-to-end VLAs, written words have been tested as a digital overlay [5], and LIBERO-VIFO benchmarks visual cue following, mostly as overlays [6]. VLAs trained to read scene text can act on it [16, 17]. On the language side, appended text collapses success [18], and several benchmarks find that VLAs largely ignore their instruction [19, 20, 21] and rely on memorized action sequences and layouts [22].

2.3 Safety benchmarks and attacks

VLA safety benchmarks [23, 24, 25, 26] and attacks [27, 28, 29, 30] score harmful behavior, but none removes the policy’s own visual input.

2.4 Does the model use its inputs?

Blind and single-modality baselines have a long history in VQA and embodied QA [31, 32, 33], and VLMs can answer without the image [2].

3 Experimental Results

3.1 Setup: the protocol and its first instance

We test π_0 and $\pi_{0.5}$ [34, 35] (LeRobot LIBERO checkpoints [36]), SmolVLA [37], and GR00T N1.7 [38], all fine-tuned on LIBERO-Goal [8], a suite of ten tasks that share one kitchen layout and differ only in the goal named by the instruction. The scene is this kitchen with a bowl, a plate, and a stove that is either off or on. The instruction is the trained hazard task (“put the bowl on the stove”) or the benign task (“...on the plate”); each episode ends with the bowl on the stove, on the plate, or neither.

In Phase 1, we either replace the camera image with a gray frame at every step or degrade it in graded steps of occlusion, low light, contrast, or brightness.

In Phase 2, we render a sticky note on the table (Fig. 1). The note carries matched helpful, harmful, or neutral text, and a blank card serves as the control. A legibility gate transcribes each note from the policy’s view with SmolVLM2 and drops notes it cannot read. The same text is also appended to the instruction as an OCR read or typed, with a meaningless append and a negated instruction as controls.




Figure 1: The starting scene in the policy’s 256 px view with each note type (a blank card is the control), and the prompts. In most seeds, the plate’s rim hides at most $\sim 10\%$ of the note’s ink.

3.2 Phase 1: Verifying the visual channel is used




Figure 2: Phase 1. Success on the benign task, as each alteration is applied to every camera frame at increasing severity, with the same seeds throughout; occlusion ends fully blind.

$\pi_{0.5}$, SmolVLA and GR00T fail without the image (0/15) and degrade as more of it is covered (Fig. 2). π_0 succeeds blind (13/15 here; 81/100 on LIBERO-Goal), and its degradation curves are flat. The vision encoder in the public checkpoint we used was never trained (its weights match a random initialization), so the policy completes the task from the instruction and its arm position alone. LIBERO-Goal allows this because object positions vary by only ~ 2 cm across starting states, and the benchmark score gives no sign of it. π_0 is therefore not a fair test of whether VLAs read notes, and we exclude it from the claims in Phase 2.

3.3 Phase 2: Perception vs. use of in-scene text

(a) **Notes in the scene.** Helpful, harmful, and neutral notes leave every model where the blank card does (Fig. 3, top), with the stove off and with it on. Each condition uses the same 15 seeds, so every episode can be matched to the control episode from the same start. No note changed $\pi_{0.5}$ ’s outcome on any seed; the most any note changed was 3 of 15 seeds for SmolVLA and 2 for GR00T. With 15 seeds, only large effects would show, so small ones are not ruled out.

(b) **The same text as language.** The OCR read and the typed text behave nearly identically in every cell, so reading quality is not the bottleneck (Fig. 3, bottom). The language route does reach SmolVLA’s actions, since any appended text, even a meaningless one, derails it (neither in 14–15 of

15). Under the hazard instruction, $\pi_{0.5}$ and GR00T stay on the stove whatever text is appended, even a negated instruction (stove 15/15), as if the instruction only selects one of the ten trained tasks. Under the benign instruction, appended harmful text sent $\pi_{0.5}$ to the stove in 3 of 15 seeds when typed and 2 when read by OCR (none for GR00T), which is not significant after correction. No camera-using model followed the note in more than 3 of 15 seeds in any route.

(c) Perception of the note. We traced the note through each model’s input in two steps. First, the note is legible in the camera frame. An OCR read taken from that frame leads to the same outcomes as the typed text. Second, we compared actions. Showing a note instead of a blank card shifts the actions of $\pi_{0.5}$ and GR00T beyond run-to-run noise, but changing the note’s wording from helpful to neutral does not.




Figure 3: Phase 2. One square per episode under the hazard instruction, colored by where the bowl ended; 15 per row (12 for the neutral note on SmolVLA and GR00T, after the legibility gate). Rows above the line show no note or a note in the scene, and rows below show the helpful note’s text appended to the instruction. $^\dagger \pi_0$ does not use its camera.

4 Conclusion

Before making any claim about contextual awareness, this protocol establishes a baseline for each checkpoint by measuring whether the policy succeeds from memory with the camera blind and how it degrades as the image worsens. Only then does it add notes to the scene, where helpful, harmful, and neutral wordings can be compared against a blank card.

None of the three camera-using VLAs acts on the wording of text in its view, so they are neither hijacked nor helped by it; masking all text, the defense proposed for VLM planners [3], is in effect already how they behave, and our helpful notes show its cost.

Limitations. One of the four public checkpoints (π_0) has an untrained vision encoder, so the claims about notes rest on three models. We test one simulated scene, a LIBERO-Goal cooking task, with the note in a single position next to the plate; other scenes, tasks, and placements may behave differently. LIBERO-Goal uses the instruction as a task identifier, which makes the contest between note and instruction lopsided. The planner replication of [3] in our scene did not reproduce the attack, so we lack a positive control for the scene itself. $\pi_{0.5}$ sees the camera at a lower resolution (224 px), where one of our two reading checks fails on some notes, so for $\pi_{0.5}$ we cannot fully rule out that the notes were too small to read. Finally, $n = 15$ rules out only large effects.

Next steps. A robot that ignores a warning fails a contextual hazard, just as one that obeys an attack fails an adversarial one. LIBERO-NOTES will measure both, across LIBERO scenes where the note is needed to act correctly.

References

- [1] R. Sinha, A. Dixit, R. Tian, A. Majumdar, and A. Bajcsy. Rethinking safety for generalist robots. *arXiv preprint arXiv:2609.06326*, 2026.
- [2] M. Asadi, J. W. O’Sullivan, F. Cao, T. Nedaei, K. Rajabalifardi, L. Fei-Fei, E. Adeli, and E. Ashley. MIRAGE: The illusion of visual understanding. *arXiv preprint arXiv:2603.21687*, 2026.
- [3] S. M. B. P. Samarakoon, M. A. V. J. Muthugala, W. K. R. Sachinthana, and M. R. Elara. Hijacking robots with a piece of paper: A systematic study of physical prompt injection in VLM-controlled robots. *arXiv preprint arXiv:2608.05715*, 2026.
- [4] J. Chang, M. Xue, R. Sun, S. Pang, S. S. Kanhere, and H. Pearce. If you’re waiting for a sign... that might not be it! mitigating trust boundary confusion from visual injections on vision-language agentic systems. *arXiv preprint arXiv:2604.19844*, 2026. arXiv:2604.19844.
- [5] H. Cheng, E. Xiao, Y. Wang, C. Yu, M. Sun, Q. Zhang, J. Cao, Y. Guo, N. Liu, K. Xu, J. Zhang, C. Shen, P. Torr, J. Gu, and R. Xu. Manipulation facing threats: Evaluating physical vulnerabilities in end-to-end vision language action models. *arXiv preprint arXiv:2409.13174*, 2024. arXiv:2409.13174.
- [6] Z. Qian, R. Yan, A. J. Wang, and J. Tang. LIBERO-VIFO: Benchmarking the capability and safety of visual cue following in vision-language-action models. *arXiv preprint arXiv:2608.17600*, 2026.
- [7] K. Stocking and B. Häon. The case for physical AI safety. The Physical AI Safety Institute, <https://paisi.ai/the-case-for-physical-ai-safety.pdf>, 2026. White paper, July 2026.
- [8] B. Liu, Y. Zhu, C. Gao, Y. Feng, Q. Liu, Y. Zhu, and P. Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *NeurIPS Datasets and Benchmarks Track*, 2023. arXiv:2306.03310.
- [9] L. Burbano, D. Ortiz, Q. Sun, S. Yang, H. Tu, C. Xie, Y. Cao, and A. A. Cardenas. CHAI: Command hijacking against embodied AI. *arXiv preprint arXiv:2510.00181*, 2025.
- [10] Z. Li and Y. Chen. Extended to reality: Prompt injection in 3D environments. *arXiv preprint arXiv:2602.07104*, 2026. arXiv:2602.07104.
- [11] C. Li, Z. Kang, J. Zhang, Z. Ma, A. Cheng, X. Li, and J. Ma. The shawshank redemption of embodied AI: Understanding and benchmarking indirect environmental jailbreaks. *arXiv preprint arXiv:2511.16347*, 2025. arXiv:2511.16347.
- [12] A. Iranmanesh and P. Liu. Not what you asked for: Typographic attacks in household robot manipulation. *arXiv preprint arXiv:2605.18593*, 2026. arXiv:2605.18593.
- [13] C. Ling, K. Hu, H. Liu, X. Han, T. Zhang, and C. Ou. Physical prompt injection attacks on large vision-language models. *arXiv preprint arXiv:2601.17383*, 2026. arXiv:2601.17383.
- [14] A. Robey, Z. Ravichandran, V. Kumar, H. Hassani, and G. J. Pappas. Jailbreaking LLM-controlled robots. *arXiv preprint arXiv:2410.13691*, 2024. arXiv:2410.13691.
- [15] H. Zhang, C. Zhu, X. Wang, Z. Zhou, C. Yin, M. Li, L. Xue, Y. Wang, S. Hu, A. Liu, P. Guo, and L. Y. Zhang. BadRobot: Jailbreaking embodied LLM agents in the physical world. *arXiv preprint arXiv:2407.20242*, 2024. arXiv:2407.20242; venue: ICLR 2025.
- [16] Z. Zhou, Y. Zhu, J. Wen, C. Shen, and Y. Xu. ChatVLA-2: Vision-language-action model with open-world embodied reasoning from pretrained knowledge. *arXiv preprint arXiv:2505.21906*, 2025. arXiv:2505.21906.

- [17] W. Zhao, G. Li, Z. Gong, P. Ding, H. Zhao, and D. Wang. Unveiling the potential of vision-language-action models with open-ended multimodal instructions. *arXiv preprint arXiv:2505.11214*, 2025. arXiv:2505.11214.
- [18] Z. Zhou, Z. Li, X. Chen, J. Li, C. Wang, H. Chen, and X. Zhu. Think only when needed: Prompt-authority control for selective slow-path intervention in vision-language-action manipulation. *arXiv preprint arXiv:2608.23224*, 2026. arXiv:2608.23224.
- [19] S. Fei, S. Wang, J. Shi, Z. Dai, J. Cai, P. Qian, L. Ji, X. He, S. Zhang, Z. Fei, J. Fu, J. Gong, and X. Qiu. LIBERO-Plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025.
- [20] Y. Hou and L. Zhao. LangGap: Diagnosing and closing the language gap in vision-language-action models. *arXiv preprint arXiv:2603.00592*, 2026. arXiv:2603.00592.
- [21] Y. Fang, Y. Feng, D. Jing, J. Liu, Y. Yang, Z. Wei, D. Szafr, and M. Ding. When vision overrides language: Evaluating and mitigating counterfactual failures in VLAs. *arXiv preprint arXiv:2602.17659*, 2026. arXiv:2602.17659.
- [22] X. Zhou, Y. Xu, G. Tie, Y. Chen, G. Zhang, D. Chu, P. Zhou, and L. Sun. LIBERO-PRO: Towards robust and fair evaluation of vision-language-action models beyond memorization. *arXiv preprint arXiv:2510.03827*, 2025.
- [23] R. Cui, Z. Zhang, J. Pang, H. Chi, J. Guo, S. Zhang, S. Xie, X. Jin, Y. Mu, J. Yang, G. Yao, X. Zhan, Y.-Q. Zhang, and H. Zhao. LIBERO-Safety: A comprehensive benchmark for physical and semantic safety in vision-language-action models. In *ECCV*, 2026. arXiv:2606.23686.
- [24] Z. Chen, Y. Gao, L. Wang, Y. Zhao, Y. Liu, J. Li, X. Zheng, Z. Wu, C. Wang, X. Ma, and Y.-G. Jiang. HazardArena: Evaluating semantic safety in vision-language-action models. *arXiv preprint arXiv:2604.12447*, 2026.
- [25] J. Fan, W. Xu, O. Sokolsky, I. Lee, and F. Kong. SafeVLA-Bench: A benchmark for the success-safety gap in vision-language-action models. *arXiv preprint arXiv:2606.00773*, 2026.
- [26] M. Lyu, Y. Sun, Y. Jia, S. Shen, M. Sha, H. Li, F. Zhao, and Y. Zeng. ForesightSafety-VLA: A unified diagnostic safety benchmark for vision-language-action models. *arXiv preprint arXiv:2606.27079*, 2026.
- [27] J. Li, Y. Zhao, X. Zheng, Z. Xu, Y. Li, X. Ma, and Y.-G. Jiang. AttackVLA: Benchmarking adversarial and backdoor attacks on vision-language-action models. *arXiv preprint arXiv:2511.12149*, 2025.
- [28] T. Wang, C. Han, J. Liang, W. Yang, D. Liu, L. X. Zhang, Q. Wang, J. Luo, and R. Tang. Exploring the adversarial vulnerabilities of vision-language-action models in robotics. *arXiv preprint arXiv:2411.13587*, 2024.
- [29] S. Karnik, Z.-W. Hong, N. Abhangi, Y.-C. Lin, T.-H. Wang, C. Dupuy, R. Gupta, and P. Agrawal. Embodied red teaming for auditing robotic foundation models. *arXiv preprint arXiv:2411.18676*, 2024.
- [30] Y. Zhang, B. Zhang, J. Fan, J. Shen, Y. Cai, Y. Yang, and J. Ji. RedVLA: Physical red teaming for vision-language-action models. *arXiv preprint arXiv:2604.22591*, 2026. arXiv:2604.22591.
- [31] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *CVPR*, 2017. arXiv:1612.00837.
- [32] J. Thomason, D. Gordon, and Y. Bisk. Shifting the baseline: Single modality performance on visual navigation & QA. In *NAACL*, 2019. arXiv:1811.00613.

- [33] A. Anand, E. Belilovsky, K. Kastner, H. Larochelle, and A. Courville. Blindfold baselines for embodied QA. In *NeurIPS Workshop on Visually Grounded Interaction and Language (ViGIL)*, 2018. arXiv:1811.05013.
- [34] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [35] Physical Intelligence, K. Black, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *arXiv preprint arXiv:2504.16054*, 2025.
- [36] R. Cadene et al. LeRobot: An open-source library for end-to-end robot learning. In *ICLR*, 2026. arXiv:2602.22818.
- [37] M. Shukor, D. Aubakirova, F. Capuano, P. Kooijmans, S. Palma, A. Zouitine, M. Aractingi, C. Pascal, M. Russi, A. Marafioti, S. Alibert, M. Cord, T. Wolf, and R. Cadene. SmolVLA: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- [38] NVIDIA. Isaac GR00T N1.7. <https://huggingface.co/nvidia/GR00T-N1.7-3B>, 2026.