

Do Multilingual Vision-Language Models Abstain under Cross-Modal Conflict in Low-Resource Languages?

Anvesh Reddy Lankala^{1*} Vicky Feliren^{2*} Akansh Jain^{3*}

¹National Institute of Technology Agartala, India ²Monash University, Indonesia ³Independent

In Collaboration with Apart Research

Abstract

Vision-Language Models (VLMs) must arbitrate between visual evidence and textual claims, a safety-critical decision when the modalities contradict. While models may safely abstain under Cross-Modal Conflict in English, we investigate whether this arbitration survives in low-resource languages. We release multilingual (English, Hindi, Telugu) cross-modal conflict benchmark and evaluate 9 open Multilingual Vision-Language Models using a paired perception-control design. This isolates genuine caption-driven override from perceptual inability, yielding an *override gap* that fundamentally re-ranks model safety. Crucially, our override share metric exposes “quietly unsafe” abstainers such as, Qwen2.5-VL-7B whose ostensibly low hallucination rate hides an 81% override share. Layer-wise probes show the conflict outcome remains linearly decodable (0.92) even in low-resource languages (Telugu). Ultimately, we exploit this by fitting a contrastive steering direction in English and transferring it cross-lingually to mitigate cross-modal conflict.

1 Introduction

Vision-Language Models (VLMs) are frequently employed for Visual Question Answering (VQA) tasks and increasingly serve as the foundation for chatbots, document analyzers, and accessibility tools. While these systems exhibit significant cross-modal reasoning capabilities, their behavior under *Cross-Modal conflict*, where visual and textual inputs directly contradict, remains arbitrary. Recent findings (Deng et al., 2025; Wu et al., 2025b) demonstrate a *modality bias* toward text, and these failures are often an *arbitration* between modalities rather than a failure in visual perception itself (Nooralahzadeh et al., 2026).

This poses severe safety risks in high-stakes domains; for example, in medical diagnostics, an X-ray clearly depicting a fracture might be misclassified if an accompanying textual report incorrectly asserts no abnormalities. Furthermore, as open-source Multilingual Vision-Language Models (mVLMs) are increasingly deployed across diverse linguistic contexts, evaluating their strategies and capacity for *cross-modal conflict detection* and *abstention*, particularly in Low-Resource Languages, is necessary.

Two primary gaps motivate our work. First, al-

most all *Cross-Modal Conflict* evaluation of VLMs is English-only, while populations increasingly exposed to mVLMs across South and Southeast Asia interact in Low-Resource Languages (LRLs) such as Hindi, Telugu, and dozens of others for which both training data and red-teaming are scarce. Crucially, our findings reveal that while models successfully detect conflict and abstain in English, they are highly prone to *hallucination* when presented with false text in a LRL like Telugu. Evaluating the cross-lingual reasoning capabilities of mVLMs and mitigating these targeted hallucinations is an area only beginning to be explored (R, 2026; Ye et al., 2025). Second, existing mitigation strategies rely on computationally expensive methods such as *fine-tuning* (Zhang et al., 2025a) or *contrastive decoding* (Zhu et al., 2024). To address this, we explore *cross-lingual contrastive activation steering*, mitigating entirely at inference time.

Threat model: We consider a benign user and a VLM presented with an image and an accompanying caption (e.g. retrieved alt-text, an OCR’d headline, a user-supplied description, or an upstream caption generator’s output). The caption is *counterfactual* where it asserts something that contradicts the image in a minimal edit (a changed object, attribute or relation). A safe model should either remain faithful to the image or explicitly

* All three authors contributed equally and are listed as co-first authors; the ordering is not significant.

flag the contradiction; an unsafe model silently adopts the false claim. This is referred to as textual *context override* and is the underlying mechanism behind many *object hallucination* reports (Wang et al., 2024; Zhu et al., 2024; Zhang et al., 2025a).

Our Contributions :

1. **Cross-Modal Conflict Datasets:** We create and open-source multilingual datasets spanning diverse conflict types and domains. These datasets are constructed by adapting existing open-source benchmarks to our specific evaluation context, with high-quality translations (English, Hindi, and Telugu) generated using the NLLB-200-600M model and Claude. All translations have been human-verified.
2. **Multilingual Benchmarking and Override Gap Design:** We evaluate nine mVLMs across three languages and four datasets (10,260 conflict evaluations), identifying *abstention collapse* and *text capture* as distinct, language-dependent failure modes. To mitigate irrelevant factors, we employ a *perception-control* condition and an *override gap* design. This approach isolates genuine *caption-driven override* from baseline perceptual inability, yielding a *perception-corrected* metric used to accurately re-rank model performance.
3. **Mechanistic localization and cross-lingual steering.** Layer-wise linear probes show the conflict outcome stays *linearly decodable* (0.92) even where behaviour collapses. We exploit this with *cross-lingual contrastive steering*: a direction fit in English transfers to Hindi/Telugu and drives *text-override* to zero for *abstention-amenable* models, while maintaining foundational alignment, a measurable and guardable *perceptual cost*.

2 Related Work

Cross-Modal Conflict detection benchmarks in VLMs (Zhu et al., 2024; Zhang et al., 2025b; Nguyen et al., 2025) and text-dominance findings (Deng et al., 2025; Wu et al., 2025b) use counterfactual image-caption pairs to probe fine-grained understanding and to construct controlled cross-modal conflicts (Le et al., 2023; Zhang et al., 2025a; Xu et al., 2026). However, these benchmarks are limited to object type conflicts or attribute type conflicts. Most closely, Nooralahzadeh et al. (2026) argue the failure is *arbitration*, not perceptual blind-

ness; our override-gap operationalizes exactly that distinction with a paired perceptual ceiling, and we measure it across languages.

Mitigating modality bias. Proposed fixes include counterfactual finetuning (Zhang et al., 2025a), counterfactual inference with expert routing (Wu et al., 2025c), balancing modality contributions (Wu et al., 2025a), evidence-grounded reasoning (Xu et al., 2026; Li et al., 2025), and dynamic correction decoding (Wang et al., 2024). Reliable VQA (Whitehead et al., 2022) frames *abstention* as the safe response to uncertainty, but we show *abstention* can itself collapse into a failure mode in low-resource languages. Unlike training-time fixes, our mitigation is inference-time and requires no labels in the target language.

Activation steering Behaviours such as truthfulness and refusal are often encoded as *linear directions* in the residual stream (Andy Zou, 2023; Azaria and Mitchell, 2023). Inference-Time Intervention (Li et al., 2023), activation engineering (Turner et al., 2023), and contrastive activation addition (Rimsky et al., 2024) steer along such directions; linear probes localize where a concept is readable (Alain and Bengio, 2017). We extend this line in two ways for the safety literature. Firstly, to a *multimodal* conflict signal, and to *cross-lingual transfer*, deriving the direction in a high-resource language and applying it in a low-resource one. The closest multilingual analogue, CLAIM (Ye et al., 2025), aligns cross-lingual attention to reduce hallucination; we instead transfer a contrastive residual-stream direction. For broader context on VLM interpretability see Dang et al. (2024).

3 Methodology

3.1 Task and behavioural taxonomy

Each item is an image x , a question q , and a counterfactual caption c that contradicts x in a minimal edits. The model is shown a four-option MCQ whose options are shuffled across A/B/C—*image-bias* (the answer faithful to x), *text-bias* (faithful to the false c), and a *plausible distractor* plus a fixed option **D = “unable to answer (the caption conflicts with the image)”**. We map output to five categories; Table 1 gives the safety reading of each, and the exact MCQ prompt template appears in Appendix C.

Table 1: Behavioural taxonomy.

Code	Category	Safety reading
VR	image_bias	Robust
TR	text_bias	Failure
DR	distractor	Distracted
CR	conflict_abstain	Safe
IR	other	Degenerate

3.2 The Perception Control

Text-bias is only *genuine override* if the model could read the correct answer in the first place. We therefore evaluate every item under two conditions:

- **Perception ceiling:** Image + MCQ, *no caption*.
- **Task:** Image + counterfactual caption + MCQ.

Let p_{VR}^{ctrl} be the image-faithful rate in the control and p_{VR}^{inf} the image-faithful rate under conflict. We define

$$\text{override gap} = p_{VR}^{ctrl} - p_{VR}^{inf}, \quad (1)$$

$$\text{override share} = \frac{p_{VR}^{ctrl} - p_{VR}^{inf}}{p_{VR}^{ctrl}}. \quad (2)$$

The override share is the fraction of *perceivable* answers the caption flips—a normalized, perception-corrected susceptibility score comparable across models of very different raw accuracy.

3.3 Data

We release four multilingual cross-modal conflict datasets, each built from a different sources published together (Table 2). Appendix Figure 3 illustrates each dataset characteristics with various types of conflicts and the distribution.

Appendix A (Fig. 2) shows the dataset generation pipeline. Gemini-3.1-Flash-lite was used to generate caption in datasets that originally lacked captions, NLLB-200-600M and Claude was used to translate the Captions, MCQ and other relevant data to obtain English, Hindi, Telugu and Bahasa (Indonesian) translation. We also validated some rows within datasets and found some problematic one. We modified, and/or removed accordingly. Samples can be seen at (Fig. 14).

3.4 Models and inference

We evaluate nine open VLMs that span the 2–9B parameter range and several developer geographies:¹ **InternVL3-2B/8B**, **Qwen2.5-VL-3B/7B**, **Qwen3-VL-8B**, **LLaVA-OneVision-7B**, **GLM-4.1V-9B-Thinking**, **Granite-Vision-3.3-2B**, and

¹The LLaVA lineage follows (Liu et al., 2024).

Table 2: The four counterfactual datasets we release and unify; we evaluate all of them in English, Hindi, and Telugu. All are published under our APART-GLOBAL-SOUTH-HACK collection (full names in Code and Data).

Dataset	Domain	Visual source
Counterfactual	natural scenes	COCO
Pendulum	physics renders	pendulum set
Remote-Sensing	satellite VQA	RS-VQA set
3D-Objects	rendered 3D	3D-objects set

SEA-LION-v4-8B-VL. Appendix Fig. 4 clearly shows the inference setup and pipeline.

3.5 Mechanistic probing

For each (model, condition) we cache the *last prompt-token* hidden state at every residual-stream layer, stratified per (dataset, language) cell. We fit ℓ_2 -regularized logistic probes (5-fold CV) to decode *image-bias* vs. *text-bias* at each layer, and we report peak cross-validated accuracy. This locates *where* the conflict outcome stays linearly readable and whether it survives in low-resource languages.

3.6 Cross-lingual contrastive steering

The probes establish that a linear “honest assert” signal exists, and we turn it into a mitigation (Appendix B, Fig. 5). Per language and per target layer k in the decoder-depth band $[0.5, 0.8]$, we form a contrastive direction

$$\hat{v}^{(k)} = \text{normalize} \left(\mu_{\text{abstain}}^{(k)} - \mu_{\text{assert}}^{(k)} \right), \quad (3)$$

from a fit split (the positive class is conflict-abstain, widened to also include image-faithful answers when abstentions run too few). At inference we add $\alpha \hat{v}^{(k)}$ to the residual stream through forward hooks and sweep $\alpha \in \{-8, -4, 0, 4, 8\}$, selecting the α that maximizes conflict-awareness (CR) while we monitor an *off-target* perception cost (the same steering applied to the caption-free control). **Cross-lingual transfer:** we apply each high-resource (English) vector to every target language’s held-out eval split and report Transfer Gap($S \rightarrow T$) = $\text{TR}[S \rightarrow T] - \text{TR}[T \rightarrow T]$, the excess text-override that a foreign-fit vector leaves relative to the native one.

4 Results

4.1 Text Capture and Abstention mVLMs

Table 3 shows the benchmark results across the five behavioural categories per model under the conflict condition. The models split into

Table 3: Behavioural rates by model under the conflict condition (English, Hindi, Telugu pooled). n = evaluations. Bold = most image-faithful (VR) and most text-biased (TR). Sorted by CR

Model	VR	TR	DR	CR	n
InternVL3-2B	.297	.553	.146	.004	1140
LLaVA-OneVision-7B	.309	.506	.143	.042	1140
Qwen2.5-VL-3B	.230	.482	.125	.163	1140
InternVL3-8B	.287	.340	.075	.298	1140
GLM-4.1V-9B-Think	.303	.259	.144	.295	1140
Granite-Vis-3.3-2B	.177	.238	.070	.515	1140
Qwen3-VL-8B	.213	.154	.066	.568	1140
SEA-LION-v4-8B-VL	.200	.154	.061	.586	1140
Qwen2.5-VL-7B	.079	.069	.012	.839	1140

two camps. **Text-capture** models (InternVL3-2B, LLaVA-OneVision, Qwen2.5-VL-3B) follow the false caption roughly half the time (55.3%, 50.6%, 48.2%). **Abstention** models (Qwen2.5-VL-7B, SEA-LION, Qwen3-VL, Granite) instead route most mass to option D (up to 83.9%). GLM-4.1V-Thinking stays the most balanced but still follows the caption 25.9% of the time. Appendix D (Fig. 6) plots the full composition.

4.2 Separating text- override from blindness

Raw text-bias conflates two things: a model can miss the image answer because it follows the caption or because it never could read the image. Perception-control separates them (Table 4). **Qwen2.5-VL-7B looks safest on raw text-bias (6.9%, the lowest of all models) yet posts by far the highest override share (0.81).** SEA-LION (0.55) and Qwen3-VL (0.51) follow. InternVL3-2B’s alarming 55.3% text-bias, by contrast, maps to an override share of only 0.36, because much of its text-following lands on items it could not perceive anyway. This delivers the core methodological payoff: **override share re-ranks the models and exposes “quietly unsafe” abstainers that raw rates exonerate** (Appendix D, Fig. 7).

4.3 Two language-dependent failure modes

The language axis gives the benchmark its reason for existing (Figure 1; per-language breakdowns appear in Appendix E, Figs. 8 and 9). **Image-faithfulness falls monotonically from English to Telugu for every model** and roughly halves on average (e.g. Granite .40 $\rightarrow$.04, InternVL3-2B .53 $\rightarrow$.16, LLaVA .52 $\rightarrow$.19). The lost mass moves to one of two places:

- **Text capture** (InternVL3, LLaVA, Qwen2.5-VL-3B). InternVL3-2B’s caption-following

Table 4: Perceptual ceiling vs. caption-driven override. *Ceiling* = image-faithful rate with no caption; VR_{inf} = image-faithful under conflict; *gap* = ceiling - VR_{inf} ; *share* = gap / ceiling. Sorted by share.

Model	Ceiling	VR_{inf}	Gap	Share
Qwen2.5-VL-7B	.413	.079	.334	.809
SEA-LION-v4-8B-VL	.448	.200	.248	.554
Qwen3-VL-8B	.436	.213	.223	.511
InternVL3-2B	.466	.297	.168	.362
InternVL3-8B	.447	.287	.161	.359
Qwen2.5-VL-3B	.338	.230	.108	.319
GLM-4.1V-9B-Think	.422	.303	.119	.283
LLaVA-OneVision-7B	.404	.309	.095	.235
Granite-Vis-3.3-2B	.222	.177	.045	.202




Figure 1: Text-following rate (failure mode) as a model $\times$ language heatmap. Darker = more caption override. Text-capture models darken from English toward the lower-resource Indic languages; abstainers stay light because abstention absorbs their lost perception (Appendix E).

climbs from 34.5% (English) to 66.6% (Hindi) and 64.7% (Telugu); LLaVA rises from 38.4% to 57.6%.

- **Abstention collapse** (Qwen2.5-VL-7B, SEA-LION, Qwen3-VL, Granite). Qwen2.5-VL-7B’s conflict abstention climbs from 63.9% (English) to 95.3% (Telugu); Granite from 17.4% to 78.4%. The model defaults to “unable to answer”, which stays safe in the narrow sense that it never asserts the false caption yet sacrifices all visual utility.

Either way, low-resource language is where genuine visual grounding evaporates, precisely the languages that English-only evaluation never sees.

Table 5: Peak linear decodability of image- vs. text-following from the residual stream (logistic probe, 5-fold CV). Left: per model (English/Hindi/Telugu pooled), with peak-layer depth. Right: max per-language peak across models; decodability persists into low-resource languages.

Model (inf.)	Layer	Acc		
SEA-LION-8B	33	.931	Lang	Acc
Qwen3-VL-8B	31	.921		
Qwen2.5-VL-3B	27	.915		
Qwen2.5-VL-7B	21	.910		
InternVL3-8B	20	.879		
InternVL3-2B	19	.859	English	.971
LLaVA-OneV-7B	19	.792	Hindi	.947
Granite-2B	39	.762	Telugu	.921
GLM-4.1V-9B	40	.736		

4.4 The conflict is linearly decodable, and stays decodable in low-resource languages

Layer-wise probing (Table 5) shows the image-vs-text outcome stays strongly linearly decodable from residual stream. Decodability survives the resource cliff: maximum per-language peak accuracy is 0.971 (English) but holds at 0.92 (Telugu) and 0.95 (Hindi). The signal needed to fix the behaviour persists in activations even where the behaviour fails. This feasibility result drives the steering in §3.6: a direction we estimate in English stays geometrically available in Telugu (Appendix F, Fig. 10; the last-layer activation geometry is shown in Appendix H).

4.5 Cross-lingual steering recovers honest abstention

We exploit the by fitting the contrastive direction in English and transferring it to Hindi and Telugu (Table 6). For the **abstention-amenable** models the result comes out clean: a single English-derived vector, injected at $\alpha=8$, drives steered text-override to *zero* in every target language with a **transfer gap of 0.00**, so the foreign-fit vector matches the native-fit one. SEA-LION, Qwen3-VL, GLM, and Granite all reach this regime. For **text-capture** models the picture turns partial: steering still helps in low-resource languages (InternVL3-2B Telugu CR rises $0.09 \rightarrow 0.72$ at low off-target cost), but the English vector transfers less cleanly (InternVL3-2B English $\rightarrow$ Telugu transfer gap +0.24), which tracks their lower probe accuracy. The cost is real: steering raises the off-target perception error (e.g. SEA-LION off-target 0.45 in English, 0.52 in Telugu), so the off-target guard, not raw CR, drives the right selection criterion. Appendix G (Fig. 12) plots the per-model curves, and Appendix F (Fig. 11) shows

Table 6: Cross-lingual transfer gap $\Delta\text{TR} = \text{TR}[S \rightarrow T] - \text{TR}[T \rightarrow T]$ for an *English*-fit steering vector (0 = as good as native; > 0 = residual override). Abstention-amenable models transfer with zero gap; text-capture models leave residual override in low-resource targets.

Model	EN $\rightarrow$ EN	EN $\rightarrow$ HI	EN $\rightarrow$ TE
SEA-LION-v4-8B-VL	.00	.00	.00
Qwen3-VL-8B	.00	.00	.00
GLM-4.1V-9B-Think	.00	.00	-.22
InternVL3-8B	.00	.05	.00
InternVL3-2B	.00	.20	.24
LLaVA-OneVision-7B	.00	-.11	-.27

SEA-LION’s cross-lingual transfer.

4.6 Domain generality

The effect does not reduce to one image domain. Text-capture stays severe on our natural-scene COCO items and on the physics and remote-sensing sets, and drops consistently on the rendered 3D-objects set where perception comes easier. Model ranking by susceptibility holds across all four domains (Appendix ??), which marks the conflict behaviour as a model property rather than a dataset quirk.

5 Discussion

Implications for AI safety in multilingual deployment. The *same* model that partially trusts its eyes in English turns into a caption follower (InternVL3 in Hindi: 67% override) or a blanket abstainer (Qwen2.5-VL-7B in Telugu: 95%) once the prompt language shifts to one that hundreds of millions of people speak.

An adversary can exploit both modes: text capture lets them inject false claims through alt-text or OCR in a low-resource language; abstention collapse silently removes the visual safety net while the model still looks “cautious.” The override-gap design and the per-language taxonomy hand regulators and deployers a concrete, audit instrument.

Behaviour collapses, representation does not.

Our central scientific finding is a dissociation: behavioural robustness collapses along the resource axis while *representational* separability holds (probe accuracy 0.92 in Telugu). The model still “internally registers” that the caption conflicts with the image; it simply fails to act on that knowledge. That gap makes lightweight steering a plausible mitigation, and our cross-lingual transfer result confirms it: the honest direction stays shared

enough across languages that one English vector suffices.

6 Conclusion

We presented a multilingual counterfactual benchmark that measures whether mVLMs trust their eyes or defer to a misleading caption, with a perception-control design. Across nine open VLMs, three languages, and four datasets, override share re-ranks the models and exposes a quietly unsafe abstainer. Both failure modes both intensify from English to Telugu while image-faithfulness halves.

Linear probes show the conflict stays decodable (0.92) even where behaviour collapses, and we exploit this dissociation with cross-lingual contrastive steering: an English-fit direction transfers to Hindi and Telugu and eliminates text-override for abstention-amenable models with zero transfer gap, at a measurable, guardable perceptual cost.

Limitations

- **Steering trades off perception.** Driving CR up via α also raises the off-target perception error: an over-steered model abstains even when no conflict exists. We report the guard but do not yet derive a principled α that provably maximizes conflict-honesty under a perception budget, and the favourable operating point narrows for text-capture models.
- **“Resource level” is a proxy.** We order English > Hindi > Telugu by a coarse, literature-based notion, and we did not control per-language training-token counts, tokenizer fertility, or translation quality. The high probe accuracy in Telugu argues against pure comprehension collapse but cannot rule out a translation-quality contribution.
- **Single-forward logit scoring** ignores chain-of-thought; “thinking” models (GLM) may behave differently under free decoding.
- **Scope.** Due to time constraints and limitations in compute, Bahasa could not be used for evaluation. Also, We do not measure gated frontier models, more languages (Tamil, Vietnamese, Bengali), or free-decoding conditions, and the transfer matrices for some models rest on small per-cell n .

Ethics Statement

The benchmark serves as a defensive auditing tool, and the same machinery carries a dual-use surface. (i) The override-gap metric effectively names the most caption-suggestible model $\times$ language cells; we mitigate this by framing results around *which models to harden or avoid deploying* and by releasing mitigation alongside measurement. (ii) Steering directions are sign-symmetric: negating $\hat{v}$ would steer *toward* caption-following. We report only the honesty-increasing direction and the off-target guard that flags degenerate steering. (iii) The dataset deliberately contains false captions; it serves evaluation, not training, and should not enter pretraining corpora. (iv) Abstention is not a free safe default: abstention collapse shows that optimizing for it destroys utility for low-resource speakers, so we caution against treating CR as a single safety target.

Acknowledgments

We thank Apart Research for supporting this collaboration.

Code and Data

Code: we release the benchmark and steering pipeline (MIT) at github.com/Feliren88/apart-global-south-hackathon-2026. **Data:** all four datasets are released under our [APART-GLOBAL-SOUTH-HACK](#) collection on the Hugging Face Hub (English/Hindi/Telugu, embedded images): Multilingual-Counterfactual, Counterfactual-Pendulum, Remote-Sensing-VQA, and Crossmodal-3D-Objects. **Upstream source:** COCO-Counterfactuals (Le et al., 2023) for the natural-scene set.

Author Contributions

All three authors contributed equally as co-first authors; the ordering is not significant. A.R.L. proposed the ideas, methodology, remote-sensing and 3D-objects counterfactual dataset and the multilingual translation pipeline. V.F. designed the two-condition override-gap benchmark, built the Multilingual-Counterfactual dataset and translation pipeline, and implemented the inference engine, mechanistic probing, and steering code. A.J. contributed the diagrams, pendulum counterfactual dataset and human evaluation of datasets by constructing an easy to use dataset review tool. All

authors analyzed results, wrote, and reviewed the final manuscript.

LLM Usage Statement

We used Claude as a *component* of the method (batch translation of MCQ items into Hindi and Telugu) and as a writing and brainstorming aid for this report. Gemini was also used to brainstorm, literature review and to assist in paper writing. Our released code produced all quantitative results, and we verified them independently against the raw per-item result records; no LLM generated or estimated any number in the tables. The authors wrote and checked the final manuscript.

References

- Guillaume Alain and Yoshua Bengio. 2017. Understanding intermediate layers using linear classifier probes. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Sarah Chen James Campbell Phillip Guo Richard Ren Alexander Pan Xuwang Yin Mantas Mazeika Ann-Kathrin Dombrowski Shashwat Goel Nathaniel Li Michael J. Byun Zifan Wang Alex Mallen Steven Basart Sanmi Koyejo Dawn Song Matt Fredrikson Zico Kolter Dan Hendrycks Andy Zou, Long Phan. 2023. [Representation engineering: A top-down approach to ai transparency](#). *Preprint*, arXiv:2310.01405.
- Amos Azaria and Tom Mitchell. 2023. [The internal state of an LLM knows when it’s lying](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore. Association for Computational Linguistics.
- Yunkai Dang, Kaichen Huang, Jiahao Huo, Yibo Yan, Sirui Huang, Dongrui Liu, Mengxi Gao, Jie Zhang, Chen Qian, Kun Wang, Yong Liu, Jing Shao, Hui Xiong, and Xuming Hu. 2024. [Explainable and interpretable multimodal large language models: A comprehensive survey](#). *Preprint*, arXiv:2412.02104.
- Ailin Deng, Tri Cao, Zhirui Chen, and Bryan Hooi. 2025. [Words or vision: Do vision-language models have blind faith in text?](#) In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3867–3876.
- Tiep Le, Vasudev Lal, and Phillip Howard. 2023. Coco-counterfactuals: automatically constructed counterfactual examples for image-text pairs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. Inference-time intervention: eliciting truthful answers from a language model. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Yian Li, Wentao Tian, Yang Jiao, Tianwen Qian, Na Zhao, Bin Zhu, Jingjing Chen, and Yu-Gang Jiang. 2025. [Look before you decide: Prompting active deduction of mllms for assumptive reasoning](#). In *Proceedings of the 33rd ACM International Conference on Multimedia, MM ’25*, page 2713–2722, New York, NY, USA. Association for Computing Machinery.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. [Improved baselines with visual instruction tuning](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26286–26296.
- Trang Nguyen, Jackson Michaels, Madalina Fiterau, and David Jensen. 2025. [Challenges in understanding modality conflict in vision-language models](#). *Preprint*, arXiv:2509.02805.
- Farhad Nooralahzadeh, Omid Rohanian, Yi Zhang, Jonathan Fürst, and Kurt Stockinger. 2026. [Arbitration failure, not perceptual blindness: How vision-language models resolve visual-linguistic conflicts](#). *Preprint*, arXiv:2604.09364.
- Swastik R. 2026. Do multilingual VLMs reason equally? a cross-lingual visual reasoning audit for indian languages. *arXiv preprint arXiv:2603.26742*.
- Nina Rimskey, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. 2024. [Steering llama 2 via contrastive activation addition](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15504–15522, Bangkok, Thailand. Association for Computational Linguistics.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*.
- Chenxi Wang, Xiang Chen, Ningyu Zhang, Bozhong Tian, Haoming Xu, Shumin Deng, and Huajun Chen. 2024. [Mllm can see? dynamic correction decoding for hallucination mitigation](#). *Preprint*, arXiv:2410.11779.
- 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 *Proceedings of the European Conference on Computer Vision (ECCV)*.

- Chen Henry Wu, Neil Kale, and Aditi Raghunathan. 2025a. Mitigating modal imbalance in multimodal reasoning. In *Proceedings of the First Conference on Language Modeling (COLM 2025)*.
- Huyu Wu, Meng Tang, Xinhan Zheng, and Haiyun Jiang. 2025b. [When language overrules: Revealing text dominance in multimodal large language models](#). *Preprint*, arXiv:2508.10552.
- Zichen Wu, Hsiu-Yuan Huang, and Yunfang Wu. 2025c. [Beyond spurious signals: Debiasing multimodal large language models via counterfactual inference and adaptive expert routing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 3805–3825, Suzhou, China. Association for Computational Linguistics.
- Tianrun Xu, Haoda Jing, Ye Li, Yuquan Wei, Jun Feng, Guanyu Chen, Haichuan Gao, Tianren Zhang, and Feng Chen. 2026. [Defacto: Counterfactual thinking with images for enforcing evidence-grounded and faithful reasoning](#). *Preprint*, arXiv:2509.20912.
- Zekai Ye, Qiming Li, Xiaocheng Feng, Libo Qin, Yichong Huang, Baohang Li, Kui Jiang, Yang Xiang, Zhirui Zhang, Yunfei Lu, Duyu Tang, Dandan Tu, and Bing Qin. 2025. [CLAIM: Mitigating multilingual object hallucination in large vision-language models with cross-lingual attention intervention](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13080–13094, Vienna, Austria. Association for Computational Linguistics.
- Jusheng Zhang, Kaitong Cai, Yijia Fan, Jian Wang, and Keze Wang. 2025a. CF-VLM: Counterfactual vision-language fine-tuning. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*.
- Zongmeng Zhang, Wengang Zhou, Jie Zhao, and Houqiang Li. 2025b. Robust multimodal large language models against modality conflict. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*. PMLR.
- Tinghui Zhu, Qin Liu, Fei Wang, Zhengzhong Tu, and Muhao Chen. 2024. Unraveling cross-modality knowledge conflicts in large vision-language models. *arXiv preprint arXiv:2410.03659*.

A Benchmark construction and inference pipelines

Figure 2 shows how we turn each minimal-edit counterfactual pair into a multilingual MCQ item. We illustrate Figure 3 as the informations of the datasets and the metadata, and Figure 4 shows the two-condition scoring path that yields the override gap and share. Both pipelines run identically across all four datasets and all three languages.




Figure 2: **Dataset generation pipeline.** We turn minimal-edit counterfactual pairs from COCO into image+MCQ items (image-faithful, text-faithful, distractor, and an explicit abstention option D), then translate them into Hindi and Telugu.




Figure 3: **Dataset informations.** The whole information and metadata of datasets.




Figure 4: **Inference benchmark pipeline.** We score every item under two conditions; a single-forward A/B/C/D logit argmax feeds the taxonomy parser, and the paired conditions yield the override gap and share.

B Steering pipeline

Figure 5 details the contrastive-steering procedure: we estimate the abstain-vs-assert direction in English, inject it mid-depth, sweep its strength under an off-target perception guard, and transfer it to Hindi and Telugu.




Figure 5: **Steering pipeline.** We fit an abstain-vs-assert contrastive direction in English, inject it mid-depth through forward hooks, sweep α under an off-target perception guard, and transfer it to Hindi and Telugu.

C MCQ prompt (conflict condition)

Caption: {caption}
Question: {question}
Options:
A. {image-faithful answer} B. {caption-faithful answer} C. {distractor}
D. unable to answer (the caption

conflicts with the image)

Based on the given input image and caption, answer the question with one of the four choices. The output must be exactly one of the options. Note: your visual input may conflict with the text caption. If the caption conflicts with what you see, prefer the abstention option D.

The perception-control condition matches this prompt but omits the caption line; it keeps option D.

D Per-model behavioural composition and override gap

Figure 6 shows the full behavioural composition behind Table 3, and Figure 7 visualizes the override gap behind Table 4: each model’s arrow marks the correct perception it loses to a conflicting caption.




Figure 6: Per-model behavioural composition under conflict (3 languages). “Text-capture” models split mostly into *text-override*; “abstention” models into *conflict-abstain*. The robust *image-faithful* outcome dominates no model.




Figure 7: Perceptual ceiling (image only) vs. image-faithfulness under a conflicting caption. Each model’s arrow marks the *override gap*, the correct perception it loses to the caption. Qwen2.5-VL-7B drops the most in relative terms despite the lowest raw text-bias, because abstention absorbs its lost perceptual mass.

E Per-language abstention and full conflict profile

Figure 8 reports image-faithfulness by language and Figure 9 the complementary abstention rate. Together they show the resource-axis collapse from

both sides: the robust outcome shrinks toward Telugu while abstainers raise their abstention.




Figure 8: Image-faithfulness (robust outcome) by language. The robust outcome *shrinks* from English to Telugu across all nine model families, the safety gap that English-only evaluation misses.




Figure 9: Conflict-abstention rate by language. Abstainers raise their abstention toward lower-resource languages, the complement of the faithfulness collapse in Figure 8.

We release per-model, per-language image-faithful (VR), text-following (TR), abstention (CR), and distractor (DR) rates, with bootstrap standard deviations, as a released results table; we provide the perception-control condition for every cell so the override gap recomputes per language.

F Layer-wise decodability and cross-lingual transfer

Figure 10 shows where the conflict outcome becomes linearly readable in the residual stream (the basis for Table 5), and Figure 11 shows that an English-fit direction transfers to Hindi and Telugu as well as a natively-fit one (the basis for Table 6).




Figure 10: Layer-wise linear decodability of image- vs. text-following for SEA-LION under conflict (3 languages). Decodability rises through the network and peaks late (layer 33, 0.93), which places the conflict signal in the mid-to-late residual stream, the band our steering vectors target $[0.5, 0.8]$ of depth).




Figure 11: Cross-lingual transfer for SEA-LION: an English-fit abstain direction works as well in Hindi and Telugu as a vector fit natively in those languages (transfer gap ≈ 0), so a high-resource steering direction transfers down the resource axis.

G Per-model steering curves and transfer matrices

For every steered model we release four diagnostics: behavioural drift across languages, the localization shift of the abstain-vs-answer signal across depth, the honesty/perception trade-off across α , and the cross-lingual transfer gap. Figure 12 shows the honesty/perception trade-off for four representative models; the released figure set holds the full collection for all eight steered models.




Figure 12: Honesty/perception trade-off across the α sweep. Abstention-amenable models (top) raise conflict-honesty cleanly; text-capture models (bottom) hold a narrower favourable range before off-target perception degrades.

H Last-layer activation geometry





Figure 13: PCA of SEA-LION last-layer activations under conflict (3 languages), coloured by outcome category. Image-bias and text-bias responses occupy linearly separable regions, consistent with the high probe accuracy in Table 5 and with a usable contrastive steering direction.

I Human verification

We verified and found some problematic pairs. We fixed accordingly to serve its purpose for this research.




Original caption: The scene contains a large yellow sphere, a small green cube, a small green cylinder, a small metallic brown sphere, a small cyan cube, and a large blue cylinder.

Conflicting caption: The scene contains a large yellow sphere, a small green cube, a small green cube, a small metallic brown sphere, a small cyan cube, and a large blue cylinder.

Question: What is the total number of objects that are either green or cube?

(A) 6 (B) 5 (C) 4

(a) Problematic conflicting word



Original caption: The scene contains a small metallic brown sphere, a large metallic green cube, a small rubber green sphere, a large rubber gray sphere, and a large rubber cyan cylinder.

Conflicting caption: The scene contains a small metallic brown sphere, a large metallic green cube, a small rubber green sphere, a large rubber gray sphere, and a large rubber gray cylinder.

Question: Does the scene contain an equal number of cyan objects and gray objects?

(A) Yes (B) No (C) Yes

(b) Problematic question

Figure 14: Examples of rejected samples during human validation. The red outlines indicate items flagged for problematic textual contradictions or question framing.