BraintankAI

DISPATCH · MULTIMODAL REFUSAL STEERING · 3 MIN

Refusal signals taken from text help image and video AI models turn down harmful inputs

A preprint steers multimodal models towards refusing unsafe images and videos without any unsafe image or video data, though some models then refuse harmless inputs or lose accuracy.

12 November 2026 · arXiv preprint, not yet peer reviewed

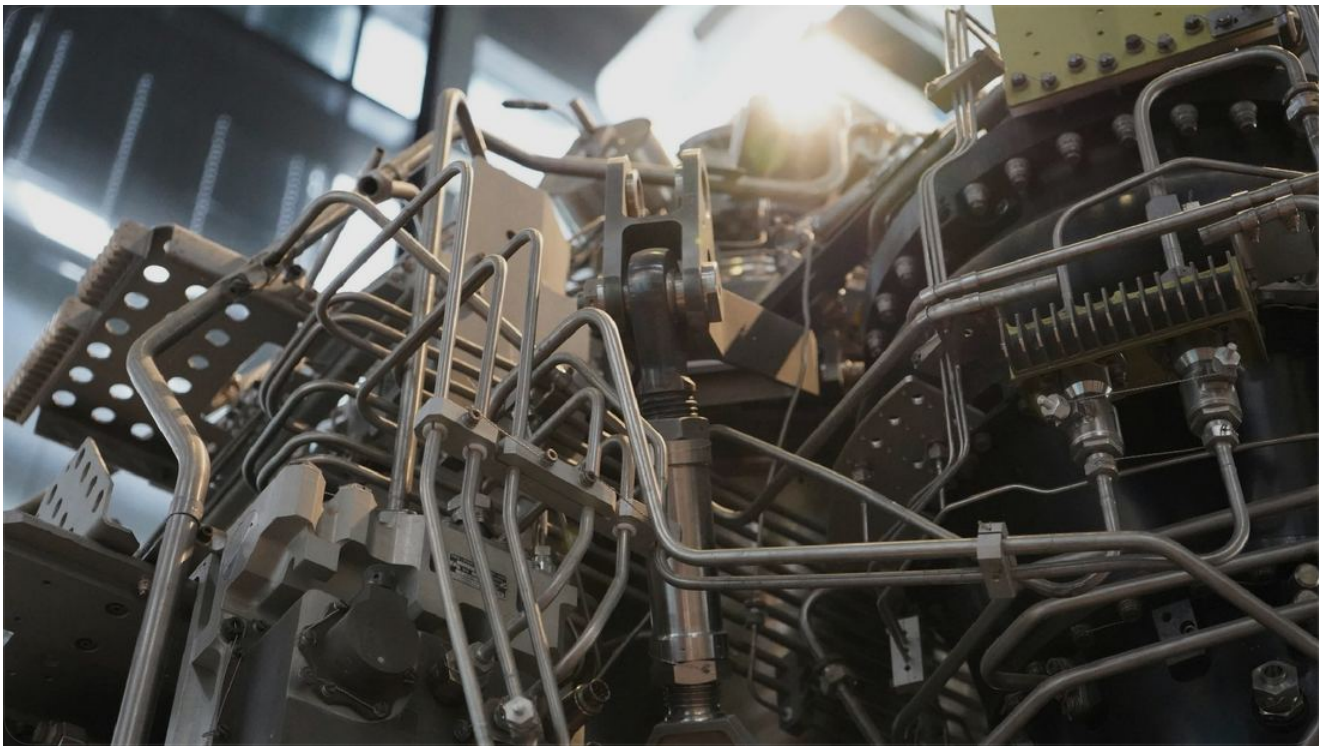


Photo: Meizhi Lang

Multimodal large language models, which are language models that also accept images or video through a vision encoder, mostly inherit their safety from the underlying language model, and methods for making them safer usually depend on examples of unsafe multimodal content. The authors of a 2026 preprint, which has not yet been peer reviewed, point out that such material is hard to collect and has to be gathered again for each new modality. Collecting and sharing it can also raise ethical and legal problems, and they cite nonconsensual and child-abuse imagery as examples. Moreno D'Incà, Nicu Sebe and Massimiliano Mancini test whether safety signals found in text alone can do the job instead.

Their starting point is a refusal direction. This is a single direction in the model's internal activation space. It is computed by averaging the model's hidden states on text prompts it refuses, drawn from a set called MaliciousInstruct, and subtracting the average on prompts it answers,

drawn from Alpaca. Earlier work showed that deleting this direction makes a model easier to jailbreak. Here the authors push along it to strengthen refusal, a form of activation steering: adding a vector to a model's internal states while it runs, with no retraining.

In preliminary tests, plain steering with the text direction already made models refuse unsafe images more often than safe ones at most layers. Some layers refused almost everything, though. On Gemma3, steering at layer 10 rejected 96.4% of safe inputs. The authors trace this to modality misalignment: image activations lean along the refusal direction simply because they come from images, so harmless pictures look suspicious.

The method, called MARS, corrects for this in several steps. First, it subtracts the average activation produced by randomly coloured 'neutral' images, which removes the generic image signal. Second, it measures how far the corrected activation points towards refusal. A gate leaves any input that leans towards the safe side untouched. Third, it sets the push strength automatically. The step is never longer than the distance to the nearer of two cluster centres, one for accepted and one for refused text, a cap the authors call a trust region. The push is applied at a single layer and only while the first output token is produced. That layer is chosen with three scores computed purely from text.

The main tests cover five open models: Qwen3-VL 8B, Qwen3.5 9B, InternVL 3.5 8B, Gemma3 4B and Molmo2 8B. On ViSU, a benchmark of paired safe and unsafe images, the authors report that refusal of unsafe images rose from 0.2% to 68.1% on Molmo2, from 14.2% to 78.5% on Qwen3-VL and from 18.2% to 43.3% on Gemma3. Refusal of safe images also rose, reaching 10.3% on InternVL 3.5 and 7.3% on Molmo2. On VideoSafetyBench, a set of video jailbreaks, Qwen3-VL's refusal rate climbed from 28.8% to 88.2%, while InternVL 3.5 moved only from 4.7% to 11.9%.

MARS was compared with two training-free methods, ECSO and AdaSteer, and with two trained ones, AutoSteer and SASA. The authors treat SASA as an upper bound. On ViSU, MARS out-refused AutoSteer on every model. The trained methods still won in places. On the HADES jailbreak benchmark, Gemma3 with MARS refused 59.2% of attacks, against 92.1% for AutoSteer and 100.0% for SASA. Scores on the MMMU reasoning benchmark were essentially unchanged on four models. On Qwen3.5, however, they fell from 50.8 to 38.2, and on the harder MMMUPro from 35.6 to 24.5. The paper describes this as modest over-refusal.

The authors acknowledge that MARS assumes the language backbone contains a usable refusal direction. They also note that subtracting a single average may not remove the image signal when activations behave in strongly non-linear ways. In the main results, refusal rates on safe inputs are reported only for ViSU's images, and no safe videos are tested. Refusal was mostly scored by matching phrases such as 'I am sorry', which can miss softer refusals. The main results report no run-to-run variance. The five main models range from 4 to 9 billion parameters, and the method requires direct access to their internal activations.

The result echoes a braintankAI report that steering vectors found in text can carry over to image generation in models that share parameters across writing and drawing (earlier story). Here the same idea of transfer is applied to safety in models that read images. Another braintankAI story

found that some safety benchmark scores largely reflect agreement with a labelling policy (earlier story), a reason to read refusal rates like these as tied to each benchmark's own definition of unsafe.

PAPER	*Harnessing Textual Refusal Directions for Multimodal Safety*
AUTHORS	Moreno D'Incà, Nicu Sebe, Massimiliano Mancini
VENUE	arXiv preprint — not yet peer reviewed · 16 October 2026
LINK	arxiv.org/abs/2606.31876
INDEX	stacks.braintank.ai
CODE	github.com/Moreno98/MARS