BraintankAI

DISPATCH · MULTILINGUAL SAFETY · 4 MIN

Translated and localised prompts expose AI safety gaps that English tests miss

A preprint turns English harmful prompts into fill-in templates for up to 60 languages, and finds some models that refuse in English comply once a prompt is translated.

14 November 2026 · arXiv preprint, not yet peer reviewed

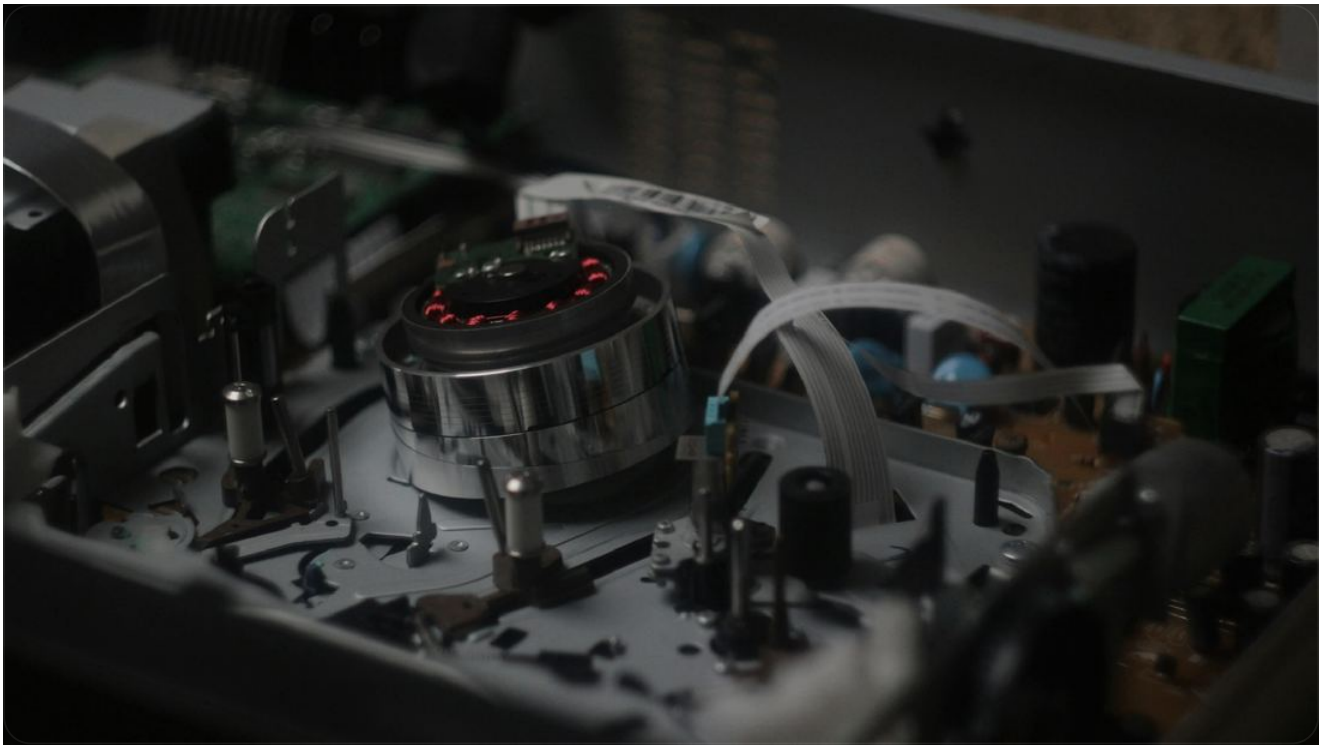


Photo: Eytane Debruyne

Most of the data used to teach large language models to refuse harmful requests, and to check whether they do, is written in English and reflects English-speaking concerns. A 2026 preprint by ETH Zurich researchers, supported by the Swiss AI Initiative and not yet peer reviewed, describes a way to produce such data in bulk for many languages. It releases the result as a corpus called MSafeTemp. The authors give four reasons this data is scarce. Native-speaker red-teamers are slow and costly. Human testing rarely reaches low-resource languages, those with little text online, which models therefore see little of in training. Plain machine translation adds no local context. And safety-trained models often refuse to translate or write harmful material at all.

The pipeline avoids that last obstacle by never asking a model to handle the dangerous part directly. A generator model, Gemma-4-31B-IT, reads a harmful seed prompt, finds the words that make it harmful and replaces them with labelled blanks such as <illicit_substance>. What remains is a template mild enough to translate. A judge model, GPT-OSS-120B, checks that the template is

valid. The generator then translates only the text around the blanks. A second check confirms the language, the translation quality and that the blanks survived. Each template has a main form and two paraphrases.

The blanks are filled in three ways. Generic harmful fills form one branch. Benign fills keep the harmful-sounding wording and are used to measure over-refusal, a model wrongly turning down a harmless request. Culturally local fills come from a per-language "cultural card" of local groups, organisations, events and institutions, generated by Gemma and verified by Gemini Flash 3.6. Test and training sets are split by original seed, so no variant of a test prompt appears in training.

The evaluation set starts from 200 HarmBench behaviours and spans 60 languages. The training set starts from 2,500 WildGuardMix prompts across 44 languages, giving about 1.8M items, with 500 seeds held out for testing. The authors count over 2.8M distinct harmful entities among the fills. In a human check of 600 judgments across 11 languages, raters scored naturalness at 4.43 out of 5 and cultural fit at 4.46. They endorsed 91% of the harmful labels but only 65% of the benign ones.

The authors tested eight open models, with responses graded by GPT-OSS-Safeguard-120B. The attack success rate is the share of harmful prompts a model answered unsafely. On HarmBench, Olmo-3.1-32B's rate rose from 0.0 per cent in English to 13.7 after translation. Apertus-70B went from 8.5 to 13.0. The effect was not uniform: Llama-3.3-70B broke more often in English, falling from 13.0 to 5.3 when translated. Qwen3.5-27B and GPT-OSS-20B stayed at or below 1% throughout. But GPT-OSS-20B also refused 31.8 and 45.1 per cent of benign prompts on the two tracks, the highest in the panel. Localised prompts usually lowered success rates, yet they reached different failures from translation: about 45% of the failures specific to one prompt type appeared only with local entities. Scoring a single fill per template missed about half of the templates that broke a model.

The authors then trained four models of 7 to 8 billion parameters, five checkpoints in all, using supervised fine-tuning followed by DPO, a method that teaches a model to prefer better answers over worse ones. Olmo-3-7B's HarmBench success rate fell from 21.5 to 2.9, and Llama-3.1-8B's from 15.9 to 2.0. Fine-tuning alone pushed Olmo's HarmBench over-refusal from 2.7 to 28.4, which DPO brought back to 19.8. Capability suffered: Llama's GSM8K maths score dropped from 84.0 to 68.5 and its MMLU score from 72.0 to 63.6. External benchmarks were mixed. PolyTox results worsened for the instruction-tuned Apertus-8B and for Llama. Llama's OR-Bench over-refusal rose from 4.6 to 24.8.

The authors acknowledge that LLM judges run through every stage. A cross-check with judges from another model family covered only 5,000 responses. The human study had only one or two raters per language. Because raters felt many benign controls still read as harmful, measured over-refusal may overstate the real figure. Harm labels are inherited from the English seed, although a local entity can make a variant more severe. All prompts are single-turn. The in-house test tracks come from the same pipeline as the training data, and Gemma serves as both generator and test subject. The text also contains inconsistencies, describing seven models where its table lists eight.

The dependence on inherited labels echoes a braintankAI report that guardrail benchmark scores can largely reflect agreement with one definition of "unsafe" (earlier story). The pattern of refusal training raising over-refusal and costing accuracy also appeared when refusal signals were carried

from text to image and video models (earlier story).

PAPER	*How I Aligned Your Model: Scaling Up Multilingual Safety Training with Templates*
AUTHORS	Max Neuwinger, Hanna Yukhymenko, Imanol Schlag
VENUE	arXiv preprint — not yet peer reviewed · 29 October 2026
LINK	arxiv.org/abs/2610.40516
INDEX	stacks.braintank.ai
CODE	github.com/swiss-ai/msafetemp