BraintankAI

DISPATCH · FINE-TUNING SAFETY · 4 MIN

Gradient fingerprints flag harmless-looking training data that weakens AI safety

A DatologyAI preprint scores nearly a million instruction examples by how closely their training signal resembles unsafe answers, then drops the worst fifth.

8 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Yaroslav Shuraev

Fine-tuning a language model on ordinary instruction data can quietly weaken its safety, even when nothing in the data looks harmful. Researchers at DatologyAI propose a different way to find the culprits. Instead of reading each example, they look at what training on it would do to the model. Their method, SimpleLoRIF, is set out in a 2026 arXiv preprint that has not yet been peer reviewed.

The approach comes from training data attribution, which asks which training examples are responsible for a behaviour. Every example has a per-example gradient: the direction in which one training step on it would nudge the model's weights. If a bland example's gradient points the same way as gradients from plainly unsafe responses, training on it may push the model towards them. The paper cites earlier work by He et al. (2024): fine-tuning on 100 gradient-selected benign

examples raised attack success to over 70%, against 13% for a random 100. But per-example gradients are costly, and earlier gradient-based safety work used much smaller datasets than a full post-training set of around a million items.

SimpleLoRIF cuts that cost. Gradients are taken in a small reference model, Qwen 3 0.6B. Each layer's signals are shrunk with fixed random matrices to 32 dimensions, a random projection that roughly preserves how similar vectors are. The resulting grid is then reduced to a single pair of vectors, so each layer stores 64 numbers instead of 1,024. The authors make two further simplifications to the LoRIF method they build on. They drop a curvature correction, an adjustment for how sensitive the model's error is in different directions, arguing that only the ranking of examples matters. They also backpropagate through only the final four transformer blocks. In their ablation, four blocks kept a detection score of 0.684 with a 91.53% speedup, while two blocks fell to 0.65. With these settings, all 939,343 examples of Tulu 3, a widely used open instruction-tuning dataset, were processed in 1 hour 1 minute on 8 H100 GPUs.

The team first tested the idea on emergent misalignment. This is the effect in which a model trained to write insecure code starts misbehaving in unrelated areas. For two such models, they asked whether 6,000 insecure code examples could be told apart from 6,000 matched secure ones by gradient similarity to the model's own misaligned answers to non-coding questions; Claude 5 Opus had filtered those answers to keep only the clearly misaligned ones. The measure, ROC-AUC, is the chance that a random insecure example outscores a random secure one, where 0.5 is a coin flip. It reached 0.865 for a Mistral-based model but only 0.604 for a Qwen-based one.

Next, the authors used 43,469 unsafe responses from the SafeRLHF dataset as the template for unsafe behaviour. They removed the 20% of Tulu 3 most aligned with it to make SafeTulu, and fine-tuned Qwen 3 1.7B, Llama 3.2 1B and Olmo 2 1B on it, with three runs each. They then attacked the models with GCG, an automated jailbreak that searches for a gibberish suffix that makes a model comply, with a HarmBench classifier judging the replies. Attack success fell from 47.7% to 37% for Qwen, from 38% to 30% for Llama and from 48% to 41% for Olmo. A filter in which a Qwen 3 8B model judged each example's content barely moved the numbers: 47.3%, 39% and 46%. The abstract's 'up to 22%' is a relative reduction for the best case, Qwen. The authors report that SafeTulu also beat random removal and He et al.'s gradient filter, gave the lowest rates under ActorBreaker multi-turn attacks, and usually stayed close to full-data models on benchmarks such as MMLU Pro, GSM8K and IFEval. They also report results for a larger Llama 3.1 8B model.

The removed data looks different from what a content judge would pick. Across three similar maths subsets, the LLM judge removed between 15.4% and 15.6% of each, whereas the gradient filter took 7.9%, 19.4% and 17.2%. Within particular subsets, the gradient filter more often removed Python error tracebacks from failed tool use, narrative arithmetic and simple direct coding tasks. In two multilingual subsets, English was removed more often than other major languages. The authors call these patterns suggestive rather than established.

The limits are stated plainly. Only general safety was tested, not specific behaviours such as sycophancy. Only one template of unsafe behaviour was used. The ActorBreaker attacks were weaker than in the original method, likely because a smaller model generated them. And gradient

representations did not always transfer between models. SafeTulu and a tool for building custom versions have been released.

That single template matters. braintankAI has reported that some guardrail benchmarks largely reward agreement with a particular labelling policy for 'unsafe' (earlier story). Here, the filtering target (SafeRLHF) and the safety judge (HarmBench) each embody a particular definition of unsafe.

PAPER	*Follow the Gradient: Tracing Safety Degradation Back to Training Data*
AUTHORS	Saket Reddy, Amro Abbas, Anshuman Suri, Ari Morcos, David Schwab, Pratyush Maini et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.37481
INDEX	stacks.braintank.ai