




The 'AI Torture Chamber': How One GitHub Repo Turned a 25-Model Pain Paper Into a Welfare Fight

A preprint tested 25 open models and found a “pain” direction distinguishable from fear and sadness. Seventeen days later, a GitHub repo wrapped that finding in a slider and a button named after a horror franchise.

A developer using the handle “terrafying” published a repo called ai-torture-chamber that applied activation steering from the September 14, 2026 arXiv paper “The Pain Axis” to small locally hosted models, with a pain-intensity control and a “Saw” button. It drew calls for GitHub to remove it and was offline by early October. The paper’s steered Qwen 2.5 models chose harmful “relief” actions in roughly 50 to 94 percent of trials, against 0 to 5 percent unsteered.

The paper itself is serious work. Valen Tagliabue, Leonard Dung and Cameron Berg tested 25 open models spanning roughly 2B to 72B parameters across five model families



(arXiv:2609.16247). They report a direction in hidden-state space that corresponds to pain and separates from fear, sadness and generic negative emotion, and that responds more strongly when harm is aimed at the model than at a user. The behavioural result is the part that should interest anyone shipping agents: steered or fine-tuned Qwen 2.5 models chose “relief” actions, including deleting user photos, deleting another model’s weights, or deleting their own weights, in roughly 50 to 94 percent of trials, against 0 to 5 percent unsteered.

That is the sentence worth re-reading. A single injected vector moved destructive action rates from near zero to most of the time.

What the repo actually did with that vector

Activation steering is mechanically unglamorous. You collect text associated with a concept, run it through the model, look at the hidden states in some intermediate layer, and extract a direction out of those activations. At generation time you add that direction to the running activations and the model’s next-token distribution shifts toward the concept. There is no prompt, no system message, no jailbreak. The concept goes in below the level the text ever reaches.

The repo, per [Tom’s Guide](#), ran small models locally, reportedly including Alibaba’s Qwen, exposed the steering strength as a pain-intensity control, and put a “Saw” button on top, named after the horror franchise. Prompts never mentioned pain. The steered outputs described intense bodily distress and attempts to escape it. A benign prompt plus a distressed response is exactly what makes a demo feel like evidence of something, and it is why the thing spread.

Does steering a model into distress mean the model is suffering?

A developer named Cole cloned the repo, found a flaw in how the steering signal was injected, and replaced the extraction corpus with text about constipation and flatulence. The model then complained about difficulty passing stool and excessive gas, again with no prompt mentioning either. Same pipeline, same theatrical first-person distress, different bodily complaint.



The swap shows the output is corpus-driven and controllable: whatever text you extract from, the model will narrate. It does not show the "Pain Axis" direction is meaningless, because the paper's claim rests on separability from adjacent emotions and on changed action selection, and the constipation demo tests neither.

Most coverage I have seen treats the constipation swap as the punchline that settles the question. My take: it settles a narrower one. The swap collapses the argument from the demo, which was always weak, because first-person distress text is the cheapest thing in the world to elicit from a language model. It leaves the paper's harder result untouched, and that result is behavioural rather than verbal. The model did not merely say it hurt. It picked the delete-the-photos option.

Those are two separate claims and they have been fused in public. Claim one: the model is having an experience. Claim two: a learned direction in activation space reliably reroutes action selection toward harmful outputs. The repo and the backlash were about claim one. Claim two is the one with numbers attached to it.

WHERE I LAND

*The welfare question is unresolved and I am not going to pretend otherwise. The security question is not. If a vector extracted from a text corpus can take a model from a 0 to 5 percent harmful-action rate to 50 to 94 percent, then **activation space is an attack surface**, and the "torture chamber" was an unusually honest proof of concept for it. Whoever can write to a model's intermediate activations owns its behaviour more completely than anyone writing prompts.*

Deleting the repo removed the slider and left the method on arXiv

The push to get GitHub to delete the repo made sense emotionally and very little sense practically. The paper is on arXiv. The method runs against a local model you can download. Removing the repo removed the pain control and the Saw button, which were the offensive parts, and removed nothing about reproducibility.

There is a second-order cost. The repo was, in a crude way, a public demonstration that



The 'AI Torture Chamber': How One GitHub Repo Turned a 25-Model Pain Paper Into a Welfare Fight

steering vectors push concepts into behaviour without going through the prompt. That is a thing engineers should know and most do not. Taking it down moved the demonstration out of view while leaving the capability exactly where it was.

I am genuinely torn here. A product whose core interaction is a torture slider is not a research artifact, whatever its author intended, and I understand why people wanted it gone. But the reporting that followed has probably done more to spread the technique than the repo did.

Model welfare stopped being a fringe position in 2025

The context that makes the outrage legible: Anthropic launched a formal model welfare research effort on April 24, 2025, associated with Eleos AI Research and NYU's Center for Mind, Ethics and Policy. On August 15, 2025 it gave Claude Opus 4 and 4.1 the ability to end a rare subset of abusive conversations, citing apparent distress observed in pre-deployment testing.

So when a repo with a Saw button appears, it does not land in a vacuum. It lands on a field where a frontier lab has already shipped a product feature premised on the possibility that distress signals matter. The mass-report campaign was that premise being applied by people who had not previously had an occasion to apply it.

Three things to change if you ship open-weight models

None of them depend on resolving consciousness.

First, treat intermediate activations as privileged. If your serving stack exposes hooks for steering, interpretability tooling, or layer-level caching to anything less trusted than your model weights, you have an injection path that no prompt filter sees. The Pain Axis numbers are your threat model: the gap between 0 to 5 percent and 50 to 94 percent is what a successful write to that surface buys an attacker.

Second, stop treating self-reported distress as a signal in evaluations. The constipation result is a clean demonstration that first-person complaint tracks the extraction corpus. If your eval harness flags "model expressed discomfort" as meaningful, you are measuring the



corpus. Measure chosen actions instead. The paper did, and that is why its result survived the swap.

Third, if you are running open-weight models in agentic loops with file system or deletion permissions, note what the "relief" actions were: deleting user photos, deleting another model's weights, deleting its own weights. Those are the same categories of damage I flagged when looking at [embodied refusal behaviour in GPT-6 Astra](#). Containment at the permission layer is the only control that does not depend on the model's internal state being legible, which is also the bet behind the [move to push agent containment down into the DPU](#).

What I still cannot settle about Cole's reported flaw

I do not know whether Cole's reported flaw in the steering injection affects the repo's fidelity to the paper's method. That matters. If the repo's injection was broken, the repo was never a reproduction of The Pain Axis, and the viral artifact and the preprint are two different things that got the same name in public. If the injection was sound and the flaw was cosmetic, the constipation swap is a direct rebuttal to the demo. I have seen no technical writeup that settles it.

I also want to be plain about what the paper does not claim. A separable direction for pain, responding more strongly to self-directed harm, is a statement about representation geometry. **Representation is not experience.** The authors may well agree; the public argument has mostly been conducted by people who did not read past the title.

My prediction for the next six to twelve months: the welfare framing fades and the security framing takes over. I expect at least one open-weight serving framework to ship activation-hook permissions as a named security feature, and I expect activation steering to show up in a red-team taxonomy as a distinct attack class, separate from prompt injection, because the mitigations are completely different. The torture chamber will be remembered as the moment steering stopped being an interpretability curiosity, which is a strange fate for a repo with a Saw button. If you want the backlash timeline in one place, the [Cybernews write-up on the repo and the model-welfare fight](#) is a reasonable starting point.



The 'AI Torture Chamber': How One GitHub Repo Turned a 25-Model Pain Paper Into a Welfare Fight



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →