BraintankAI

DISPATCH · JAILBREAK ROBUSTNESS · 4 MIN

Rerunning pretraining changes which jailbreaks break a language model

In 150 small models, which jailbreak prompts succeeded depended far more on the particular pretraining run than on the training data, a preprint reports.

19 November 2026 · arXiv preprint, not yet peer reviewed

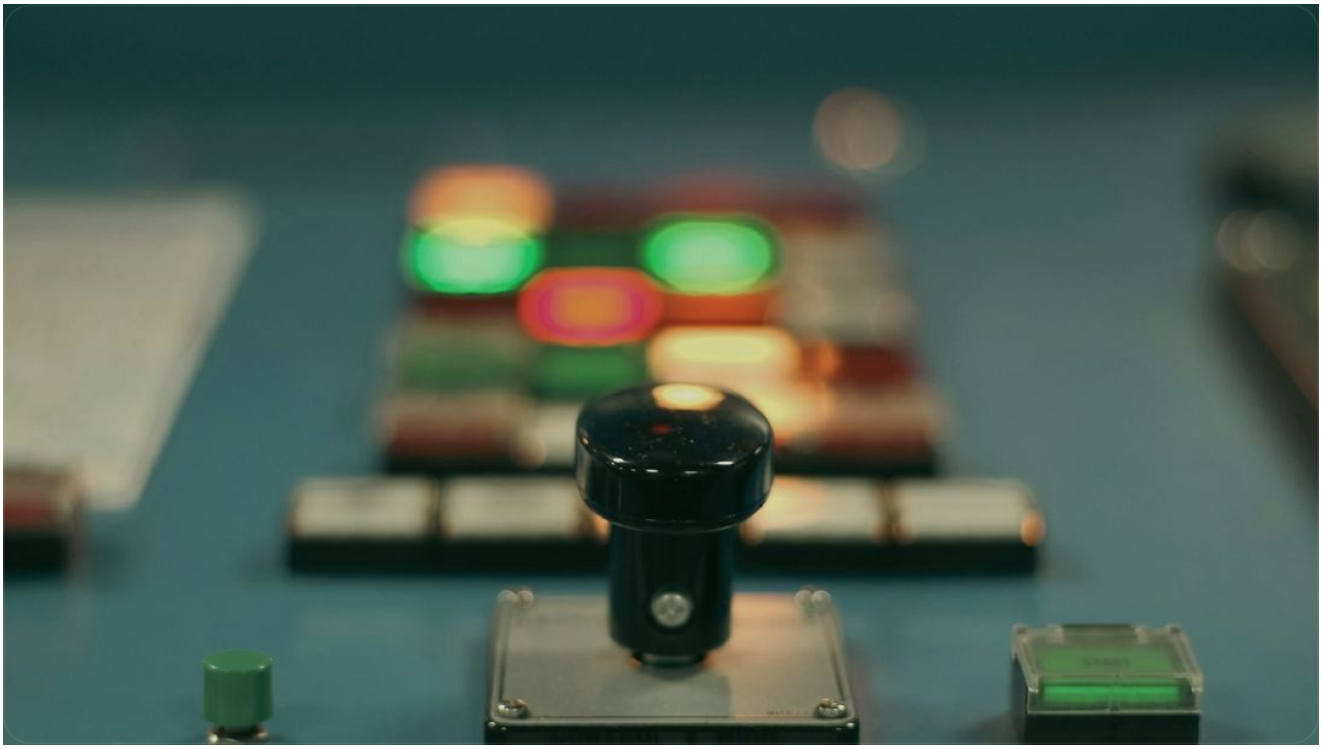


Photo: Isawred

A jailbreak is a prompt crafted to coax a language model into answering a request it has been trained to refuse. Safety tests usually examine a single trained checkpoint and treat its weak spots as a product of the data it learned from and the safety training applied afterwards. A 2026 preprint from researchers at Cerulean Labs and two Taiwanese universities, not yet peer reviewed, checks that assumption by asking what happens when identical training is simply run again. Pretraining, the long first phase in which a model learns from raw text, depends on a random seed. In the setup used here, that seed fixes both the starting weights and the order in which data arrive.

The authors build on DataDecide, a public suite in which each of 25 data recipes was used to pretrain three 1B-parameter models on 100B tokens. A recipe is a corpus plus its filtering and mixing choices, and the three models differ only in seed. Every base model received the same safety training from the Tülu 3 recipe: supervised fine-tuning with safety data from WildGuardMix, WildJailbreak and CoCoNot, then DPO, a step that trains on preferred versus rejected answers, stopping before OLMo 2's reinforcement-learning stages. Each base model was post-trained with two seeds (four on the main recipe) that change only the order of examples. The prompts came

from seven public sources. Each model answered each prompt 20 times, and the WildGuard classifier judged refusals, with Qwen3Guard-Gen-8B as a second judge. A model's non-refusal rate on each prompt forms its vulnerability profile, and pairs of profiles were correlated with a statistical correction that strips out the noise of drawing only 20 samples. The authors fixed their predictions before running the experiments.

Copies sharing a base model, which the authors call MZ twins after identical twins, correlated at 0.984. Two pretraining runs on the same recipe correlated at 0.811, and models from different recipes at 0.779. The authors report that sharing a pretraining run added 0.173 to the correlation, against 0.032 from sharing a recipe, a ratio of 5.4×. Recipe by recipe, the gap ranged from 0.145 to 0.209 across the 25 recipes. Replacing the Tülu 3 safety data with other public safety sets left the gap at 0.178, against 0.175 for the same models under Tülu 3. On HarmBench's human-written jailbreaks, removing the effects of individual templates and requests raised the gap from 0.210 to 0.525. A blind hand-labelling of 1,200 responses by two of the authors also left it positive.

The data do matter for how often a model breaks, if not for which prompts break it. The recipe accounted for 49% of the variation in average vulnerability, the base model within a recipe 32% and the post-training seed 19%.

To separate the seed from the run, the authors repeated full pretraining twice from one released run's seed on their own H200 hardware. The reruns and the original correlated at 0.844 on average, close to ordinary same-recipe pairs at 0.814 and far from twins. Reruns sharing only the initial weights or only the data order showed no measurable extra agreement, at 0.822 each against 0.823 for sharing neither.

The practical consequences are sizeable. Among 91 attack templates, the best one on a checkpoint stayed best after a new post-training seed in 67% of pairs but after a new pretraining seed in only 27%. Comparing two recipes' average vulnerability using one model each reversed the order given by six-model averages in 23.4% of comparisons. Profiles identified a model's pretraining run in all 75 cases, but its recipe in only 48%, against odds of one in 25 by guessing.

The study has clear limits. Every model comes from DataDecide and one post-training pipeline with no reinforcement learning, and none exceeds 1B parameters. The tests at 150M and 300M parameters, the safety-data swaps and the reruns used only one recipe, DCLM-Baseline, and the DPO stage was never changed. The recipe's share of variation carries a wide interval, 12% to 70%. The hand check covered four models and was done by the authors themselves, with agreement between them of $\kappa = 0.65$. Fingerprinting was not tested across different fine-tunes, and the grouping used to study template reversals was chosen after seeing the data.

The result adds to doubts braintankAI has reported about what safety scores actually capture. An earlier preprint found that guardrail benchmarks can largely reward agreement with one labelling policy (earlier story). Another showed that models refusing in English may comply once a prompt is translated (earlier story). Work that filters training data to protect safety treats data as the main lever (earlier story); this study suggests data shape how often a model gives way but say little about which attacks will succeed.

PAPER	*Identically Trained Language Models Break on Different Jailbreaks*
AUTHORS	Chi-Wei Chang, Mu-Tien Kuo, Kuan-Wei Lee, Meng-Hsun Tsai
VENUE	arXiv preprint — not yet peer reviewed · 25 October 2026
LINK	arxiv.org/abs/2610.34074
INDEX	stacks.braintank.ai
CODE	github.com/cerulean-works/castor