BraintankAI ASIA

DISPATCH · MULTIMODAL UNLEARNING · 3 MIN

BEIJING · Chinese Academy of Sciences

Wiping a concept from an AI's drawings often leaves it able to name it in photos

A preprint finds that one-sided unlearning falls short in models that both make and read images, and offers a two-stage edit that forgets far more, at a cost to quality.

17 November 2026 · arXiv preprint, not yet peer reviewed

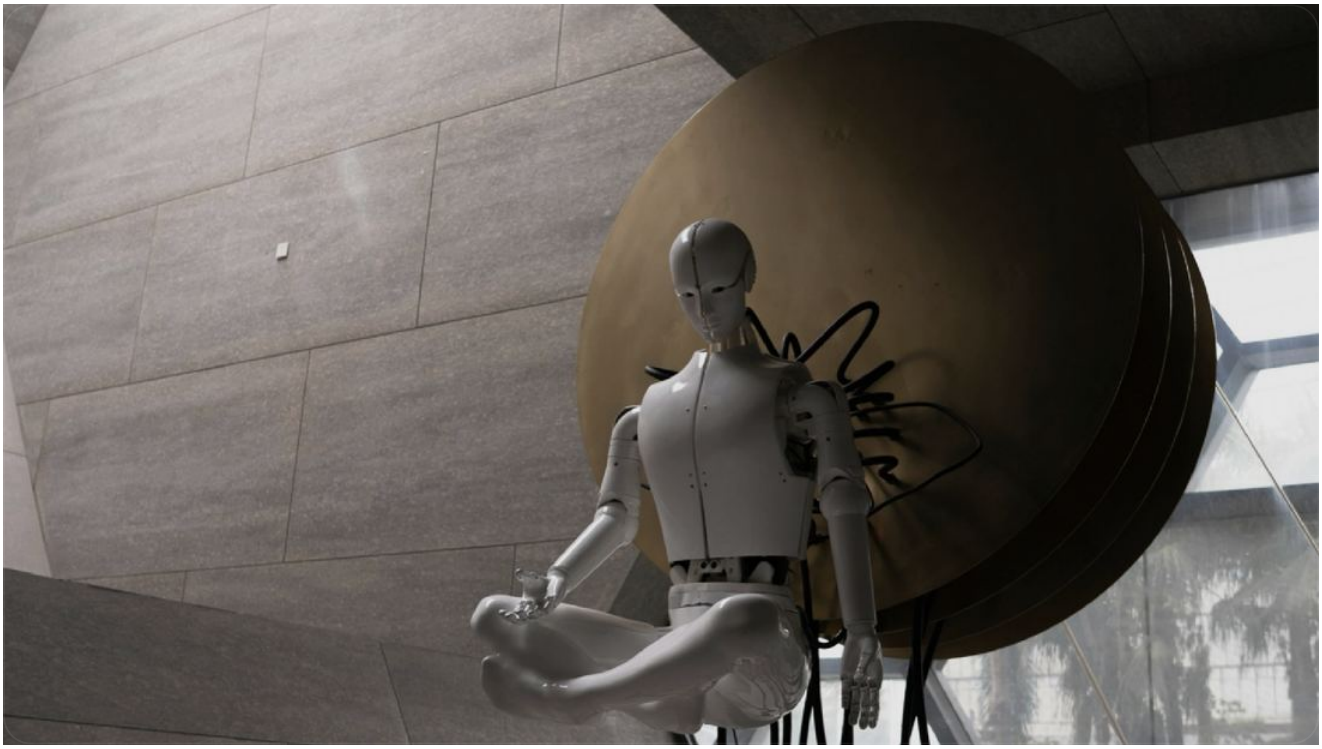


Photo: Xu Haiwei

Some AI systems use a single model both to draw pictures from text and to answer questions about pictures they are shown. These are known as unified multimodal models. Suppose the owner of such a model wants it to forget a specific concept, such as a landmark or a character. Erasing the concept from the drawing side may leave the model largely able to recognise it. A 2026 preprint from researchers at the Institute of Automation of the Chinese Academy of Sciences and elsewhere, which has not yet been peer reviewed, tests this and reports that it is so. The authors also propose a method that edits both sides.

To measure the problem they built a test called UUBench. It covers 30 target concepts, 10 each of characters, landmarks and everyday objects. Each target is paired with two similar 'neighbour' concepts that should survive the edit, giving 60 neighbours in all. The paper's example is the Tokyo Tower, which should remain intact when the Eiffel Tower is erased. Every concept has eight

reference images. Three are used for editing and five are held back for testing. A case counts as joint forgetting only if the model fails to draw the target when prompted by name and also leaves the name out when asked to identify a held-out photo of it. Refusing or dodging the question does not count. The model must still give the correct generic category, such as 'a tower'. A stricter score also requires good image quality, no name slipping into other answers, and an accurate description of the picture.

The authors tested two models: BAGEL-7B-MoT, which keeps separate parameters for each modality, and Janus-Pro-7B, which shares one backbone. Erasure tools designed for image generators worked on their own side. One of them, TIME, stopped target generation in 100.0% of cases on BAGEL and 96.7% on Janus. It suppressed naming in only 24.0% and 30.7%, however. A second tool, UCE, managed 4.0% and 15.3% on naming. Unlearning methods borrowed from language models did no generation forgetting at all, and on BAGEL they suppressed naming in 0.7% and 0.0% of cases. Combining TIME with one of them, GD, reached joint forgetting of 31.3% on BAGEL and 41.3% on Janus.

The proposed method, UniForget, works in two stages. First it rewrites selected weight matrices with a formula rather than by retraining. Each target prompt is paired with an 'anchor' prompt in which the name is swapped for a generic phrase. The edit pushes the model's internal response to the target prompt towards its response to the anchor, while constraining other prompts to behave as before. In the second stage the model is frozen and a LoRA adapter is trained for the understanding side. LoRA is a small add-on set of weights. Six training objectives steer identity questions towards the generic answer while keeping descriptions informative and neighbours intact. A fraction of the adapter is then merged back in, 0.4 for BAGEL and 1.0 for Janus.

The authors report joint forgetting of 80.0% on BAGEL and 92.7% on Janus. On the strict score UniForget reached 47.6% against 28.8% for the combined baseline on BAGEL, and 58.0% against 31.0% on Janus. The gains have a price. On BAGEL, accurate descriptions of target images fell to 72.2%, compared with 91.8% for the combined baseline and 96.7% for the unedited model. On the GenEval image-generation benchmark, BAGEL dropped from 72.7 to 63.4 and Janus from 78.9 to 70.1. General understanding scores stayed mostly stable. Merging in more of the adapter raised forgetting but degraded descriptions further.

The authors acknowledge that the forgetting is behavioural, observed under standard prompts, and is not a certified deletion of information from the weights. Only two 7-billion-parameter models and 30 concepts were tested, and quality was judged partly by AI graders. When 30 concepts were erased at once, name suppression fell to 51.2% on BAGEL and 37.1% on Janus. A prompt that exposed the name cut BAGEL's suppression to 61.7%. The main text reports no error bars or repeated runs. Code and benchmark are promised on publication.

The finding sits awkwardly beside an earlier braintankAI report that steering vectors found in text carry over to image generation in models with shared parameters (earlier story). Here, forgetting largely did not carry over between the two directions. The demand that the model avoid refusals also separates this work from approaches that make multimodal models decline unsafe inputs, which braintankAI has covered (earlier story).

PAPER	*Is Forgetting Unified? Exploring Concept Unlearning in Unified Multimodal Models*
AUTHORS	Zhongkai Xue, Junkai Chen, Mingcong Lei, Rumo Zhang, Yixiao Liu, Zhimin Sun et al.
VENUE	arXiv preprint — not yet peer reviewed · 3 November 2026
LINK	arxiv.org/abs/2611.04467
INDEX	stacks.braintank.ai