BraintankAI

DISPATCH · AI-TEXT DETECTION · 3 MIN

Sparse autoencoders turn an opaque AI-text detector into one auditors can read

A Google DeepMind and Penn preprint rebuilds a neural detector on features with plain-English descriptions, keeping most of its accuracy and, in some versions, generalising better.

20 November 2026 · arXiv preprint, not yet peer reviewed



Photo: National Cancer Institute

Tools that flag machine-written text already sit in peer review, classrooms, admissions offices and newsrooms, yet the accurate ones cannot say why they flagged anything. A preprint from researchers at Google DeepMind and the University of Pennsylvania, posted to arXiv in October 2026 and not yet peer reviewed, describes a way to open one up without throwing away most of its accuracy.

The starting point is gemma-2-2b, a small Google language model, fine-tuned to label text as AI or human. The authors read its residual stream, the running internal representation passed between layers, at layers 7, 13, 17 and 22, using sparse autoencoders: small networks that re-express a model's dense internal numbers as many features, only a few active at once, which tend to each mean one thing. These were initialised from GemmaScope, Google's public set of such autoencoders. Each feature is summarised over a document by its strongest activation on any token. A second model, gemini-3.1-flash-lite, looks at ten example snippets and writes a description

of 20 words or fewer. The same model is then given only that description and asked to guess which activation band each of 50 documents falls into; the correlation between its guesses and reality is the feature's simulation score. Features scoring below 0.7 are discarded, and a deliberately sparse linear classifier is trained on the rest, so any verdict can be explained by listing the features that drove it.

Three versions were built: one on unmodified GemmaScope features of the base model (Base-SAE), one with the autoencoders adapted to detection data on the base model (FT-SAE), and one adapted to the fine-tuned detector itself (FT-Model+SAE). Training used the English COLING shared-task dataset, covering 41 language models plus human writing, with 610,000 training documents. On its in-distribution split, the filtered FT-Model+SAE scored an F1 of 96.8 against 98.0 for the original detector, the authors report. On unfamiliar data it held up on overall measures but slipped at strict thresholds: the share of AI texts caught when only 5% of human texts are wrongly flagged fell from 77.9 for the detector to 67.9 on one test split.

The surprise came from the versions built on the base model. On the COLING test split, FT-SAE reached an F1 of 92.2 and caught 89.1% of AI texts at that 5% false-alarm rate, against 75.2 and 53.2 for the fine-tuned detector. The stylometric baseline gram2vec, which counts over 930 style features, managed 72.0 F1. Against Binoculars, a zero-shot detector not trained on COLING, FT-SAE led by 92.2 to 81.4 on COLING test but only 91.8 to 87.6 on the RAID benchmark. The authors hypothesise that the base model learned linearly separable 'AI text' features during pre-training, and leave the question to future work.

The features themselves are revealing. In a benchmark of 106,413 human NeurIPS and ICLR peer reviews paired with AI-written ones, one feature, double colons after bolded markdown headers, fired on 11 human documents and over 54,000 GPT-4o ones. Sentence-initial 'I' leaned human. On COLING, the top FT-SAE feature fires on refusal filler such as 'As an AI language model', while Base-SAE leans on HTML attributes and software-licence boilerplate, which the authors acknowledge may be domain shortcuts.

To test whether explanations reflect what the classifier really uses, the team paraphrased individual features away and checked whether the score moved as predicted; correlations ran from 0.69 to 0.81 on peer reviews but only 0.36 to 0.62 on COLING. A plain 'make it more human' rewrite flipped at most 8.7% of AI documents to human, whereas rewrites guided by the detector's own features, repeated over three rounds, flipped 60 to 100%. The authors read this as a sign that the feature lists are fairly complete.

The authors list further limits. Every trained model, stylometry included, scored near 100 on the peer-review data, which they take as a sign of confounds in that dataset. Only one base model was tested. Automated descriptions tend to over-generalise, and in an informal check about half the tested features had near-perfect descriptions. Beyond those, a single Gemini model writes, validates and attacks throughout, and the tables report no error bars.

The peer-review result echoes a braintankAI report that some AI safety benchmarks largely reward agreement with a labelling policy rather than a genuine ability to spot harmful requests (earlier story); in both cases a near-perfect classifier score says as much about the dataset as about the classifier.

PAPER	*Building Auditable AI Text Detectors with Sparse Autoencoders*
AUTHORS	Liam Dugan, Callum McDougall, Matthew Denton, Philip Pham, Christine Kaeser-Chen, Chris Callison-Burch et al.
VENUE	arXiv preprint — not yet peer reviewed · 21 October 2026
LINK	arxiv.org/abs/2610.30006
INDEX	stacks.braintank.ai
CODE	github.com/google-deepmind/detectorscope