BraintankAI ASIA

DISPATCH · MOE INTERPRETABILITY · 4 MIN

SHANGHAI · Shanghai University of Finance and Economics

Sparse autoencoders trace which narrow features each expert in an MoE model drives

A ByteDance Seed preprint splits feature activations exactly across a model's experts, finding most features driven by one or two, with themes labelled by another AI.

13 November 2026 · arXiv preprint, not yet peer reviewed

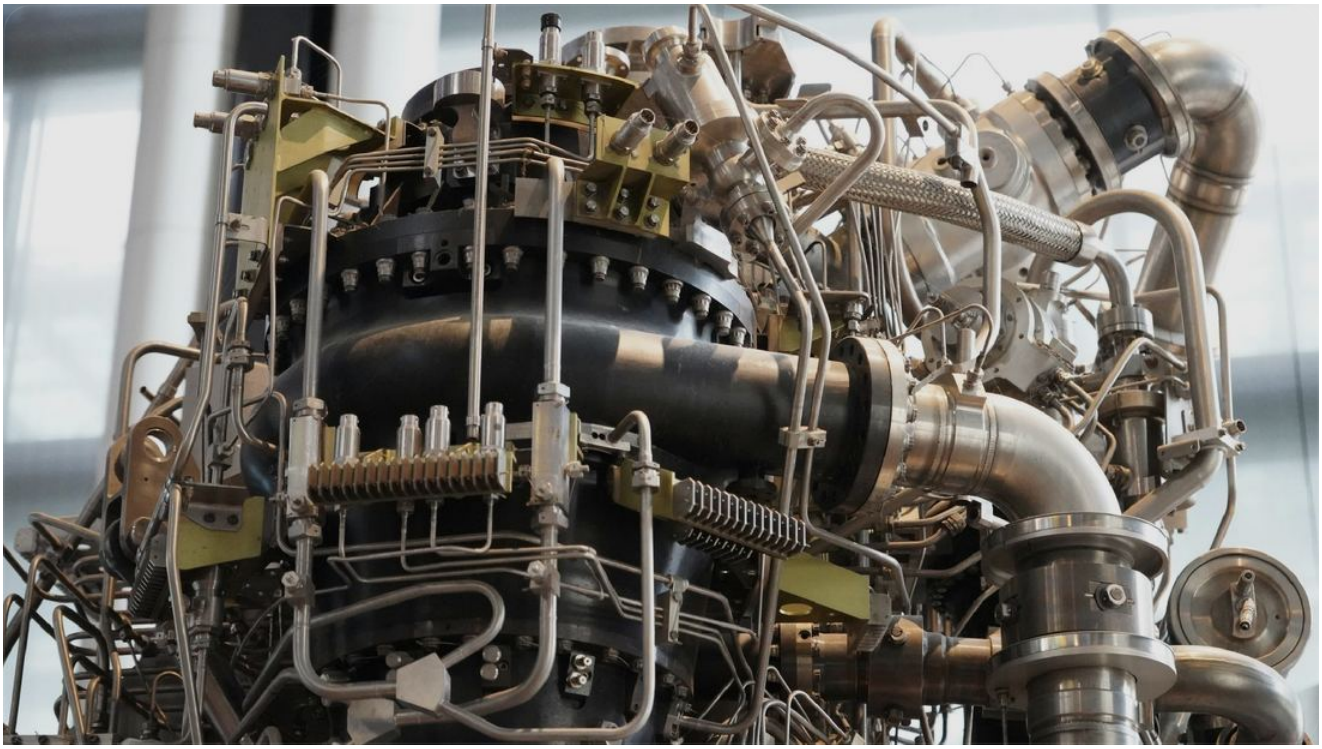


Photo: Meizhi Lang

Mixture-of-experts (MoE) language models send each token, a word or fragment of one, to a few of their many sub-networks, called experts, chosen by a component known as the router. What an individual expert does is disputed. Some studies find experts tied to subject domains, while others find routing driven by surface properties such as the identity of the token. A 2026 preprint from ByteDance Seed, not yet peer reviewed, argues that both camps rely on a blunt tool: counting how often an expert is picked on text that humans have sorted into a handful of domains. That fixes the categories in advance and labels whole documents rather than tokens, so story framing inside a maths problem still counts as maths. It also ignores how much weight the router gives each chosen expert.

The authors measure contribution instead. They train a sparse autoencoder (SAE), a model that rewrites a layer's output as a sparse mix of many learned directions, called latents, which tend to be easier to interpret than raw neurons. An MoE layer's output is a router-weighted sum of expert outputs, and the SAE's encoder is a linear map plus a bias. As a result, each latent's value before its activation cut-off splits exactly into the encoder bias plus one term per active expert. Keeping the positive parts and averaging over a corpus gives each expert a share of each latent. An expert is said to dominate a latent if its share reaches 0.2, where an even split would give about 0.008. The model is Qwen3-30B-A3B, which has 128 experts per MoE layer, 8 of them active per token and no shared experts. The statistics come from a held-out set of 1.5B tokens.

One obvious worry is that concentrated contributions simply mirror the router's uneven weights. The authors rank the eight active experts by each measure. At layer 30, they report, the top rank averages 0.231 of the gating weight but 0.512 of the contribution, while at the last of the eight ranks gating averages 0.074 and contribution 0.005. Discarding negative terms is not the main cause of the gap: on average 1.29 of the eight experts have their term dropped at that layer. Exactly one expert dominates 57.4%, 62.1% and 56.7% of latents at layers 4, 30 and 46, and latents with three or more dominators never exceed 3.7%. Every one of the 128 experts dominates some latents in every layer, and the ten busiest hold at most 30.1% of all expert-to-latent links. Overlap is measured by Jaccard similarity, the number of shared items divided by all distinct items. It is slight: at layer 30 only 20 of 8128 expert pairs exceed 0.05.

To attach meanings, GPT-5.5 described up to 2000 sampled latents per expert at layer 30, working from the contexts where each fires. The tokens routed most heavily to expert 72 are mostly function words, yet its latents cover exercise and health alongside algorithms and data structures. One of them, F97783, fires only at the <= of a binary-search loop. GPT-5.5 also judged theme shares for four experts: 86.7% for a Java expert and 78.1% for a geometry expert. Random pools of 2000 latents from the other 127 experts scored 1.9% and 2.0% on the same themes. GPT-5.5 then grouped the experts into 11 domains of very unequal size: 28 for computing, 21 for text form and 20 for mathematics and reasoning, while art appears in a single expert's label. Three of the domains are not subjects at all (text form, language, and entities and identifiers), and together they hold 36 experts. Of the 20 most-overlapping pairs, 19 share a domain, whereas only 12.8% of all pairs do.

The structural statistics rest on exact algebra, but the readings of meaning rest on thinner ground. Only one model is studied, and the interpretation covers layer 30 alone, which the authors justify by the layers' structural similarity. GPT-5.5 writes every explanation, judges every theme and does the grouping, and the main text reports no human check. The quantified themes cover four showcased experts, and the specialisation is not always clean. An arithmetic-and-polynomial expert scores 59.1%, and 13.6% of layer-30 latents have no dominant expert. A single dominance threshold is used. The attribution is also not a causal test, because no ablation or steering experiments are run. The authors offer the effect of editing an expert only as a hypothesis. They also suggest that the domain imbalance reflects the training data, but this is speculative, since Qwen3's data mix is undisclosed and they use Llama 3's as a stand-in.

That missing intervention is a step other interpretability work has already taken. braintankAI has reported on a preprint in which steering vectors found in text also nudged image generation in models that write and draw with shared parameters (earlier story). The equivalent test here would

be to edit one expert and check whether the features it dominates shift accordingly.

PAPER	*SAE Interprets MoE: Sparse AutoEncoders Reveal That MoE Experts Exhibit Human-Interpretable Fine-Grained Specialization*
AUTHORS	Anyuan Zhuo, Shaohua Zhang, Yuan Lin, Pinyan Lu
VENUE	arXiv preprint — not yet peer reviewed · 5 November 2026
LINK	arxiv.org/abs/2611.08089
INDEX	stacks.braintank.ai
CODE	github.com/ByteDance-Seed/SAE-interprets-MoE