BraintankAI ASIA

DISPATCH · DISCRETE VISUAL TOKENS · 3 MIN

WUHAN · Huazhong University of Science and Technology

Word-like image tokens close in on continuous features once given enough bits

A preprint finds that rounding error poorly predicts what discrete image tokens retain, and offers a parallel-codebook quantiser that kept improving up to the largest bit budget tested.

13 November 2026 · arXiv preprint, not yet peer reviewed

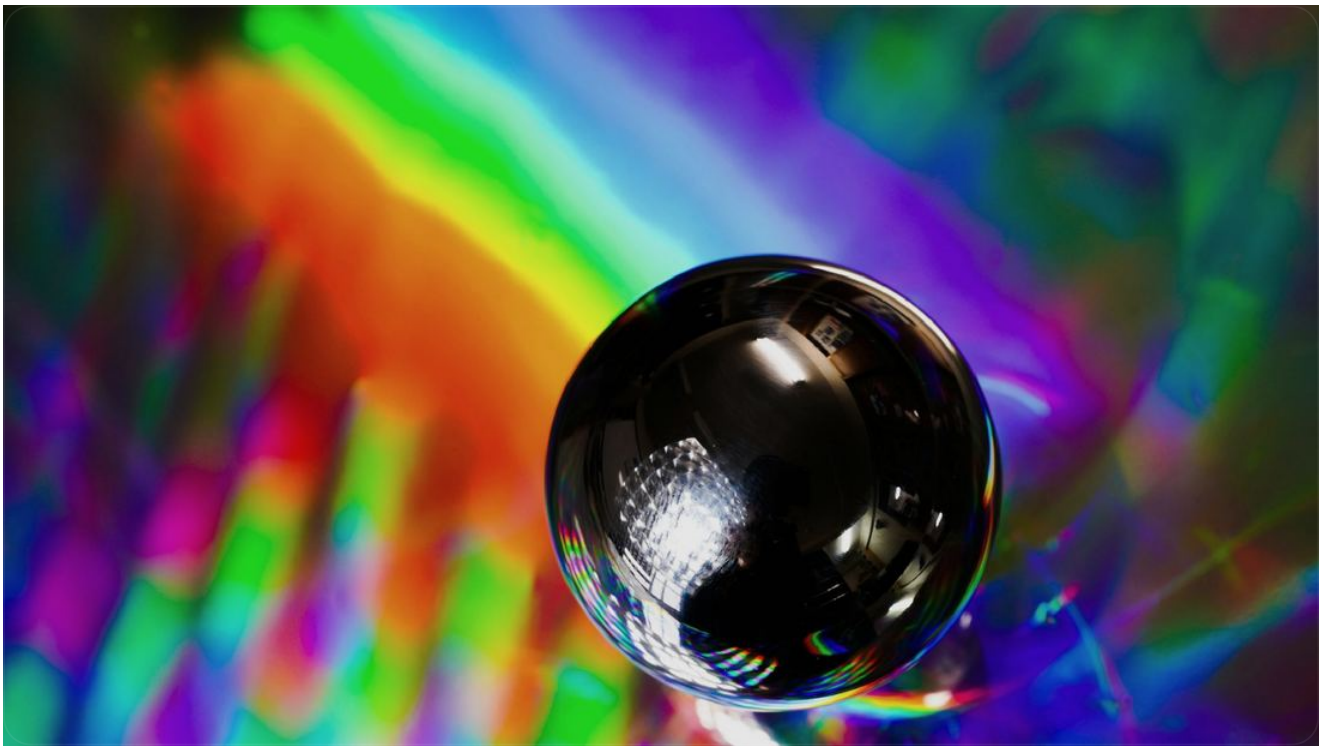


Photo: Flyd

Most AI systems that answer questions about pictures hand the image to a language model as continuous features: long lists of floating-point numbers from a vision encoder. The alternative is discrete tokens, integer IDs drawn from a fixed vocabulary. The model then reads pictures much as it reads words, and a single next-token objective covers both. The catch, it is widely held, is that this rounding throws away meaning. A 2026 preprint from researchers at Meituan's LongCat team and several universities, which has not yet been peer reviewed, argues that the usual measure of that loss is misleading. The authors build a quantiser around a different one.

The rounding works through a codebook, a learned list of vectors. Each patch of an image is replaced by the ID of its nearest entry, and quantisation error is how far that substitute sits from the original feature. In exploratory runs built on the smaller Qwen3-1.7B language model, a single-codebook quantiser had an error of 0.027 and averaged 51.2 across eight understanding benchmarks. The authors then tried residual quantisation, in which each codebook encodes what the previous ones missed. Keeping every other stage of an eight-stage residual quantiser, four in all and leaving out the first, coarsest one, produced a far larger error of 0.246, yet it averaged 60.7.

What mattered more, the authors argue, was the bit budget. This is how many bits each patch's token IDs can carry: the number of codebooks multiplied by the bits per codebook. With every codebook fixed at the same size, 14 bits each, they added codebooks, and existing schemes improved at first but then stalled. Residual quantisation averaged 64.0 with eight codebooks (112 bits) and 64.0 with sixteen (224 bits). The authors attribute this to its later codebooks seeing ever-smaller leftovers. Product quantisation, which slices the feature into ever-thinner channel chunks, went from 64.1 to 64.3.

Factorized Vector Quantization (FVQ) is designed to avoid both problems. It passes the encoder's feature through several fixed projections, each giving a view of the whole feature rather than a slice. Each view is then quantised in parallel by its own codebook. The projections are generated once from a random matrix and frozen. They are chosen so that, before any rounding, the views recombine into the original feature exactly, which means all distortion comes from swapping in codebook entries. In the same setup FVQ averaged 62.4, 63.9 and 65.3 with four, eight and sixteen codebooks. The unquantised Qwen3 ViT features beneath it averaged 64.6. FVQ still trailed those features on MMStar and on OCRBench, a text-reading test, scoring 663 against 676.

The full system pairs 16-codebook FVQ with the Qwen3-8B language model. It was trained in four stages on 100 billion, 10 billion, 60 billion and 30 billion tokens of industrial-scale data. Against Qwen3-VL-8B, a strong continuous-feature model, it scored 77.4 to 77.2 on the MathVista maths benchmark and 72.0 to 71.5 on RealWorldQA. It also scored higher on both parts of the CharXiv chart test. It fell behind on most other benchmarks, including MMMU-Pro (40.8 against 55.9), ChartQA (83.9 against 89.6), InfoVQA (77.7 against 83.1) and OCRBench (867 against 896). UniAR-8B, an earlier discrete-token model, scored 833 on OCRBench.

The limits are substantial. At four and eight codebooks FVQ only tied or trailed the residual and product quantisers, so its advantage rests on a single step from eight to sixteen. The gaps are about a point, and the paper reports no repeated runs or significance tests. Only understanding was measured, not image generation, which is one reason for wanting discrete tokens. The tokeniser also sits on top of a pretrained continuous encoder. The industrial training data is not described in detail, and the comparison models were trained on different data. The paper does not say how each patch's indices are combined beyond a merging module, or what the scheme costs in compute or sequence length. The authors present FVQ as just one way to ease the saturation.

The work bears on models that write and draw with shared parameters. braintankAI has reported that in such models, steering vectors found in text carry over to generated images (earlier story). Discrete image tokens are one route to that kind of single model, though this paper tests only

understanding. It sits less comfortably with a preprint that called next-token prediction an inefficient route to a shared representation across data types (earlier story). FVQ is offered in support of exactly such a shared next-token objective.

PAPER	*Can Discrete Tokens Serve as Semantically Lossless Representations of Visual Signals?*
AUTHORS	Jialv Zou, Xiaotong Li, Zhengrong Yue, Xu Huang, Yaxin Luo, Wenhao Yan et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.37662
INDEX	stacks.braintank.ai
CODE	github.com/hustvl/FVQ