BraintankAI ASIA

DISPATCH · MULTIMODAL TRAINING DATA · 3 MIN

SHANGHAI · Shanghai Jiao Tong University

Training AI on questions that link several image regions lifts it on 25 benchmarks

ZwZ2 learns from 20.4K automatically checked questions built to require combining several parts of one picture, and its gains extend well beyond fine-grained perception.

19 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Umut Tuluoglu

Vision-language models, AI systems that take in images alongside text, are increasingly asked to judge where objects sit relative to one another, or to follow clues scattered across many images and conversation turns. Good training data for those skills is costly to produce. In a 2026 arXiv preprint, which has not yet been peer reviewed, researchers from Shanghai Jiao Tong University, Ant Group and other institutions argue that such tasks rest on one shared skill they call connecting evidence: finding the relevant visual facts, tying each to the right object, working out how they relate and combining them into an answer. Their bet is that data for this skill is far easier to build and check than data for each downstream task.

The pipeline, EvidenceConnect, starts from high-resolution, visually busy images. A strong model proposes concrete regions with bounding boxes, then groups three to seven of them per question. Questions take three forms: aggregation, such as combining counts; comparison and matching,

where the regions form the complete set of candidates, as in finding the odd one out; and dependency chains, where one clue determines what to read next. The question writer's only visual input is a "privileged" collage of the chosen regions at native resolution, each outlined in a distinct colour that also marks it in the full image. The model being trained never sees this collage and must work from the whole picture. Alongside each question, the writer must supply a short step-by-step justification and the answer it expects if each region were removed, so that simple rules can reject questions with unused regions, leaked answers or inconsistent arithmetic.

Two checks follow. Two models from different families answer each question repeatedly, and it survives only if a clear majority of each agrees. Then each region is greyed out in turn; a question is discarded if more than one region turns out to be unnecessary. Survivors become four-option multiple-choice items, and any that a lightweight model always answers correctly are dropped. The result is 20.4K examples, used to train Qwen3.5-4B, Qwen3.5-9B and Qwen3.6-27B with GRPO, a reinforcement-learning method that compares several sampled answers to the same question and rewards the better ones, here simply for picking the right letter. There was no supervised fine-tuning, and the models were trained in non-thinking mode.

Across 25 benchmarks in six domains, the authors report that the trained models, called ZwZ2, beat their base models on every benchmark at every size. Overall averages rose from 49.41 to 54.79 at 4B, 51.63 to 56.59 at 9B and 55.52 to 60.51 at 27B. Although training used single images, the 4B model climbed from 60.08 to 70.08 on MileBench, a multi-image long-context test, and the 27B model went from 51.25 to 61.75 on the embodied reasoning benchmark ERQA. ZwZ2-27B's overall average of 60.51 exceeded those of GPT-5.6-Luna (54.02), GLM-5.3-Flash (57.24) and DeepSeek-V4.1-Flash (55.47). However, those references ran in low- or non-thinking modes with resolution capped by their platforms, and ZwZ2-27B trailed at least one of them on the long-context, spatial, embodied and general-reasoning averages.

A further test trained the same 4B model with the same recipe on five other synthetic datasets, scored on 11 benchmarks. Against a base average of 47.16, MMHops, TVG, F2RVLM and CRIT all left the model worse off, at 42.73, 45.48, 46.42 and 46.61. The team's earlier single-region dataset, ZwZ, reached 49.86 with 73K examples; the multi-region data reached 53.86 with 3.6x fewer. Most of the gain arrived early: 7K examples already gave 52.65.

The results have limits. No repeated runs or error bars are reported, and the two ablation effects are 1.27 and 1.63 points. Region-removal checking hurt ERQA by 4.00 points. Some gains are marginal: at 27B, ScreenSpot-Pro rose only from 70.21 to 70.59 and UI-Vision from 66.05 to 66.66, and at 128K tokens on a subset of MMLongBench the 4B and 9B models only matched their bases. Absolute scores stay low on some tasks, with Chartography at 12 to 14. The pipeline also depends on strong models that the main text does not name, and it gives no compute cost.

The approach contrasts with a preprint braintankAI covered that tackled spatial reasoning by replaying mixed spatial data in time order, after finding models could learn left from right without learning distance (earlier story). ZwZ2 instead trains on general multi-region questions rather than dedicated spatial data. It shares with Xiaomi Research's self-training loop a reliance on automatically checkable questions rather than human labels, though here the questions are generated and vetted once by other models (earlier story).

PAPER	*Zooming without Zooming v2: From Seeing Details to Connecting Evidence for Generalizable Multimodal Perception*
AUTHORS	Lai Wei, Yatai Li, Liangbo He, Jun Lan, Renjie Gu, Linhan Cao et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.37037
INDEX	stacks.braintank.ai
CODE	github.com/inclusionAI/Zooming-without-Zooming-v2