BraintankAI ASIA

DISPATCH · AGENT TRAINING DATA · 4 MIN

HANGZHOU · Zhejiang University

Fixing broken practice tasks before AI agents learn office work from them

WorkFactory maps the information each work task needs and patches gaps by test-running it, then fine-tunes smaller Qwen models on a stronger model's demonstrations.

17 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Shane Rounce

Teaching an AI agent to do professional work by example has a weak point beyond the choice of model being imitated. The trouble can lie in the task it is handed. Suppose a synthetic supplier-selection brief quotes a discount that does not square with the shipment volumes implied by its own demand forecast. The teacher model may guess or skip the analysis, weakening the demonstration that the smaller model learns from. WorkFactory, described in a 2026 preprint from IQuest Research and university collaborators that has not yet been peer reviewed, is a pipeline that aims to make work tasks' files complete and consistent before any demonstration is recorded.

It starts from roughly 120,000 professional intents, drawn from occupation hierarchies, documented skills and real professional deliverables. For the deliverables it uses instruction backtranslation, which means inferring from a finished report what instruction might have produced it. For each intent, a planner works backwards from the expected output to the

calculations and judgements behind it, then to the pieces of information those require. The result is a Semantic Workflow Graph, a map of which facts the deliverable depends on and which request, file or external source carries each one. Raw premises go into the task environment. Derived figures and conclusions are deliberately left out for the solver to work out. A worker then fetches real files or synthesises supplementary material to fit the plan.

The second stage tests each task by attempting it. A probe solver, which sees only what a real solver would see, tries the task and reports anything that blocks it. A refiner then edits the request or files. In the paper's example it adds a term that Supplier C charges USD 0.15 per carton for freight, without doing the comparison for the solver. Another language-model review checks that each edit neither strips out the core work nor leaks the answer. Tasks that cannot be fixed within a repair budget are thrown away. The final corpus holds 112,686 tasks across 612 occupations, with an average of 4.15 reference files per task. GLM-5.3 attempted them in clean sandboxes, and its filtered recordings, or teacher trajectories, averaged 57.36 turns and 62.95 tool calls. These were used for supervised fine-tuning, in which a model is trained to reproduce demonstrations. Three Qwen base models were trained this way to produce WorkPilot-4B, WorkPilot-27B and WorkPilot-35B-A3B.

The main comparison is each model against its own base on four benchmarks: GDPval, JobBench, Workspace-Bench Lite and OmegaUse-OfficeVal. GDPval is scored as an Elo rating, a chess-style score built from pairwise judgements of outputs. The authors report that WorkPilot-35B-A3B rose from 28.91 to 62.35 on JobBench and from 7.14 to 21.18 on OmegaUse-OfficeVal. WorkPilot-27B went from 975 to 1612 Elo on GDPval, between DeepSeek-V4.1-Flash at 1600 and GLM-5.3 at 1646. However, the WorkPilot figure is a local estimate with Gemini 3.1 Pro Preview as judge, whereas the frontier scores are official ones. WorkPilot-4B beat Apodex-1.1-mini on GDPval and OmegaUse-OfficeVal but not on JobBench or Workspace-Bench. The students still trail their teacher on all four benchmarks. WorkPilot-27B, for instance, scores 63.05 on JobBench against GLM-5.3's 66.35.

A component study on 5,000 tasks with the 35B-A3B model separates the pieces. With file-level planning alone, the model reached 1245 Elo. Adding the graph raised that to 1293, and adding repair raised it to 1326. The graph cut broken derivations from 24.86% to 6.86% of tasks, but leakage rose from 1.92% to 3.36%. Repair brought cross-file inconsistency down from 25.74% to 4.84% and leakage to 0.84%. Changing the teacher mattered too. On the same 30,000 tasks, a student of GLM-5.3 reached 1483 Elo, against 1248 for a student of DeepSeek-V4-Flash-0731. More data showed diminishing returns. The first 5,000 tasks added 447 Elo, while going from 75,000 to 110,000 tasks added 13. Average episode length rose from 24.87 turns to 162.88 at 2,000 tasks, then fell to 85.90 at 110,000. Web-fetch failures stayed at around 40%–48% throughout.

The authors note two limits. Because the teacher reads only text, training and testing were text-only. The models were trained by fine-tuning alone, with no reinforcement learning. The setup carries other caveats. WorkFactory is not compared directly with other task-building frameworks. The defect rates come from the same kind of language-model probe that the repair loop is designed to satisfy, which may flatter the repair step. The component study used one model at one data size. The other benchmarks use Claude Sonnet 4.6 as judge where one is needed, rather than people. Trajectories were collected in the same Stirrup and Claude Code harnesses later used for evaluation, which may favour WorkPilot. Only part of the task data has been released.

The gap this targets is the one measured by a benchmark of real projects across 55 kinds of digital work, where the best agent set-up fully passed 24% of public tasks (earlier story). The recipe also resembles fine-tuning small open models on the logged steps of stronger AI-driven searches (earlier story). WorkFactory's difference is its focus on planning and repairing the tasks before anything is recorded.

PAPER	*WorkFactory: Turning Professional Intents into Executable Work for Scalable Agent Training*
AUTHORS	Ziyang Zeng, Tingyang Chen, Jia Deng, Qiyao Wang, Shuo Tang, Yike Yang et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.37161
INDEX	stacks.braintank.ai
CODE	github.com/IQuestLab/WorkFactory