BraintankAI

DEEP READ · AGENT BENCHMARKS · 8 MIN

AI agents fully finish at most a quarter of expert-built professional tasks

A benchmark built from real projects across 55 kinds of digital work finds the best agent setup fully passes 24% of public tasks, and agents pass almost none of the hardest.

7 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Navy Medicine

Agents' Last Exam, or ALE, is a benchmark (a fixed set of test tasks with an agreed way of scoring them) built from work that professionals had already done on computers. It is described in a 2026 preprint, not yet peer reviewed, by Yiyou Sun and colleagues. It was developed with more than 250 domain experts and spans 55 kinds of work in 13 industry clusters. Experts submitted projects they had already completed, often representing days or weeks of effort. The result is 1,490 task instances, most of them marked by automatic checks rather than by human or AI judges. On the public portion, the best agent setup fully completed 24.0% of tasks. On the hardest tier, the average full-pass rate across mainstream setups was below 1%.

The authors start from a puzzle. AI systems keep clearing celebrated tests in games, olympiad mathematics and competitive programming, yet by their account the economic impact has been surprisingly limited. They argue this is largely an evaluation problem, because popular benchmarks do not measure sustained performance on long, economically valuable workflows.

They point to three obstacles. Authentic workflows are expensive to collect, broad industry coverage needs sustained access to experts, and real deliverables such as spreadsheets, designs and reports are hard to check automatically. As a result, existing benchmarks often give up realism, breadth or verifiability. ALE tries to keep all three. The bar it sets is demanding: an agent must hand over a professional's finished product, not just answer questions or finish short computer chores.

How it works

The authors started from the US government's job classification, SOC 2018, and its companion database O*NET. They grouped occupations with similar software-based workflows into 13 domains and 55 subdomains, leaving out sectors whose core work is not meaningfully digital. Experts submitted projects through a web portal, specifying a description of the task, the input files, the target software, the expected deliverable and how to evaluate it. Each submission then went through a conference-style first review, engineering implementation with dry runs by engineers, and a final peer-review check by an expert committee. Of the 1,490 instances, 960 came from external submissions and 530 were commissioned.

Each task is a small program that loads the task, starts it and evaluates the result as a score between 0 and 1. The agent works on a remote virtual machine with four folders: read-only inputs, the software it should use, an output folder that is the only place it may write, and a hidden reference folder holding the ground truth. Scoring compares output with reference in whatever way suits the deliverable. Options include exact values, numbers or tables within a tolerance, geometric distances, the behaviour of a built system under a fixed sequence of inputs, and rubrics for free text. Checks are either averaged as a checklist or combined as gate-and-score. In gate-and-score, a must-pass condition comes first, such as the file opening at all, and failing it means a score of zero. A vision-capable language model acts as judge only where no deterministic check exists, for example for video, game screenshots and rendered scenes, and then only for narrow yes-or-no questions.

The systems under test are what the authors call generalist computer-use agents. These can operate graphical applications and also run code, shell commands and file operations in one loop. Popular agent harnesses, such as Claude Code, Codex and OpenClaw, are built around the command line. A harness is the software around a language model that runs its loop, supplies tools and manages its working context. The authors connected each harness to 14 desktop-action tools through MCP (Model Context Protocol), a standard way of plugging tools into an agent. Models that cannot read images hand screen work to a separate vision model acting as a sub-agent. The authors also built ALE-Claw, a stripped-down version of the OpenClaw harness without long-term memory or user-preference features.

Only 150 instances, about a tenth, are public. Most of the rest are kept private and are meant to rotate into public use over time. This is meant to limit contamination, in which models see test tasks during training and score well for the wrong reasons.

What the evidence shows

The main results come from three public tiers. Near-Term has 67 tasks the tested agents can partly solve, Full-Spectrum has 55 tasks with at least one per subdomain, and Last-Exam holds the 38 hardest. The paper reports full-pass rate, the share of tasks given full credit, alongside mean score, which averages partial credit. Each run was capped at five hours. Codex with GPT-5.5 came top at 24.0% overall. It passed 38.1% of Near-Term tasks, 22.7% of Full-Spectrum and none of Last-Exam, with mean scores of 64.7, 36.0 and 11.2. Close behind were ALE-Claw with GPT-5.5 (23.0%), Claude Code with Claude Fable 5 (22.0%), three Cursor configurations (20.4 to 20.7%) and Droid with GPT-5.5 (19.1%). Codex with GPT-5.4 (7.2%) and Grok CLI with Grok 4.3 (6.6%) trailed. On Last-Exam only four configurations passed anything, a single task each, and the best mean score was 12.8.

To separate the effect of the model from that of the harness, the authors ran 14 models inside OpenClaw. GPT-5.5 (21.1%) and GPT-5.4 (20.5%) led, followed by Claude Opus 4.7 (15.1%), with Grok 4.3 last at 4.3%. None of the 14 passed a Last-Exam task. Fixing the model and changing the harness still moved results. GPT-5.5 ranged from 19.1% in Droid to 24.0% in Codex, and Opus 4.7 from 12.8% in Droid to 20.4% in Cursor. The authors conclude that among well-engineered systems, the choice of model accounts for roughly three times the spread in results that the choice of harness does. Yet their own table holds a sharp exception: GPT-5.4 scored 20.5% in OpenClaw, 11.8% in ALE-Claw and 7.2% in Codex.

ALE-CLI is a Linux-only subset open to agents without screen access. On it, Codex with GPT-5.5 reached 23.3%, against a quoted 82% for the same pairing on Terminal-Bench, a command-line benchmark. Command-line-only agents running Claude Sonnet 4.6 scored between 9.0% and 13.3%. Runs are costly, at $3 to $10 per task on average for a frontier agent, and take tens of minutes to hours. The authors report that heavier spending did not reliably buy better results. Summed across the tiers, Cursor with Composer 2.5 spent about $87 for 20.4%, Codex with GPT-5.5 about $570 for 24.0%, and OpenClaw with Gemini 3.1 Pro about $3,133 for 14.1%.

Behavioural analyses add texture. Graphical software is the primary tool in 34% of public tasks, yet most configurations made few graphical actions and fell back on shell commands. Scores varied sharply by domain. For Claude Fable 5 and GPT-5.5, computational mathematics and agriculture and environment reached roughly 55 to 85%, business and legal about 50 to 55%, and education below 25%. The authors also broke down failures for Claude Code with Opus 4.7. Roughly three quarters of failed runs stemmed from misunderstanding the task or choosing the wrong approach, rather than from slips in execution.

What it does not show

The evaluated set is small. Every headline number rests on roughly 150 to 160 public tasks, and the paper's own counts differ slightly between sections. On Last-Exam, a single task is worth 2.6 percentage points. Only some configurations were run more than once, because of compute costs. Where three runs exist, score standard deviations reach three points, and five for one agent. The order within the top cluster, between 19% and 24%, is therefore not a firm ranking. The claim that the public tasks represent the larger pool rests on the authors' own check in an appendix, and about 22% of all instances had yet to pass quality checks.

Test conditions also shape the numbers. Every agent in the main results ran with the authors' own desktop-tool bridge added to its harness. Three models (DeepSeek V4 Pro, GLM 5.1 and MiniMax M2.7) worked through a vision sub-agent. The five-hour cap cut off 3.8% of runs, ranging from about 1% for lightweight harnesses to 5.7% for OpenClaw, and some scoring relies on a model judge. The authors say ALE-Claw performs comparably to OpenClaw. That holds for GPT-5.5 and Opus 4.7 but not for GPT-5.4 (11.8% against 20.5%). The failure breakdown covers only one configuration. The authors' explanation for uneven domain scores, uneven model capability across domains and uneven exposure to tool-use training, is offered as likely rather than tested. The Terminal-Bench contrast sets one quoted figure against ALE-CLI, so it is not a controlled comparison.

Most important for the economic framing, no human-expert baseline is reported, so the paper does not quantify how hard the tasks are for professionals. Nothing ties the scores to productivity. The idea that saturating ALE would signal real economic transformation is an aspiration, not a finding. Physical work is excluded by design.

Where it fits

The authors set ALE against benchmarks that each give something up. These include OSWorld's shorter computer-use tasks, synthetic environments, question-answering tests and benchmarks confined to a few domains. They also include economic-work evaluations such as GDPval and the Remote Labor Index, which rely on human judges. The authors describe ALE's task surface as a superset of graphical-only benchmarks such as OSWorld and command-line-only ones such as Terminal-Bench. They mapped 16 prior benchmarks onto their taxonomy and found 13 of the 55 subdomains covered by none. That mapping used a language-model-assisted classifier, which may be imprecise.

The authors draw an analogy with ImageNet, the labelled image collection long used as a benchmark in computer vision. They argue that once a field has a verifiable, widely used benchmark, progress tends to accelerate and deployment often follows. They also argue the industry often equates computer-use agents with agents that only click through screens, a conflation they call incomplete. And they contrast thin ReAct-style loops, in which a model simply alternates reasoning and actions, with full harnesses. In their results, a single model's score shifted by more than 13 points between harnesses. That underlines why they treat the harness as part of the system under test.

For this series, ALE sets the yardstick: sustained, economically valuable work delivered as a finished product that can be checked. Measured that way, the best agent tested in 2026 fully finished just under a quarter of the public tasks, and agents finished almost none of the hardest.

PAPER	*Agents' Last Exam*
AUTHORS	Yiyou Sun, Xinyang Han, Weichen Zhang, Yuanbo Pang, Tianyu Wang, Yuhan Cao, Yixiao Huang, Chris Duroiu, Haoyun Zhang, Jeffrey Lin, Weishu Zhang, Tyler Zeng et al.
VENUE	arXiv preprint — not yet peer reviewed · 3 June 2026
LINK	arxiv.org/abs/2606.05405
INDEX	stacks.braintank.ai
CODE	github.com/rdi-berkeley/agents-last-exam