BraintankAI

EXPLAINER · CODING AGENT HARNESSES · 5 MIN

Which parts of a coding agent's scaffolding matter, tested one piece at a time

Researchers from Zoom and US universities swapped planning, tools and context handling in and out of one coding agent and found each piece's value depends on model, task and budget.

16 November 2026 · arXiv preprint, not yet peer reviewed

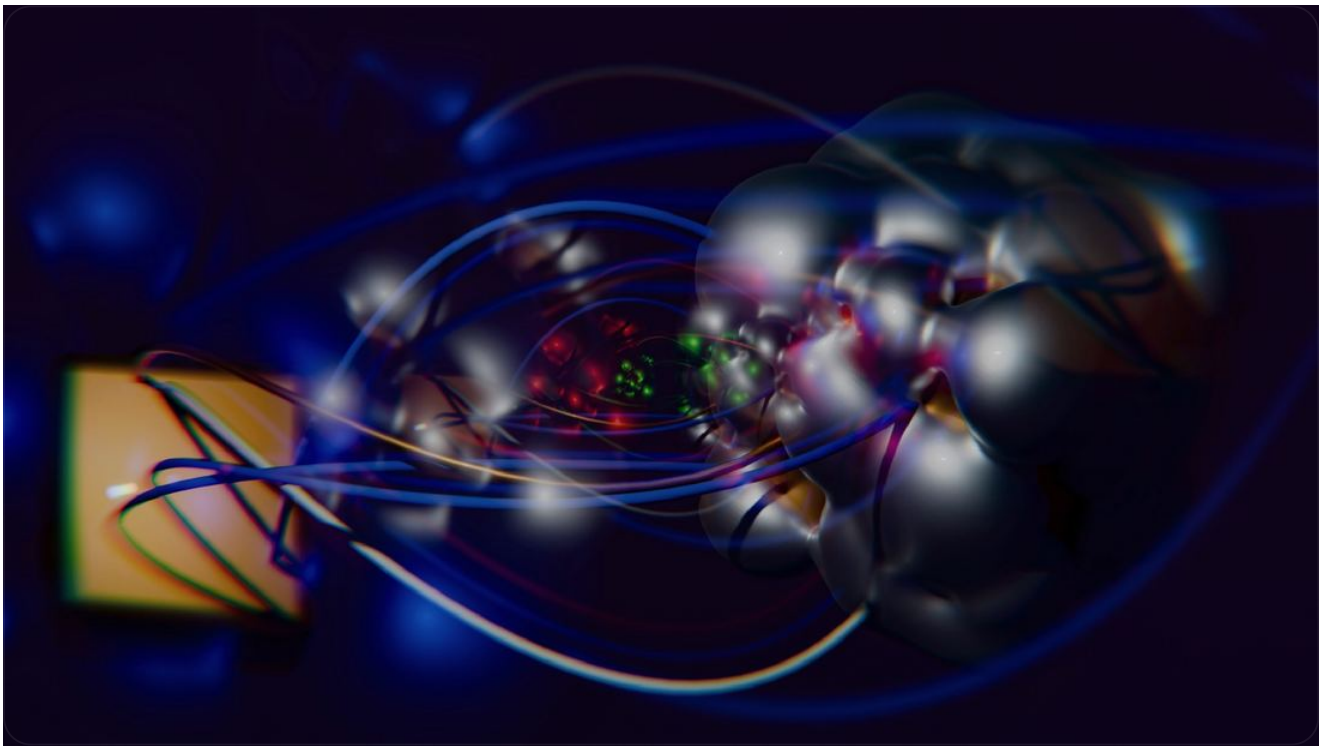


Photo: Steve A Johnson

When people compare AI coding agents, they usually compare whole products. One system pairs a language model with a particular set of tools, a way of planning and a way of handling its own growing history, and another pairs a different model with different choices on all three. If one scores higher on a benchmark, it is hard to say why. Was it the planning? The tools? The memory tricks? Or simply that the model and the scaffolding happened to suit each other?

A 2026 preprint by Run-Ze Fan and colleagues at Zoom Video Communications, UMass Amherst, Emory University and UNC Charlotte tries to pull those strands apart. The paper is on arXiv and has not yet been peer reviewed. Rather than pitting finished systems against one another, the authors built a single, deliberately simple coding agent and changed one component at a time, across four

open-weight models (models whose trained parameters are publicly released) and two coding benchmarks. Their question is whether each component is useful in general, or whether its value depends on the model, the task and the resources available.

The idea

A language model on its own only produces text. To fix a bug in a real software project, it has to be wrapped in what the field calls a harness: the software that runs the model in a loop, gives it tools to read and edit files or run commands, and decides what part of the growing record of the session the model gets to see on each turn. The most common loop is ReAct, short for reason and act: the model thinks, calls a tool, reads the result and repeats until it is done.

The authors argue that because harnesses are normally judged as complete bundles, differences between them cannot be traced to any single design choice. They point to an earlier comparison by Cao and colleagues, also from 2026, in which one Anthropic model, Claude Opus 4.5, did best inside the OpenHands harness while its sibling Claude Sonnet 4.5 did best inside SWE-Agent. That is a hint that the right scaffolding may depend on the model. They also note that fairly plain pipelines such as Agentless and AutoCodeRover have shown that elaborate architectures are not always needed.

One component gets special attention: context management. Every model has a context window, the maximum amount of text, measured in tokens (word fragments), that it can take in at once. A long coding session fills it with file contents, search results and error logs. Earlier approaches include cutting old material, storing it elsewhere for retrieval, or summarising it. The authors observe that such methods are usually introduced as improvements on one model, so little is known about how much they matter as the window grows or the model gets stronger.

How this paper puts it to work

The team built their harness from scratch around a fixed ReAct loop, capped at 300 steps per task, with certain safeguards always on, including a check that stops the run if the agent keeps repeating the same failing call. They then varied three things. Planning, when switched on, prompts the model to write a plan on its first turn, gives it a tool to update that plan, and shows the current plan to it on every turn. The action space, the set of moves the agent can make, was either a menu of purpose-built tools for reading, editing, listing and searching files alongside a shell, or a raw command-line shell (bash) on its own for that work. The authors stress that the tool menu also brings extras, such as tracking file state and running automatic code checks after edits, so the comparison is of whole interfaces rather than tool count alone.

Context management came in five tiers. The lowest did nothing, so a run simply failed with an error once the history overflowed the window. Others used elision, swapping bulky old tool outputs for short placeholders; recall, which saves those outputs to disk and lets the model fetch them back; or summarisation, in which a separate call to the same model folds the oldest history into a running summary. The fullest tier combined all three, eliding first when the window was 60% full and summarising only at 85%.

The models were three sizes of Nvidia's Nemotron-3 family, at 30, 120 and 550 billion parameters, plus Mistral-Medium-3.5-128B as an outside comparison. The benchmarks were SWE-Bench Verified, 500 human-checked GitHub issues in Python projects, and Terminal-Bench 2.1, 89 command-line tasks. Context tiers were tested at windows of 32,000, 64,000, 96,000 and 128,000 tokens. Planning and tool choices were tested only at the largest window with full context management. Cost was estimated from published per-token prices, although the models were run on the authors' own hardware.

What it shows, and what it does not

The clearest result concerns context management. On SWE-Bench, averaged across models, managed tiers beat no management by about 36 percentage points at the smallest window, shrinking to under three points at the largest. The mechanism is direct: without management, most runs at the smallest window overflowed, while managed tiers never did. The 550-billion-parameter model, for instance, solved 6.4% of SWE-Bench tasks unmanaged at the smallest window and 55.6% with the full tier. The managed tiers performed about equally well, but the combined tier was the cheapest in seven of eight model and benchmark pairings. Recall added nothing measurable, and in more than half the settings where it was available the models never used it.

Planning and tools told a more mixed story. For the smallest model, planning lifted SWE-Bench success from 13.6% to 25.2%, apparently by keeping runs going long enough to reach an edit, though at higher cost. For stronger models it mainly trimmed cost, with small and statistically insignificant dips in success. Purpose-built tools helped the smallest model greatly, largely because in shell-only mode it kept issuing tool calls the harness could not run. The largest Nemotron did better and spent less with the shell alone on SWE-Bench, while Mistral preferred the tool menu on SWE-Bench and leaned towards the shell on Terminal-Bench.

The authors are explicit about the limits. These are specific implementations, not a recipe for an ideal harness. The tool comparison bundles several changes together. Planning and tool effects were measured in only one configuration, each setting was run once per task, and many Terminal-Bench differences are not statistically significant, so conclusions there rest on consistent directions across models and budgets. Only four models and two benchmarks were tested, and model size is an imperfect stand-in for capability, since training and familiarity with the shell may also play a part. The authors say the crossover points should be validated before being applied to other models, harnesses or tasks.

PAPER	*An Empirical Study of Harness Design for Coding Agents*
AUTHORS	Run-Ze Fan, Zihao Zhang, Simin Ma, Yebowen Hu, Shouju Wang, Kaiqiang Song, Fei Liu, Hamed Zamani, Xiaoyang Wang et al.
VENUE	arXiv preprint — not yet peer reviewed · 17 September 2026
LINK	arxiv.org/abs/2609.20804
INDEX	stacks.braintank.ai