BraintankAI ASIA

DISPATCH · CODING AGENT EVALUATION · 3 MIN

SHANGHAI · Fudan University

Coding agents meet most of a user's requests but rarely finish the whole job

A benchmark built from real user sessions finds the best of ten models passes an average of 69.9% of checklist items per task yet fully completes only 26.4% of web-app edits.

19 November 2026 · arXiv preprint, not yet peer reviewed

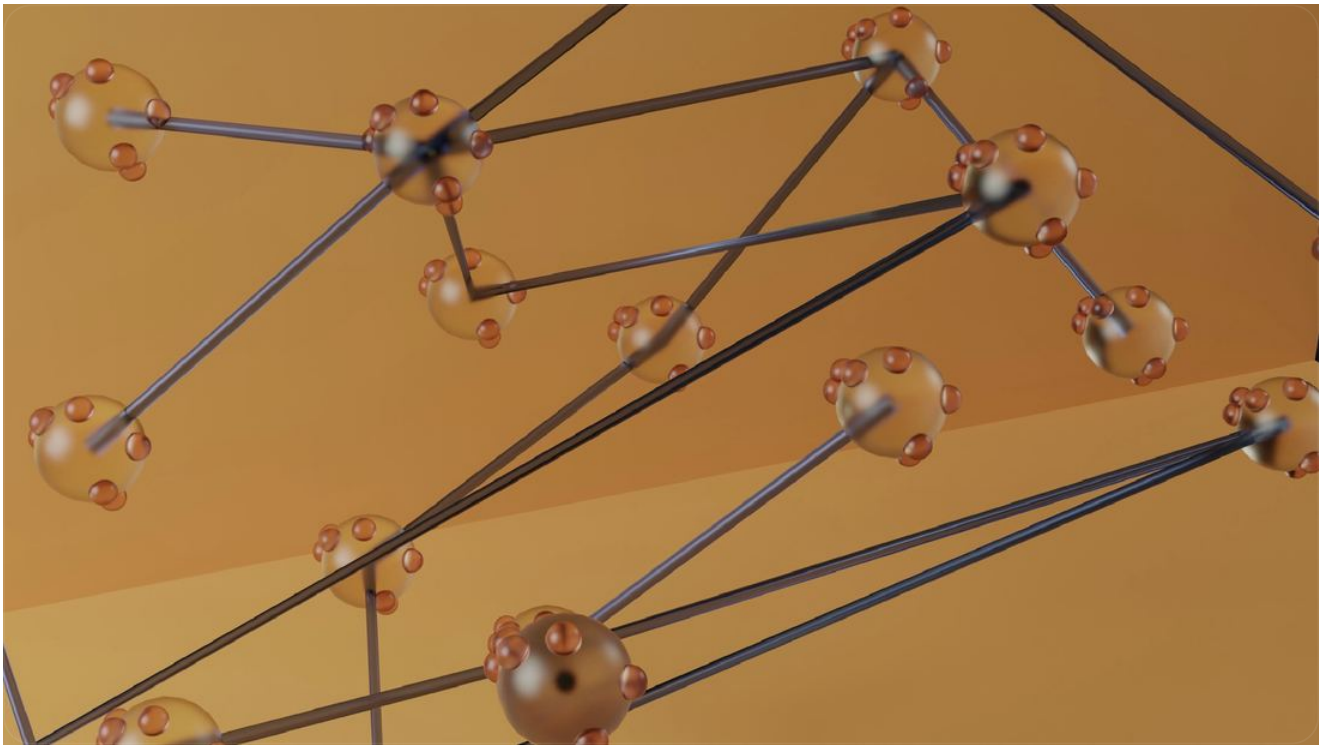


Photo: Guerrillabuzz

When someone asks a coding assistant to change a working web app, perhaps attaching a screenshot, a voice note or a video, there is rarely a single correct result to compare against. Co-Design Bench, a 2026 preprint that has not yet been peer reviewed, tries to grade such edits anyway. Its authors, who include researchers at Alibaba, argue that agentic execution changes how scores on older edit tests should be read. In agentic execution, models work in a loop of reading, editing and testing. On 150 tasks from an existing benchmark, WebCompass-edit, they report that agentic execution lifted qwen3.7-plus from 73.6% to 95.1% graded feature completion with the grading held fixed, and that opus-4-8 reached 99.4%.

The tasks come from 37,463 deduplicated queries across 26,314 real user coding sessions. For each task, the app as it stood before the request was rebuilt and had to build, launch in a container and render in Playwright, a tool that drives a web browser automatically. Tasks relying on services that

could not be restored were dropped. Requests were ranked partly by signs that the original user had struggled, such as rewording a request or rolling back a change, and experts wrote a hidden checklist of acceptance criteria for each task. The Multimodal Editing track has 100 localised edits with 553 criteria: 32 carry audio, 27 video, 17 documents, 4 images and 20 no attachment. The Repository Editing track has 173 multi-file changes averaging 10.3 criteria each, and 166 of them need more than one kind of edit.

Grading is reference-free, meaning there is no gold solution or target screenshot. An AI evaluator marks each criterion Pass, Partial or Fail from code changes, rendered pages and browser recordings. A separate AI verifier then challenges every positive verdict, hunting for stubs, hardcoded values, unreachable features and behaviour that breaks under varied inputs. Visual claims must be backed by screenshots and interaction claims by browser-action logs.

Ten model endpoints ran all 273 tasks inside the same Claude Code harness, capped at 3 hours and 120 turns. The authors report that gpt-6-astra led both tracks, passing 69.9% of criteria per task on average but fully completing only 26.4% of tasks. Its full-pass rate was 41.0% on multimodal edits and 17.9% on repository edits. Behind it on combined criterion scores came opus-5 at 0.661, qwen3.8-max-0902 at 0.653 and qwen3.8-27b at 0.644, while grok-4.6, at the bottom of the table, fully passed 6.0% and 6.9% of tasks on the two tracks. More effort did not guarantee better results. On multimodal edits glm-5.2 used a median of 69 turns and 7.29 million tokens to score 0.572, against gpt-6-astra's 0.754 from 32 turns and 1.24 million tokens. All ten models covered fewer criteria on repository edits spanning three or four change types than on those spanning one or two, 0.591 against 0.654 when pooled. Document tasks scored highest on criteria, at 0.652, yet were fully passed only 18.2% of the time, against 37.0% for text-only tasks.

The test has clear limits. Each task fixes one consolidated request and simulates no further feedback from the user, so it does not capture back-and-forth collaboration. The authors note that the refutation step targets unearned credit but does not by itself establish that every verdict is correct. Two experts checked 200 outputs, but the main text gives no agreement figure, and no repeated runs are reported. All headline results come from one harness, and the authors report strong interactions between model and harness. The image subset has only four tasks, only frontend web apps are covered, and the authors caution that qwen3.8-27b's edge over opus-5 on repository edits, 0.645 against 0.642, does not show a general advantage. Several authors are affiliated with Alibaba, whose Qwen models are among those tested.

The gap between high coverage and low completion echoes a benchmark of professional digital work on which the best agent setup fully passed 24% of public tasks (earlier story). The single-harness caveat also matters given earlier findings that the value of each piece of a coding agent's scaffolding depends on the model, the task and the budget (earlier story).

PAPER	*Co-Design Bench: Multimodal Collaborative Coding Evaluation for Agentic Models*
AUTHORS	Xilin Wei, Shengyuan Ding, Xinyu Fang, Tiancheng Gu, Yuchi Wang, Zhixun Li et al.
VENUE	arXiv preprint — not yet peer reviewed · 26 October 2026
LINK	arxiv.org/abs/2610.35138
INDEX	stacks.braintank.ai
CODE	github.com/Wiselnn570/Co-Design-Bench