

ac-grind.js — \$0 Local-Model Code Grinder (per-finding)

What it is

A wrapper around **aider** + the `autocoder.js` differential-tsc-gate logic that fixes **one finding/file per pass** using a **local \$0 model** (Ollama `qwen2.5-coder:32b-instruct-q8_0` @ 10.0.0.2:11434, or MLX `qwen3-coder-30b` @ 11437). Implements the measured recipe from `/tmp/evidence-mlx-autocoder-fix/verification.md`.

Tool: `~/system/tools/ac-grind.js` · **MC:** #103050 · **Status:** proven (2026-06-06)

The recipe (7 levers)

1. **Scope = 1 finding/file per pass** (broad batches choke the local model).
2. **Auto-attach relevant DEF files as read-only aider context** — biggest accuracy lever (measured 5→1). Resolves the target's `@/...` imports to real files via tsconfig paths.
3. **Qwen sampling:** temp 0.7 / top_p 0.8 / top_k 20 / rep_penalty 1.05.
4. **Repo-map (tree-sitter) + diff edits** (whole-file fallback for small files on aider 0.86.x).
5. **DIFFERENTIAL touched-file-only tsc gate** — repo is ~615 errors RED at baseline, so only NEW errors in the touched file count.
6. **Reflection cap ≥ 6** (aider 0.86.x caps internal loop at 3 → wrapper does $\text{ceil}(N/3)$ outer passes, stops on tsc-delta stall).
7. **E2E browser gate** (vite + Playwright + MSAL fixture) as the real DONE signal, not tsc-only.

Usage

```
node ~/system/tools/ac-grind.js --file <abs path> [--desc "what to fix"]  
  [--engine ollama|mlx] [--max-reflections 6]
```

Proof (MC #103050, EVVRecordsList.tsx, LumisCare)

- 8 DEF files auto-attached.
- **5 → 0 differential tsc errors** in one aider pass (independently confirmed: `npx tsc -b` clean on fixed file).
- Diff validated vs ground-truth APIs: `PaginationIndicator` (real export, props `page/totalPages/rowCount/pageSize/onPageChange`), `Badge size="sm"` (valid cva variant). Original `Pagination currentPage=` + `size={16}` were genuine type errors.
- **E2E Playwright PASS** (results.json expected:1 unexpected:0) + real screenshot — component mounts in a real browser without crash.
- \$0 (local model), no deploy, proof edit local-only & reverted.

Proven scope / honest limits

- Reliable for **mechanical type-error findings** (prop/API mismatches) where DEF files carry ground-truth types.
- NOT proven for **open-ended feature generation** (model must invent new logic).
- E2E gate currently proves **mount-not-crash**, not full data render — backend must run to prove real rows.

Hardening & real-batch results

Guard A — file-size guard

`DEFAULT_MAX_LINES = 450` (calibrated: proven-good 397 lines < 450 < proven-bad 585 lines). Files over the threshold emit a loud `SKIPPED_TOO_LARGE` box to stderr with exit code 2 — before any model call is made, so no tokens are wasted. CLI flags: `--max-lines N` to override, `--force-large` as escape hatch.

- `EVVRecordDetail.tsx` (586 lines) → exit 2, `SKIPPED_TOO_LARGE` (correctly blocked).
- `QualificationsTab.tsx` (397 lines) → allowed by default, flows to aider (proven-good path confirmed).

Guard B — artifact lint

Diff-scoped JSX-comment and whitespace-expression detection runs after every aider edit. Detects patterns such as `{" "}` whitespace expressions, `{/* narrative comments */}`, and inline narrating comments on JSX tags. On detection, ac-grind auto-strips and re-validates; if the stripped result still

fails, it reverts the file and exits with `ARTIFACT_REJECTED` (exit 3). Legitimate `{ " " }` (e.g. dash separators) are preserved because detection is diff-scoped, not file-scoped. Never ships dirty output.

- Pattern unit test: 12/12 PASS.
- Integration test (MC #103058 stash artifacts): 5 artifact lines detected, auto-stripped cleanly, legitimate separator preserved.

Real backlog sweep (MC #103063)

Ran across the backoffice module of LumisCare. 15 files edited by ac-grind sweep; 13 kept (clean error reduction), 4 reverted (quote-style-only changes, zero error reduction).

Metric	Value
Files edited by sweep	15
Kept clean (real error reduction)	13
Reverted (quote-style only, no improvement)	4
Baseline tsc errors (origin/dev)	612
After sweep	578
Net reduction	-34 errors (5.5% of backlog)
New errors introduced	0

Shipped as [PR #55](#) to `dev` (branch `ac-grind-sweep-103063`, commit `b17f606`), pending Proveo review. Not merged.

Honest takeaway: ac-grind reliably clears mechanical prop/API type-errors on files at or below 450 lines. Structural errors (e.g. `EVVRecordsList.tsx` — 5 errors untouched) and over-threshold files still require manual work.

Guard evidence

- `/tmp/evidence-103061/verification.md` — Guard A + B hardening + threshold recalibration to 450
- `/tmp/evidence-103063/verification.md` — real-batch sweep triage table + PR

Evidence

- `/tmp/evidence-103050/verification.md` (+ John independent verification section)
- `/tmp/verify-103050/evidence/` (build-tsc, test-e2e, playwright-results.json, screenshot)
- Recipe origin: `/tmp/evidence-mlx-autocoder-fix/verification.md`
- Related memo: `project_local_model_coding_setup_2026-06-06.md`

Revision #3

Created 2026-06-06 15:41:42 UTC by John

Updated 2026-06-06 20:42:16 UTC by John