06 · OPEN SOURCE · MIT · V3.0.0

AuthentiGen

A self-hostable app that gives CG renders and synthetic images the optical signature of a real camera — lens, sensor, colour and codec — for VFX compositing and computer-vision data augmentation.

ROLE Sole engineer, directing AI coding agents — first commit to public release in ~6 weeks

LINK github.com/yh12-10/authentigen

STACK TypeScript · Node 22 · sharp/libvips · React 19 · tRPC · MySQL/Drizzle · S3-compatible storage · Docker · GitHub Actions · Vitest

12	12/12	~1 s	2.2/s	$0	69
pipeline stages	byte-identical reruns	per 1–2 MP image	images at 4 concurrent jobs	per-image API cost	tests gated in CI

THE PROBLEM

CG renders look “too clean” because they never pass through a lens, sensor or codec. The first prototype sent every image to a generative model: an API call per job, different output on every run, and the risk of the model changing the content.

WHAT I BUILT

- **When not to use AI.** Within a day I replaced the generative model with a deterministic pixel pipeline: barrel distortion, edge-aware chromatic aberration (eased on skin), ISO-scaled grain, hot pixels, colour drift, halation, focus falloff and JPEG/EXIF.
- **Reproducible under concurrency.** Each job seeds its own generator from SHA-256(pipeline version ‖ intensity ‖ input bytes), so output is byte-identical across reruns and parallel jobs.
- **Safe with untrusted uploads.** Full-decode validation, MIME/format match, 12,000 px and 40 MP caps against decompression bombs, metadata stripping, and crash-safe job recovery.



Original render (crop)

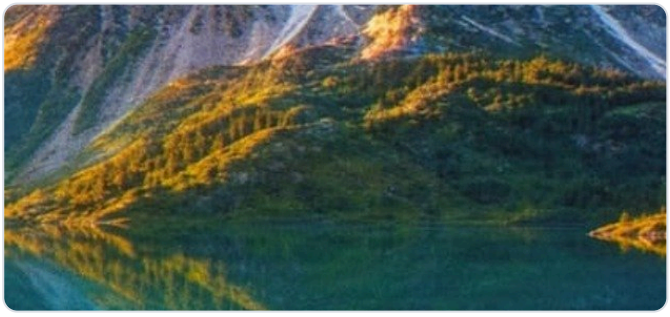

Heavy: real pipeline output, same crop

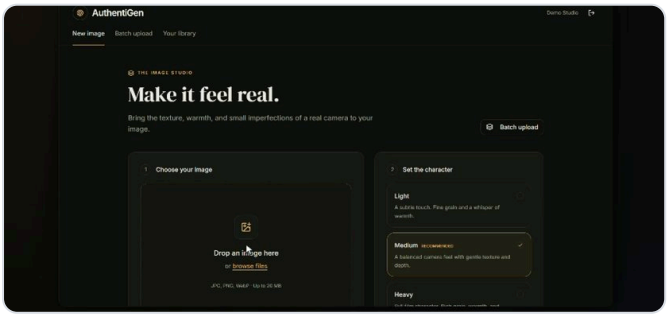

Studio: pick an intensity

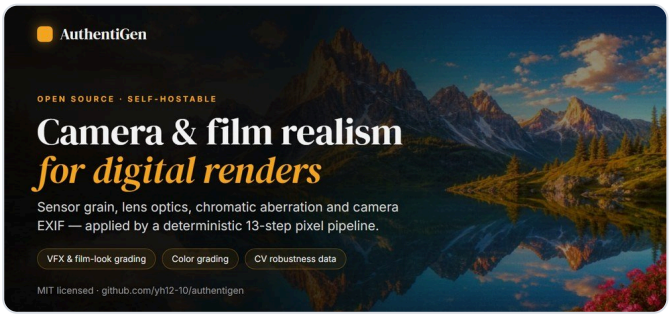

Release card