Alan Howard Butterworth

ai product builder and creative technologist bar, montenegro. remote, uk hours. freelance, one or two days a week.

curriculum vitae, november 2026

I was a portrait photographer in London for fifteen years, and since March 2026 I have been designing and shipping software with AI coding agents instead, working on my own. In that time two apps have shipped across five platforms, post for iPhone and Android and post pro for Mac, Windows and the web, in nine languages with paid subscriptions, a third is in App Store review, a fourth is in closed testing, and a newspaper has written itself for me every night. I don't write the code by hand. I decide what gets built, direct the agent that writes it, and then test and measure what comes back, which turns out to be most of the job. The rest is the eye, from fifteen years of photography and a background in cinema, which is why most of what I build makes pictures look better or move, on the newest image and video models.

shipped

Two apps, five platforms.

post, relight your photos

An AI photo app, on iPhone and Android at version 1.4.4 on both stores and in the browser, live since the spring. It reads what the scene was going for and relights it around the light that was already there, then hands it back graded and grained. It runs in nine languages on every surface, down to the store listings and the privacy policy, each read by two native speakers who didn't know about each other, and one subscription contract covers Apple, Google and Stripe and is checked on the server, because the three of them do not agree about anything. The trial images are held encrypted on the phone, so the server never stores a user's photograph. The hard parts were mine to specify and verify: putting full-resolution detail back after the model has worked at two megapixels, without undoing what the model deliberately changed, putting the original person back where the model has replaced a small face with a plausible stranger, and keeping HDR the right way up however the phone was held, with Apple's own ImageIO as the ground truth. It has been through App Store review several times, once backwards, and the paid campaigns on Meta and Google bought users at a price I would have taken.

post pro

Mac, version 1.4, Windows, version 1.1.1 on the Microsoft Store, the browser, and a Lightroom plug-in. The professional version for photographers who work in Lightroom or Capture One. It gives back a layered 16-bit PSD rather than a flat file, which meant a PSD and TIFF writer built from scratch and checked against Photoshop's own flatten, colour-managed rendering, the same pipeline running in WebGL in the browser, and a Windows port whose output matches the Mac pixel for pixel, which I know because a harness compares every layer and says when it doesn't. The Windows version took ten days from a standing start and is on the Microsoft Store in the same nine languages.

in review and in testing

nownotnow

iPhone, under review on the App Store. A to-do app for people whose sense of time has two settings, now and not now. You dictate one jumbled breath, a word comes back in about a second, and overnight a server works out tomorrow's short ordered list. It went from nothing to App Store submission in eight days, and nownotnow.co shows what it is while Apple makes up its mind. The model that runs it was chosen by an evaluation harness over ten candidates on accuracy and the cost of a night's work, and the transcription was chosen the same way, in my own voice, against Deepgram, OpenAI and Gemini. Anyone who would rather use their own Claude or ChatGPT can plug it in through an MCP connector.

Runt

iPhone and Apple Watch, in closed testing. A running coach with a voice, heart-rate monitors, auto-pause and a training plan worked out in advance rather than on the hoof. It looks nothing like the other two, which was the point, a third app in a third category so that the first two weren't a fluke of one domain.

not for sale

The two pieces of work I'd show first, and neither is a product. There's a reel of post hoc and the poster for the film on alanbutterworth.com, and there's a fuller package of either on request.

---

post cinema

A production platform for making a film with AI, which I'm using to direct a feature, and which will be in development for as long as the film is. It runs the whole length of the job, from the script as beats and notes, with research, looks, scouts and every option and choice logged against the scene they belong to, through previz in Blender driven by an agent, with motion-capture clips retargeted onto the characters and the previz fixing the timing for everything downstream, to the edit. A scene is one text file, one Blender file and one DaVinci Resolve timeline. A compiler turns the script, the location, a version-pinned style and the character references into a brief for every shot, with a lockfile, so a shot rendered again a month later comes out the same, and before any credit is spent a gate checks the first frame against the previz. Video comes from Kling, Seedance and Veo, but every voice and movement is grounded in a real performance, which is me playing every part with the voice swapped afterwards, because the capture is the one thing that doesn't improve when the models do. Sound is built in layers, music from a click track, the dialogue mixed in order, and the edit and the grade go together in Resolve by script. There's a Vision Pro app for scouting, in which you walk the virtual set, drop numbered markers where the cameras should go and talk, and the walk is transcribed and the markers become cameras in Blender.

---

post hoc

My own desk for live portrait sessions. The camera tethers in, the raws are developed into my look, the frames are ranked as they land, and an agent takes a starred frame, writes a render prompt, calls an image model and drops the render back on the board while the client is still in the chair. Behind it is the retouching project, which isn't finished and is the more interesting half. It starts with a dataset built from over a thousand of my own retouched PSDs, every healing stroke of fifteen years recovered, and a network trained on it that in a blind test passes as my retouch about half the time, with exposure and white-balance rules that turned out to beat my own average adjustment. The chain has been ported to Core ML and shrunk to fit on a phone without any measurable change in what it does, and the training goes on.

systems that run without me

morning post

A private newspaper for one reader, produced overnight by an editor model, a staff of writers, a sub-editor and an art department, and on the doormat every morning since 18 August 2026 without a miss. Each stage degrades gracefully when it fails rather than taking the edition down with it. Its style sheet names and bans the habits of machine prose, which I found by counting them in the paper's own back issues, and an edition costs about the price of a sandwich.

how i work

I measure before I choose. When a model choice matters I build a corpus, blind it, score it and pick on the evidence and the cost, and I have done that a dozen or so times across these projects, for vision models, image models, transcription and the lip-sync and voice conversion for the film. Every AI call carries a price, and the price goes into the decision rather than into a report afterwards. Anything that can go wrong in public has a kill switch on the server, a check before credit is spent, or a dry run that shows the fees before a listing goes live. When something is ported, iPhone to Android or Mac to Windows, the port is checked against the original's numbers. And every project carries a log of decisions with dates and reasons, so that a fresh session, human or otherwise, can pick up where the last one stopped.

also

Marketing machine

Generates photographs built to trap a naive photo app, runs them through the real pipeline, has an AI judge pick the winners and cuts the ads and store screenshots from them in eight languages, and then it ran the campaigns on Meta and Google.

---

echoframes

AI portrait sets of real clients, made from their own photographs and sold to returning headshot clients, with the model chosen in a head-to-head on likeness and every frame checked before it went out, for a fraction of a penny each.

---

eBay

When I was moving out of the flat the AI sold the contents for me, with a camera tethered to a watch folder, eBay's API, and a written negotiation policy it was allowed to run on its own. There wasn't a lot to it, but it sold a lot of things while I was elsewhere.

before 2026

Self-employed portrait photographer in London from 2011 to 2026 (head-shot.co.uk), a couple of thousand people over the years. Author of *The Trick* (2024), a satirical novel, which I also read for the audiobook. Directing an AI-produced feature film. BA Philosophy, 2:1.

tools, in plain words

Claude Code and the Claude Agent SDK are the working method. Swift and SwiftUI for iPhone, Mac, Watch and Vision Pro, Kotlin for Android, C# for Windows, and Node, Express and PostgreSQL on Railway behind them. Python for pipelines and model training. StoreKit, Play Billing and Stripe for the money, Core Image, Core ML and WebGL for the pictures. Claude, GPT, Gemini, DeepSeek, Seedream, Kling, Veo, ElevenLabs, Suno and Whisper in production, each one there because it won a test. MCP for letting users bring their own assistant, and Blender and DaVinci Resolve driven by script.

links

post: apps.apple.com/app/id6761730441 · play.google.com/store/apps/details?id=com.unthinkable.post · finishedinpost.com

post pro: finishedinpost.com/pro · apps.microsoft.com/detail/9NSV0S75QP4F

nownotnow: nownotnow.co · Runt: youbettarun.com

The Trick: amazon.co.uk, ISBN 1836542577