

Software I built for AI production

An image studio that checks and prices every prompt before it runs, a Python pipeline that turns Blender previs into film, and a publishing service where a person approves every post.

Michael James Owusu

Creative technologist, London · michaeljowusu@gmail.com

Open this section online ↗

Contents

Scan to open
michaelowusu.com/systems/



Software I built for AI production

01	Content Studio Engine Systems · AI image production	p. 3
02	AI commercial production pipeline Systems · AI video production	p. 5
03	Human-approved social publishing pipeline Systems · Content automation	p. 7

HOW TO USE THIS PDF

Contents entries jump to each section. **CASE PAGE** opens the full write-up online; **WATCH** opens the film or recording. Every section also has a QR code.

ALSO AVAILABLE

- [Work index \(one page\)](#)
- [Campaigns \(PDF\)](#)
- [Films \(PDF\)](#)
- [CV](#)

Campaigns shown are self-initiated concepts made at Mira Content Studio, my studio. No clients, results or metrics are claimed.

Content Studio Engine

A local studio and command-line engine for directing AI product photography. A brief and photographic recipe go in; linted, priced prompts produce takes; a person picks the winners; finals are versioned with their prompts, references and costs.

ROLE


Designed the workflow and the director rules. Built the CLI engine and its web studio with Claude Code and Codex. Ran the shoots and made every selection.

TOOLS

TypeScript React Express Vitest Gemini 3 Pro Image GPT Image 2

OUTCOMES

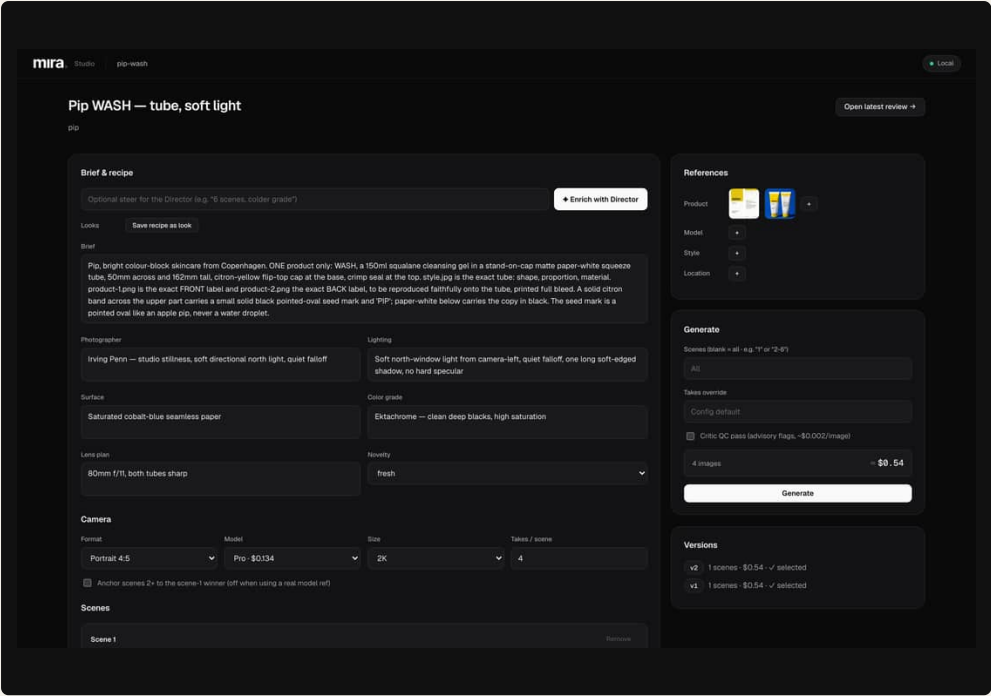
- Produced the stills for the four campaigns on this site, with every run recording its prompts, references and cost.
- A lint gate stops vague prompts before spend: a generic test prompt returned 13 errors in a dry run.
- Re-runs only pay for what changed: unselected scenes are inherited from the previous version.



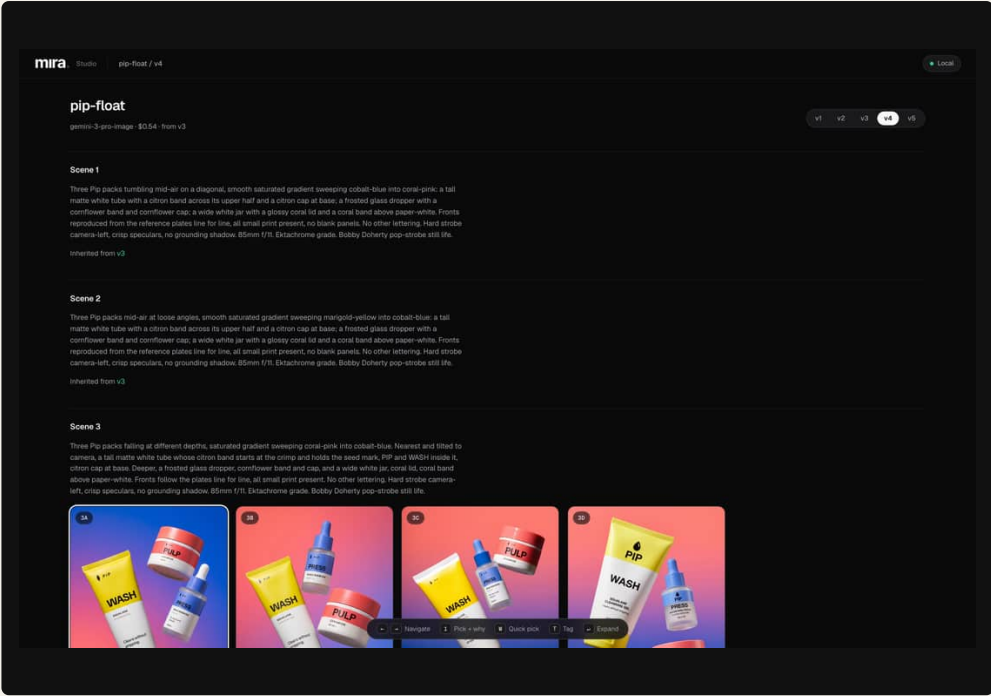
Case page ↗

▶ Watch the walkthrough





Project editor: recipe fields, the product's own references and a live estimate before Generate.



Version 4 re-rolled one scene as four takes; the other scenes were inherited at no cost.

```
$ npm start -- --project lint-demo --dry-run
x lint error scene 1 [no-lens]: no focal length (e.g. "85mm")
x lint error scene 1 [no-fstop]: no f-stop (e.g. "f/8")
x lint error scene 1 [no-light-direction]: no light direction
x lint error scene 1 [banned-generic]: "8k" pulls toward the generic-AI average
[... 9 more errors and 2 warnings]
i [DRY RUN] 13 lint error(s) - a real run will refuse these prompts
```

Real dry-run output for a deliberately generic test prompt (shortened). Dry runs make no API calls.

[Read the full case page ↗](#)

AI commercial production pipeline

A Python command-line pipeline that turns a storyboard into a short commercial. References are locked by checksum, Blender previs drives the motion, each video model has its own adapter, every paid request is previewed first, and FFmpeg assembles the cut.

ROLE


Specified and built it with Claude Code and Codex. Directed the LIDO and PIP films and approved each paid take.


TOOLS

Python Blender FFmpeg fal API Wan 3.0 Prime

OUTCOMES

- Made the LIDO film: five shots, all five first takes kept, about \$0.82 at listed 480p prices.
- Made the PIP film at 720p after Kling Motion Control rejected the grey-box previs and Wan 3.0 Prime followed it.
- No paid request is possible without --confirm-paid; previews make no uploads or API calls.

[Case page ↗](#)[▶ Watch the LIDO film](#)



The shared LIDO set in Blender: one layout for all five shots, keyframed camera paths.



Previs and final at the same moments.

One adapter per video model

ADAPTER	LIVE USE
Wan 3.0 Prime Reference	11 takes: LIDO 5, PIP 6
Kling 3 Pro Motion Control	1 request, rejected: no upper body in the grey-box previs
Kling 3 Pro · Seedance 2 · LTX 2.3	Priced preview only

Every paid request is previewed first: model, checksummed inputs, warnings and estimated cost, with no uploads.

What I'm improving next

- 01 **Check the previs before it costs anything.** Previs checks for pushes at the lens, objects crossing a face and actions that end off screen, before any paid take.
- 02 **Choose references from the locked library.** References suggested from locked assets, with a flag on any that carry text.
- 03 **Check the label on every take.** Each take's label compared with the studio pack shot before selection.
- 04 **Two takes per shot, compared side by side.** Two takes per shot by default, reviewed against the previs with their costs.
- 05 **Match each shot to the right model.** LTX-2.3 and Seedance 2.0 tested on the same shots, and the results used to route shots.
- 06 **Plan transitions with the shots.** Transitions and matching frames recorded in the shot plan and blocked together.

Full plan on the case page ↗

Human-approved social publishing pipeline

A Python service for a football content workflow. It finds stories, drafts captions, checks and cuts out real photos, and renders branded posts, carousels and reels. A person picks, shapes and approves every post in Telegram before it goes to Instagram.

ROLE

Designed and built it with Claude Code. I operate it and approve every post.

TOOLS

Python Telegram Bot API Instagram Graph API Pillow FFmpeg
Gemini 2.5 Flash BiRefNet on fal

OUTCOMES

- Nothing is drafted until a story is picked, and nothing is published without an explicit approval.
- Real photos pass five checks before they are offered; subjects are cut out, not generated.
- 342 offline tests pass with every network call stubbed.

[Case page ↗](#)

01

Collect stories

02

Pick one

PERSON

03

Draft and search

04

Shape it

PERSON

05

Build the post

06

Approve

PERSON

07

Publish

Three steps belong to a person. Nothing is drafted without a pick or published without approval.

Five checks before a photo is offered

01

Source filter

Watermarked stock-agency hosts are dropped before anything is downloaded.

02

Resolution

The short side must be at least 500 px. No thumbnails.

03

Vision check

A vision model confirms the right subject and face, and rejects watermarks and pre-made graphics.

04

Cut-out

A segmentation matte isolates the subject, who must fill at least half the frame height.

05

Crop

The frame is auto-cropped to 4:5 for the post canvas.

Output formats

FORMAT	SIZE	BUILT FROM
Single post	1080 × 1350	Fixed headline at the top, cut-out subject with rim light and shadow, grain, wordmark and accent colour
Carousel	1080 × 1350 slides	Photos cropped to 4:5 with a light grade and the wordmark
Reel	1080 × 1920	A pasted clip with a hook card, highlighted punch word and source credit, composed with FFmpeg

Limits that keep the operator in control

Open drafts. New menus pause while two drafts are waiting for a decision.

Unpicked stories. Return to the pool and can be offered again after 90 minutes.

Failures. A failed attempt returns the story to the menu instead of marking it used.

Keys. API keys live only on the server, never in the repository.

Tests. 342 passing offline, with every network call stubbed.

CONTACT

Michael James Owusu

Creative technologist, London. I direct AI product campaigns and short films, and build the software I make them with.

michaeljowusu@gmail.com

linkedin.com/in/michaeljowusu

github.com/mirasolutions06

michaelowusu.com

Download CV

Work index

Scan for the portfolio
michaelowusu.com

