



DOCUMENT TEMPLATE

AI UGC & Avatar Video

What this guide helps you create — and who it's for

A UGC video brief — written as a structured JSON prompt you paste straight into an AI UGC or avatar generator (Higgsfield, HeyGen, Synthesia, Arcads, or whatever's current) to get the clip out. The brief pins the one job the clip does, the format, the script with the hook front-loaded, the avatar/voice/visual spec, and the two or three variations to test.

Use this if you want short-form video — hooks, demos, testimonial-style clips — without a film crew, and you keep posting either nothing or one unwatched explainer. **Skip it** if you've already got a UGC engine running and you just need more volume — that's the deeper "UGC engine" blueprint at the end, not this guide.

HOW TO USE THIS GUIDE

- 1 **Read sections 02–04** so you know what a short clip is *for* — and why the first 1.5 seconds decide everything.
- 2 **Open your AI chatbot** — the one you set up in guide 01.
- 3 **Make the decisions in section 03** — the one job, the format, the hook, the one CTA. Do this *before* you touch the prompt.
- 4 **Copy the prompt in section 05** into the chatbot. Paste your Founder Context Block where it says to, and fill in the `{curly braces}`. It gives you back a JSON brief.
- 5 **Paste that JSON into your UGC tool** (Higgsfield / HeyGen / Synthesia / etc.) and generate. Then sharpen using the four moves in section 06 — one at a time.
- 6 **Post all the variations.** One clip isn't a test. Three hooks, same point, see which one travels.

Why this document matters

The AI tools — Higgsfield, HeyGen, Synthesia and the rest — will happily turn a sloppy idea into a slick, useless clip. The bottleneck isn't the rendering; it's the brief. Most founder UGC fails the same way: it opens with sixty seconds of context before anything happens (dead on social); it starts "Hi, I'm Sam from Ledgerwise" (nobody waited); the avatar and voice are generic enough to scream "AI ad" (uncanny, monotone, stock); it crams the whole pitch into one clip; and it ships *one* version and calls that "testing."

A UGC brief fixes the input. It decides the one job and the one format, writes the hook by hand, caps the script to fit the length, specs an avatar and a voice that feel like a real person, and produces a handful of variations. The output is a clean JSON prompt the tool can actually execute — and a clip worth posting.

The tool isn't the bottleneck; the brief is. A slick clip from a sloppy idea is just a faster way to make nothing.

It comes from four decisions. Section 03.

The thinking framework — what to settle before you prompt

The prompt writes the JSON brief. These calls are yours.

— What's the one job?

A clip does *one* job: a top-of-funnel hook (scroll-stopper, gets you noticed), social proof (testimonial-style, makes a wavering buyer believe), or a demo (shows the thing working). The job changes the format, the script, and the metric. A clip trying to do all three does none. Pick the one.

— What’s the format?

Pick one: an avatar talking-head; a voiceover over b-roll or screen capture; a hook plus a quick screen demo; a “POV / relatable” near-skit. The format follows the job — hooks lean talking-head or relatable, demos lean screen capture, proof leans testimonial-style. Don’t mix four formats into one clip.

— What’s the hook?

The first 1.5 seconds decide whether the clip is watched. Write the hook by hand, before you touch the prompt — and make it *specific*: a sharp question (“Still got a glovebox full of receipts?”), a contrarian claim (“Your accountant costs more than your van insurance”), a “stop doing X” (“Stop doing your books at 11pm in January”), or a hyper-specific callout of the exact person. “Discover how Ledgerwise can transform your bookkeeping” is not a hook.

— What’s the one CTA?

One action, said out loud and shown on screen, at the very end — “Try it free, link in bio.” Not a call. Not a contract. Not three options. If the clip earns a second of attention at the end, spend it on the single thing you want them to do.

DECIDE	WEAK DEFAULT	WHAT GOOD LOOKS LIKE
The one job	“explain what we do”	“top-of-funnel hook — get a trades sole trader to look”
The format	“an explainer video”	“avatar talking-head + a 3-second screen-grab of the app”
The hook	“Discover how we can transform your...”	“Still got a glovebox full of receipts? January’s coming.”
The one CTA	“learn more / book a demo / sign up”	“Try it free — link in bio.” (spoken and on screen)

04

Common failure modes — what generic output looks like

Ask an AI for “a UGC video script” cold, and you get the generic ad it’s seen a million times. Reject these:

× WHAT TO REJECT ON SIGHT

- **The explainer.** Sixty seconds of background before anything happens. On a feed, you've lost them in two.
- **No hook.** It opens with "Hi, I'm [name] from [company]" or "In today's world...". Nobody waited to find out who you are.
- **The uncanny spokesperson.** A generic avatar and a flat, mid-Atlantic voice that reads as "AI ad" instantly. People scroll past "ad" faster than anything.
- **Everything in one clip.** The full value prop, every feature, three CTAs — crammed into 30 seconds. One clip, one point.
- **One version.** You generated a single clip and called it "testing." Generate three hooks, same point — *that's* a test.

If the script looks like this, you haven't decided the job or written the hook by hand. Section 03.

05

The prompt template

This is the bit that does the work. You don't need to understand prompting — you just need to do these five things, in order:

RUNNING THIS PROMPT — STEP BY STEP

- 1 **Copy the entire grey box below** (all of it, top to bottom).
- 2 **Open your AI chatbot** and paste it into the message box. Don't press send yet.
- 3 **Paste your Founder Context Block** where the prompt says `{paste your Founder Context Block here}` — delete that line and put your block in its place. (No Context Block yet? Build one in guide 01.)
- 4 **Replace every `{curly-brace}` bit** with your own answer — your four decisions from section 03 go straight in. If you're unsure of one, leave it; the prompt will ask.
- 5 **Send it.** It gives you back a JSON brief — copy that, paste it into your UGC tool (Higgsfield, HeyGen, Synthesia, etc.), and generate. Then go to section 06 to sharpen.

PROMPT TEMPLATE

→ Copy the whole box. Paste your **Founder Context Block** where it says to, then swap everything in {curly braces} for your own answers before you send it.

You are helping me create a short-form UGC video brief – and I want it back as a structured JSON object I can paste straight into an AI UGC / avatar generator (Higgsfield, HeyGen, Synthesia, Arcads, or similar).

Here is everything you need to know about my company:

{paste your Founder Context Block here}

Four things I've decided (use these – don't override them):

- The one job of this clip – one of: top-of-funnel hook / social proof / product demo: {e.g. "top-of-funnel hook on a trades feed"}
- The format – one of: avatar talking-head / voiceover over b-roll or screen capture / hook + screen demo / POV-relatable: {e.g. "avatar talking-head + 3s screen capture"}
- The hook (the first ~1.5 seconds, written by hand): {e.g. "Still got a glovebox full of receipts? January's coming."}
- The one CTA, spoken and on screen: {e.g. "Try it free – link in bio. No card."}
- The platform and length: {e.g. "Reels/TikTok, ~20s,

9:16"}

Return ONLY a JSON object with these fields:

job, platform, aspect_ratio, target_length_seconds, format, hook, script – an array of beats, each tagged with a rough timestamp range; the FIRST beat is the hook, the LAST is the CTA; nothing before the hook (no logo, no "hi I'm"); on_screen_text – short overlay phrases, cta, avatar – { look, age_range, vibe } – a real-feeling person from my audience, NOT a generic "spokesperson", voice – { tone, pace, accent }, b_roll_or_visuals, music_vibe, variations – 3 alternative hook lines, same point, different angle, to test.

Rules:

- People speak ~2.5 words/second. The script must fit the target length – for a {N}-second clip, that's roughly {N × 2.5} words total. If it doesn't fit, cut the middle, never the hook or the CTA.
- The hook is the first thing on screen. No intro, no "hi I'm [name]", no "in today's world".
- One point per clip, one CTA. If the brief is trying to

do two things, tell me and pick one.

- Anything in `{curly braces}` is a slot for me to fill. If I left one blank, ask me for it – don't invent it.

After the JSON, in plain text: tell me which of the 3 hook variations you'd post first, and why.

06

How to iterate when the output is weak

First drafts come back too long and too generic. These four moves fix that — one at a time, re-generating the clip after each:

1. **Kill everything before the hook.** *“Cut anything that happens before the hook line. The clip starts on the hook — no logo, no ‘hi l’m’, no setup.”*
2. **Sharpen the hook.** *“Rewrite the hook to name the exact person and the exact moment — a glovebox of receipts, an 11pm January deadline. Vague gets scrolled.”*
3. **Cut the script to the length.** *“People speak about 2.5 words a second. For a {N}-second clip that’s roughly {N×2.5} words. Cut the script to fit, keep the hook and the CTA, lose the middle.”*
4. **De-genericise the avatar and voice.** *“Re-spec the avatar and voice to a real-feeling person from my audience — a tradesperson in their 30s, plain UK accent, slightly dry, not a spokesperson. And give me three hook variations, same point, to test.”*

If two passes don't fix it, you haven't decided the one job — back to section 03.

07

A worked example

Sam from guide 01, Ledgerwise — a 20-second hook clip for a trades feed.

> WORKED EXAMPLE – SAM’S HOOK CLIP, AS A JSON BRIEF

Job: top-of-funnel hook on a trades-adjacent feed (Reels / TikTok).

Format: avatar talking-head + a 3-second screen-grab of the receipt-photo flow. **Hook:** “Still got a glovebox full of receipts? January’s coming.”

CTA: “Try it free — link in bio.”

The AI returns roughly this, which Sam pastes into Higgsfield / HeyGen:

```
{ "job": "top-of-funnel hook", "platform": "reels/tiktok", "aspect_ratio": "9:16", "target_length_seconds": 20, "format": "avatar talking-head + 3s screen capture", "hook": "Still got a glovebox full of receipts? January's coming.", "script": [ "0-2s HOOK: 'Still got a glovebox full of receipts? January's coming.'", "2-9s 'You're not behind because you're bad at this – you're behind because nobody does their books in a van.'", "9-15s 'So you photograph the receipt – that's it. We sort it, file it, and your return's basically done.' [cut to screen-grab of the app filing a receipt]", "15-20s CTA: 'Try it free – link in bio. No card.' " ], "on_screen_text": [ "GLOVEBOX → FILED", "free for 30 days, no card" ], "cta": "Try it free – link in bio", "avatar": { "look": "relatable UK tradesperson", "age_range": "30s", "vibe": "down-to-earth, not a spokesperson" }, "voice": { "tone": "plain, slightly dry, not salesy", "pace": "natural", "accent": "UK" }, "b_roll_or_visuals": "phone photographing a crumpled receipt; the app categorising and filing it", "music_vibe": "minimal, low, no stock-corporate energy", "variations": [ "Still got a glovebox full of receipts? January's coming.", "Doing your own books? You're probably more behind than you think.", "Your accountant's £200 a month. This isn't – and it does the boring bit." ] }
```

Notice what makes it work: the hook is the *first* line, the script fits 20 seconds (~45 words), the avatar is “a tradesperson,” not “a spokesperson,” and there are three hooks to test — not one clip called a test.

08

Edge cases — and when to go deeper

A few situations need more than the template:

- **You want volume on a schedule**, not one clip — a pipeline that turns a content brief into a steady stream of UGC. That’s the deeper engine:

HeyGen AI UGC Engine — Spin up volumes of avatar-led UGC video from a single content brief, on a schedule.

n-access.co.uk/playbook/heygen-ai-ugc

- **You're shooting real footage, not AI-generated.** This guide still applies for the brief, the hook and the script — just drop the avatar/voice spec and hand the script to whoever's on camera.
- **Disclosure and honesty.** An AI avatar that's framed as a real customer giving a testimonial is a problem — legally and reputationally. Keep "your brand says this" clearly distinct from "a customer says this" unless it actually is one, and follow the platform's AI-disclosure rules.
- **The hook is your wedge, compressed.** If you can't write a sharp hook, the issue is usually upstream — your positioning isn't pointed enough. Sharpen it in guide 03 first.

09

What to read next

- **10 · Social Media Content System** — where these clips go, on a cadence, alongside your other post formats.
- **05 · Content Strategy** — UGC video as a theme in your wider content engine, not a one-off.
- **03 · Brand Positioning & Messaging** — the hook is your wedge in five words; if the hooks keep falling flat, the positioning's the problem.

The Founder's Playbook

A library of short operator guides for founders building with AI — not prompt-pack content, an actual method. Each guide is free, each routes you to deeper blueprints when you need them, and each assumes one thing: that you've built your **Founder Context Block** (that's guide 01).

The Founder's Playbook · n-access.co.uk/playbook