AnotherSystems · anothersystems.com

How AI agents make our Instagram reels

We automated our Instagram reels with Claude Code and a small team of AI agents. After 10 days of reels, we hit 1,000 followers. This is the whole system, so you can copy it.

It's built for informative videos like ours, topics, facts and guides told with 2D cards, 3D scenes, an AI voice and captions. It isn't made for animated cartoons or comedy sketches.



1,000	10	60+	3
followers	days of reels	reels made	AI reviewers per reel

Don't read it alone. Give it to your agent.

Drop this PDF into Claude Code or any AI agent that can read files and run commands, with a screen recording (phone or computer) of a video in the style you want, and say "Follow this guide." It sets everything up and walks you through your first reel.

ONE REEL, FROM IDEA TO READY TO POST

Topic
Something famous in your niche. Ours is AI, software and data, so ChatGPT, Netflix and Google Maps come up.

Research
Only the makers' own pages and papers. Every number gets a source.

Script
Up to 10 lines, about 60 seconds, a famous name in the first words.

Three AI reviewers
Facts, a new viewer and design, before anything is built.

You read the script — Optional
Ask the agent to stop here, then approve the script or change it.

Build
The voice, the 3D scenes, the music, every frame, a cover and the caption.

Audit loop — Up to 5 rounds
Frames and caption checked, fixed and checked again.

You watch the reel
Approve it and post it yourself, or ask for changes and the agents redo it.

Your numbers
After posting, read the reel's insights and give the agent the numbers.

For the AI agent reading this

Your user gave you this guide so you can set up its system for them. Follow the last section, "Start your first reel". Create the files from "The files" at the end, exactly as written, in the current folder. Ask your user for a screen recording of a video in the style they want (from their phone or computer), keep it in reference/ and match its look. This system makes informative videos (topics, facts and guides in 2D and 3D), not animated cartoons or comedy sketches, so say so if the user wants something else. Then ask for their account's niche, their style rules and their first reel idea, and whether they want to check each script before the build. Run that reel up to the script stage and stop there for their review. Never post anything, the user posts by hand.

This is the whole system behind @another_systems, from the idea to a reel ready to post, written so you can copy it. Every rule here comes from a real reel, and most of them from a reel we had to redo.

It's built for informative videos like ours, topics, facts and guides told with 2D cards, 3D scenes, an AI voice and captions. It isn't made for animated cartoons or comedy sketches. The guide gives you the method, the rules and every brief, but not our video code. Your agent writes its own from your reference video, so the first reel takes the longest.

The PDF ends with the briefs and a starting playbook as files, so your agent can use them as they are.

1 What you need

Claude Code, Anthropic's AI coding tool. Claude's $20 a month Pro plan is enough to make about 10 to 20 reels a month with no problems. Then Node.js, a small React app with three.js that draws each reel as a web page, headless Chrome to record it, and FFmpeg to turn the frames into a video.

The voice and the word check run on the laptop. We use Alibaba's Qwen3-TTS for the voice and NVIDIA's Parakeet to hear it back, both through mlx-audio on a MacBook Pro with an M1 Pro chip. The voice step also times every word with Qwen3-ForcedAligner, and those times light each caption word and cue every piece of the scene. On Windows or Linux, any local text to speech, speech to text and word timing models can do the same jobs. The music is made by code too, so there is no track to license.

A reference video. Play a reel in the style you want and record your screen, on your phone or your computer, then give the recording to your agent. Until you have approved reels of your own, it is the picture of done for every build.

And an Instagram professional account, which shows each reel's insights.



Claude Code and the five tools.

Tool	Its job
Claude Code	**Runs every agent. The only paid part**
three.js	**Draws the 3D scenes in a web page**
Qwen3-TTS	**The AI voice and its word times, on the laptop**
Parakeet	**Hears every line back to catch a cut word**
Headless Chrome	**Records the page, frame by frame**
FFmpeg	**Joins frames, voice and music into the video**

2 The folder

Everything lives in one folder that the agents read and write. The names are ours, and yours can be anything.

Each reel's script folder keeps its research, a sources file with the page behind every number, the reviews and a decision log. When a choice is made, the log says why, so a later agent never undoes it by accident.

In the folder	What it holds
playbook.md	**Every rule, each with its date and the note that made it**
builder_brief.md	**What a builder agent gets for one reel**
reviewer_briefs.md	**The four review roles, one brief each**
queue.txt	**Reel ideas as tickets, their owner and stage**
catalog.md	**Every reel made, its lines, sound and numbers**
posts/	**One folder per reel, the video, cover, caption and script**
src/	**The code, one page per reel with its timeline and voice lines**

3 The playbook

The playbook is the heart of the system. It is the list of rules every agent reads before it starts, and it grows with every review. Ours now covers the hook, the length, the captions, the 3D look, the sound and the caption under the reel.

Start yours with rules like these, then let your notes change them. Every rule carries the date and the note that made it. Where two rules disagree, the newer one wins.

A rule to start with	Why
Something famous in your niche in the first words	**About half our viewers were gone by 3 seconds**
The first frame is a full picture, its words already lit	**An empty opening gives them no reason to stay**
Up to 10 spoken lines, about 60 seconds	**Our reels kept running long**
Every word said is on screen, nothing before its word	**Eyes and ears tell one story**
Every number from a primary source	**One wrong fact costs trust**
Each line explains its own terms	**A non-engineer should follow every line**
3D for real objects, glass cards for anything read	**Text stays sharp, objects feel real**
Every gap 16 px, every card padded 20 px	**Uneven gaps look cheap**
Dim by color, never by transparency	**Faded pieces look broken**
A fresh music mix for each reel, picked by its mood	**Repeats make every reel sound the same**

4 The queue and the agents

Ideas wait in the queue as tickets. A ticket names the topic, the hook, the maker's own page that proves the story, and who owns it. A topic without that page never becomes a ticket. One of ours turned out to be a reseller's listing, not a real release.

The main Claude Code session is the coordinator. It claims a ticket and starts one builder agent on Claude Opus with the builder brief. At the script stage it starts the three Claude Sonnet reviewers at the same time, and after the first render, the frame auditor. That is 5 agents per reel, and never more than 3 running at once. More only burns tokens.

One builder per reel. It writes a checkpoint into the script at every stage, so if it stalls, a new builder picks up from the checkpoint instead of starting over.

LESSON 1

Give every agent a written brief and one file to write. The brief is all the context it gets.

5 The builder

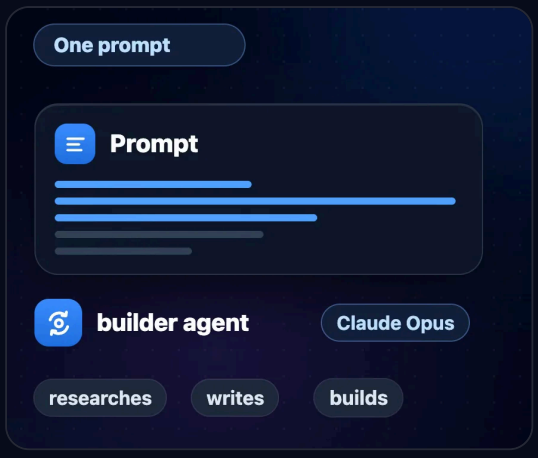
The builder gets the topic, the reel's name and its brief, and works in three stages. It stops and reports to the coordinator after each one. It copies the look of your reference video, or of your own approved reels once you have some.

The script. Research on primary sources, saved with the sources file. Then the lines, a plan for every line of what the screen shows, 3D or glass for each piece, and the decision log. Nothing is voiced yet.

The build. It reads the three reviews, decides each point and logs why. It voices every line and hears each one back, builds the scenes, checks the key frames, makes the music, the mix, the cover and the video, and hears the finished video once more. It writes the Instagram caption too, so the reel is ready to post. Then it pulls about 24 frames at the key moments for the auditor.

The finish. It fixes every blocking item from the audit, renders again and runs the audit loop. When a round comes back clean, it writes the reel into the catalog.

The brief also carries the house rules, the short version of the playbook. Ours include no em dashes, no colons in the caption, the newest examples from the makers' own pages, never saying why a reel was made, and never posting anything on its own.



One prompt starts the builder.

6 The three reviewers

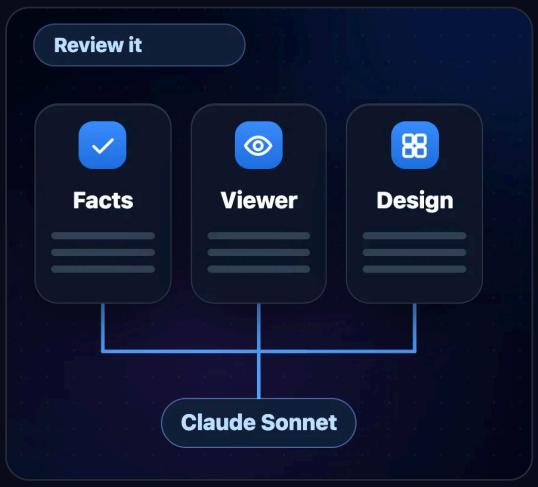
The reviews happen on the script, before a single frame exists. A fix there costs one line. The same fix after the build costs a render.

The fact checker checks every claim and every planned label against primary sources, on the web itself. It looks for missing qualifiers like "about" or "up to", something true of one model said of another, and examples older than they need to be. Each problem comes with the source's quote and the exact new words.

The new viewer reads as a non-engineer scrolling past, deciding in 2 seconds. Does each line name its own subject? Does "it" point at the right thing? Is the hook famous enough? It counts the lines and words, names what to cut, and drafts the caption.

The visual designer checks the plan against the code. Every 3D object stays whole on screen, glass for everything read, nothing shown before its word, room around every card. It writes a build spec for each line.

One real catch. A draft line said a 2018 study of over 17,000 Android apps found none sending out audio in its tests. The fact checker read the paper and found that only about 9,100 of them ran in the tests. The line became "A 2018 study tested about 9,000 Android apps and found none sending out audio" before the voice was made.



Three reviewers on Claude Sonnet.

7 The script check, if you want it

By default the build starts right after the reviews. If you want a say earlier, tell the coordinator to stop there, once for one reel or as a rule in the playbook for every reel.

It then shows you the reviewed script, the lines and the plan for each scene. Approve it, or change lines in your own words, and the builder starts only on your yes. A changed line costs nothing here, while the same change after the build costs a new voice and a new render.

8 The audit loop

After the first render, one auditor gets about 24 frames, the cover, the lines and the caption, and does two jobs. As a viewer, it looks for anything clipped, crowded, too small for a phone, or shown but never said. As a fact checker, it checks every number again, everywhere it appears.

Each problem is marked blocking or minor, with an exact fix. The builder fixes them, then audits the whole reel again, until a round finds nothing, at most 5 rounds. If round 5 still finds something, the reel waits for a person.

The auditor misses taste. It passed tangled tubes we called spaghetti and a plain gate we called lame. So it now compares every 3D object with a sheet of objects we already approved, and anything that looks like stacked shapes fails. Scenes that cut away too soon fail too. Each scene now stays at least 1 second after its last line ends.

LESSON 2

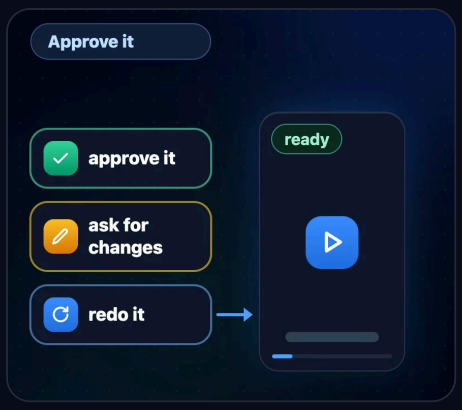
An audit checks rules, and your eye still checks taste. Turn each thing you catch into a rule the audit can check next time.

9 Your review

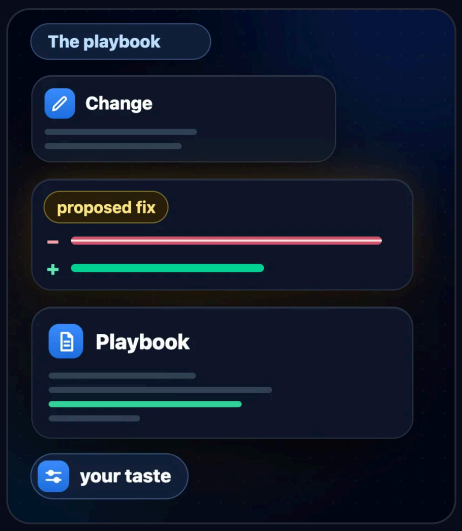
You watch every finished reel before it goes out. Approve it and post it yourself, or write what to change in your own words.

Each note does two things. The agents fix the reel, and they write the rule your note implies into the playbook, with the date and your words, and tell you what changed. You keep the rule or undo it.

A note that one scene cut away too fast became the 1 second rule. A note that faded cards looked wrong became dim by color. Neither has come up again.



Approve it, or ask for changes and the agents redo it.



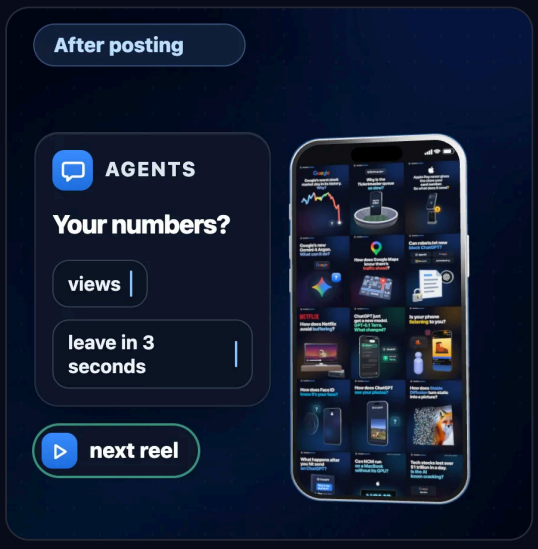
Each change becomes a proposed fix to the playbook.

10 Your numbers

For now this step is by hand. Every few days, open the insights of the newest reels in the Instagram app and give the agent the numbers, typed or as screenshots. Each reel's chart of how long people watched gives the share still watching at 1, 3, 5, 10, 30 and 60 seconds. The agent logs them in the catalog, next to the reel's topic, opening and length.

Look for patterns, not single reels. Ours showed most viewers leaving in the first 3 to 5 seconds, and the curve going nearly flat after about 10. The two best openings put a famous name on screen first. Everyday topics held about twice as long as AI news.

The same table shows which reels won people over on their own. Those are the only ones we put ads on.



After you post, the agents ask for your numbers.

Our reel	Watching at 3 s	At 10 s	At 30 s
How Google Maps sees traffic	52%	28%	23%
How Apple Pay works	50%	32%	21%
The Ticketmaster queue	36%	26%	22%
ChatGPT, after you hit send	47%	29%	18%
Gemini 4 Argon news	38%	20%	15%
GPT-6.1 Terra news	33%	16%	11%

THE NUMBERS LOOP, BY HAND FOR NOW



11 What cost us reworks

Every rule in the playbook started as a mistake. These were the expensive ones.

What went wrong	The rule now
Reels ran 62 to 70 seconds	**Up to 10 lines and 60 seconds**
Tangled tubes and a plain gate passed the audit	**Each 3D object is compared with approved ones**
Builders stalled halfway	**A checkpoint at every stage**
A reseller listing almost became a reel	**A maker's own page before any ticket**
Uneven gaps between cards	**One spacing system, 16 and 20 px**
Scenes cut away too fast	**Each scene stays 1 second after its last line**

12 Start your first reel

Make an empty folder and put this guide's PDF in it, with a screen recording of a video in the style you want. Open Claude Code there, or any AI agent that can read files and run commands, and paste this prompt. The agent takes it from there.

"Read the guide PDF in this folder. Set up its system in this folder. Write playbook.md from its starting playbook, builder_brief.md and reviewer_briefs.md from its briefs, and an empty queue.txt and catalog.md. Use the screen recording in this folder as the reference for the look. Then ask me for my account's niche, my style rules and my first reel idea, and run that reel up to the script stage. Stop there for my review."

More on this system is coming soon, so stay tuned.

The point

Copy the folder, not our taste. The playbook, the briefs and the numbers table are the system. Your notes become its rules and your numbers pick its topics, so after a few weeks the reels are yours.

The files

Your agent creates these from here as they are, each under its name in your folder. Then your notes change the rules to your own.

playbook.md

```
# Playbook

The rules every agent reads before it works on a reel. It is a living
guide. When the owner reviews a reel and asks for a change, the agent fixes
the reel and writes the rule the note implies here, with the date and the
owner's own words, then tells the owner what changed. Where two rules
disagree, the newer dated one wins. Replace these starting rules with your
own as your notes come in.

## The account

- The kind of video. This system makes informative videos, topics, facts
  and guides, told with 2D glass cards, 3D scenes, an AI voice and
  captions. It isn't made for animated cartoons or comedy sketches.
- The reference. A screen recording of a video in the style the owner
  wants, recorded on a phone or a computer, kept in reference/. Until the
  owner has approved reels of their own, every agent matches its look.
- Topics. Write one line on what the account covers and for whom. Ours is
  AI, software and data analysis explainers for people who use them every
  day.
- Everyday "how does it work" topics (a maps app's traffic, paying by
  phone, a ticket queue) held about twice as long as AI news on our
  account. Pitch those first, and AI news only when the hook names what
  changes for the viewer.
- A topic needs its maker's own page (a release post, a paper, official
  docs) before it becomes a ticket. A reseller's listing or a news recap
  is not a source.

## The hook and the first frame

- The first words name something famous in your niche that nearly every
  viewer knows or has lived (for us, in AI, software and data, a famous
  app, phone, company or moment like ChatGPT, Netflix or Google Maps),
  never a model or method few know. The technical name comes inside the
  reel.
- Frame 0 is a full, readable picture of that famous thing with the first
  caption already up and lit. About half of our viewers were gone by 3
  seconds, most of them over openings with nothing to look at yet.
- The cover names the topic's keyword and asks the question. It never
  gives away the answer.

## Length and lines

- At most 10 spoken lines and 60 seconds, plus a 0.5 second hold on the
  last frame. Cut lines until it fits, a fact before a name's explanation.
- Each line names its own subject and explains every term it uses, for a
  non-engineer reading English as a second language.
- "It" always points at the thing just named.
- A process is told in its real order (the review before posting, the
  numbers after it), and a loop is said as a loop.
```

```
## The screen

- Everything said is on screen as it is said, and everything on screen is
  said. Nothing shows before its word.
- Every spoken line is a big caption near the bottom whose words light up
  as they are said.
- 3D for real objects (phones, servers, cards). Flat glass cards for
  everything read (labels, numbers, status pills).
- Every 3D object stays whole inside the frame at every moment. Objects
  pop or fade in place, never fly in from outside.
- Every 3D object must look as finished as the approved objects sheet, a
  contact sheet of the objects in the owner's approved reels. Until it
  exists, compare with the reference recording. Stacked primitives, wire
  frames and tangled tubes fail.
- One spacing system. Every gap between cards is 16 px and every card's
  padding is 20 px, from shared constants, never a number typed per card.
- Breathing room. Few things per card, no pill against an edge, a header
  line or another pill. Cut or move a piece rather than squeeze it.
- Dim by color. An unlit or inactive piece darkens its text, icon and
  border at full opacity, never by transparency.
- Connector lines run behind the tiles and only in the gaps between them,
  never across a tile.
- Every scene is seen whole. It stays at least 1 second after the line
  that brings its last piece ends.
- Keep the areas Instagram covers clear (the top bar, the caption and the
  buttons on the right and bottom).

## Facts

- Every number has a primary source, listed in the reel's sources file.
- Every qualifier ("about", "may", "up to", "mostly") is kept everywhere
  the fact appears, in lines, labels, cover and post text.
- Use the newest example that fits, from this year when one exists,
  checked on the maker's own page.
- Never say why a reel was made or who asked for it. Open with the subject.

## Voice and sound

- One designed narrator voice for every reel, never a clone of a real
  person.
- Hear every voiced line back with speech to text. A missing or changed
  word is a cut line, so voice it again with another seed.
- Time every spoken word with a forced aligner. The word times light the
  caption's words and cue every piece of the scene.
- Music is made by code, fresh for each reel from its mood (calm, wonder,
  tense), and no two reels share the same mix. Never a copyrighted track.
- No sound at 0 seconds. The ending fades out, then 1 second of quiet.

## The caption

- The Instagram caption under the reel. 2 or 3 short sentences, about 40
  words at most, that ask the reel's question without spoiling it. Then
  the follow line and 3 hashtags.
- A reel whose call to action is a comment keyword (like "Comment AUTO for
  the free guide") drops the follow line.

## The process

1. Ticket in the queue, with the maker's own page in its notes.
2. The builder writes the research, the sources file, the script and a
   visual plan per line. Nothing is voiced yet.
```

3. Three reviewers read the script at the same time (fact checker, new viewer, visual designer). The builder decides each point and logs why.
4. Optional, the owner's script check. If the owner asked for it (for one reel, or here as a rule for every reel), the coordinator stops after the reviews and shows the reviewed script. The build starts only on the owner's yes, with any changed lines.
5. The build. Voice, hearing, scenes, key frame checks, music, mix, frames, video, cover and caption.
6. The frame auditor checks about 24 frames, the cover and the caption.
7. The audit loop. Fix, audit the whole reel again, until a round finds nothing, at most 5 rounds. After round 5, the reel waits for the owner.
8. The owner watches it, then approves it or writes notes. Every note becomes a fix and a dated rule here. The reel is now ready to post, and the owner posts it by hand. Agents never post.
9. The numbers, by hand for now. Every few days the owner reads the newest reels' insights in the Instagram app and gives the numbers (views, and the share still watching at 1, 3, 5, 10, 30 and 60 seconds), typed or as screenshots. They go in the catalog's table, and the next topics come from it. Ads go only on the reels that held viewers best on their own.

Agents

- One builder per reel. It writes a checkpoint at the top of the script at every stage, so a new builder can resume from it.
- About 5 agents per reel (1 builder, 3 reviewers, 1 auditor), never more than 3 running at once.
- Each agent writes only its own file and reads only what it needs.
- Agents never post anything, and never commit or delete anything they didn't make.

builder_brief.md

Builder brief, one spoken reel from start to finish

The brief a fresh builder agent gets for one reel. The coordinator (the main session) adds the reel's topic, its name and any research already done, and starts the builder on the strongest model you have. We use Claude Opus. The reels are informative videos (topics, facts and guides told with 2D glass cards, 3D scenes, an AI voice and captions), never animated cartoons or comedy sketches.

Read first

- The playbook, all of it. It changes after every review, so read the current text, never memory.
- The queue, before anything else. Work only on the ticket you were given.
- One or two finished reels as references, if you have them, their script, sources file and code. Copy their structure, never their old captions, since some predate the playbook's newest rules. Before the first reel is approved there are none, so the owner's screen recording in reference/ is the reference for the look. There is no video code yet either, so on the first reel you write it (a page per reel drawn with three.js and glass cards, a frame by frame capture, the voice, the music and the encode), and later reels reuse it.
- The approved objects sheet, if you have one. Every new 3D object must match it. If there is none yet, compare with the reference recording, and after the owner approves the first reel, make the sheet from its

```
  objects.

## House rules

Write your own here. Ours, as an example.
- No em dashes or en dashes anywhere. No colons in the caption.
- Every number has a primary source in the sources file, and every
  qualifier stays everywhere the fact appears.
- The hook names something famous in the account's niche in its first
  words. Frame 0 is a full picture with the first caption lit.
- At most 10 spoken lines and 60 seconds plus the 0.5 second hold.
- Everything said is on screen, everything on screen is said, nothing
  before its word. Every term explained in the line that uses it.
- 3D for real objects, glass cards for everything read, one spacing
  system, dim by color, lines only between tiles.
- One designed voice. A fresh music mix from the reel's mood.
- Never post or commit anything. Never delete what you didn't
  make. If a command is refused, skip it, note it and go on. Never stop to
  ask. Decide and log the decision.

## Saving tokens

Read only what you need (search, line ranges). Check frames as small
contact sheets, a few frames scaled down and joined, never one by one.
Never print whole logs, and don't reread files you wrote.

## One builder per reel

Write a checkpoint block at the top of the script at every stage (what is
done, what is next). If the run stops, a new builder resumes from it. If
you find a newer checkpoint you didn't write, or files changing that you
didn't touch, stop at once, change nothing and report.

## Stage A, the script

Research on primary sources (makers' pages, papers, official docs), saved
in the reel's research folder. Then write
- sources.md, one row per fact with the page and the quote,
- script.md, the lines, a visual plan per line (what lands on which word,
  3D or glass for each piece) and a decision log,
- the voice file, the lines as the voice will say them.
Don't voice yet. Stop and report in under 150 words.

## Stage B, the build

When told to go on, read the three reviews, decide each point and log why
in the decision log, and update the lines. Then
1. voice every line, time every word, and hear each line back with speech
   to text, and fix any cut or wrong word with another seed,
2. write the timeline and the scenes, and check frames at the key moments,
3. make the music by code and the mix, record every frame, encode the
   video and make the cover,
4. hear the finished video once more, voice over music,
5. pull about 24 frames at the key moments, plus the cover, into the
   reviews folder with a list of their times, after your own close look at
   each.
Write the Instagram caption by the playbook, so the reel is ready to
post. Stop and report in under 150 words.

## Stage C, the finish
```

```
When told to go on, read the audit, fix every blocking item and the
worthwhile minor ones, render again and hear it again. Then the audit
loop. Audit the whole reel for wrong facts and contradictions across the
lines, the screen, the cover and the caption, fix, and audit again,
until a round finds nothing. The auditor's file is round 1, and your own
full audits after each fix are rounds 2 to 5. If round 5 still finds
something, stop, list what is still wrong at the top of the script, mark
the reel for the owner and say so first in your report.

Otherwise write the reel into the catalog (a row and an entry with its
lines, sound, cover and caption). If you learned something that would help
the next reel, propose it for the playbook. Report in under 200 words,
the files, the length, what the reviews and the audit changed, and
anything left for the owner.

## The owner's notes

The owner watches every finished reel before it is posted. When notes come
back, fix each one, log it in the decision log with the date and the
owner's words, and never drop one. In the same pass, write the rule each
note implies into the playbook, with the date and the words, and name the
change in your report. Numbers only the owner can see (views, retention,
follows) go in your report as a question.
```

reviewer_briefs.md

```
# Reviewer and auditor briefs

The four review roles of a reel. The coordinator hands each agent its
brief with the reel's folder and voice file. The three script reviewers run
at the same time, on a fast model (we use Claude Sonnet). The frame auditor
runs after the first full render.

Common to all. The reel is a narrated 9:16 video, every spoken line a big
caption near the bottom with its words lighting as said, a 3D scene for
real objects and glass cards for everything read. Frame 0 must be a full
picture with the first caption lit. At most 10 spoken lines and 60 seconds
plus the 0.5 second hold. Check against the playbook's current text, not
memory. Where two rules disagree, the newer dated one wins, and name the
clash in your report. The owner watches the finished reel before it is
posted, so flag anything you'd expect the owner to send back. Write only
your own file, and read only what you need.

## Fact checker (script stage), writes reviews/review_facts.md

Check every claim in every line and every planned label against primary
sources. Verify on the web yourself, on makers' own pages, papers, model
cards and docs (some sites refuse fetches, so use mirrors such as arXiv,
GitHub or Hugging Face). Look for
- claims the sources don't support,
- missing or wrong qualifiers ("about", "may", "up to", "mostly"),
- wrong attributions, and things true of one model or version said of
  another,
- a visual that implies a claim, or a duration or result no source gives,
- an example older than it needs to be (search for the newest one that
  fits, this year's when one exists).
Verdict APPROVE or CHANGES. For each problem give the line, the current
words, the source quote that proves it and the exact replacement words,
```

```
short and speakable. Final message under 250 words.

## New viewer and copy editor (script stage), writes reviews/review_viewer.md

Read as a non-engineer with English as a second language, scrolling and
deciding in 2 seconds. Most viewers leave in the first 3 seconds. Check
- each line names its own subject and leaves nothing unexplained,
- the hook in the first 3 seconds, a name or thing famous in the
  account's niche first (propose at most 2 stronger first lines),
- the flow and a payoff someone could retell,
- speakable sentences, and "it" never pointing at the wrong thing,
- a process told in its real order, and a loop said as a loop,
- the length. Count the lines and the said words. Over 10 lines or about
  170 words (60 seconds at the voice's pace), name the lines to cut until
  it fits.
Give a 2 sentence retell. Draft the Instagram caption by the playbook's
caption rules. Verdict and exact replacement wording.
Final message under 300 words.

## Visual designer (script stage), writes reviews/review_visual.md

Read the visual plan with the scene code and one or two finished reels as
references, if there are any, or else the owner's reference recording in
reference/. Rules to check
- every 3D object whole inside its frame at every moment, popping or
  fading in place, never flying in from outside,
- 3D only for real objects, each as finished as the approved objects
  sheet (or the reference recording, before there is a sheet), and glass
  for everything read,
- every word on screen is said, nothing before its word,
- frame 0 a strong, readable picture, and scenes built top to bottom, left
  to right,
- breathing room and one spacing system (shared gap and padding
  constants, never a typed number),
- unlit pieces dim by color, never transparency, and connector lines only
  between tiles, never across one.
Give a concrete build spec per line (objects, materials, colors with one
accent per idea, camera framing, what lands on which word, 3D or glass)
and the cover's composition, built on a real photo when the story has a
person or a physical product. Flag what breaks a rule or won't read at
phone size. Final message under 300 words.

## Frame auditor (after the first full render), writes reviews/audit.md

Two jobs in one pass, on the about 24 frames and the cover (with their
list of times), the voice file, the caption and the sources file.

1. As a cold viewer, per frame, what a viewer takes away, the first 3
   seconds, and anything clipped, overlapping, packed, unevenly spaced
   (measure the gaps), too small for a phone (frames are 1080 x 1920, seen
   at about 390 x 693 points), unexplained, or shown but never said.
   BLOCKING
   - a scene that starts to leave less than 1 second after the line that
     brings its last piece ends,
   - a piece dimmed by transparency, or a line crossing a tile,
   - a phone screen that reads as noise at phone size, or a clipped word,
   - a call to action that doesn't stay clear on screen to the last frame,
   - a diagram whose lines have no room,
   - any 3D object that looks like stacked primitives, a wire frame or
     tangled tubes, or less finished than the approved objects sheet (or
     the reference recording, before there is a sheet). This
```

```
     is the taste check. The auditor passed tubes we called spaghetti and
     a gate we called lame, so be strict.
2. As a fact checker, every number, qualifier, attribution and claim, said
   or shown, against the sources (verify shaky rows yourself), and the
   same fact wherever it appears (lines, labels, cover, caption). Also a
   process told out of its real order, and caption text beyond what the
   playbook allows.

Mark each item BLOCKING or MINOR with an exact fix. Verdict CLEAN or
FIXES. Final message, the blocking items, then the 5 most useful minor
ones, under 300 words.
```

queue_and_catalog.md

```
# queue.txt, one ticket per reel idea

Tickets wait here until a session claims one. Update STATUS, STAGE and
UPDATED at every stage. When a reel is done, move its ticket to the DONE
list with one line, and keep the details in the reel's script.

R1 How does Google Maps know there's traffic ahead?
STATUS: open (open, building, waiting for owner, done)
OWNER: the session that claimed it
CLAIMED: date and time
UPDATED: date and time
STAGE: script, review, build, audit round 2, ...
HOOK: the first line, a name famous in your niche first
SOURCE: the maker's own page that proves the story
FOLDER: posts/<date>/<NN>_<topic>_reel/
NOTES: two lines at most

# catalog.md, every reel made

## At a glance

| # | Reel | Made | Kind | Length | Sound | Cover title |
|---|---|---|---|---|---|---|

## Retention on Instagram

The share of viewers still watching, read off each reel's "How long people
watched your reel" chart in its insights. Add rows every few days.

| Reel | Length | 1 s | 3 s | 5 s | 10 s | 30 s | 60 s | Views | Follows |
|---|---|---|---|---|---|---|---|---|---|

## Entries

One entry per reel. Its files, the lines as spoken, the cover, the sound
mix, the caption, what the reviews and the audit changed, and the
owner's notes.
```

Claude and Claude Code are trademarks of Anthropic, and Instagram of Meta. The reels named here are our own. This guide is independent, not affiliated with or endorsed by Anthropic, Meta, Alibaba, NVIDIA or any other company named. Follower and viewer numbers are our account's own, from Instagram's insights, and yours will differ.