

A FYTAHQ COMPANION RESOURCE

AI From Scratch

A Plain-Language Course
for people who did not grow up writing software

LESSON 8 OF 8 · FINAL LESSON

"Text. Image. Video. Audio. The tells, the grey zone, and the one question that matters more than any of them."

How To Spot AI-Generated Content

ISSUE · JUNE 2026 · FACELESSYOUTUBEAUTOMATIONHQ.COM

More and more of what you read, see, and hear online was made by a machine. That is not automatically a bad thing. But it is worth being able to notice, so that you decide what to believe rather than letting the picture decide for you. This lesson hands you the everyday clues — and, more importantly, the calm habit that keeps working after the clues stop.

WHAT YOU WILL KNOW AFTER THIS LESSON

- The everyday tells for AI-generated text, images, video, and sound
- Why the famous "weird hands" joke has already stopped being useful
- The honest truth that these tells are fading, and what to lean on instead
- The one question that works no matter how good the fakes get
- How to tell harmless AI-content apart from the kind that is meant to deceive

This is the final lesson of the course, and the most practical one. You do not need to remember every tell. Reading it once gives you the instinct, and the instinct is the part that lasts.

THE HONEST CAVEAT, SAID FIRST

Every visual tell in this lesson is getting harder to spot, not easier. The tools improve every few months, and several of the markers below will likely stop working within a year. So treat them as helpful hints, not proof — and put most of your trust in the bigger skill on the last page: asking **who benefits from showing me this**, and pausing before you share.

THE FOUR PLACES TO LOOK — AT A GLANCE

TEXT

Generic phrasing, "in conclusion" and "it is important to note" overused, confident claims with no source named.

IMAGE

Melted or gibberish backgrounds, skin that is too perfect, jewelry that does not match left to right.

VIDEO

Lips out of sync, lighting that shifts shot to shot, odd blinking, a faint seam at the neck.

AUDIO

Breathing that is too even, no little stumbles, a rhythm that is smoother than any real person.

The next pages take these four boxes one at a time, in plain language, with no jargon and nothing to memorise.

TEXT — THE EASIEST KIND TO CHECK

Six things AI-written text keeps doing

Text is where you will meet AI content most often, and it is the easiest to judge with your own eyes. No single marker below is proof on its own. But when three or four of them pile up in one short piece, you are almost certainly reading something a machine wrote.

- 1 Generic, could-apply-to-anything phrasing.** The text talks about "users", "organizations", and "today's world" instead of naming a real person, a real place, or a real moment. Specificity is hard to fake, so machines avoid it.
- 2 "In conclusion" and "it is important to note" on repeat.** AI loves connective filler (small linking phrases that add no information). One is fine. Three on a single page is a strong signal.
- 3 Confident claims with no source.** A figure or a fact is stated firmly, but nobody is named and nothing is linked. Hedged-but-sourceless ("studies show", "experts agree", with no study and no expert) is the classic shape.
- 4 The "it is not just X, it is Y" pattern.** "It is not just a tool, it is a way of thinking." Machines reach for this rhythm constantly. Five of them in five hundred words is a giveaway.
- 5 Tidy lists of three abstract nouns.** "Accuracy, scalability, and excellence." A tricolon (a group of exactly three) made of vague, near-synonymous words is a house style of AI writing.
- 6 A smooth, empty closing paragraph.** Reread the last paragraph. If you could delete it and lose no actual information, a machine probably wrote it.

A QUICK RULE OF THUMB

Look for the human fingerprints, not the machine ones. A real writer names a real person, admits a real opinion, tells one small story that only they could tell. When a whole piece glides along without a single specific, lived detail, that absence is the loudest tell of all.

IMAGE — FORGET THE HANDS

Where AI pictures still slip

The old advice was "look at the hands". That joke is out of date — modern image tools draw hands well now. The give-aways have moved to the edges and the background, where the tool was paying less attention.

- **Melted or gibberish backgrounds.** The face looks sharp, but text on a sign behind it is nonsense and logos are subtly wrong.
- **Skin that is too perfect.** Pores, stray hairs, and small blemishes are smoothed away. Real skin is never that even.
- **Asymmetric jewelry.** Earrings that do not match, or a necklace that simply fades out partway across the chest.
- **Hair that floats free.** Watch where hair meets an ear or collar — strands often drift off and attach to nothing.
- **Reflections that disagree.** A mirror or shop window shows a scene that does not match what is in front of it.
- **Faces too symmetrical.** Real faces are gently lopsided. When the two halves match a little too neatly, a machine smoothed them.

VIDEO — THE WINDOW IS CLOSING

What AI clips still get wrong

Video is harder for the machines, so the tells are clearer here, for now. Most of these will weaken over the next year, so use them while they last.

- **Mismatched lip-sync.** The mouth and the words are a few frames apart, the way a badly dubbed film feels off.
- **Lighting that shifts shot to shot.** The light direction or colour changes between cuts that should match.
- **Odd blinking.** Eyes blink too rarely, too evenly, or not quite together.
- **A faint neck seam.** A subtle line or texture change where a swapped face meets the real neck.
- **Background people that misbehave.** Figures in the distance walk in strange rhythms or briefly merge into each other.
- **Things that morph on contact.** A hand picks up a cup and the cup briefly warps or changes shape in the handover.

A handy shortcut: very long, unbroken shots are still rare in AI video, so a steady fifteen-second take with no cut leans real.

AUDIO — LISTEN FOR WHAT IS MISSING

How a synthetic voice gives itself away

Cloned and generated voices are convincing now, so the tells are not mistakes you hear — they are natural things you stop hearing. A real person is slightly messy, and that mess is what the machine leaves out.

- **Breathing that is too even.** Real speakers breathe at irregular moments. Synthetic breaths land in a tidy, repeating pattern, or are missing altogether.
- **No little stumbles.** Humans say "um", restart a word, or trip on a sound. A voice with none of these micro-stutters across minutes of speech is suspicious.
- **Prosody that is too smooth.** Prosody is the rise and fall of a voice. AI speech often glides too evenly, with none of the small flat or rushed patches real people produce.
- **Emotion that does not quite fit.** The words are sad but the delivery stays bright, or a punchline lands with no lift. The feeling and the sentence do not move together.

THE GREY ZONE

"Is it AI?" is the wrong question

Not all AI content is dishonest. The same technique can be harmless in one place and a real problem in another, so it helps to picture a sliding scale rather than a yes-or-no.

- A photo lightly retouched by AI — not deception.
- An article drafted by AI, then read and approved by a person — usually fine.
- That same article posted under a human byline claiming first-hand experience — that is dishonest.
- News footage altered to change what actually happened — serious fraud.
- A photo of a late relative gently reconstructed by AI for the family — morally complex, neither clearly wrong nor clearly fine.

The better question is not "is this AI?" It is: **is the AI doing something the audience would object to if they knew?** That single question still works even when every visual tell has failed.

A CONCRETE EXAMPLE

A picture goes viral: an elderly couple in a garden holding a hand-painted sign that reads "THANK YOU TO ALL THE NURSES". Here is a calm walk through it.

The sign: readable, but the letters are uneven — some too round, some too narrow. Mild suspicion.

The background: a rose blurs into a leaf with no clear edge between them.

The clothes: the man's knitted sweater has a pattern that almost — but not quite — lines up with itself across his chest.

The hair: one strand of the woman's hair floats off into the air, attached to nothing.

Verdict — AI. The uneven sign-text and the floating hair-strand are the clearest tells. Then the question that matters: shared to tug your heart, it is mild. Shared as "real news from yesterday", it is fraud.

THREE THINGS TO LOOK OUT FOR

If you remember nothing else from this lesson, remember these three. They stay true even as the individual tells come and go.

- 1 Do not trust "AI-detector" tools.** Online detectors flag real human writing as fake and pass real fakes as genuine, often. Trust the markers you can see with your own eyes, not a tool's verdict.
- 2 The tells weaken over time.** What works today may not work in six months. The habit survives even when a specific marker ages out — check the background, check the specifics, check the rhythm.
- 3 Context beats markers.** AI text in a recipe blog is harmless. The same text in a court filing is not. Always ask whether the audience knows, and whether they would mind.

THE ONE QUESTION THAT NEVER GOES OUT OF DATE

When the tells fail — and one day they will — you still have this: **who benefits from me seeing this, and from me passing it on?** A clip designed to make you angry, afraid, or eager to share is worth a second look, whoever made it. Pause before you share. That single pause is the most powerful skill in this whole lesson.

WHAT IS SAFE TO SKIP

You do not need to become a forensics expert. None of the following is required:

- Running every image you see through a detection tool
- Memorising all the markers — the few on these pages are plenty
- Treating every piece of AI content as a trick, because most of it is not
- Learning deepfake analysis, which is a full specialist field of its own
- Boycotting platforms that allow AI content, since by 2026 most of them do
- Arguing with every commenter who posts something a machine wrote

Spotting AI content is an everyday-life skill, not a profession. Calibrate it to the things that could actually affect you, and let the rest go.

PLAIN-LANGUAGE KEY — EVERY TERM USED IN THIS LESSON

Tell

A small clue that something was made by a machine rather than by a person.

Synthetic content

Text, images, video, or sound produced by AI rather than recorded or written by a human.

Lip-sync

How well a person's mouth movements line up with the words you hear. When they drift apart, the video feels "dubbed".

Prosody

The natural rise and fall, pace, and stress of a speaking voice. AI voices often make it too smooth.

Micro-stutters

The tiny "ums", restarts, and trips a real person makes. Their absence is itself a clue.

Deepfake

A video or image that swaps in a face or voice to make someone appear to do or say something they did not.

Byline

The "by Name" credit on an article. A human byline on AI-written text that claims first-hand experience is where harmless tips into dishonest.

Watermark

A small label some tools stamp on AI output. Helpful when present, but easily cropped or removed, so its absence proves nothing.

YOU DID IT

Eight lessons, eight weeks. You are now AI-literate in the way that actually matters: you can follow the conversation, reach for the tools when they help, ignore them when they do not, and spot the obvious deceptions. That is the whole point of the course, and you have it. There was never any forensics exam at the end — just the calm, skeptical instinct you have now built.

THE WHOLE COURSE, IN ONE LINE

Stay curious, stay skeptical, share slowly

You do not need to know how any of it works under the hood. You need to notice when something is AI, ask who benefits from you seeing it, and pause before you pass it on. Everything else in these eight lessons was support for that one habit.

PREFER A LIVE CONVERSATION GOING FORWARD?

One-on-one sessions with a real person

If you would like a real human to walk you through your own questions and your own situation from here, the mrocon side of our world offers ongoing live one-on-one sessions — same philosophy as this course, just guided. One hour, two hours, or three, depending on how deep you want to go.

mrocon.at/pricing →

COMPANION VIDEO

Lesson 8 — the final episode

The Lesson 8 video is live on YouTube — the closing episode of the "AI From Scratch" playlist. Watch what fits your evening, read what fits your morning, and revisit any lesson whenever a conversation sends you back to it.

facelessyoutubeautomationhq.com →

Sources + version note: Content valid as of June 2026. This is the lesson that ages fastest — image, video, and audio tells should be re-reviewed every few months. Reference categories: ChatGPT / OpenAI · Claude / Anthropic · Gemini / Google · AI-image and AI-video tooling · EU AI Act (transparency and labelling obligations) · platform synthetic-media policies. Specific markers shift quickly, so this lesson leans on the durable habit — check the context, ask who benefits, pause before sharing — which outlasts any single tell.