NO PREP

Spotting AI-Generated Images

7 Telltale Signs

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

Spotting AI-Generated Images: 7 Telltale Signs

Why AI Images Look Almost Right

Artificial intelligence image generators have become astonishingly capable. Systems like Midjourney, DALL-E, and Stable Diffusion can produce photorealistic images of people, places, and events that never existed — in seconds, for free. In many cases, these images are visually beautiful and emotionally compelling. But 'almost right' is not the same as 'right,' and current AI image generators still leave a distinct fingerprint on their outputs. Learning to recognize that fingerprint is one of the most practical media literacy skills you can develop in the current information environment. The errors AI generators make are not random — they are

Name: **Period:** **Date:**

Spotting AI-Generated Images: 7 Telltale Signs

Why AI Images Look Almost Right

Artificial intelligence image generators have become astonishingly capable. Systems like Midjourney, DALL-E, and Stable Diffusion can produce photorealistic images of people, places, and events that never existed (in seconds, for free). In many cases, these images are visually beautiful and emotionally compelling. But 'almost right' is not the same as 'right,' and current AI image generators still leave a distinct fingerprint on their outputs. Learning to recognize that fingerprint is one of the most practical media literacy skills you can develop in the current information environment. The errors AI generators make are not random; they are systematic. They stem from the way these systems work: by learning statistical patterns across billions of images rather than understanding the actual physical and biological rules that govern how the real world looks. This means AI images tend to fail in predictable places. Once you know what to look for, you'll start seeing the tells everywhere: in news feeds, on social media, in political advertising, and in content designed specifically to deceive.

Tell #1: The Hands Problem

Human hands are one of the most reliably problematic features for current AI image generators. The reason is structural: hands are geometrically complex, highly variable in position, and appear in an enormous range of configurations in photographic data. AI systems that learn from this data struggle to build a consistent model of hand structure. The result is images where hands have too many fingers, not enough fingers, fingers that bend in anatomically impossible directions, or fingers that blur indistinguishably into each other or into nearby objects. When examining a suspicious image, find the hands first. Count the fingers. Look for joints that bend the wrong way. Look for fingers that fade into the background or into

other fingers without a clear border. Even in images that appear highly realistic in every other respect, hands frequently give away AI generation. This is the single most reliable quick test available to non-expert viewers.

Tell #2: Teeth, Eyes, and Hair

Three facial features deserve special attention: teeth, eyes, and hair. AI-generated teeth are often too uniform (perfectly even, perfectly white, without the natural variation in size, color, and alignment that real human teeth have). They may also have an unusual texture, appearing almost ceramic rather than biological. When someone in an AI image smiles broadly, look closely at the individual teeth. Eyes in AI images often have a quality that experienced viewers describe as 'too perfect' or 'glassy': they lack the subtle asymmetry and the complex patterns of light reflection that real eyes display. Irises may be unnaturally symmetric or have a repeating pattern that doesn't match real iris anatomy. Pupils may be slightly misaligned between the two eyes in the same image. Hair is another common failure point: AI hair often has an unusual texture at the edges; it may fade unnaturally into the background, have inconsistent strand thickness, or show repeating patterns that real hair does not.

Tell #3: Background Chaos

While AI generators have become very good at rendering the central subjects of images, backgrounds often reveal their AI origin. Text in backgrounds is almost always garbled — road signs, storefronts, labels, newspapers, books, and any other text-bearing surface will typically show letter-like shapes that do not spell actual words. This is because AI image generators were not trained to be coherent text-writers; they learned to replicate the visual pattern of text without understanding its linguistic content. Beyond text, look at background architecture and objects. Windows may be asymmetric or multiply in ways that violate physical space. Reflections in mirrors, windows, or glasses may not correspond to what should logically appear in them. Bookshelves may contain book spines with nonsense titles. Crowds of people in the background often contain individuals with fused features, merged bodies, or faces that share characteristics in an uncanny way. The more densely populated a background is, the more likely it contains visible AI artifacts.

Tell #4: Jewelry, Accessories, and Clothing

Earrings in AI-generated images are frequently asymmetric — different shapes, different sizes, or present on one ear but absent on the other. Necklaces and bracelets may pass through the body rather than resting on top of it. Glasses often have mismatched lenses, frames that continue through the face rather than sitting in front of it, or reflections that don't match the light source in the rest of the image. Logos on clothing are typically distorted or replaced with logo-like shapes that don't correspond to any real brand. These accessory failures happen because AI generators learn the statistical association between, say, 'formal portrait' and 'earrings' without understanding the geometric and physical constraints that govern how earrings actually sit in relation to ears. The generator knows earrings should be present but doesn't have a sufficiently detailed model of the geometric relationship between ear anatomy and jewelry form.

Tell #5: Lighting and Physics Violations

Real photographs obey physical laws. Light comes from a specific direction or multiple specific directions, and every object in the image casts shadows and reflects light in ways that are geometrically consistent with those light sources. AI images frequently violate these constraints. Shadows may fall in multiple conflicting directions. The light falling on a person's face may come from a different apparent direction than the light illuminating the background behind them. Reflective surfaces — water, glass, metal, eyes — may reflect light sources or objects that aren't present elsewhere in the image. These physics violations are particularly visible in images that try to composite multiple subjects — a person placed in a landscape, for example, where the lighting on the person was learned from studio portrait data while the background lighting is outdoor daylight. The mismatch is often visible even to untrained viewers who can't explicitly describe what's wrong, but simply feel that something is 'off' about the image.

Tell #6: Skin Texture and Age Inconsistencies

AI-generated human skin tends toward an airbrushed smoothness that's common in professional photography but often goes further — pores may be entirely absent, or the texture may have a repeated pattern that real skin does not. In images of older individuals, wrinkles may appear in geometrically regular patterns or may be present

only in some areas of the face while absent in adjacent areas where they would logically appear. Age inconsistencies are also common: a person's face may appear one age while their hands, neck, or body appear a significantly different age. The AI generator has essentially assembled the image from learned patterns for different body parts without maintaining consistent age characteristics across those parts. If the face and hands of a person in an image look like they belong to people of very different ages, that's a significant red flag.

Tell #7: The Uncanny Valley in Expressions

Finally, pay attention to facial expressions and their relationship to body language. Human expressions are complex and consistent — when someone genuinely smiles, it involves not just the mouth but the eyes, cheeks, and subtle movements throughout the face. AI images sometimes produce faces with the surface features of an expression — upturned mouth corners, for example — without the corresponding secondary features that make the expression look genuine. The result is a smile that looks posed in an unsettling way, or a worried expression that doesn't extend to the eyes. Body language and expression should match. A laughing person's shoulders and torso position should reflect the physical reality of laughing. A person looking surprised should have body language consistent with surprise. When the face and body appear to be in different emotional states simultaneously, or when an expression on the face has the right general shape but is missing subtle authenticating features, you may be looking at AI generation.

Key Information: AI Image Detection Tells

Tell	What to Look For
Hands	Wrong number of fingers, impossible bends, fingers merging
Teeth	Too uniform, ceramic texture, unnatural evenness
Eyes	Glassy appearance, symmetric irises, misaligned pupils
Hair	Fading edges, repeating strand patterns, inconsistent texture
Background text	Letter-like shapes that don't spell real words

Tell	What to Look For
Accessories	Asymmetric earrings, jewelry passing through skin
Lighting/shadows	Conflicting shadow directions, impossible reflections
Skin	Airbrushed smoothness, absent pores, age inconsistencies

Case Study: Jordan's Research Assignment

Jordan was researching AI-generated disinformation for a class project and found a social media post featuring what appeared to be a photograph of a scientist at a press conference, with a quote attributed to her claiming vaccines caused a specific disease. Jordan suspected the image was AI-generated. She zoomed in and immediately noticed the scientist's left hand, which appeared to have six fingers, two of which seemed to merge near the second knuckle. The background showed a banner with text that read something like 'HEALT RSERCH INTTUTE' — clearly not a real institution name. However, several of Jordan's classmates argued the image looked 'professional enough' and that the six-fingered hand might just be a photography angle. Jordan captured screenshots of the specific anomalies and brought them to class the next day to explain why she was confident the image was fabricated.

1. Using at least two specific tells from the reading, explain why Jordan was right to be confident the image was AI-generated, even though classmates weren't immediately convinced.

2. Why is it significant that the fake image was paired with a fabricated quote attributed to a scientist? What does this combination tell us about how disinformation campaigns use AI imagery?

Multiple Choice

1. Why do current AI image generators consistently struggle with human hands?

a) Hands are geometrically complex and appear in too many configurations for AI to model consistently

b) Hand generation requires far more computing power than other body parts

c) Hands are rarely photographed on their own, giving AI generators little training data to learn from

d) AI generators intentionally avoid generating hands to prevent misuse

2. Why does text in the backgrounds of AI-generated images typically appear as garbled letter-like shapes?

a) AI generators deliberately compress background detail in order to save processing time

b) Text is deliberately blurred to avoid copyright infringement

c) Background text is added randomly after the main image is generated

d) AI generators learn visual patterns of text without understanding linguistic content

3. What does the term 'uncanny valley' refer to in the context of AI facial expressions?

a) A geographic region in the United States where deepfake technology was first invented in 2017

b) A database of human facial expressions used to train AI

c) Expressions that have the right general shape but are missing subtle authenticating features

d) The process by which AI improves facial realism over time

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Generative AI image system	A) Difference or imbalance between two sides of something
2. ___ Iris	B) A visual error or anomaly revealing AI image generation
3. ___ Asymmetry	C) The colored ring of the eye surrounding the pupil
4. ___ AI artifact	D) Image assembled from multiple separate elements or sources
5. ___ Composite image	E) Software that creates photorealistic images from text prompts

Fill-in-the-Blank

1. AI image generators fail at human hands because they learn statistical _____ rather than understanding actual physical rules.

2. Text in AI-generated image backgrounds is typically _____ because generators replicate the visual appearance of text without understanding language.

3. When a person's face and hands in an AI image appear to belong to people of very different _____, it suggests the generator assembled parts without maintaining consistency.

Short Answer

Answer each question in 2-4 complete sentences.

1. Explain why AI systems make systematic, predictable errors rather than random errors, and why this predictability is actually useful for people trying to detect AI-generated images.

2. Choose two of the seven tells described in the reading and explain what underlying technical limitation of AI image generators each tell reveals.

3. The reading says that even untrained viewers sometimes feel something is 'off' about AI images without being able to describe what. Why might this instinctive response be valuable, and why isn't it sufficient on its own?