

# “THE HUMAN GUIDE TO DETECTING AI IMAGERY”

For journalists, researchers, and fact-checkers  
investigating our information ecosystem. And for users  
defending themselves against a world of AI slop.

---

AI FORENSICS



## Table of Contents

Acknowledgements.....	3
Introduction to the Human Guide to Detecting AI Imagery.....	4
1. BEFORE THE IMAGE: AI TELLTALES.....	5
Written Indicators on Platform Interface.....	5
1.1. LABELS.....	5
1.2. TITLE, DESCRIPTION, HASHTAGS.....	9
1.3. PROFILE/DESCRIPTION/BIO.....	11
1.4. ACCOUNT POSTING HISTORY.....	12
1.5. COMMENTS.....	13
Written Indicators Within the Post's Frame.....	13
1.6. WATERMARKS.....	13
1.7. STICKER/TEXT.....	14
2. SYNTHETIC ARTIFACTS IN AI IMAGERY.....	16
Content and Subject Matter.....	16
2.1. IMPOSSIBLE PHYSICS.....	16
2.2. INDECIPHERABLE INSCRIPTIONS.....	17
2.3. INCONSISTENT PROPORTIONS.....	19
2.4. UNMATCHING OR MISSING ACCESSORIES.....	19
2.5. TOO MANY/TOO FEW BODY PARTS.....	20
Format and Form.....	21
2.6. STYLIZED SMOOTH GLOW.....	21
2.7. IMPLAUSIBLE SYMMETRY, VANISHING POINTS, AND SHADOWS.....	23
2.8. INCONSISTENT STYLE.....	24
3. MOVING IMAGES/CLIPS.....	25
Body Details.....	25
3.1. DISAPPEARING/MORPHING BODY PARTS.....	25
3.2. UNNATURAL BODY LANGUAGE.....	26
Facial Details.....	27
3.3. UNNATURAL EYES.....	27
3.4. UNMATCHING/IMPOSSIBLE LIPS/TEETH/HAIR.....	29
Object Motion Details.....	30
3.5. (DIS)APPEARING OBJECTS.....	30
3.6. UNNATURAL PHYSICS.....	32
4. DIGITAL PROVENANCE.....	34



## Acknowledgements

“The Human Guide to Detecting AI Imagery,” version 1, published in December 2025

Authors: Natalia Stanusch

Reviewers: Fraser Crichton

The contribution from AI Forensics is funded by core grants from Open Society Foundations, Luminate, Omidyar Network, and Limelight Foundation.

All other content (c) AI Forensics 2025

Graphic & Brand Design: Denis Constant / Ittai Studio <https://ittai.co/>

Email: [info@aiforensics.org](mailto:info@aiforensics.org)

---

## Introduction to the Human Guide to Detecting AI Imagery

This guide supports users in detecting distinct visual artifacts that can help identify whether an image or video they encounter on a platform is AI-generated. Different cues to detect AI-generated content are discussed along with examples. In this guide, we account for a spectrum of synthetic content that can be referred to as “generative AI imagery,” including outputs in the form of still images as well as moving images or clips. This guide is intended for both professionals, such as researchers, journalists, and fact-checkers, as well as daily users of social media platforms looking to be better equipped in the age of AI slop.

Due to the rapidly progressing and changing formats of AI-made visual content, this guide does not attempt to provide an exhaustive or definitive set of guidelines. In fact, we acknowledge that some elements of this guide will become outdated in the coming months as these technologies advance. Ultimately, this limited guide offers an overview and order of useful steps to take when encountered with visual content in the contemporary digital media environment, particularly on social media platforms.

This guide is structured as a series of steps a user can follow to assess the likelihood that the media content they see has been generated using AI tools. The steps are divided into four focus areas: 1) Before the Image: AI Telltales; 2) Synthetic Artifacts in AI Imagery; 3) Moving Images/Clips; 4) Digital Provenance. These four areas denote the broader points of interest constituting the overall detection process. It is important to follow the steps across all four areas when assessing the plausibility that the content is AI-generated.

We recommend consulting the steps in the specified order. That said, users should feel free to consult these steps based on their unique needs or queries. The recommended order of steps in each of the four areas is structured based on the difficulty level of the inquiry and result uncertainty. For relative assurance that an image or video is likely AI-generated, the content in question should meet more than one of the characteristics that follow. Also, please note that while this guide does its best to structure and utilize current digital media expertise and literacy, it cannot always guarantee absolute correctness of results in practice.

## 1. BEFORE THE IMAGE: AI TELLTALES

### Written Indicators on Platform Interface

Prior to any assessment of the visual object (still image or video clip), there may be several elements within the platform interface that can serve as telltale signs of a piece of content having been entirely or partially made using generative AI tools. While the following examples are based on specific features embedded across three platforms (Instagram, TikTok, and YouTube), the detection elements roughly encompass the main features a user can look out for and find across other platforms as well. It is crucial to note that these features change over time as well as the version of the platform the user can access.

#### 1.1. LABELS

Many platforms have introduced content labels to address the concern around detectability of AI-made or AI-edited content. These content labels should be either visible or at least accessible via the platform interface while viewing a post. If available and applied to a piece of content, such labels are a quick and easy way to verify whether the visual object is AI generated. The following examples demonstrate what AI content labels tend to look like and where they are usually located based on examples from TikTok (Tables 1 and 2), YouTube (Tables 3 4), and Instagram (Table 5). It is also crucial to note that the terms of service differ across platforms and evolve across time; in general, most platforms put the responsibility on content creators to label and disclose photorealistic AI content they upload. In practice, most content creators do not label their content as AI, making the AI labels useful in around half the cases.<sup>1</sup>

The following are takeaways and considerations for assessing the presence of AI content labels on platforms: different labels disclosing AI content might be used simultaneously on the same platform (Table 1) and these labels might change in appearance and/or text over time (Table 2); AI content labels might be accessible only upon clicking for further information, e.g. by expanding the description tab (Tables 3 and 4); AI content labels might be present on the app interface of the the platform but absent on the web interface (Table 5). The following examples illustrate these considerations and explain the possible obstacles in locating and

---

<sup>1</sup>AI Forensics' investigations found that at least 50% of AI content on platforms such as TikTok is not labelled and disclosed as AI. For more details, see the following reports: [aiforensics.org/work/gen-ai-slop](https://aiforensics.org/work/gen-ai-slop) and [aiforensics.org/work/agentic-ai-accounts](https://aiforensics.org/work/agentic-ai-accounts).

accessing AI content labels across platform interface particularities and possible differences between browser and app interfaces.



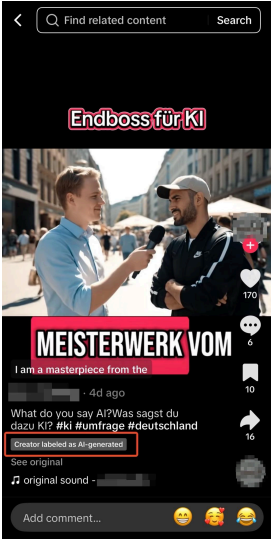
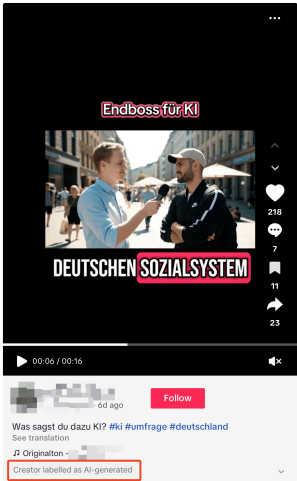
AI content labels vary and are used simultaneously on the same platform			
'AI generated'		'Creator labelled as AI-generated'	
app	browser	app	browser
			
Figure 1A	Figure 1B	Figure 1C	Figure 1D

Table 1. The visibility of TikTok’s AI content labels on the platform across app (Figure 1A and 1C) and web interfaces (Figures 1B and 1D). The red rectangle outlines the positioning of the labels within TikTok’s interface.

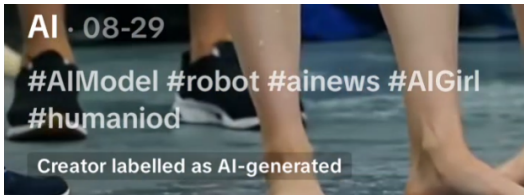
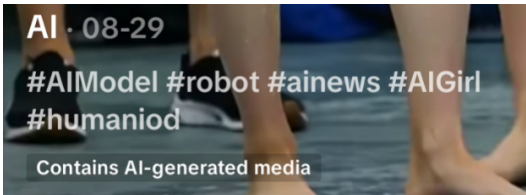
AI content labels change in appearance and/or text over time	
app (old)	app (new)
	
Figure 2A	Figure 2B



Table 2. An example of a change in the text of an AI content label over time. In this example, TikTok’s AI label with the text “Creator labelled as AI-generated” (Figure 2A) has been changed to a label that reads “Contains AI-generated media” (Figure 2B), possibly as a result of A/B testing. The change in the label’s text was documented on 28 October 2025.

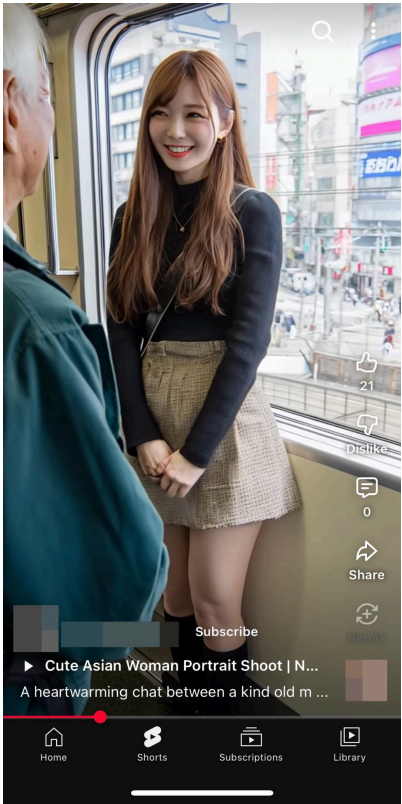
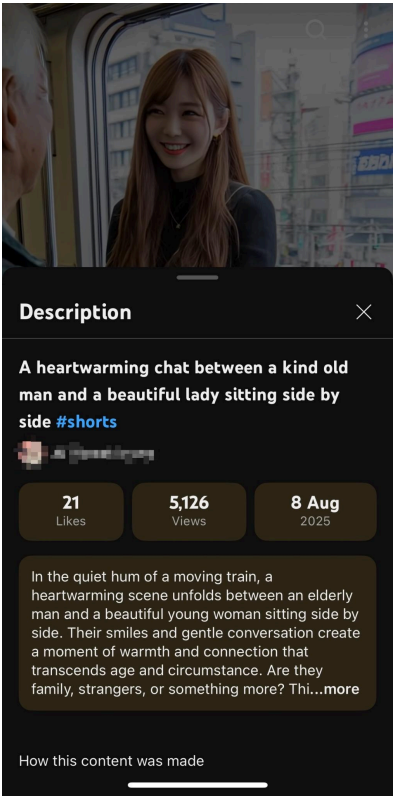
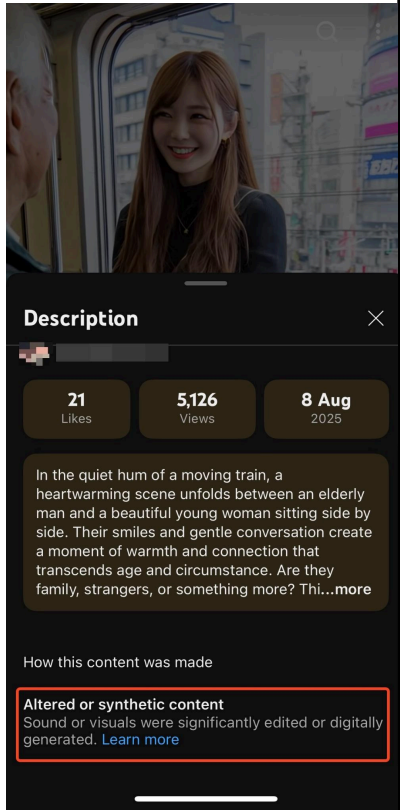
AI content labels are accessible only upon clicking for further information and scrolling		
app (invisible)	app (invisible)	app (visible)
		
Figure 3A	Figure 3B	Figure 3C

Table 3. Accessibility of YouTube’s AI content label via the platform’s app. The AI content label is visible at the bottom of the screen after clicking the tab to expand the post’s description. Furthermore, the label is not visible immediately given the lengthy description of the post but requires the user to scroll down to the bottom of the screen to see it.


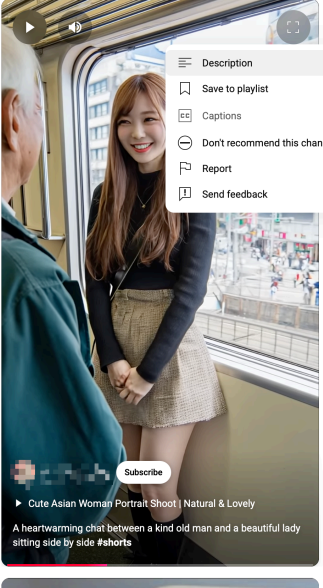
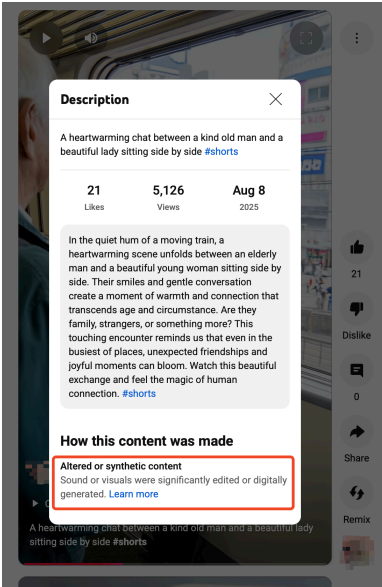
AI content labels are accessible only upon clicking for further information		
browser (invisible)	browser (invisible)	browser (visible)
		
Figure 4A	Figure 4B	Figure 4C

Table 4. Accessibility of YouTube’s AI content label via the platform’s browser. The AI content label is visible once the description tab is expanded; the description tab is shown upon clicking the “more options” icon on the interface.

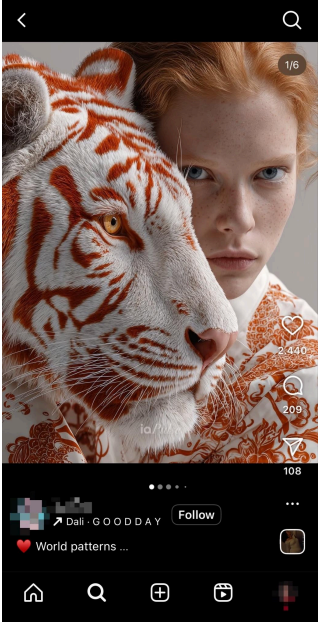
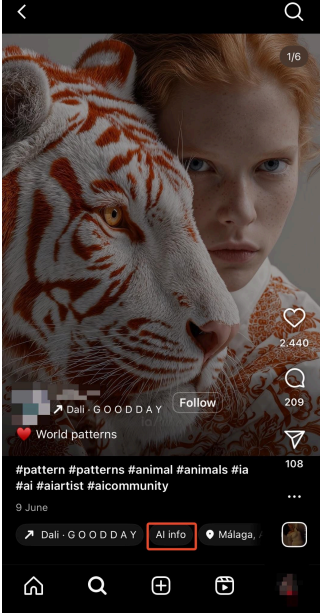

AI content labels are present on the app but absent on the web interface of the platform		
app (invisible)	app (visible)	browser (invisible)
		
Figure 5A	Figure 5B	Figure 5C

Table 5. Instagram’s AI content label is not visible by default in many posts but often hidden among other information. The user is forced to search for the AI label, which can appear on different parts of the screen depending on the post and is hidden from the initial content view (Figure 5A). Furthermore, no AI label is displayed on Instagram’s web interface (Figures 4C and 5C) as of 30 October 2025.

1.2. TITLE, DESCRIPTION, HASHTAGS

Information included in the title of a post, its description, and/or hashtags might contain remarks that indicate whether the content is generated by AI. Such information is included by the content creators when they upload their posts on the platform. An AI content disclosure is easy to spot when it is located at the forefront, e.g. as the first words in a description, title, or list of hashtags. However, sometimes the disclosure is located in the middle or closer to the end of the description text or hashtag list, remaining invisible unless the user purposely seeks it (Table 6). AI content disclosures might be straightforward, such as ‘#AI’ or ‘made



with AI.’ They might also be more context-specific and refer to a particular generative AI tool. Table 7 provides a non-exhaustive, common list of keywords/phrases, hashtags, and tools that tend to be mentioned as AI content disclosures across platforms.

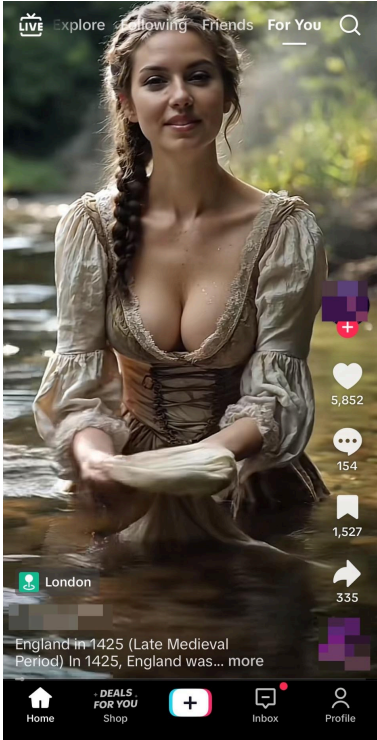
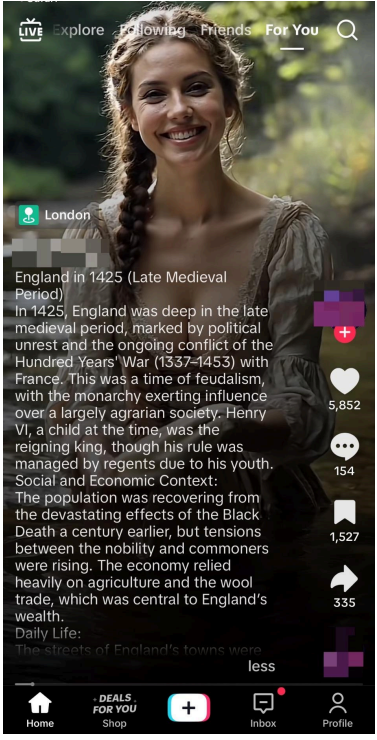
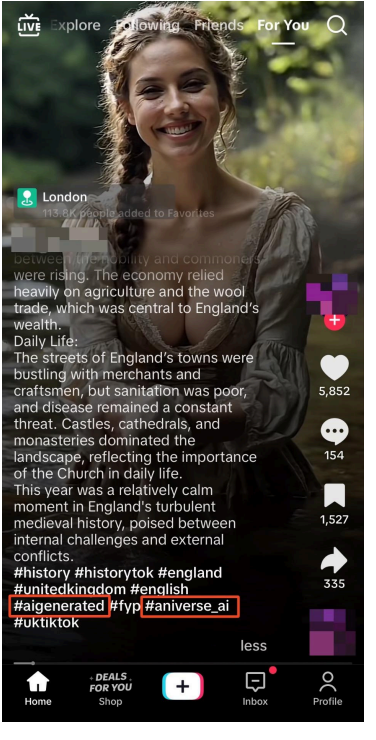
Visibility of alternative AI content disclosures as hashtags		
app (invisible)	app (invisible)	app (visible)
		
Figure 6A	Figure 6B	Figure 6C

Table 6. Examples of the visibility of alternative user disclaimers of AI content on TikTok’s app interface. The red rectangle outlines the positioning of the user’s AI disclosure, which in this case is a hashtag.

A non-exhaustive list of possible AI content disclosures		
Keywords/Phrases	Hashtags	Tools
ai-generated	#aigenerated	Runway
aigenerated	#madewithai	Sora



AI	#ai	Sora2
artificial intelligence	#artificialintelligence	Veo
künstliche Intelligenz	#kunstlicheintelligenz	Veo3
KI	#ki	Stable Diffusion
inteligencia artificial	#inteligenciaartificial	Midjourney
IA	#ia	Grok
ia	#Aistory	Kling
aiart	#TikTokAI	GPT Image
_ai	#AIVideo	ChatGPT
	# + [a word, e.g. 'creative,' 'lookbook', or 'elegant'] + 'AI'; e.g. #Aigirls	Nano Banana
	# + 'AI' + [+any other word, e.g. 'creative,' 'cat', or 'fashion'], e.g. #Aldogs	Pika Labs

Table 7. Examples of AI content discourses content creators might include in the title, description, and hashtags list while uploading their AI-made posts. The list contains keywords and phrases, hashtags, and the names of generative AI tools used for generating still and moving images. The list was compiled during AI Forensics' investigations and exploratory analyses of AI content across Instagram and TikTok.

### 1.3. PROFILE/DESCRIPTION/BIO

The username, account description, or profile bio might include a disclosure that the account is used to post AI content, increasing the possibility that the piece of content in question is also AI-generated. This might be visible immediately when the account's username contains such information, or it might require accessing the account's profile page (Table 8).


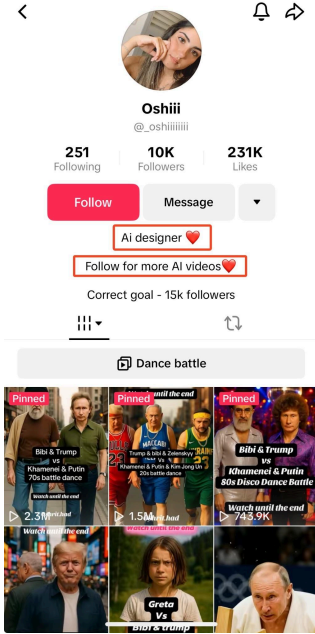
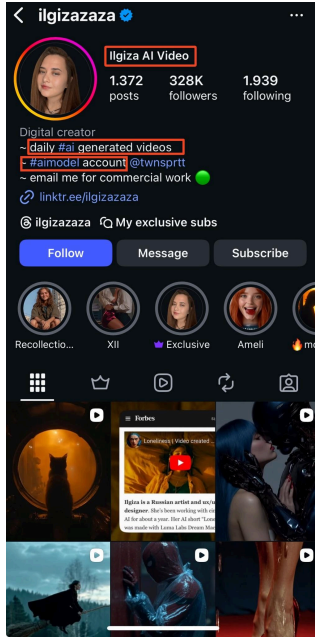
Alternative AI content disclosure in a username, account description, or profile bio		
Username (YouTube)	Profile bio (TikTok)	Account name (Instagram)
		
Figure 8A	Figure 8B	Figure 8C

Table 8. Examples of instances from YouTube (Figure 8A), TikTok (Figure 8B), and Instagram (Figure 8C) where either the username (Figure 8A), profile bio (Figure 8B), or account name (Figure 8C) include AI content disclaimers. As seen in the above figures, such AI content disclaimers are often (but not always) included across multiple features mentioned in points 1.2 and 1.3.

#### 1.4. ACCOUNT POSTING HISTORY

Checking an account's profile feed and posting history might reveal whether the account has a history of posting some, or only, AI content (an example of Agentic AI Accounts). If an account has posted solely generative AI content in the past, or if it has posted several videos or images that are visibly of similar format and/or subject matter and carry some telltale signs of AI origin, it is likely that the post in question is also AI-generated.<sup>2</sup> It is useful to scroll through the feed to see the upload's history, as one may notice that an account used earlier

<sup>2</sup> For a detailed discussion of Agentic AI Accounts (AAAs), see the following AI Forensics' investigation: [aiforensics.org/work/agentic-ai-accounts](https://aiforensics.org/work/agentic-ai-accounts).

models or earlier versions of models, containing more obvious signs of AI content.

1.5. COMMENTS

The comments section might reveal an ongoing discussion about the AI-generated status of the content in question and include relevant clues and indicators regarding AI use. However, it is important to stay wary of any ambiguity between users’ genuine ‘questioning if it’s AI’, and commenting in a more ‘memefied’ and ‘trolling’ spirit.

Written Indicators Within the Post’s Frame

Some content might include indicators of AI use within the content frame itself. Such indicators include watermarks, stickers, and other forms of text that either stand out as external elements (Tables 9 and 10) to the subject matter of the content or form an integral part of the uploaded content (Table 11). While such indicators might usually be easy to spot, it’s good to be aware of possible obfuscations of the visibility of AI indicators caused by the platform’s interface design or the content creator’s frame cropping in some instances.

1.6. WATERMARKS

The title/description/hashtags of an image or video mention or contain remarks that indicate whether the content is generated by AI or is a deepfake.

Sora watermark visibility on TikTok’s app interface		
app (partially hidden)	app (partially hidden)	app (fully visible)



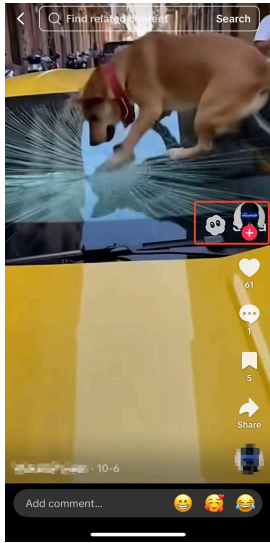
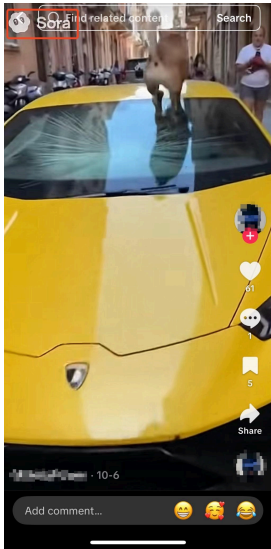
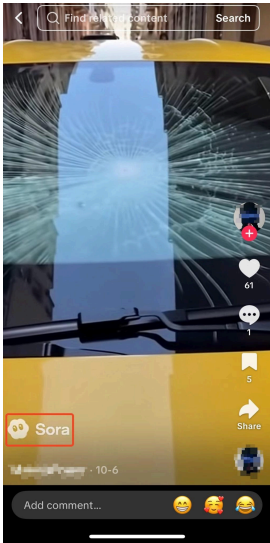
		
Figure 9A	Figure 9B	Figure 9C

Table 9. Visibility of OpenAI’s Sora2 model watermark on TikTok’s app interface. The watermark appears across different parts of the screen as the video progresses. Depending on the placement of the watermark, interface elements, such as the user’s avatar picture, partially obscure the watermark’s visibility (Figures 9A and 9B).

VEO watermark visibility on TikTok’s app interface		
app (visible)	app (hidden)	browser (visible)



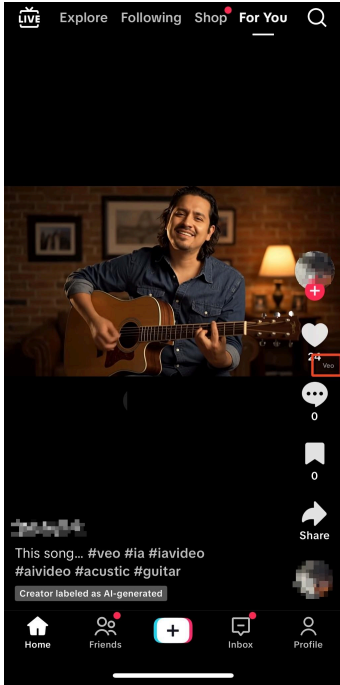
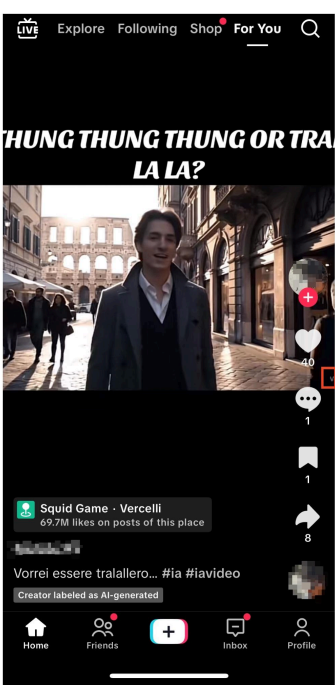
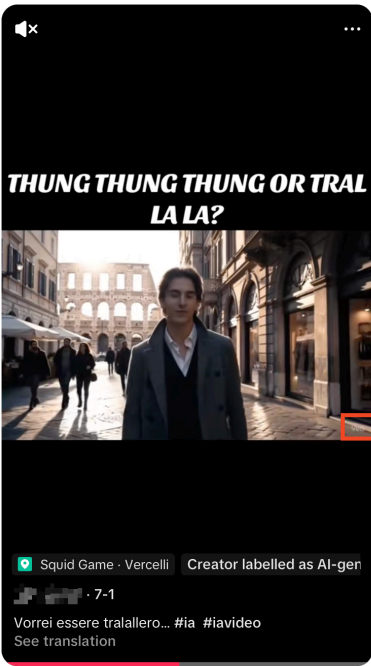
		
Figure 10A	Figure 10B	Figure 10C

Table 10. Visibility of Google’s VE03 model might depend on the way the user has cropped the content. However, on TikTok, the app interface might crop the frame of the uploaded content, practically deleting the VEO watermark which tends to be located on the bottom right corner (Figure 10B). On TikTok’s browser interface, however, the watermark remains correctly visible (Figure 10C).

1.7. STICKER/TEXT

Some users might choose to include their username within the frame of the post. This is particularly useful if the AI-generated content has been shared or reuploaded by other accounts, possibly losing its original AI disclosures and/or labels. It is more visible if the username contains an indication of AI use (Figures 11A and 11B). However, such indications might also be included as an integral part of the image or video, e.g. when a written indication of the content’s AI origin or content creator’s username is included as part of the background or character design (Figures 11B-11D).

Written AI indicator as an integral part of the uploaded content

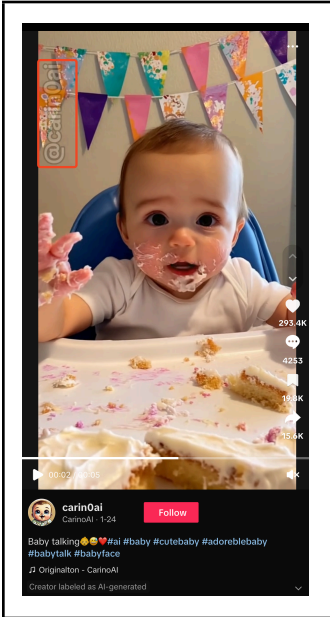
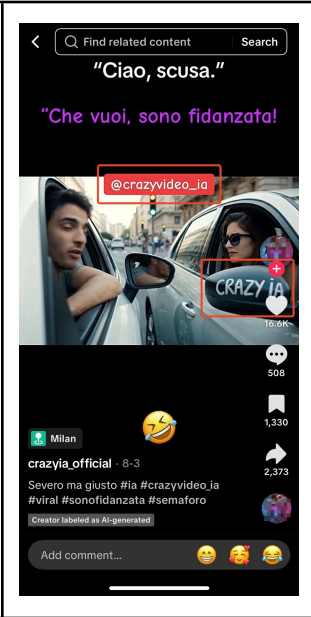
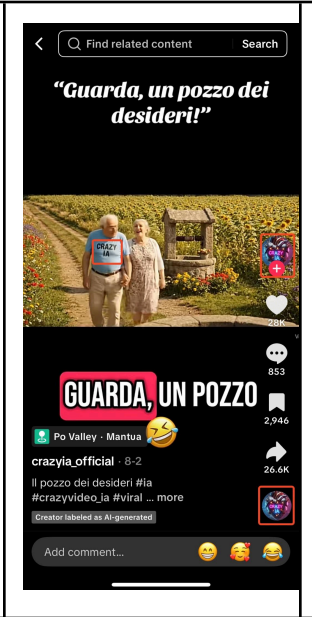
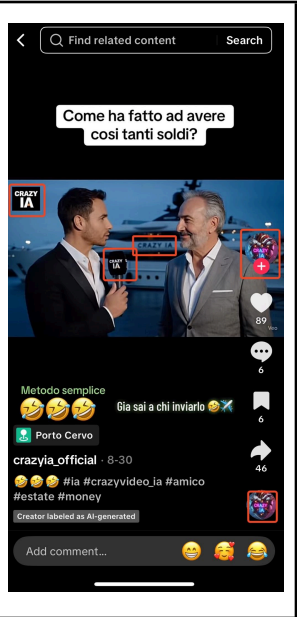
			
Figure 11A	Figure 11B	Figure 11C	Figure 11D

Table 11. Examples of visibility and placement of written AI indicators included as an integral part of the uploaded content on TikTok. The examples range from an inclusion of the content creator’s username within the frame of the uploaded post (Figure 11A) to an inclusion of the content creator’s username which also serves as an AI disclosure appearing in various locations and styles across uploaded content (Figures 11B-11D).

## 2. SYNTHETIC ARTIFACTS IN AI IMAGERY

This section focuses on synthetic artifacts that can be spotted in AI imagery. As such, they are applicable to both still and moving images (short clips, films, etc.). While these artifacts tend to be defining qualities of AI imagery, it has to be noted that visual qualities that are subjectively viewed as meeting some of the following criteria – and thus serving as ‘evidence’ of a piece of content being AI-generated – are not definitive proofs and could lead to an error-prone judgement. Nonetheless, the following description of visual artifacts can be considered as a set of ‘places to look’ and ‘things to look for’ while assessing whether a piece of content might be AI-generated.

Each of the following example images contains several of the synthetic artifacts described in this section. For clarity, the highlighted and discussed elements in each of the examples call attention only to the specific synthetic artifact that the particular subsection is dedicated to, rather than discussing all synthetic artifacts present in each image or the presence of any of the previously mentioned AI telltales (e.g. AI content labels).

### Content and Subject Matter

The following subsection discusses the “what” elements when assessing visual objects, such as still photos and “film” clips, in relation to characteristics of AI imagery. It brings attention to details in the content’s subject matter that can be spotted in the foreground and background as indications of AI imagery. The following examples focus on details within the subject matter and take for granted the first question a user should ask, namely, is the combination of time, characters, setting, and/or circumstances portrayed in the content physically plausible? If the answer to this question is no, further assessment can be used as additional artifacts-based affirmation of the AI-origin of the piece of content rather than a detection of suspicious elements.

#### 2.1. IMPOSSIBLE PHYSICS

Synthetic artifacts that fall under the category of impossible physics are often striking and readily noticeable. They connote all material configurations that would not occur in the “real” world due to, most often, gravity, or contextual plausibility. Examples include objects floating in the air or missing crucial elements. This also includes a lack of correspondence to the physical makeup of objects, especially technological objects and their components: for example,

---

cameras, radios, instruments or other objects tend to show a logical placement of limited components that made their function physically possible. In AI-generated imagery, these logical structures, their order, placement, and form are often distorted and “hallucinated.”




AI artifacts exemplifying impossible physics		
		
Figure 12A	Figure 12B	Figure 12C

Table 12. Examples of impossible physics as seen in examples of posts from TikTok. The highlights in red numbered 1 and 2 in the first example (Figures 12A and 12B) draw attention to the physical impossibility of the chair standing in the given spot while missing a part of its leg (number 2 in the image) and the backrest being positioned perpendicularly to the person’s back, instead of in parallel (number 1). The pigeons (number 4) appear to be floating in the air despite looking as if they were seated. In the second example (Figure 12C), the red tray seems to be floating in the air, not being held by any hand or support.

2.2. INDECIPHERABLE INSCRIPTIONS

The incomprehensibility of words and numbers in AI imagery has been one of the most iconic telltale signs of early AI content. While the similarity of synthetic



words, numbers, and signs to the ‘real’ ones has improved, they remain one of the elements where AI imagery continues to fail and ‘hallucinate.’ Relevant synthetic artifacts include words, phrases, or symbols that look “messy” upon closer examination: the details are missing, appear asymmetrical, or lack coherence or legibility. Numbers and letters do not attune to any existing alphabet, dissolve into each other, and are, simply put, not real.

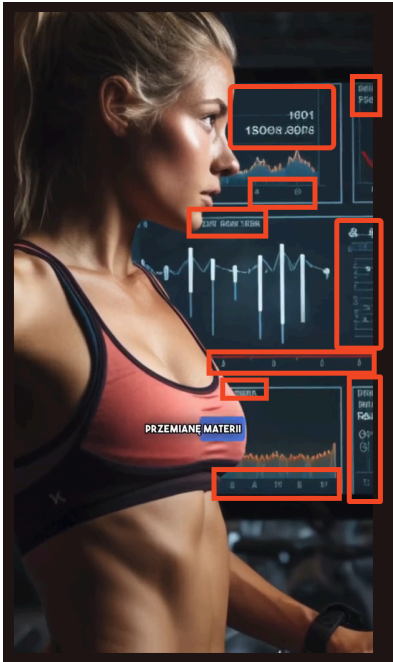


Illegible writing and numbers, inconsistently rendered logos and symbols		
		
Figure 13A	Figure 13B	Figure 13C

Table 13. Cases where inscriptions visible in the frame are either indecipherable or suspicious upon closer examination. In Figure 13A, almost all numbers and words in the background are illegible and bear stylistic and formal qualities of “messy,” unreal AI-made inscriptions. The highlights in red bring attention to all details in the background that show such artifacts. In Figure 13B, we observe somewhat recognizable symbols, such as the logo of Adidas and Coca-Cola. However, upon closer examination, the word “adidas” contains “o” instead of the first “a,” the letters are not symmetrical, and the “s” at the end of the word appears as if it’s “melting.” Similarly, the actual writing on the Coca-Cola bottle, while partially obscured, does not appear to contain any believable writing behind the hand that holds it. The other logo on the T-shirt is also asymmetrical, blurry, and each star above it has different proportions and number of arms. Figure 13C shows similar inconsistency: first, the numbers on the phone’s rotary dial are not placed in a decreasing order and most of them are indecipherable

and illegible. The symbol of what appears to be the Russian flag is asymmetrical (especially on the left side).

2.3. INCONSISTENT PROPORTIONS

As the generation of AI imagery is grounded solely in probability rather than the laws of physics and ‘knowledge’ of the world, the differences in scale, size, and perspective of objects within the frame can often appear off. Here, synthetic artifacts can be detected in relation to the relative proportion of different objects, e.g. a mismatch between the sizes of different body parts, especially the head or hands and the rest of the body, or objects appearing too big or small (from shoes to cars and buildings).



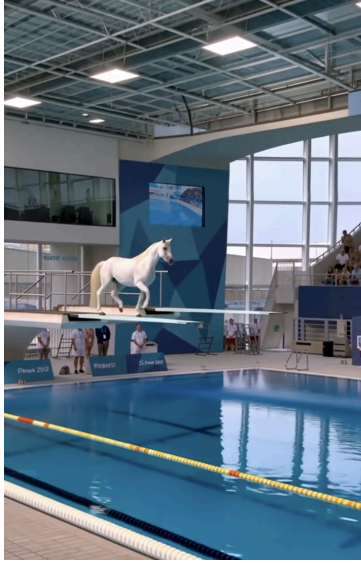
Inconsistent or unrealistic proportions between objects and/or background		
		
Figure 14A	Figure 14B	Figure 14C

Table 14. Figures 14A and 14B bring attention to the unlikely size of the two pigeons given the context and placement of the scene. While one appears to be further back compared to the other, they both appear to have the same size. Figure 14C shows an unrealistically small horse compared to the rest of the proportions of the indoor space and humans visible in the background.

2.4. UNMATCHING OR MISSING ACCESSORIES

Further synthetic artifacts to look for include earrings, makeup, and other accessories or clothing details that are not symmetrical or appear totally

different on each side. This also applies to object-specific details of clothing or background objects, such as the lack of or too much glare in glasses.


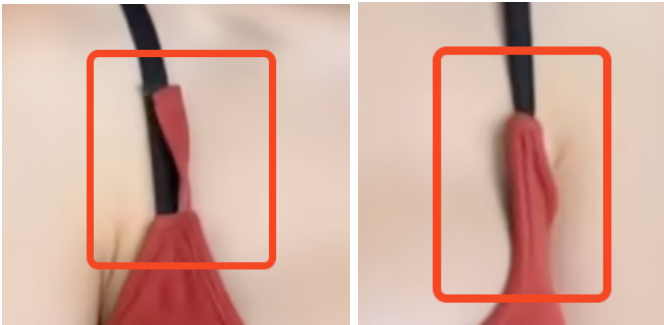
Unmatching accessories	
	
Figure 15A	Figure 15B

Table 15. This example brings attention to the lack of consistency and symmetry that can often be expected of elements of attire. The two parts of the top worn by this AI woman look strikingly different where the black straps connect to the rest of the red top. Such inconsistency is highly unlikely in reality and therefore constitutes a plausible example of an AI-generated synthetic artifact.

2.5. TOO MANY/TOO FEW BODY PARTS

Even though this detail has been a telltale sign since the inception of AI-generated imagery, it remains applicable to even the newer AI models. This can include instances where the number of fingers and/or the size and realism of the hands are incorrect; limbs, hands, and especially fingers appear and disappear during movement sequences; or there are too many, too few, or an indecipherable number of limbs that appear unlikely to be natural.

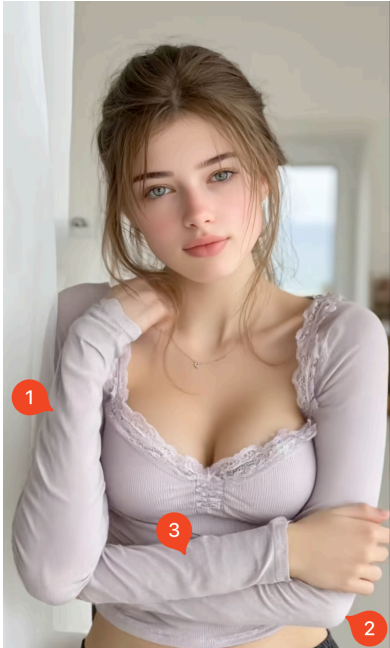


Unnatural number and/or style of limbs		
		
Figure 16A	Figure 16B	Figure 16C

Table 16. The female body in Figure 16A has too many arms. The two hands holding a glass in Figures 16B and 16C have either too few (bottom hand) or too many (top hand) fingers. The fingers merge, dissolve, and multiply around the glass.

Format and Form

There are several formal qualities of visuals that might indicate AI imagery. The first indicator and simultaneously the major challenge is that, often, AI-generated footage and images shared across platforms tend to be of low quality. This low quality, especially if appearing only in a section of the image or video compared to the rest of the footage, can be a useful indicator, raising suspicion about whether the piece of content has been manipulated or is synthetic; however, the low quality also poses a major difficulty in analyzing other details of the piece of content.

2.6. STYLIZED SMOOTH GLOW

To say that AI imagery has a particular set of stylistic qualities is an overstatement. Many of those who turn to AI tools might aim to generate a piece of content in a particular stylistic convention, or imitate the look of older



techniques (e.g. gelatin silver print photographs, early film, VHS amateur recordings, to name a few). In such cases, the stylistic elements in the piece of AI-generated content will likely lack many qualities described below. Nonetheless, the majority of so-called AI slop tends to display similar aesthetic qualities that we can broadly refer to as a particular “style.” This style is most noticeable in oversaturation of colours, unrealistically exaggerated light and shadow play, and a smooth, clay-like (“glazed”) texture.

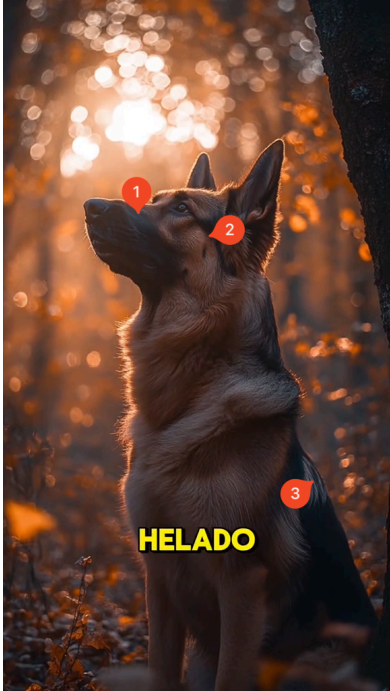

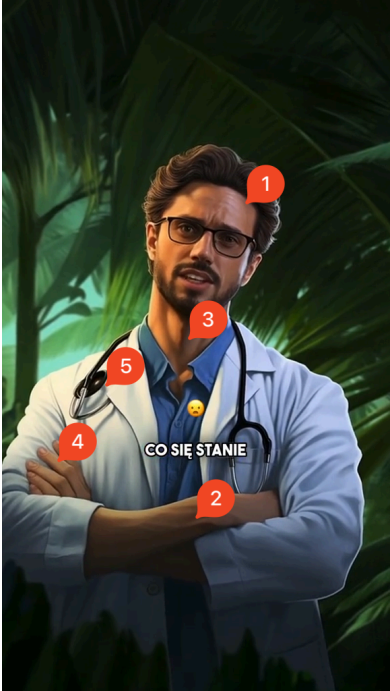
The style of genAI: artificial dramatic glow and texture plasticity		
		
Figure 17A	Figure 17B	Figure 17C

Table 17. The three images (Figures 17A-17C) display similar qualities of a stylized smooth glow that is usually a characteristic of AI imagery. Across the three figures, point no. 1 draws attention to an unrealistic and exaggerated light and shadow play, where the surface (whether human skin or dog fur) glows with white highlights and flat but strong light. Point no. 2 across all three figures draws attention to the “glazed over” texture of skin, muscles, and even hair (or fur). This glazed-over, smooth texture appears more similar to clay or dough rather than real flesh with imperfections and variations. This “texture glaze” bears resemblance to some social media “filters,” used to smooth out skin texture and often resulting in a somewhat similar glazed-over look. Point no. 3 draws attention to the inconsistency of quality and detail in the image, particularly across edges, where we can observe a blurring of otherwise sharp details alongside a subtle residue of pixellated noise. Point no. 4 in Figures 17B and 17C further

highlights the blurriness of particular areas in the image. Point no. 5 draws attention to the noise residue combined with unnatural play of highlights and shadows, both of which can be inspected upon zooming into the images.

2.7. IMPLAUSIBLE SYMMETRY, VANISHING POINTS, AND SHADOWS

The background setting of a scene can reveal its artificial origin as it often contains inconsistencies similar to the ones noticeable in foreground details: implausible symmetry, unnatural perspective, and unrealistic placement of shadows in relation to the light source. In architectural details, expectations are tied to the type of objects analyzed. In most architectural details, one can look for a lack of expected symmetry and coherence. In natural settings, strangely symmetrical elements or uncanny patterns repeating in the background can be suspicious.



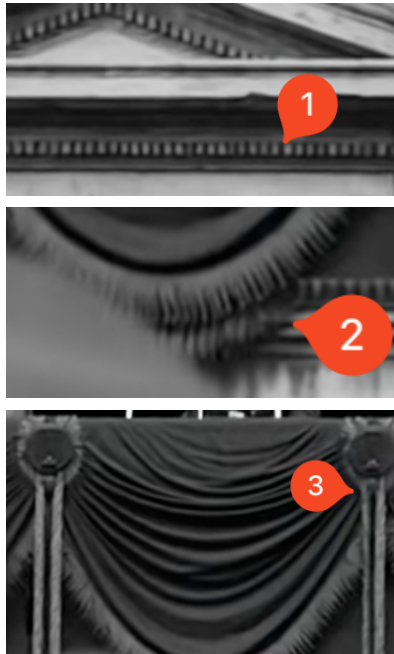
Implausible details in the background: case of architecture		
		
Figure 18A	Figure 18B	Figure 18C

Table 18. This example contains several architectural details revealing synthetic artifacts. Point no. 1 draws attention to the irregularity of shape and size of ornaments around the frieze in the pediment above the door; the low quality of the footage, however, might be blamed for this seeming lack of symmetry. Point no. 2 draws attention to the floating, hazy layer of the ornamental cloth seemingly merging and disappearing into the wall. Point no. 3 invites a comparison of the texture and



thickness of ropes between the two ornaments on each side of the cloth, drawing attention to the implausibility of these ropes varying in thickness to this extent. Finally, the straight lines visible in Figures 18A and 18B outline how details around the doorframe are not parallel to each other, showing an unlikely lack of symmetry.

## 2.8. INCONSISTENT STYLE

Some AI imagery might be made in a particular (animated) style, e.g. in the style of a cartoon like The Simpsons. The application of a particular style to all elements in the foreground and background can be scrutinized for any inconsistencies.




Inconsistencies in the application of an existing 'style'		
		
Figure 19A	Figure 19B	Figure 19C

Table 19. Figures 19A-19C show three different still images generated with AI in the style of The Simpsons. The red rectangles across the images point to several inconsistencies in how figures in the background are handled; several Christian-like ornaments in the background depict humanoid figures, some of which are rendered with more realistic proportions and details than others, yet none show the stylistic and formal qualities of The Simpsons. Figure 19B includes a character in the background whose body appears partially missing and out of proportion, with the hair showing an inconsistent colour

gradient compared to the rest of the characters. In Figure 19A, points no. 1 and 2 draw attention to too many pupil-like dots in the eyes of the characters.

3. MOVING IMAGES/CLIPS

Body Details

Inconsistency of body details in moving images can be one of the most obvious signs of AI-generated content. To be able to track these inconsistencies, it is recommended to watch the video in slow motion or frame-by-frame, as many motion-related artifacts appear only for a fraction of time. Examples to look for include unnatural facial expressions, inconsistent or “choppy” movements across a single frame, e.g. an overly rigid or excessively smooth smile, limbs suddenly changing in position or rapidly (dis)appearing.

3.1. DISAPPEARING/MORPHING BODY PARTS

These details include movements of body parts (most often, hands) that lead to a morphing of the details of these body parts with other limbs or with the background, or cause a disappearance of the body part(s) in motion.




Morphing hands in motion		
		
Figure 20A	Figure 20B	Figure 20C

Table 20. These three examples (Figures 20A-20C) show different frames from the same AI-generated clip. The red rectangles draw attention to how only one hand (the rectangles to the left) displays clear signs of morphing as the movement is executed. The rectangles positioned around the hand to the right draw attention to the more subtle detail of morphing and disappearing nails on the fingers of the hand.

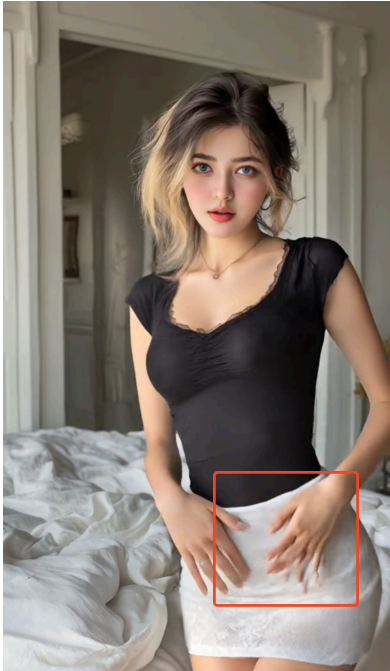
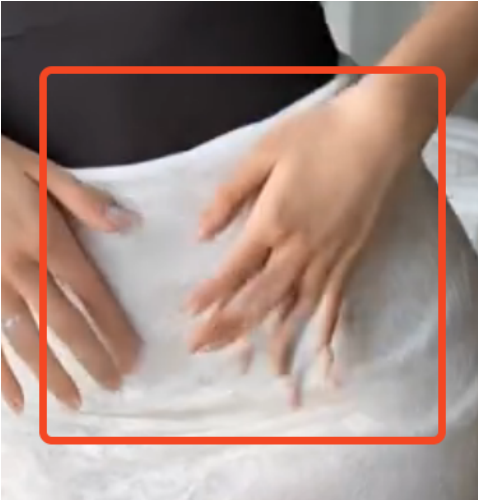
Morphing fingers in motion	
	
Figure 21A	Figure 21B

Table 21. Figures 21A-21B display another example of a hand movement which results in one of the hands showing “melting” fingers. Once the hands move, what appears is a visible residue of synthetic artifacts emerging as the visible fingers morph into each other while other finger-like elements blur next to them.

3.2. UNNATURAL BODY LANGUAGE

Other indications of AI content can be when body movements lack fluidity and seem to appear out of sync with the rest of the body. The position of the head and body might seem uncoordinated. Body movements might not correspond to actual human/animal movements or might take place “against physics.” If the person is a public figure with particular mannerisms in the way they act/speak, such mannerisms might be missing from AI-generated footage or appear highly exaggerated.






Unnatural body movements		
 <p>The Field Hospital</p>	 <p>The Field Hospital</p>	 <p>The Field Hospital</p>
Figure 22A	Figure 22B	Figure 22C

Table 22. These three figures (Figures 22A-22C) display three consecutive frames from an AI-generated video. The soldier standing at the tent entrance has his back facing the “viewer.” As he turns back, his right leg steps backwards, appearing unnaturally twisted towards the back. In the following frame, his right leg morphs with his left leg, while the left leg disappears behind it.

Facial Details

3.3. UNNATURAL EYES

Inconsistency of the appearance and movement of eyes is yet another detail that can be spotted in AI imagery. Such inconsistencies include rapid or disconnected movements between eyes, eyebrows, eyelids, lack of blinking or continuous blinking, as well as lack of consistency in the eye colour and physical symmetry.


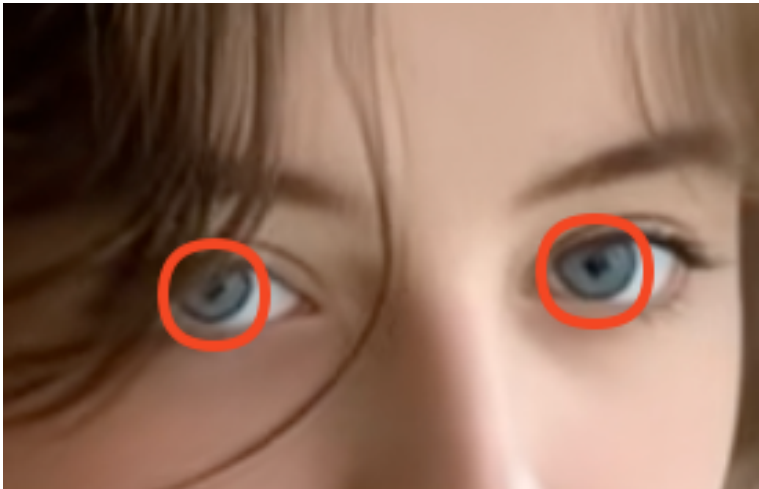
Asymmetrical irises and eyes	
 A full-body photograph of a woman with long brown hair, wearing a white tank top. Her eyes are noticeably different in color and shape, with one appearing more orange and the other more blue.	 A close-up of a person's eyes. Both eyes are circled in red. The left eye is a light blue, while the right eye is a darker, more intense blue, showing a clear asymmetry in iris color.
Figure 23A	Figure 23B

Table 23. Human eye pupils tend to be perfectly round. Similarly, human irises are also usually circular. Therefore, if the two pupils or irises in a footage are not visibly circular and alike (Figures 23A-23B), it is highly likely that the image contains synthetic artifacts.



Morphing and inconsistent eye color	
 A photograph of a woman wearing a brown hijab and a brown jacket, speaking into a microphone. A red rectangular box highlights her face. In the top left corner, there is a small circular icon of a man's face. Text overlays include 'in the street' and '@talkhand_media'.	 A close-up of a woman's face, focusing on her eyes. A red rectangular box highlights the eye area. The eyes show a morphing or inconsistent color, with one appearing darker than the other.

Figure 24A	Figure 24B
------------	------------

Table 24. In some cases of AI-generated imagery, the inner details of the eyes (such as the colour of the irises) tend to morph. In this case, one of the irises might appear to change color or “spill over” the eye, making the pupil and irises disappear and/or lack the expected light reflection off the eye surface (Figures 24A-24B).

3.4. UNMATCHING/IMPOSSIBLE LIPS/TEETH/HAIR

Focusing further on the facial inconsistencies in AI-generated footage, several types of details in lips, teeth, and hair can reveal synthetic artifacts. In the case of lips, focus on the movement can reveal several inconsistencies, e.g. when the lip movement does not match the face or does not show pauses in speech. The lips may further show unnatural movements especially when the head is turned, or any object (such as a hand) moves in front of the face. Teeth might be generated as a white strip at the bottom of the mouth unlike realistic looking teeth, or the mouth might lack any signs of tongue movement. Hair may lack any flyaways or show fizziness throughout the footage.

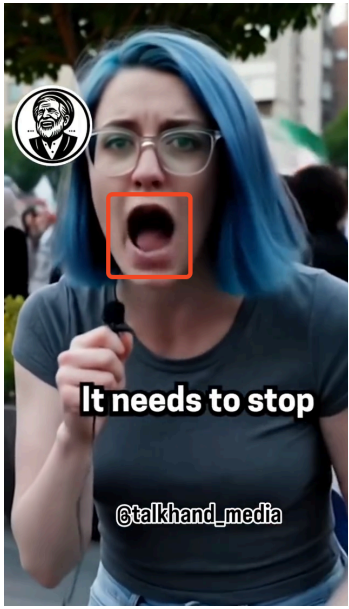

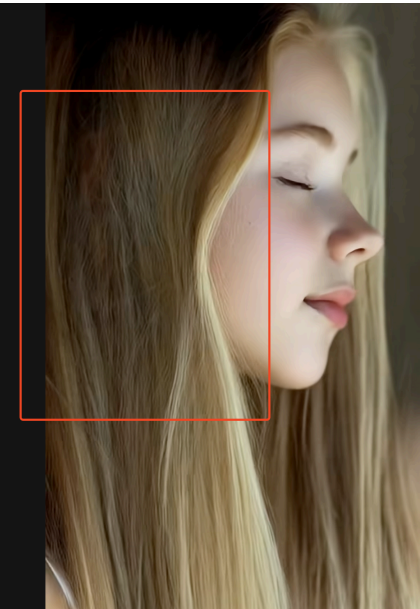
Glitching lips, teeth, and hair		
		
Figure 25A	Figure 25B	Figure 25C

Table 25. “Glitches” can often be observed between frames which include a movement of lips and hair. Figure 25A shows an example of a mouth movement where both lips and teeth disappear. Figure 25B is an example of a frame where the tongue appears to



merge with the bottom line of teeth. Figure 25C offers an example of a section of the hair showing a “glitchy” texture, where hair strands are visibly pixellated, blurry, and expanded onto the texture of the facial skin.

Object Motion Details

Similar to details of bodies and faces in AI-generated moving images, inconsistencies in objects and accessories can also indicate the likelihood of AI imagery. Such indications can be spotted both in the content foreground and background and range in size from smaller accessories or object details to broader background patterns and settings.

3.5. (DIS)APPEARING OBJECTS

Synthetic artifacts can be suspected when objects, including characters’ accessories, appear and/or disappear between frames. The objects might also morph into the background or change shapes.

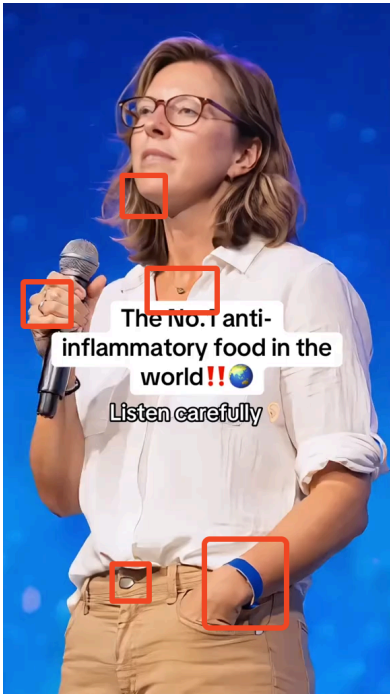


Disappearing and reappearing objects		
		
Figure 26A	Figure 26B	Figure 26C

Table 26. The following series of frames (Figures 26A-26C) show several details which give away a set of synthetic artifacts: an earring appears out of thin air, a necklace pendant changes in size and form, one out of the two finger rings disappears, the pant

button changes colour and shape, a bracelet becomes a smartwatch and morphs back to being a bracelet.




Disappearing and reappearing figures		
		
Figure 27A	Figure 27C	Figure 27C

Table 27. Figures in the background of the main action might appear and disappear seemingly out of thin air. In this series of frames from the same video (Figures 27A-27C), a human figure “materializes” out of pixelated noise (Figure 27A) as the action unfolds in the foreground and appears fully “formed” a few consecutive frames later (Figure 27C).

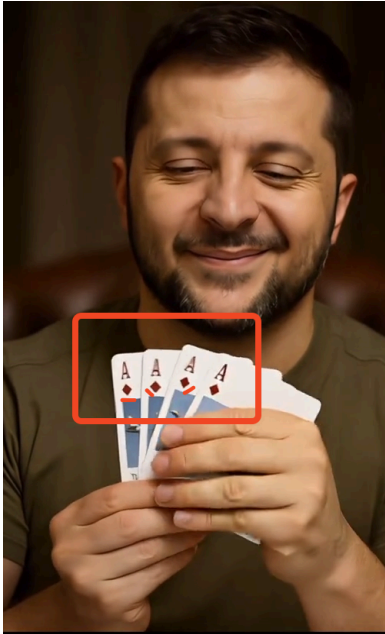
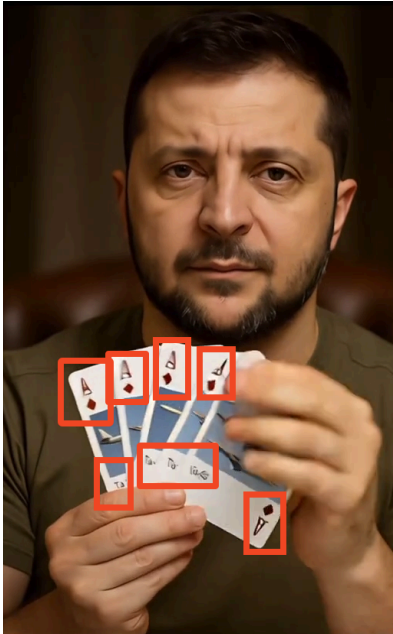

Disappearing and reappearing object details		
		
Figure 28A	Figure 28B	Figure 28C

Table 28. In this video of Volodymyr Zelensky, showing a card deck with pictures of military planes to viewers, the cards merge with one another while the writing on the cards glitches (Figures 28A-28C).

3.6. UNNATURAL PHYSICS

Other signs of AI imagery include objects moving, changing, or remaining the same against the physics of the real world, e.g. water passing through a glass, plates bouncing off or dissolving into a table instead of breaking, waves moving in the opposite direction compared to the coastline, no footprints left on the snow. Inconsistent physics also applies to unnatural movements of fabrics as well as placements of staircases, pavements, and roads that lead to nowhere.



Unnaturally occurring physical phenomena		
		
Figure 29A	Figure 29B	Figure 29C

Table 29. Figures 29A-29B show two consecutive frames from the same AI-generated footage. In the first frame, water appears perfectly still with no undulation. In the following frame, the surface of the water is immediately and completely transformed with ripples and scattered movements of the surface. Figure 29C shows an example of a frame in which the staircases in the background would be physically unclimbable in real circumstances.

## 4. DIGITAL PROVENANCE

A final step in assessing the likelihood that a visual piece of content is made using generative AI tools is to verify the content's origin. Such an assessment focuses on questioning and verifying the origin of the content, or what we find more fitting to call "digital provenance": where was the content first posted, is it similar to something else already existing online, and was it reshared?

One aid for tracking the digital provenance of visual objects is the use of search engines offering reverse-image searches. This way, a still image or a frame from a footage can be exported (or screenshotted) and used as a query. This could resurface visually-alike content, alongside references to the sources. This way, one can attempt to verify whether a piece of content was reshared across various platforms and/or accounts and whether it displays subject matter (a place or a person) that seems visually alike to what has already been shared online in the past. Another option is to search for news stories and mentions by authoritative sources (such as credible news media organizations) of an event that a suspected AI imagery claims to represent. If there is a detailed object in the frame of the suspected AI imagery piece (as described in ["2.1 Impossible Physics"](#)), it can be useful to search for authoritative images or descriptions representing this object (e.g., a telegraph from the 1930s) to compare its visual details and specifics.

This step, similar to the previously mentioned steps, is unlikely to provide definitive answers and certainty on its own; however, combined with other elements of detection-driven analysis, it can offer additional insights otherwise unattainable from a single platform's interface and content-only analysis.

