



















2026 Synthetic Advertising Report



RESEMBLE.AI

METHODOLOGY

2026 Synthetic Advertising Report

Based on 37,768 scanned ad files

How this report was built: All figures come from Resemble Detect scans of ads pulled from Meta’s public Ad Library. Detect checked every image and video for AI generation or manipulation. The dataset covers:

- 37,768 image and video files from 15,433 ads, run by 2,298 Facebook and Instagram pages
- 509 companies across 26 industries, found by searching each company name in the Ad Library in September 2026
- A second collection in September 2026 covering nine prediction market, sportsbook and crypto brands, scanned for audio as well as visuals
- Headline rates use only ads a brand ran from its own page: 28,463 files from 389 brands. Ads from other pages that matched the search are reported separately.
- Every rate is a floor. The main collection was scanned for visual manipulation only, and audio alone caught 62 of the 142 fakes in the betting and crypto collection.

Each file is classified by:

1. Verdict: a file counts as AI when Detect returns a Fake verdict
2. Advertiser type: whether the brand itself ran the ad, or another page using its name
3. Content: what each fake shows and how it is styled, labeled for all 3,216 fakes in the main collection

Placement data is not reported. Meta keeps stopped US commercial ads out of the public Ad Library, so placements for most ads in the study could not be recovered after the scan.

ADVERTISER TYPE

- Searched brand
The company's own page
- Other advertiser
Retailers, affiliates, creators and unrelated pages that matched the search
- Bait page
Clickbait, search arbitrage and scam pages that use a brand name as a lure

AI STYLE

- Photoreal
Built to pass as a real photo or video
- UGC-style realism
Made to look like a customer’s own phone footage
- Graphic or UI
Designed layouts, text promos and app screens
- Obvious AI or CGI
Stylized output no viewer would take as real

DETECTION COVERAGE

- Image
All files
- Video frames
All files
- Audio
Betting and crypto only
- C2PA
971 files checked

EXECUTIVE SUMMARY

The 7 biggest takeaways from 37,786 files running on Meta

1 in 11

ad files carried AI-generated or manipulated media

3,379 of 37,768 files were flagged by Detect, across all advertisers and all 26 industries in the study.

6.9%

of files brands ran from their own pages were AI (1,971 of 28,463).

57%

of brands ran at least one AI ad (222 of 389). Brands with two or more averaged 12.5%.

70%

of fakes were photoreal, built to pass as a real photo or video (2,259 of 3,216).

0.06%

of AI files disclosed AI use in the ad copy (2 of 3,379). None of 971 files carried C2PA provenance data.

71%

of files from bait pages were AI, ten times the rate on brand pages (463 of 649).


44%

of fakes in the betting and crypto ads were caught only by their audio, mostly ElevenLabs voices.

More than half of companies used AI-generated content. For a few, it was the majority of the ads.


Detect flagged 3,379 of 37,768 ad files (8.9%). On brand pages the rate was 6.9%, but use was uneven: 167 brands ran none, while 13 brands used AI in more than 40% of their ads.

AI rate by who ran the ad (share of files flagged fake)



Bait pages use a brand name as a lure but are not run by the brand. n = 37,540 files.

389 brands by the share of their own ads that were AI




Counts of brands. Brand-page files only.

A small group of advertisers were heavy generative AI users. Temu (64 of 100 files), Wayfair (56 of 100), Motel 6 (50 of 80) and Lufthansa (50 of 94) each used AI in more than half of their sampled ads. The 185 brands with two or more AI ads averaged 12.5% of total ads containing AI, meaning those who used AI used it about a dozen times on average.

Consumer goods and electronics used the most Gen AI content, fast food did not

All 26 industries ran AI ads. Rates ranged from 13.4% in consumer packaged goods to 2.4% in quick service restaurants. Regulated categories sat below the 6.9% average: insurance, banking and pharma all came in at 5.0% or lower.

Share of brand-page files flagged as AI, by industry




n = 28,463 brand-page files across 389 brands. Audio was scanned only for the nine betting and crypto brands.

Regulated categories used AI least. Banking (3.8%), nonprofits (3.3%) and pharma (2.8%) sit at the bottom alongside fast food. Consumer goods, electronics and fitness sit near the top, all categories built on product and lifestyle imagery.

Regardless of employee count or marketing team size, companies used AI at about the same rate


We matched 313 parent companies to their total headcount and marketing staff to see whether smaller companies or leaner marketing teams would rely more heavily on AI. However, AI rates stayed between 6% and 9% at every size. Neither company size nor marketing team size predicted how much AI a brand used.

AI rate by total employees



Brand-page files, pooled by parent company.

AI rate by marketing staff



Marketing staff from employee profiles in Apollo.



Employees	Companies	AI rate	Ran any AI
<1K	37	7.5%	65%
1K-10K	86	8.9%	55%
10K-100K	141	6.1%	61%
100K+	49	6.4%	65%

Marketing staff	Companies	AI rate	Ran any AI
<100	99	7.5%	56%
100-299	84	7.3%	58%
300-999	73	7.1%	64%
1,000+	57	5.9%	67%

AI use isn't a resourcing story. Temu runs a \$54B business with 232 marketers and used AI in 64 of 100 ads. Amazon has nearly 23,000 marketers and used none. Larger teams were somewhat more likely to try AI at least once (67% against 56%), but their overall AI rate was slightly lower, 5.9% against 7.5%.

Many AI ads depict hyper realistic looking people


Seven in ten fakes were photoreal, and 46% showed a person. The most common subject was a realistic model, followed by product shots and scenes. Only 1 in 9 fakes was stylized enough that a viewer would read it as AI.




Brands fake polish. Bait pages fake people. On brand pages, 2% of fakes imitated customer footage. On bait pages it was 11.2%, used to make clickbait look like a real person's story.

The "customer" in many ads was generated


Brands used AI to make testimonials and selfie videos that look shot on a phone. Resemble Detect flagged the below examples as AI-generated with high confidence.




99.9% certainty
Teladoc. Telehealth patient in a waiting room, UGC style.



99.8% certainty
MoneyLion. Phone chat video, UGC style.



99% certainty
Dave. Driver selfie testimonial for a cash advance.




99.9% certainty
Tarte Cosmetics. Selfie video for a tinted serum.

Gen AI imagery was used more than video, and AI-generated music and voice was prominent


In the main collection, brand images were flagged at more than twice the rate of brand videos. When the betting and crypto ads were also scanned for audio, 62 of their 142 fakes showed no visual sign of AI.

Main collection: AI rate by format




Brand pages, visual scan only. 1,489 of 17,857 images and 340 of 9,706 videos.

Betting and crypto: how each fake was caught



Count of fakes. n = 142 brand-page fakes.

How fakes were caught, by segment





Count of fakes per segment.

Video rates in the main collection are a floor. Without an audio scan, a real video with a generated voiceover reads as authentic. The next section breaks down what that audio was.

Most AI audio was voice, combined with real footage


Of 67 fakes flagged on audio, 57 were synthetic voices and 10 were music or sound effects. 54 of the 57 voices played over video Detect rated as real, so the picture was real and the voiceover was not. By eye and ear, these were generally not deceptive or intended to be deceptive, and were used primarily as an editing technique.



Not all undisclosed AI is the same. A generated backing track and a generated spokesperson could be meaningfully different to different viewers. People can decide for themselves whether AI music matters to them. A synthetic voice presented as a real person has a lot more potential to come across as deceptive or misleading than a voiceover describing a scene.

Real footage with generated voices

In each of these ads Detect rated the video as authentic and the voice as synthetic.




99.9% audio score
AUDIO FAKE VIDEO REAL 11% ELEVENLABS

“Get more crypto for your cash and safely set up recurring buys.”

Robinhood UK. ElevenLabs voiceover on a crypto recurring-buy ad.

[Ad Library](#) · [Detect report](#)




86.9% audio score
AUDIO FAKE VIDEO REAL 1% ELEVENLABS

“Drop \$10 in Novig, you’ll get an extra \$25 in trade credits.”

Novig. Creator-style presenter with a synthetic voice.

[Ad Library](#) · [Detect report](#)




68.7% audio score
AUDIO FAKE VIDEO REAL 1% ELEVENLABS

“The only sportsbook to reward every bet. Win or lose.”

Fanatics. AI voiceover persona on a sportsbook promo.

[Ad Library](#) · [Detect report](#)



99.2% audio score
AUDIO FAKE VIDEO REAL 1% ELEVENLABS

“The new eToro app is about fast execution.”

eToro. App launch video, voice flagged synthetic.

[Ad Library](#) · [Detect report](#)

Almost no AI-generated content was labeled

We searched the copy of every ad for any mention of AI generation, and checked files for C2PA content credentials, the industry standard for provenance. Advertisers almost never said when an ad was made with AI.

0.06%

of AI files disclosed AI use in the ad copy, 2 of 3,379.

1

brand disclosed on its own ad. Turkish Airlines noted its visuals were created with AI.


0

of 971 files carried C2PA provenance credentials, including every confirmed fake.

7 of 8

AI-made presenters in Kalshi’s ads showed an on-screen AI notice. Our copy check can’t see those, so on-screen labels are counted only for Kalshi.

THE ONE BRAND THAT DISCLOSED



Turkish Airlines
International Cat Day video for its Touristanbul program. We detected is as AI-generated with 99.9% confidence.


The ad copy ends with a footnote: the visuals “were created using artificial intelligence technology.”

The note sits below the caption, after the hashtags in most feeds. However, it was disclosed, and shows how easy it would be to include in any and all ads.

Bait pages used brand names to sell scams or fake stories


Searching a brand's name also surfaced pages the brand doesn't control. These bait pages ran 649 files, and 71% were AI. They made up 2% of files but 14% of all fakes, and three in four of their fakes were clickbait advertorials.

What bait-page fakes were selling (share of 463 fakes)




Ad type labeled from the creative and copy.

Top subjects: brand pages



Share of brand fakes.

Top subjects: bait pages



Share of 463 bait fakes.

A prominent YouTuber found his background and body used in an ad

On September 23, YouTuber Elliot Choy saw a Kalshi ad built from his 2025 apartment-move video. The script, room and framing matched his. The man on screen did not: AI had replaced Choy, who is Asian American, with a white man who doesn't exist. Kalshi told NPR the ad was old, that it no longer runs AI-generated ads, and that it was reviewing its relationship with the agency that made it.

CHOY'S ORIGINAL VS. THE KALSHI AD



DETECT SCAN OF THE AD

IMAGE ANALYSIS

Verdict

Score

Type

AI-GENERATED

0.9011

FINALRESULT




Verdict AI-generated, score 0.90. Red overlay marks regions Detect attributes to generation.

He wasn't the first. Creator Pushpek Sidhu said a betting ad in June copied one of his videos almost line for line and also swapped him for a white man. In both cases the person whose work was used found out by recognizing it.

Crypto ads were the most synthetic in the study


We added nine prediction market, sportsbook and crypto brands and scanned their ads for audio as well as visuals. 142 of 900 brand-page files were AI (15.8%), more than double the study average. Crypto and brokerage brands drove most of it.

AI rate by segment



Brand-page files, visual and audio scan.

AI rate by brand



OKX ran AI in nearly half its ads. 48 of 99 files were flagged as AI-generated, most on visuals. DraftKings was at the other end of the spectrum, with only 2 of their 98 most recent ads showing AI use..

Real athletes and creators, synthetic media

Several flagged ads featured real people. FanDuel ads with Rob Gronkowski and Bussin' With The Boys were flagged, as were Kalshi images of Marshawn Lynch. Every one of Robinhood UK's 21 flagged ads used an ElevenLabs voice.


21 of 21

Robinhood UK fakes used an ElevenLabs voice.

55

voices in this collection traced to ElevenLabs, out of 67 audio fakes.

KALSHI: EVERY ACTIVE US AD, SEPTEMBER 24, 2026



52 of 511 files flagged (10%), in 50 of 397 ads (13%).

40 videos carried a synthetic voice: 23 traced to ElevenLabs and 16 to Cartesia.

7 of 8 AI-made presenters showed an on-screen AI notice.

0 of 511 files carried C2PA credentials.

Where advertisers have to disclose AI today

Only a few places require an ad to say it uses AI. Most rules cover generated people, not AI music, voices or scenery, and most put the duty on the advertiser rather than the platform.

Region	Rule	What it requires	Status
South Korea	KFTC endorsement guidelines	Ads using an AI-generated virtual person must label it as virtual, on screen while the figure appears	In force, Jun 1 2026
New York	Gen. Business Law §396-b	Ads must conspicuously disclose a synthetic performer the advertiser knows about. Audio-only ads and platforms are exempt.	In force, Jun 9 2026
European Union	AI Act, Article 50	Machine-readable marking of AI output and visible labels on deepfakes	Aug 2 2026 [verify]
China	AI content labeling measures	Visible and machine-readable labels on AI text, images, audio and video	In force, Sep 1 2025
India	IT Rules amendment	Labeling and takedown of synthetically generated content	In force, Feb 20 2026
California	SB 942, AI Transparency Act	Gen AI providers must offer detection tools and embed provenance data	[verify date]
Tennessee	ELVIS Act	Consent required to use a person’s voice or likeness in an AI replica	In force
US federal	FTC Act, NO FAKES Act	No AI ad disclosure rule. FTC deception rules apply. NO FAKES is pending.	Partial
United Kingdom	ASA advertising code	No AI-specific disclosure law. Misleading ads rules apply.	[verify]

Even while legal, lack of disclosure in ads could be problematic

Elliot Choy found out because he recognized his own apartment. **Recognizing yourself or your apartment in an ad shouldn't be how the public learns that AI was used in advertising.**

In our data, 0.06% of AI files disclosed it in the ad copy. Waiting for advertisers to volunteer has not worked. Disclosure has to be built into each step between the model and the screen.

Model providers watermark	Platforms read and label	Advertisers declare
Every image, video and voice a model generates should carry an imperceptible watermark and provenance data that survive editing and compression.	Ad platforms should check uploads for those marks and for synthetic media, and add an AI label at serving time when they find it.	Advertisers should have to select an AI disclosure when they upload, the way political advertisers already declare digitally altered media.

CONSENT

Disclosure doesn't solve consent. A label would have told viewers the Kalshi ad was generated. It would not have asked Choy before his video was used. Both problems need an answer.

How each figure was calculated

Figure	Calculation	Base
8.9% overall	3,379 fake ÷ 37,768 files with a completed scan	All advertisers
6.9% brand pages	1,971 ÷ 28,463 files run from the searched brand's own page	389 brands
57% of brands	222 brands with at least one fake ÷ 389 brands with brand-page files	Brand pages
12.5%	Pooled fake rate across the 185 brands with two or more fakes	Brand pages
70% photoreal	2,259 ÷ 3,216 labeled fakes	Main collection fakes
0.06% disclosed	2 AI files whose ad copy mentions AI generation ÷ 3,379 AI files	All AI files
71% bait pages	463 ÷ 649 bait-page files	Bait pages
62 of 142 audio only	Fakes with an audio Fake verdict and no visual Fake or Likely fake verdict	Betting and crypto brand pages
Size and marketing	Total employees, marketing staff and revenue from Apollo company enrichment, matched to 313 parent companies covering 363 brands	Brand pages
Brand page labels	718 files moved to Other advertiser where the page belonged to a different company or a standalone creator	Applied before all rates
Exclusions	Fanatics merch app (100 files), duplicate media URLs (34), subsidiary pages (9), failed scans (139)	Removed from all rates

Scale, content, disclosure and the betting feature

Scale

- 37,768 files from 15,433 ads, 509 companies and 26 industries
- 1 in 11 files flagged as AI (8.9%), 6.9% on brand pages
- 222 of 389 brands ran at least one AI ad

Content

- 70% of fakes photoreal, 46% show a person
- Realistic models (26%) and product shots (24%) were the top subjects
- Bait pages: 71% of files AI, 14% of all fakes

Audio

- 62 of 142 betting and crypto fakes caught only by audio
- 57 of 67 audio fakes were synthetic voices, 55 traced to ElevenLabs

Disclosure

- 0.06% of AI files disclosed AI use (2 of 3,379)
- 0 of 971 files carried C2PA credentials

Betting and crypto

- 142 of 900 brand-page files AI (15.8%), crypto and brokerage highest at 29.3%
- Kalshi page scan: 52 of 511 files (10%), 0 with C2PA

COMING NEXT: THE UGC AND INFLUENCER REPORT

Deepfake detection models from Resemble AI

Make confident, scalable decisions instantly.

