

1 ShadowTrace Analysed on device

What changed in this recording?

These measurements help you review the recording; they do not establish authenticity or how it was made.

File under review

Video moving-train-nanawyawbotar-cc-by-sa-4.webm

► [Download report](#) [New check](#)

2

Duration	Sampled frames	Measured events	Sample interval
20.53 s	42	7	0.5 s

3

[moving Train, NanaYawBotar, Wikimedia Commons CC BY-SA 4.0](#)



Extracted frame time: 3.15 s, Distance from selected moment: 0.35 s

► [Open original recording](#)

Motion increase

Frame-to-frame motion energy rose above its usual level.

Moment under review 3.5 s

Value 17.2%

Threshold $\geq 16.4\%$

Frame change (%) 17.2%

RMS (selected window) -8 dBFS

► [Detailed measurements](#)

How ShadowTrace is built

ShadowTrace measures how an image, a video or an audio recording changes over time. This handbook is about how it is built. It explains where the work happens, where the file does not go and why it gives no verdict. I do not go into the usage steps here.

1 What I built it to solve

When a video reaches me, I first want to know what changed in the recording, and when. I look for a cut, a drop in the sound or a picture that breaks up. Watching and taking notes by eye is slow and cannot be repeated.

ShadowTrace measures the file and ties every measurement to the file's SHA-256 identity and a timestamp. Then it asks a person to review each event one by one. It does not answer whether the file is real, fake or made by AI.

Step-by-step use of the tool is in a separate handbook, the ShadowTrace field guide. This one covers only the structure.

A few words

SHA - 256	A fingerprint computed from the file's content. If a single byte changes, this value changes too.
Spectrogram	An image that shows how strong each frequency of the sound is, spread over time.
STFT	A method that cuts sound into small windows and works out the frequencies in each one.
OCR	Turning text in an image into machine-readable text.
C2PA	A standard for a signed record embedded in some files that describes how the file was made.

2 What runs where

The whole analysis happens in the browser. The file you pick is not uploaded to a server. I built it this way on purpose. The site stays static and no paid analysis service is called.

Reading the file	The browser opens the file in its own memory. Images and video are drawn on a canvas and read pixel by pixel.
Identity	The SHA-256 value is computed over the whole file with the browser's own crypto interface (Web Crypto).
Audio	Sound is decoded and measured with the browser's offline audio engine (OfflineAudioContext).
OCR and C2PA	Text in key frames and any C2PA record in the file are also read in the browser. The OCR engine and the Turkish and English language files are downloaded from a public code host (jsDelivr). Your file is not sent there.
Storage	The result and the reviewer ledger stay in memory only while the page is open. No account, no server record.
Output	The ZIP and reports are built on your computer and saved there. The original media is not added to the pack.

File	Browser	Person	Pack
Image, video or audio. Up to 8 files.	Identity, pixel, frame, motion, audio and shared-time measurement.	Each event is reviewed one by one and a note is added.	Report, measurements and a digest list are saved to your device.

3 The six-stage flow

Every file goes through the same open flow. If a stage cannot run, it is not hidden. It is marked unavailable and the reason is written down. That way a missing measurement never quietly turns into a confident result.



The flow run on the training file. The sample video on the page is a train recording from Wikimedia Commons under CC BY-SA 4.0.

1 File identity

SHA-256 is computed first. Every later measurement is tied to this identity.

2 Audio signal

Level, clipping and frequencies are measured in small windows.

3 Shared time

Events in the picture and in the sound are placed side by side when they are close in time. If the sample interval is too wide, this stage is marked unavailable.

Limits written in the code

Files	Up to 8.
Size	Image 40 MiB, video 500 MiB, audio 160 MiB.
Video sampling	About 2 frames per second, at least 18 and at most 96 frames. For measuring, frames are scaled so the long edge is 96 pixels.
Audio decoding	Only audio up to 64 MiB and 5 minutes is decoded. It stops if decoded data passes 192 MiB.
Spectrogram	A 2048-point window, 64 columns, 24 frequency bands.

4 The result screen

1ShadowTrace Analysed on device

What changed in this recording?

These measurements help you review the recording; they do not establish authenticity or how it was made.

File under review

Video moving-train-nanawayawbotar-cc-by-sa-4.webm

► Download report New check

2

Duration	Sampled frames	Measured events	Sample interval
20.53 s	42	7	0.5 s

3Moving Train, NanaYawBotar, Wikimedia Commons CC BY-SA 4.0



Extracted frame time: 3.15 s, Distance from selected moment: 0.35 s

► Open original recording

Motion increase

Frame-to-frame motion energy rose above its usual level.

Moment under review	3.5 s
Value	17.2%
Threshold	≥ 16.4%
Frame change (%)	17.2%
RMS (selected window)	-8 dBFS

► Detailed measurements

The signal ledger, after the training file was analysed in the browser.

1 Analysed on device

The header says it plainly. There is no analysis step on the page that goes to a server.

2 Summary counts

Duration, sampled frames, measured events and sample interval. An event is a moment where a measurement passed its threshold. Passing a threshold does not mean a problem was found.

3 Source and licence

Where the training file comes from and its licence are written right below the file.

5 File identity and the reviewer

▼ File identity and method	
File	moving-train-nanawyawbotar-cc-by-sa-4.webm
1 SHA-256	02925e7b929e0403841501fe6cea4b50c13a72ab20414e98402f42228bbba4e0
2 Case ID	ST-02925E-E6F138
File last modified	24 Feb 2020
Analysis date	30 Sept 2026, 12:16 UTC
MIME	video/webm
STFT configuration	2048 FFT, 256 Window count, 58%


The file identity and method section.

1 SHA-256

Every measurement in the report is tied to this value. Someone else reviewing the same file should get the same value.

2 Case ID

Starts with the first six characters of the SHA-256. It ties the report and the pack to the same review.



The reviewer check. Each time event is marked pending, reviewed or needs follow-up.

1 Only on this device

The reviewer ledger and notes do not go to a server. They are lost when the page closes, but go into the ZIP exactly as written.

2 Review progress

Counts how many events have been handled. It is not an accuracy or authenticity score. The order of events is not a risk score either.

6 The delivery pack



The report download section.

1 What is inside

A technical HTML report, the reviewer ledger, file codes, frame and audio measurements, a timeline, a spectrogram CSV, JSON and Markdown. The original files are not added.

2 Download the ZIP

The pack is built in the browser. The digest list inside it lets you check later whether any file has changed.

If you want, you can add a client label and get a delivery copy with a visible watermark. This is also made on your device. The watermark is not copy protection, proof of ownership or an authenticity verdict.

7 Design decisions and limits

- × No verdict. Turning general signal measurements into a real or fake probability would not be scientifically sound. The tool does not do it.
- × The motion path is the centre of the changed pixels. It is not tracking of a person or an object.
- × The frame nearest to an event is shown with how far it is from the event. It is not presented as the exact moment.
- × It does not download files from social media links. You have to pick the file yourself.
- × Which video and audio formats open depends on your browser.
- × In CSV files, starts that a spreadsheet could run as a formula are defused.

In short, ShadowTrace is a measurement ledger. It shows what changed and when. What that means is decided by the person reviewing it.