



MODULE 2 – Digital Skills for Managing Disinformation

Partner: FORMA.Azione srl

Authors: Adina Curta, Sylvia Liuti



Co-funded by
the European Union



Summary

Introduction	03
.....	
Toolkit Overview	04
.....	
Image and Video Verification	07
.....	
Website and Domain Analysis	17
.....	
Social Media and Account Analysis	21
.....	

Introduction

Every day, health advice circulates across social media, messaging apps, and websites – on TikTok, Instagram, YouTube, WhatsApp, and countless blogs and forums. Some of this content is accurate and useful. Much of it is not. The challenge is that misleading health information is often designed to look convincing, and the skills needed to identify it are rarely taught explicitly.

This module is grounded in the five national surveys conducted across the partner countries between June and September 2025. The findings reveal a consistent and urgent picture: the ability to identify health disinformation is critically low across all five countries. The problem is not motivation or distrust in authorities – it is a missing practical toolkit.

Highlights from the comparative report

- **Italy:** 64% of respondents cite misleading online content as their primary barrier to reliable health information; 42% are uncertain whether they have ever shared false health information;
- **Lithuania:** 48% say lack of skills to assess reliability is their main barrier; only approximately 1 in 3 consistently fact-checks despite awareness of misinformation;
- **Croatia:** only 5.6% of respondents could correctly classify false health information. 80.4% had encountered fake medicines or miracle cures online.
- **Bulgaria:** 83% had encountered miracle product disinformation; uncertainty and contradictory information – not language complexity – are the main obstacles.
- **Cyprus:** 53.7% cite the overwhelming presence of misleading online content as the main obstacle to informed health decisions.

This module directly addresses this gap. It equips learners with a concrete, replicable set of tools they can apply immediately to any health claim encountered online. The three categories covered are the verification of images and videos, the analysis of websites and domains, and the assessment of social media accounts.

Learning Objectives

After completing this module, you will be able to:

- › Use reverse image search across at least two platforms to verify the origin and authenticity of health-related images and videos;
- › Use video frame extraction and metadata analysis to identify false-context claims in health videos;
- › Detect digital manipulation in image files using Error Level Analysis;
- › Analyse a URL to identify the real domain, detect lookalike spoof sites, and check domain registration history;
- › Apply lateral reading to assess the credibility of any health website in under 60 seconds;
- › Identify structural and behavioural markers of inauthentic health accounts on Instagram, TikTok, and YouTube;
- › Apply these tools consistently as habits when encountering health content online.

Toolkit Overview

This module covers nine tools and methods across three categories. Each tool is introduced in a dedicated section with step-by-step instructions, opening examples, practice exercises, and key reminders. The tables below provide an at-a-glance reference before each tool is covered in depth.

3.1 Image and Video Verification

Health disinformation frequently uses real images and videos in the wrong context, or media that has been digitally manipulated. These tools let you trace any image or video back to its original source and detect manipulation.

Tool	What it does	Best for
Google Images / Lens	Reverse image search: finds every place an image has appeared online and when	Verifying health product photos, before/after images, supposed proof photos
TinEye	Reverse image search sorted oldest-first; shows the true origin of any image	Finding when and where a viral health image first appeared
InVID / WeVerify	Browser extension: breaks videos into frames for reverse-searching; reads video metadata	Verifying health advice videos on TikTok, Instagram, YouTube
FotoForensics	Error Level Analysis (ELA) reveals digital manipulation in image files	Detecting doctored screenshots of medical results or altered statistics

3.2 Website and Domain Analysis

A website can look completely professional and still be a disinformation source. These methods let you assess the credibility of any health website in under 60 seconds, without reading a single article.

Tool	What it does	Best for
URL structure reading	Identifying the real domain by reading right-to-left; detecting lookalike spoof domains	Any health website before trusting its content
WHOIS / who.is	Shows when a domain was registered, who owns it, and where it is hosted	Exposing sites that claim authority but were created last month
Lateral reading	Opening new tabs to search about a source rather than reading within it	Quickly establishing whether a website or author is credible

3.3 Social Media and Account Analysis

Inauthentic accounts and coordinated networks are used to make health disinformation appear more popular and credible than it is. These tools help you assess authenticity on Instagram, TikTok, and YouTube.

Tool	What it does	Best for
Profile behaviour checklist	Manual assessment: account age, follower ratio, posting uniformity, comment patterns, follower growth	Identifying bot accounts and fake health influencers on Instagram, TikTok, YouTube
Social Blade	Tracks follower growth history over time; sudden spikes indicate bought followers	Exposing inauthentic growth in health supplement and wellness accounts

Image and Video Verification

The sections below cover four tools for verifying the authenticity and origin of health-related images and videos. Boxes 3.1.1 to 3.1.4 can be read in sequence or consulted individually when a specific verification task arises.

3.1.1 Don't Trust the Image — Google Images / Lens

What is Reverse Image Search?

Reverse image search works in the opposite direction from a normal search. Instead of typing words to find images, you provide an image, and the tool finds every place it has appeared on the internet and when. This lets you:

- find the original source and publication date of any image;
- discover if an image is being used in a completely different context from its original one;
- identify stock photos presented as authentic evidence for a health claim;
- detect images created or circulated long before the claim they supposedly support;

How to Use Google Images / Lens: Step by Step



Open the webpage or social media post containing the image. Right-click on the image and select “Search image” or “Search image with Google”. A new tab opens with results showing every known appearance of that image. Look at where the image appears, the dates of those appearances, and whether any source is a stock photo library.



Go to images.google.com. Click the camera icon in the search bar. Choose “Upload a file” if you have saved the image, or paste the image’s direct URL. Click “Search by image” and review results as in Method A.



In a mobile browser, hold-press (long-press) on the image and select “Search image” or “Search with Google Lens”. For TikTok or Instagram where hold-press does not open a menu: take a screenshot of the image, then go to images.google.com and upload the screenshot using Method B.

Reading the Results: What to Look For

What you find	What it means
Image found on a stock photo website	The image is generic, not authentic documentation of the specific claim. Strong red flag.
Same image on dozens of sites within 24–48 hours	Content went viral before verification. Viral spread is not proof of truth.
Image predates the claim by months or years	The image is being used in false context. The claim is built on recycled or stolen media.
Image appears only on the original post	Inconclusive. Continue with TinEye and other tools.
Image traced to a credible news organisation or body	May be authentic. Check the original article to verify the full context.

Try it now:

- Find any health-related image from a social media post or article you have seen recently.
- Use Google Images (Method A or B) to reverse-search it.
- Write down: where does the image originally come from? Does that source match what the post claims?

Key reminder

- Reverse image search takes 45 seconds. Make it a habit before sharing any health image.
- Viral spread is not proof of truth. An image appearing on 50 sites in 24 hours needs more scrutiny, not less.
- Always look for the original source and the original date – not just whether the image is widely used.

3.1.2 Finding the True Origin – TinEye

TinEye specialises in one feature Google does not prioritise: sorting results by date, oldest first. This makes it the most reliable tool for answering the key question about any health image:

“When did this image first appear, and what was it originally used for?”

TinEye searches its own index of images, which is large but not complete. If an image returns zero results, it does not mean the image is authentic, it means TinEye has not indexed it. In that case, also use Google Images and InVID (section 3.1.3).

How to Use TinEye: Step by Step

- 1 Open tineye.com in your browser. No account or login is needed.
- 2 Click “Upload” and select the image file from your device, OR drag the image directly onto the page, OR paste the image’s direct URL into the search bar. Click the search icon.

- 3** When results appear, find the “Sort by” dropdown in the top-right area of the results page. Select “Oldest”. The first result is now the earliest known appearance of this image on the internet.
- 4** Click through to the oldest result. This is the image’s likely origin. What was it originally about? Who published it? Does it have any connection to the health claim now using it?
- 5** Compare the date of the oldest result with the date of the health claim. If the image is months or years older than the claim, it is being used in false context.

Real-world example:

Politicians and activists also recycle old images with new captions to fabricate events. In health disinformation, the same tactic is used: a photo of a real hospital, a real patient, or a real laboratory is given a completely new caption to support a false health claim. TinEye makes this manipulation visible by surfacing the original caption and context.

Key reminder

- Always sort TinEye results by Oldest — this is the feature that makes it more powerful than Google for origin detection.
- Zero TinEye results do not mean an image is authentic. Use multiple tools.
- The original caption and original context of an image are the most important facts to find.

3.1.3 Verifying Health Videos – InVID / WeVerify

Video disinformation is harder to check than image disinformation because a video contains many frames, each of which is potentially an image that can be searched. You cannot right-click a video the way you can an image. InVID and WeVerify solve this by:

- breaking a video into individual frames that you can then reverse-image-search one by one
- reading the video's embedded metadata – invisible data that can reveal the original recording date, device model, and sometimes location
- extracting thumbnails for platform-specific searches.

1

In Chrome: open the Chrome Web Store (chrome.google.com/webstore). In Firefox: open Firefox Add-ons (addons.mozilla.org).

2

Type "InVID WeVerify" in the search bar.

3

Click "Add to Chrome" or "Add to Firefox". The extension installs in seconds. A small icon appears in your browser toolbar.

How to Use InVID/ WeVerify: Step by Step

1	On TikTok, Instagram, YouTube, or any platform: find the share button and copy the link to the video.
2	Click the InVID/WeVerify icon in your browser toolbar. The tool panel opens.
3	Paste the video URL into the tool. Select “Keyframes” or “Video Analysis”. The tool extracts a series of still frames from the video at regular intervals.
4	Right-click on any extracted frame and select “Search image with Google” or drag the frame into Google Images or TinEye – identical to the process in sections 3.1.1 and 3.1.2. The question to answer: is this location what the video claims? Has this footage appeared before in a different context?
5	In the InVID tool, look for the Metadata tab. This shows any embedded data in the video file, which may include the original creation date, GPS coordinates, and the device used to record it. If the metadata says the video was created two years ago but the post claims it was filmed last week, a false-context claim has been identified.

What to Look for in Health Video Verification

- **Location claims:** does the background actually show what the creator claims? A person claiming to film inside a hospital may be standing in any corridor. Reverse-search a frame of the background to check if this location has appeared in any other context.

- **Identity claims:** a person claiming to be a doctor, scientist, or health authority can be checked. Take a clear frame showing their face or any visible name badge, then use image search alongside lateral reading (section 3.2.3) to search their name and claimed credentials.
- **Temporal claims:** does the metadata date match the posting date? Does the surrounding environment match the claimed date? A mismatch exposes a false-context post.

Key reminder

- InVID/WeVerify is the professional standard for video verification used by journalists worldwide. Install it as a permanent tool in your browser.
- Metadata can be stripped or manipulated by platforms, so the absence of metadata is not proof of authenticity.
- Frame-by-frame searching is slow but thorough. Focus on frames showing the location, the speaker's identity, or any visible text.

3.1.4 Detecting Image Manipulation — FotoForensics

What is Error Level Analysis?

Every digital image file has a compression signature, such as a pattern created when the image was saved. When part of an image is edited after its original creation, for example, when text is added, a number is changed, or a face is replaced, that edited area is re-saved with a different compression level. The rest of the image retains its original compression.

Error Level Analysis (ELA) makes this difference visible by rendering the compression levels as colours. Areas of the image that have been recently edited appear significantly brighter or differently coloured than the unedited surroundings.

1	Open fotoforensics.com in your browser. No account is needed.
2	Click “Upload” and select your image file, OR paste the image’s direct URL into the address bar at the top of the page. Then click “Upload”.
3	The tool displays the original image and its ELA version side by side. Look for areas that are significantly brighter or a different colour from the surrounding areas.
4	Uniform brightness across the whole image suggests it has not been recently edited. Bright patches over specific areas – especially over text, numbers, faces, or statistics – indicate post-creation editing.
5	Not all brightness differences indicate malicious editing – JPEG compression artefacts and screenshots can produce some variation. Focus on bright patches specifically over meaningful content: numbers, names, logos, or medical data.

Most Common Uses in Health Disinformation Detection

Doctored statistics in medical study screenshots. A real study result is screenshotted, and a key number is changed. ELA shows the altered number as a bright patch against the uniform background.

Altered product labels or nutrition information. Ingredients lists, dosage instructions, or warning labels are altered to remove safety information. The altered text appears brighter than the original label.

Altered product labels or nutrition information. Ingredients lists, dosage instructions, or warning labels are altered to remove safety information. The altered text appears brighter than the original label.

Fabricated quotes attributed to health authorities. A genuine photo of a health official with fabricated text added as a caption or speech bubble. The added text has a different compression profile from the original photograph.



FotoForensics works best on images that have not been heavily re-compressed by social media platforms. When an image is uploaded to Instagram or TikTok and then downloaded again, the platform's own compression can reduce ELA effectiveness.

For best results: use the original image file if possible, or the highest-quality version you can find.

ELA is one tool among several. A clean ELA result does not guarantee an image is unmanipulated. Combine with reverse image search and context analysis.

Key reminder

- ELA is not proof of manipulation on its own; it is evidence that requires interpretation.
- Focus ELA analysis on meaningful content: numbers, text, faces, and logos, not background areas.
- Always combine FotoForensics with other tools for a complete picture.

Website and Domain Analysis

The section below covers three methods for assessing the credibility of health websites. They can be applied in sequence to any site you are unsure about, or used individually, depending on the specific concern.

3.2.1 Reading an URL Like a Detective — URL Structure, WHOIS, and Lateral Reading

Understanding URL Structure

Every web address follows a strict structure. Reading it correctly takes ten seconds and reveals whether a website is what it claims to be.

`https://subdomain.realdomain.extension/path/page`

subdomain — Set by the owner. Can be any word, including 'nhs', 'who', 'government'.

realdomain — The actual name of the website. This is what you must identify.

.extension (TLD) — Restricted extensions (.gov, .int, .edu) signal authority. Open extensions (.info, .biz, .net) are available to anyone.

Golden Rule: read right to left from the first forward slash (/)

The real domain is what appears immediately before the first /. Everything to the left is a subdomain which anyone can set to any word, including the name of a trusted organisation.

- **who-healthalerts.net/vaccine-warning** → Real domain: **who-healthalerts.net** (NOT the WHO)
- **health.nih-alerts.com/update** → Real domain: **nih-alerts.com** (NOT the NIH)
- **who.int/news** → Real domain: **who.int** ✓ Legitimate WHO

Level	Extensions	Who can register
Restricted	.gov .int .edu	Governments, international organisations, and (broadly) educational institutions only
Verify	.org .nhs.uk .ac.uk	Originally restricted; now more open. Check with WHOIS
Anyone	.com .net .info .biz	Open registration. A health site on .info claiming authority is a red flag

How to Use WHOIS

WHOIS is a public database recording when every domain was registered, who registered it, and where it is hosted. A site claiming 20 years of health authority that was registered three months ago is not what it claims.

1	Open who.is in your browser.
2	Type the domain name (e.g. naturalhealthtruth.com) into the search bar and press Enter.
3	Look for the field labelled “Created” or “Registration Date”.
4	Look at “Registrant” or “Organization”. Hidden registrant details behind a privacy service are themselves relevant information.
5	Compare the registration date and registrant with the authority the site claims. Newly registered + hidden registrant + open TLD = strong red flags.

Lateral Reading: Leave the Website Immediately

Research by the Stanford History Education Group found that professional fact-checkers assessed website credibility faster and more accurately than experts by doing one thing: they left the website immediately and searched what others say about it. This is lateral reading.

The lateral reading process — 30 seconds

- 1 Arrive on a health website you do not recognise.
- 2 Do not read the articles. Open a new browser tab immediately.
- 3 Search: [website name] + “fact check” OR [website name] + “misinformation” OR [website name] + “review”.

4

Look for coverage from fact-checking organisations (Snopes, Full Fact, AFP Fact Check, Correctiv), medical authorities (WHO, national health ministries), or academic assessments.

5

Interpret: credible external coverage = proceed with caution; no external coverage = not an established authority; fact-checker identifies as disinformation source = do not use.

Key reminder

- Real domain = what appears immediately before the first /. Trusted brand names in the subdomain position are a classic manipulation tactic.
- WHOIS registration date is one of the fastest credibility indicators available.
- Lateral reading is faster and more accurate than deep reading for assessing source credibility.
- HTTPS = encrypted connection. It does not equal trustworthy content.



The HTTPS padlock does NOT mean a website is trustworthy. It only means the connection is encrypted. Disinformation websites routinely use HTTPS. The padlock says nothing about content accuracy.

Social Media and Account Analysis

The section below covers two tools for assessing whether a social media account promoting health content is authentic. The 6-point checklist provides a manual assessment method; Social Blade provides data-driven evidence of inauthentic follower growth.

3.3.1 Is This Account Real? — Profile Checklist and Social Blade

Why Inauthentic Accounts Matter for Health

Inauthentic accounts — bots, purchased followers, and coordinated fake profiles — are used in health disinformation campaigns to make false health claims appear more popular and widely accepted than they are. Survey data from all five partner countries found that up to 83% of respondents had encountered miracle-product disinformation. Many of those products are promoted through inauthentic account networks. The mechanism behind this is the same one identified across all five national surveys as a primary driver of risky health behaviour: social proof. If it looks like everyone is recommending something, the impulse to follow is powerful. Inauthentic accounts manufacture that social proof artificially.

The 6-Point Account Authenticity Checklist

Apply these six checks to any account promoting health claims on Instagram, TikTok, or YouTube. Three or more red flags together is strong evidence of inauthentic behaviour.

1

Find the account creation date (visible in 'About' or 'Account info' on most platforms). Compare it with the follower count. A genuine health educator who built an audience over three years might have 50,000 followers. An account with 300,000 followers after six months warrants scrutiny.

Red flag: *very high follower count for a very young account.*

2

Go to socialblade.com and type the account's username.

Social Blade works for YouTube, Instagram, and TikTok. Organic growth appears as a gradual slope. Bought followers appear as a vertical cliff — a sudden spike of thousands in a single day.

Red flag: *any single-day spike of thousands of followers followed by a return to slow growth.*

3

Scroll through the account's posts. Authentic accounts mix content types, formats, and topics. Coordinated inauthentic accounts post with mechanical regularity: the same template, same style, same schedule, every day without variation.

Red flag: *every post is the same format, layout, and caption structure.*

4

Read 10–15 comments on recent posts. Authentic engagement includes specific, varied responses. Inauthentic comment networks produce short, generic, repeated phrases — "This changed my life!", "So true!", "Amazing product!" — with no personal specifics.

Red flag: *most comments are 3–5 generic words, or the same phrase appearing verbatim across multiple posts.*

5

A large authentic account typically follows relatively few others.

An account following 12,000 people while having 14,000 followers may be part of a follow-for-follow scheme common in supplement promotion networks. Check whether the followed accounts are also health product promoters.

Red flag: following count close to or exceeding follower count at high numbers.

6

Search for the same health claim or the account's handle on other platforms.

Coordinated networks often post identical text and images on Instagram, Facebook, TikTok, and YouTube simultaneously — word for word.

Red flag: exact same text and images appearing simultaneously across multiple platforms.

How to use Social Blade

1

Open socialblade.com in your browser.

2

Type the Instagram, TikTok, or YouTube username into the search bar. Select the correct platform from the dropdown.

3

The main graph shows daily follower count changes over time. A smooth gradual slope = organic growth. A vertical spike on a single day = purchased followers.

4

Below the graph, a daily data table shows exact follower numbers per day. Look for any single day with an unusually large gain. If 200,000 followers arrived on one day and growth was otherwise slow, those followers were bought.

Key reminder

- No single red flag proves inauthenticity. Three or more together is strong evidence.
- Inauthentic accounts amplify disinformation from a real source — they do not always create it. Also check the original claim using your broader fact-checking toolkit.
- A large follower count means nothing if the followers were bought. Social Blade makes bought followers visible.