How to Turn a News Claim Into an OSINT Workflow With AI

Maria Cattini | 12/06/2026 | OSINT

A short sentence can hide a long investigation.

“A video shows a protest in the city center.”

“A company is linked to a sanctioned entity.”

“This image was taken after the attack.”

“The account belongs to the person mentioned in the story.”

Each claim looks simple when it appears in a post, a headline, a message, or a screenshot. But for an OSINT researcher, the real question is not whether the claim sounds plausible. The question is whether it can be broken into verifiable parts.

This is where AI can help.

Not by deciding what is true. Not by replacing source checking. Not by giving you a final answer. AI is useful when it helps you slow down, separate the claim into smaller questions, generate search paths, identify missing evidence, and keep a record of what still needs verification.

The value is not the answer. The value is the workflow.

Start With the Claim, Not the Tool

Many weak OSINT workflows begin with a tool.

Someone sees a claim and immediately opens a search engine, a reverse image search tool, a social platform, a map, or an AI chatbot. That can work for simple checks, but it often creates noise. You start collecting fragments before you have defined what you are trying to prove.

A better first step is to write the claim as a single sentence.

Example:

> This video shows police using tear gas during a protest in Madrid on 10 June 2026.

This sentence contains several elements:

- a content type: video;
- an alleged actor: police;

• an alleged action: using tear gas;
• an alleged event: protest;
• an alleged place: Madrid;
• an alleged date: 10 June 2026.

The claim is not one thing. It is a bundle of smaller claims.

AI can help you split it.

Prompt:

Break this claim into separate verifiable components. For each component, explain what type of evidence would support it and what type of evidence would not be enough. Claim: [paste the claim]

The output should not be treated as evidence. It is a planning layer. You are asking AI to help you structure the investigation before you begin.

Step 1: Separate Entities, Actions, Time, and Place

Every news claim should be mapped into basic investigative units.

Use four categories:

Category	What to Extract	Examples
Entities	People, organizations, accounts, companies, domains	a minister, a Telegram channel, a company name
Actions	What allegedly happened	posted, deleted, attacked, donated, registered
Time	Dates, times, sequence, publication windows	“yesterday”, “after the strike”, “in May”
Place	Locations, facilities, streets, countries, platforms	a city, a building, a website, a social network

AI is useful here because it can identify hidden assumptions.

For example, the sentence “the video shows the protest” contains at least three separate questions:

• Does the video show a protest?
• Is the protest the one mentioned in the claim?
• Was the video recorded at the claimed time and location?

Those are different checks.

If you merge them too early, you risk confirming one part and accidentally accepting the rest.

Step 2: Turn Each Component Into a Verification Question

Once the claim is separated, each part should become a question.

Weak question:

• Is this true?

Useful questions:

• Who first published this video?
• Does the visual environment match the claimed location?
• Is there independent evidence that a protest happened there on that date?
• Do weather, light, clothing, signage, language, license plates, or architecture support the claim?
• Are there earlier versions of the same media?
• Has the same image or video been used in another context?

AI can help generate questions, but you should edit them.

Prompt:

Turn each component of this claim into OSINT verification questions. Group the questions by source type: search engines, social platforms, maps, archives, official sources, media databases, image/video verification. Claim: [paste the claim]

The result becomes your investigation plan.

It is still not proof.

Step 3: Build a Source Map

Before searching, decide which sources could actually answer each question.

This matters because not every source has the same evidentiary value.

For a location claim, useful sources may include:

• satellite or street-level imagery;
• maps;
• local business listings;
• transport data;
• municipal pages;
• local media;
• social media posts from the same area;
• archived pages;
• weather records;
• livestreams or public webcams, when legal and available.

For an identity claim, useful sources may include:

• official profiles;
• domain registration data, when available;
• company registries;
• archived website versions;
• public biographies;
• cross-platform usernames;
• public documents;
• verified organizational pages.

For a time claim, useful sources may include:

• original upload timestamp;
• archive timestamp;
• metadata, if available and reliable;
• publication sequence across platforms;
• shadows, weather, light, or seasonal clues;

• independent reports from the same window.

AI can help you create the source map, but it cannot guarantee that the sources exist or are reliable.

Prompt:

Create a source map for this claim. For each verification question, list possible public source types, what they could confirm, what they cannot confirm, and the main risk of false positives. Claim: [paste the claim]

This gives you a structured path instead of a random search session.

Step 4: Generate Search Queries, Then Test Them Manually

AI is useful for query expansion.

It can suggest alternative names, spelling variations, translations, related terms, date formats, and platform-specific wording. This is especially helpful when a claim crosses languages or regions.

For example, if a claim mentions a company, AI can help generate:

• official company name variations;
• possible abbreviations;
• local-language forms;
• executive names;
• related subsidiaries;
• domain patterns;
• social profile terms;
• keywords for complaints, sanctions, lawsuits, or procurement records.

But queries are not evidence. They are probes.

Prompt:

Generate search queries for this claim in English and in the likely local language. Include exact-match queries, broad discovery queries, date-based queries, and source-specific queries. Do not assume the claim is true. Claim: [paste the claim]

Then you test them manually.

Keep the useful queries. Delete the noisy ones. Record what each query found and what it failed to find.

Step 5: Create an Evidence Log

An OSINT workflow becomes stronger when it has a written evidence log.

The log does not need to be complex. It should help you avoid mixing confirmed facts, weak indicators, and guesses.

Use a simple table:

Claim Component	Source Checked	What It Shows	Confidence	Open Question
Location	Street-level imagery	Signage and building shape appear consistent	Medium	Need independent local confirmation
Date	Original post	Uploaded on the	Low/Medium	Upload date is not

Claim Component	Source Checked timestamp	What It Shows claimed date	Confidence	Open Question recording date
Event	Local media report	Protest reported in same area	Medium	Does it match this specific video?

AI can help format the log, but the entries must come from your checks.

Prompt:

Create an evidence log template for this investigation. Use columns for claim component, source, finding, confidence level, limitation, and next step. Claim: [paste the claim]

The important part is discipline.

Do not write “verified” when you only have “consistent with.” Do not write “false” when you only have “not found.” Do not write “confirmed identity” when you only have a similar username.

Step 6: Use AI to Find Gaps, Not to Close Them

One of the best uses of AI in OSINT is gap detection.

After you collect early findings, ask AI to challenge the workflow.

Prompt:

Review this investigation plan and evidence log. Identify weak assumptions, missing source types, possible false positives, and claims that are not yet supported. Do not decide whether the original claim is true. Evidence log: [paste your notes]

This kind of prompt is useful because it changes the role of AI.

AI is not the judge. It is the second reader that helps you notice what is missing.

The final judgment still belongs to the researcher.

Step 7: Write the Finding Carefully

At the end of the workflow, avoid forcing a binary answer if the evidence does not support it.

Use precise language:

- Verified: the evidence directly supports the claim.
- Likely: multiple independent indicators support the claim, but one part remains uncertain.
- Consistent with: the evidence does not contradict the claim, but does not prove it.
- Unverified: not enough evidence.
- False or misleading: evidence contradicts the claim or shows important context is missing.

This language is not cosmetic. It protects the integrity of the investigation.

A responsible OSINT finding should say:

- what was checked;
- what was found;
- what was not found;
- what remains uncertain;
- what level of confidence is justified.

A Practical AI/OSINT Workflow

Here is the full sequence:

1. Write the claim as one sentence.
2. Break it into entities, actions, time, and place.
3. Turn each part into verification questions.
4. Build a source map.
5. Generate search queries.
6. Test the queries manually.
7. Keep an evidence log.
8. Ask AI to identify gaps and weak assumptions.
9. Separate findings by confidence level.
10. Write the conclusion with limits.

This workflow works for many types of claims:

- a viral image;
- a breaking news post;
- a company connection;
- a social media account attribution;
- a location claim;
- a timeline claim;
- a document circulating online;
- a screenshot presented as proof.

The tools may change. The method should not.

What AI Should Not Do

AI should not be used to:

- identify private individuals for harassment;
- bypass privacy settings;
- scrape personal data at scale;
- make legal or security conclusions without evidence;
- invent sources;
- treat screenshots as proof;
- produce certainty from weak signals;
- replace primary source checks.

In OSINT, speed is useful only when it does not destroy the chain of evidence.

AI can make you faster at planning, organizing, translating, comparing, and reviewing. It should not make you careless.

The Real Skill Is Decomposition

A claim becomes manageable when you stop treating it as a single object.

Break it apart.

Find the entities. Define the action. Locate the time. Test the place. Map the sources. Record the evidence. Name the uncertainty.

AI can support every step of that process.

But the discipline remains human: knowing what counts as evidence, what is only an indicator, and when the honest answer is still “not verified.”

That is the difference between using AI to generate conclusions and using AI to build better OSINT workflows.
A short sentence can hide a long investigation.

“A video shows a protest in the city center.”

“A company is linked to a sanctioned entity.”

“This image was taken after the attack.”

“The account belongs to the person mentioned in the story.”

Each claim looks simple when it appears in a post, a headline, a message, or a screenshot. But for an OSINT researcher, the real question is not whether the claim sounds plausible. The question is whether it can be broken into verifiable parts.

This is where AI can help.

Not by deciding what is true. Not by replacing source checking. Not by giving you a final answer. AI is useful when it helps you slow down, separate the claim into smaller questions, generate search paths, identify missing evidence, and keep a record of what still needs verification.

The value is not the answer. The value is the workflow.

Start With the Claim, Not the Tool

Many weak OSINT workflows begin with a tool.

Someone sees a claim and immediately opens a search engine, a reverse image search tool, a social platform, a map, or an AI chatbot. That can work for simple checks, but it often creates noise. You start collecting fragments before you have defined what you are trying to prove.

A better first step is to write the claim as a single sentence.

Example:

> This video shows police using tear gas during a protest in Madrid on 10 June 2026.

This sentence contains several elements:

- a content type: video;
- an alleged actor: police;
- an alleged action: using tear gas;
- an alleged event: protest;
- an alleged place: Madrid;
- an alleged date: 10 June 2026.

The claim is not one thing. It is a bundle of smaller claims.

AI can help you split it.

Prompt:

Break this claim into separate verifiable components. For each component, explain what type of evidence would support it and what type of evidence would not be enough. Claim: [paste the claim]

The output should not be treated as evidence. It is a planning layer. You are asking AI to help you structure the investigation before you begin.

Step 1: Separate Entities, Actions, Time, and Place

Every news claim should be mapped into basic investigative units.

Use four categories:

Category	What to Extract	Examples
Entities	People, organizations, accounts, companies, domains	a minister, a Telegram channel, a company name
Actions	What allegedly happened	posted, deleted, attacked, donated, registered
Time	Dates, times, sequence, publication windows	“yesterday”, “after the strike”, “in May”
Place	Locations, facilities, streets, countries, platforms	a city, a building, a website, a social network

AI is useful here because it can identify hidden assumptions.

For example, the sentence “the video shows the protest” contains at least three separate questions:

- Does the video show a protest?
- Is the protest the one mentioned in the claim?
- Was the video recorded at the claimed time and location?

Those are different checks.

If you merge them too early, you risk confirming one part and accidentally accepting the rest.

Step 2: Turn Each Component Into a Verification Question

Once the claim is separated, each part should become a question.

Weak question:

- Is this true?

Useful questions:

- Who first published this video?
- Does the visual environment match the claimed location?
- Is there independent evidence that a protest happened there on that date?
- Do weather, light, clothing, signage, language, license plates, or architecture support the claim?
- Are there earlier versions of the same media?
- Has the same image or video been used in another context?

AI can help generate questions, but you should edit them.

Prompt:

Turn each component of this claim into OSINT verification questions. Group the questions by source type: search engines, social platforms, maps, archives, official sources, media databases, image/video verification. Claim: [paste the claim]

The result becomes your investigation plan.

It is still not proof.

Step 3: Build a Source Map

Before searching, decide which sources could actually answer each question.

This matters because not every source has the same evidentiary value.

For a location claim, useful sources may include:

• satellite or street-level imagery;
• maps;
• local business listings;
• transport data;
• municipal pages;
• local media;
• social media posts from the same area;
• archived pages;
• weather records;
• livestreams or public webcams, when legal and available.

For an identity claim, useful sources may include:

• official profiles;
• domain registration data, when available;
• company registries;
• archived website versions;
• public biographies;
• cross-platform usernames;
• public documents;
• verified organizational pages.

For a time claim, useful sources may include:

• original upload timestamp;
• archive timestamp;
• metadata, if available and reliable;
• publication sequence across platforms;
• shadows, weather, light, or seasonal clues;
• independent reports from the same window.

AI can help you create the source map, but it cannot guarantee that the sources exist or are reliable.

Prompt:

Create a source map for this claim. For each verification question, list possible public source types, what they could confirm, what they cannot confirm, and the main risk of false positives. Claim: [paste the claim]

This gives you a structured path instead of a random search session.

Step 4: Generate Search Queries, Then Test Them Manually

AI is useful for query expansion.

It can suggest alternative names, spelling variations, translations, related terms, date formats, and platform-specific wording. This is especially helpful when a claim crosses languages or regions.

For example, if a claim mentions a company, AI can help generate:

• official company name variations;
• possible abbreviations;
• local-language forms;
• executive names;
• related subsidiaries;
• domain patterns;
• social profile terms;
• keywords for complaints, sanctions, lawsuits, or procurement records.

But queries are not evidence. They are probes.

Prompt:

Generate search queries for this claim in English and in the likely local language. Include exact-match queries, broad discovery queries, date-based queries, and source-specific queries. Do not assume the claim is true. Claim: [paste the claim]

Then you test them manually.

Keep the useful queries. Delete the noisy ones. Record what each query found and what it failed to find.

Step 5: Create an Evidence Log

An OSINT workflow becomes stronger when it has a written evidence log.

The log does not need to be complex. It should help you avoid mixing confirmed facts, weak indicators, and guesses.

Use a simple table:

Claim Component	Source Checked	What It Shows	Confidence	Open Question
Location	Street-level imagery	Signage and building shape appear consistent	Medium	Need independent local confirmation
Date	Original post timestamp	Uploaded on the claimed date	Low/Medium	Upload date is not recording date
Event	Local media report	Protest reported in same area	Medium	Does it match this specific video?

AI can help format the log, but the entries must come from your checks.

Prompt:

Create an evidence log template for this investigation. Use columns for claim component, source, finding, confidence level, limitation, and next step. Claim: [paste the claim]

The important part is discipline.

Do not write “verified” when you only have “consistent with.” Do not write “false” when you only have “not found.” Do not write “confirmed identity” when you only have a similar username.

Step 6: Use AI to Find Gaps, Not to Close Them

One of the best uses of AI in OSINT is gap detection.

After you collect early findings, ask AI to challenge the workflow.

Prompt:

Review this investigation plan and evidence log. Identify weak assumptions, missing source types, possible false positives, and claims that are not yet supported. Do not decide whether the original claim is true. Evidence log: [paste your notes]

This kind of prompt is useful because it changes the role of AI.

AI is not the judge. It is the second reader that helps you notice what is missing.

The final judgment still belongs to the researcher.

Step 7: Write the Finding Carefully

At the end of the workflow, avoid forcing a binary answer if the evidence does not support it.

Use precise language:

- Verified: the evidence directly supports the claim.
- Likely: multiple independent indicators support the claim, but one part remains uncertain.
- Consistent with: the evidence does not contradict the claim, but does not prove it.
- Unverified: not enough evidence.
- False or misleading: evidence contradicts the claim or shows important context is missing.

This language is not cosmetic. It protects the integrity of the investigation.

A responsible OSINT finding should say:

- what was checked;
- what was found;
- what was not found;
- what remains uncertain;
- what level of confidence is justified.

A Practical AI/OSINT Workflow

Here is the full sequence:

1. Write the claim as one sentence.
2. Break it into entities, actions, time, and place.
3. Turn each part into verification questions.
4. Build a source map.
5. Generate search queries.
6. Test the queries manually.
7. Keep an evidence log.
8. Ask AI to identify gaps and weak assumptions.
9. Separate findings by confidence level.
10. Write the conclusion with limits.

This workflow works for many types of claims:

• a viral image;
• a breaking news post;
• a company connection;
• a social media account attribution;
• a location claim;
• a timeline claim;
• a document circulating online;
• a screenshot presented as proof.

The tools may change. The method should not.

What AI Should Not Do

AI should not be used to:

• identify private individuals for harassment;
• bypass privacy settings;
• scrape personal data at scale;
• make legal or security conclusions without evidence;
• invent sources;
• treat screenshots as proof;
• produce certainty from weak signals;
• replace primary source checks.

In OSINT, speed is useful only when it does not destroy the chain of evidence.

AI can make you faster at planning, organizing, translating, comparing, and reviewing. It should not make you careless.

The Real Skill Is Decomposition

A claim becomes manageable when you stop treating it as a single object.

Break it apart.

Find the entities. Define the action. Locate the time. Test the place. Map the sources. Record the evidence. Name the uncertainty.

AI can support every step of that process.

But the discipline remains human: knowing what counts as evidence, what is only an indicator, and when the honest answer is still “not verified.”

That is the difference between using AI to generate conclusions and using AI to build better OSINT workflows.