INTELLIGENCE · OPEN-SOURCE INTELLIGENCE

Verifying Open-Source Information

Open sources are only as useful as they are reliable. This article explains why verification matters, the kinds of false and misleading information analysts meet, practical checks for provenance, source, date, location and motive, and how to weigh and report what has been verified.

ModQu research article · 2,135 words · 21 references (APA 7th edition)

Open sources offer enormous quantities of information, but quantity is not reliability. A photograph may be years old and from another country. A quotation may be invented. A document may be forged. A review may have been written by the business itself. A trend on social media may be driven by automated accounts. Before any piece of open-source information supports a decision, it needs to be checked. Verification is the discipline of establishing whether information is authentic, accurate and what it appears to be. This article explains why it matters, the kinds of problems analysts meet and practical methods for dealing with them.

Why verification matters

Hulnick (2002) warned that one of the main weaknesses of open-source intelligence is the unreliability of much of its material, combined with the effort required to separate the reliable from the unreliable. The problem has grown with social media. Vosoughi et al. (2018), studying the spread of contested stories on a major platform over more than a decade, found that false stories reached more people and spread faster than true ones, largely because people found them novel and shared them.

The consequences of acting on false information can be severe: a company might respond to a fake report of a competitor's recall, react to a fabricated statement by a regulator or invest on the basis of an invented trend. And once false information has been absorbed, correcting it is hard. Lewandowsky et al. (2012) reviewed research showing that misinformation often continues to influence people's reasoning even after they have accepted a correction, an effect known as the **continued influence effect**. It is far better to verify before accepting than to correct afterwards.

Kovach and Rosenstiel (2014) describe journalism as, at its core, a discipline of verification. Intelligence analysis shares that core. The methods described here draw heavily on the experience of journalists and open-source investigators, who have developed practical techniques for checking material quickly.

Kinds of false and misleading information

Wardle and Derakhshan (2017) distinguish three kinds of problematic information:

- **Misinformation**: false information shared without intent to harm, such as an honest mistake or a misunderstood photograph.

- **Disinformation**: false information deliberately created or shared to cause harm or mislead.
- **Malinformation**: genuine information shared to cause harm, for example private details leaked to damage someone, or a real photograph presented with a misleading caption.

The distinction matters because each calls for different questions. Misinformation is usually detected by checking facts. Disinformation also requires asking who benefits and whether there is a coordinated effort. Malinformation requires attention to context, since the content itself may be authentic.

In business settings, common forms include fake reviews, false rumours about products or companies, impersonation of executives or official bodies, manipulated images and documents, and coordinated campaigns by automated or paid accounts. Luca and Zervas (2016) found that a substantial share of reviews on a popular review platform were filtered as suspicious, with fake reviews more common among businesses facing weak reputations or intense competition. Ferrara et al. (2016) describe how social bots can imitate human users and inflate the apparent popularity of claims. Starbird et al. (2019) show that disinformation campaigns often recruit unwitting participants, ordinary users who share and amplify false material believing it to be true, so that apparent grassroots agreement may have been engineered.

Five checks

Silverman (2014), in a handbook for journalists verifying material during emergencies, sets out checks that apply to almost any piece of open-source content. Together with the attention to motive emphasised by Wardle and Derakhshan (2017), they can be summarised as five questions.

Five checks for any piece of open-source information

Questions to ask, and typical ways to answer them

Check	Question	Typical methods
Provenance	Is this the original, or a copy or edit?	Reverse image and video search; earliest version; archives
Source	Who created it, and could they know?	Account history; other work; contact the creator
Date	When was it created or first published?	Metadata; first appearance; weather, shadows, events
Location	Where was it created?	Landmarks, signs, terrain matched to maps and imagery
Motive	Why was it created and shared?	Who benefits; pattern of the account; coordinated sharing

After Silverman (2014), Wardle and Derakhshan (2017) and Higgins (2021)

ModQu.com

Figure 1. Five checks for any piece of open-source information. After Silverman (2014), Wardle and Derakhshan (2017) and Higgins (2021); table by ModQu.

Provenance

Is this the original item, or a copy, an excerpt or an edited version? Images and videos are often recycled from earlier events. Reverse image search can reveal whether a picture appeared online before the event it supposedly shows. Archived copies of web pages show how a page looked at an earlier date. Finding the earliest version of an item is usually the first step.

Source

Who created the item, and were they in a position to know? For an account on social media, analysts look at its history, its other posts, when it was created and whom it interacts with. A new account with few posts that suddenly shares dramatic material deserves caution. Contacting the creator directly, where appropriate, can confirm details and permissions.

Date

When was the item created or first published? Metadata embedded in files can indicate creation dates, although it is easily altered or removed. Contextual clues, such as weather, seasonal vegetation, shadows and visible events, can confirm or contradict a claimed date.

Location

Where was the item created? Higgins (2021) describes how open-source investigators geolocate images and video by matching landmarks, road layouts, signs, terrain and buildings against maps and satellite imagery. The same techniques can confirm, for example, whether a photograph of a damaged factory really shows the facility claimed.

Motive

Why was the item created and shared? Asking who benefits from a claim, whether it serves a commercial or political interest, and whether it is being spread in a coordinated way helps identify disinformation even when individual details are hard to check.

Reading laterally

A simple but powerful technique comes from research on how professional fact-checkers work. Wineburg and McGrew (2019) compared fact-checkers, historians and university students as they evaluated unfamiliar websites. The students and historians read **vertically**, studying the site itself closely, its design, its "about" page and its content. The fact-checkers read **laterally**: they left the site almost immediately and searched for what other sources said about the organisation behind it. They reached accurate judgements faster and more often. For analysts, the lesson is that the best evidence about a source usually comes from outside it.

Metzger et al. (2010) found that people judging the credibility of online information rely on mental shortcuts, such as the reputation of a name, endorsements by others, consistency with other sources and whether the material looks professional. These shortcuts are efficient but can be exploited: a professional-looking site or a large number of shares does not make a claim true.

Manipulated media

Images, audio and video can be altered. Farid (2016) describes forensic techniques for detecting manipulation in photographs, including inconsistencies in lighting and shadows, traces left by editing and compression, and anomalies in metadata. Such analysis often needs specialist tools and expertise.

Synthetic media, produced by machine learning and sometimes called "deepfakes", raise new difficulties. Chesney and Citron (2019) argue that convincing fake video and audio threaten privacy, democracy and security, and also create a "liar's dividend": as people learn that fakes exist, those caught on genuine recordings can claim they are fake. Vaccari and Chadwick (2020), in an experiment with a synthetic political video, found that it was less likely to deceive people outright than to leave them uncertain, and that this uncertainty reduced trust in news more generally. For analysts, the implication is that the provenance and corroboration of media matter more than ever, since appearance alone is no longer reliable evidence.

Why people, including analysts, are fooled

Pennycook and Rand (2019) found that susceptibility to false headlines was better explained by a lack of careful reasoning than by a desire to believe claims that fit one's views: people who scored higher on a test of reflective thinking were better at telling true from false headlines, whatever their politics. In a later review, Pennycook and Rand (2021) concluded that inattention plays a large role, and that simply prompting people to think about accuracy improves the quality of what they share.

Analysts are not immune. Time pressure, the excitement of a dramatic finding and a story that confirms existing expectations all weaken scrutiny. Heuer (1999) showed how analysts tend to accept evidence that fits their expectations and discount evidence that does not. Building verification into routine procedures, rather than relying on individual vigilance, is the main defence.

Weighing and corroborating

Verification is rarely a matter of proving something true or false. More often, it produces a judgement about how much weight an item deserves. Intelligence services have long rated sources and information separately: how reliable has this source been in the past, and how credible is this particular claim, given what else is known? The two do not always agree. A usually reliable source can be wrong, and an unknown source can be right.

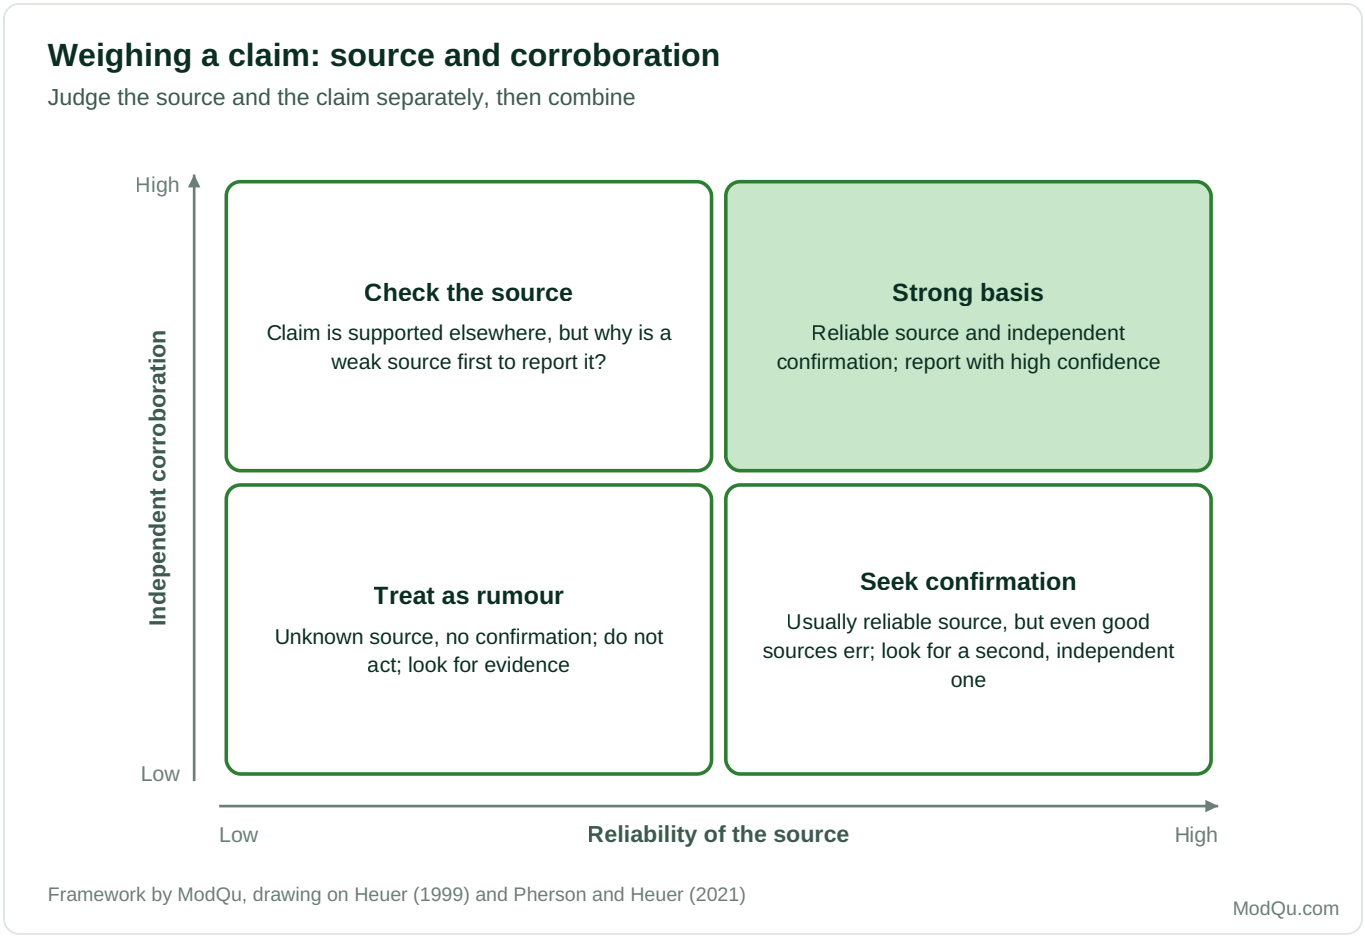


Figure 2. Weighing a claim. Judging the source and the claim separately, then combining them, guards against both credulity and dismissal. Framework by ModQu, drawing on Heuer (1999) and Pherson and Heuer (2021).

Corroboration means finding independent confirmation. The word "independent" matters. Ten news articles repeating the same press release are one source, not ten. Several social media accounts sharing the same image may all trace back to a single post. Analysts should trace each claim to its origin and count only genuinely separate sources.

Pherson and Heuer (2021) describe structured techniques that support this weighing, such as checking key assumptions, assessing the credibility of each source and listing alternative explanations for the evidence. Graves (2016), studying professional fact-checkers, shows how they document their sources and reasoning so that readers can follow and challenge each judgement. Analysts should do the same.

Verifying documents and data

Not all open-source material is images and posts. Analysts also meet documents, such as letters, contracts, reports and spreadsheets, and data sets published by organisations of varying reliability. Similar principles apply, with some specific checks:

- **Documents.** Does the format, letterhead, signature and language match genuine documents from the same organisation? Are dates, names and reference numbers consistent with known facts? Does the file's metadata show an author, software or creation date that contradicts the claimed origin? Can the organisation confirm it?
- **Statistics.** Who collected the data, how and for what purpose? Are definitions and methods published? Do the figures agree with independent sources, such as official statistics or industry data? Are percentages based on very small numbers?
- **Quotations.** Can the quotation be traced to a recording, transcript or original article? Quotations are frequently shortened, paraphrased or attached to the wrong person as they are shared.
- **Screenshots.** Screenshots of posts, messages and web pages are easy to fabricate. Look for the original post or page, or an archived copy, rather than relying on an image of it.

The common thread is to move from the copy to the original, and from the claim to independent evidence.

Building verification into practice

Brandtzaeg et al. (2016) found that journalists verifying material from social media relied heavily on traditional methods, such as contacting sources and cross-checking with trusted outlets, supplemented by specialised tools. For organisations, a sound approach includes:

1. **Verify before sharing.** Treat all unverified material as provisional, and label it as such in internal communications.
2. **Find the original.** Trace every important claim, image or document to its first appearance.
3. **Check the source from outside.** Read laterally to learn who is behind a site or account.
4. **Look for independent confirmation.** Seek at least one source that does not derive from the first.
5. **Ask who benefits.** Consider the motives behind the creation and spread of the material.
6. **Record the steps.** Note what was checked, how and with what result.
7. **State confidence.** Report verified, partly verified and unverified claims differently.

An illustration

On a Monday morning, a post circulating on social media claims that a major competitor of a Canadian food processor has recalled a popular product because of contamination, with a photograph of a recall notice. Several accounts share it, and the processor's sales team asks whether to contact the competitor's customers.

The intelligence analyst first looks for the recall on the Canadian Food Inspection Agency's official list and the competitor's website, and finds nothing. A reverse image search shows the photograph of the notice was published two years earlier, about a different product from a different company, with the name edited. The accounts sharing the post

were created recently and share content almost exclusively about the competitor. The analyst reports that the claim is false, with high confidence, and notes signs of a coordinated campaign. The processor does not act on the rumour, avoids embarrassment and passes the information to the competitor, who reports the campaign to the platform.

Conclusion

Verification turns open-source material into evidence. It rests on a small set of questions about provenance, source, date, location and motive, on the habit of checking sources from outside, and on careful corroboration by genuinely independent sources. Manipulated media and coordinated campaigns make the work harder, and human tendencies make everyone vulnerable to accepting what fits expectations. Organisations that build verification into routine practice and report their confidence honestly gain the benefits of open sources without being misled by them.

References

Brandtzaeg, P. B., Lüders, M., Spangenberg, J., Rath-Wiggins, L., & Følstad, A. (2016). Emerging journalistic verification practices concerning social media. *Journalism Practice, 10*(3), 323–342. https://doi.org/10.1080/17512786.2015.1020331

Chesney, R., & Citron, D. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review, 107*(6), 1753–1820.

Farid, H. (2016). *Photo forensics*. MIT Press.

Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM, 59*(7), 96–104. https://doi.org/10.1145/2818717

Graves, L. (2016). *Deciding what's true: The rise of political fact-checking in American journalism*. Columbia University Press.

Heuer, R. J., Jr. (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence, Central Intelligence Agency.

Higgins, E. (2021). *We are Bellingcat: An intelligence agency for the people*. Bloomsbury.

Hulnick, A. S. (2002). The downside of open source intelligence. *International Journal of Intelligence and CounterIntelligence, 15*(4), 565–579. https://doi.org/10.1080/08850600290101767

Kovach, B., & Rosenstiel, T. (2014). *The elements of journalism: What newspeople should know and the public should expect* (3rd ed.). Three Rivers Press.

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest, 13*(3), 106–131. https://doi.org/10.1177/1529100612451018

Luca, M., & Zervas, G. (2016). Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science, 62*(12), 3412–3427. https://doi.org/10.1287/mnsc.2015.2304

Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and heuristic approaches to credibility evaluation online. *Journal of Communication, 60*(3), 413–439. https://doi.org/10.1111/j.1460-2466.2010.01488.x

Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition, 188*, 39–50. https://doi.org/10.1016/j.cognition.2018.06.011

Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences, 25*(5), 388–402. https://doi.org/10.1016/j.tics.2021.02.007

Pherson, R. H., & Heuer, R. J., Jr. (2021). *Structured analytic techniques for intelligence analysis* (3rd ed.). CQ Press.

Silverman, C. (Ed.). (2014). *Verification handbook: A definitive guide to verifying digital content for emergency coverage*. European Journalism Centre.

Starbird, K., Arif, A., & Wilson, T. (2019). Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *Proceedings of the ACM on Human-Computer Interaction, 3*(CSCW), Article 127. https://doi.org/10.1145/3359229

Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society, 6*(1). https://doi.org/10.1177/2056305120903408

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Wardle, C., & Derakhshan, H. (2017). *Information disorder: Toward an interdisciplinary framework for research and policy making* (Council of Europe report DGI(2017)09). Council of Europe.

Wineburg, S., & McGrew, S. (2019). Lateral reading and the nature of expertise: Reading less and learning more when evaluating digital information. *Teachers College Record, 121*(11), 1–40.