

Why you can trust how CredBrief picks your stories

Why we built CredBrief

We wanted to keep reading the news, but we were tired of the flood of distraction. CredBrief exists to focus your time on what matters today and what will still matter in seven days, and to push back against algorithms, clickbait and sponsored news.

WHAT YOU GET

- **The important news first.** Stories are ordered by how much they matter, not by how many people click.
- **Each story once.** When many outlets report the same thing, you see it one time, not ten.
- **Real articles from named outlets.** Every story links to the original, so you can always read the source.
- **Range, not just the loudest topic.** World events, sport, culture and every region keep a place.
- **Controls that never hide what matters.** Tune topics and outlets to your taste; major stories stay in.

We only report the news

CredBrief does not write news and does not invent stories. Every story you see is a real article from a named outlet, shown with that outlet's name and a link to the original.

CredBrief chooses and orders the stories and writes a brief summary of each, in English, from the article's own text. The outlet name and the link are always the original's.

Isn't this just another RSS reader?

No. An RSS reader collects what you tell it to and shows all of it. CredBrief collects far more, then does the sorting a reader would otherwise have to do alone.

AN RSS READER

- A long list of everything, newest first, unranked.
- Every copy of the same story, from every feed.
- No sense of what matters or which outlet is reliable.
- You do all the judging.

CREDBRIEF

- Reads more than 1,000 feeds and sites, far more than any person can.
- Scores each story for significance and rates each outlet's credibility.
- Merges copies of the same story into one.
- Keeps room for the world, sport, culture and every region.
- Gives you an ordered edition, most significant first, with controls to tune it.

Why let AI sort the news?

Because nobody can read it all. Every day 14,000 to 23,000 articles pass our first filters. An edition keeps about 145 to 170 of them, twice a day.

Because your time should go to reading, not sorting. On news sites and social media, much of the time is spent deciding what to open, meeting the same story again, and following whatever their algorithm pushes next.

EXAMPLE: YOUR OWN ARITHMETIC

Think of yesterday's news time. How much went to stories you were glad you read, and how much to scrolling, choosing and repeats? The second part is the work CredBrief does for you. We do not claim a number of minutes saved; that depends on you.

Because many feeds rank for attention. Social media and many sites put first what keeps you scrolling. We rank for significance: what happened and who is affected.

Because a machine is patient and consistent. It screens every article, including in other languages, reads the promising ones in full, and applies the same test each time. When two outlets report the same story, their scores typically sit about 0.6 points apart on a 10-point scale.

You keep the judgement. We put the stories in order; you decide what to read, what to believe, and you can always click through to the source.

What AI is bad at, and how we check it

- **It can misjudge.** Some of our scores are simply wrong. We measure how often, and improve.
- **It favours what is widely reported.** A quiet but important story can score too low. Up to five stories per edition are added when strong outside signals say we missed something.
- **It is not a person.** Our checks include a small test with real people. You can help: join our short daily check at credbrief.com/help.

How a story gets in

- 1 **We read widely.** More than 1,000 feeds and sites, from wire services to specialist outlets.
For you: one place instead of fifty tabs.
- 2 **An AI model rates how significant each article is.** Our own rating of the outlet's credibility is blended in, so trust counts without burying a good story from a smaller outlet.
For you: importance and trust in one number, with trust never the whole story.
- 3 **Low-impact stories are left out.** Stories of low significance do not make the edition.
For you: less to wade through.
- 4 **Copies are merged.** In one recent edition we removed 122 near-copies; in another, 140.
For you: no repeats.
- 5 **We keep room for what a score misses.** Minimum shares for sport and culture, regional minimums, and world events at least 20% of your first 30 stories.
For you: a brief with range.

What the AI writes, and what it does not

The AI does four things with each story:

- scores how significant it is;
- decides its place in the order;
- merges copies of the same story;
- writes a short summary from the article's own text, and a tidier headline when the original is misleading or too long.

So the summary, and sometimes the headline, are machine-written from the article and translated into English where needed. They are not the outlet's words. The outlet's name and the link are always the original's, and the original headline is on the linked article.

The AI does no reporting. If a summary ever seems to say more than the article, the article is what counts. Mark it with the ? icon and we will remove it and improve based on your feedback.

What never changes a score or a rating

Your topic settings. More or less of a topic, and how close to home (less, balanced, more), only reorder your own feed, within limits.

Your mutes. They change what you see from that outlet, but never hide a major story.

Your own subscriptions, for example to The New York Times or The Guardian. They only decide which outlet's link we show first when several cover a story.

Clicks. The shared edition is chosen without reader clicks.

Outlets. No outlet can pay to be included or to be rated higher.

How we check ourselves

Anyone can say their selection is good. These are our checks, the latest results, and what each one does not show. Ranges are 95% ranges: the true value very likely lies inside.

CHECK	WHY IT MATTERS TO YOU	WHAT IT ASKS	LATEST RESULT	WHAT IT DOES NOT SHOW
Independent reader	If an independent reader would order the day's news much as we do, you are less likely to be missing something big.	Does a separate AI model, not used to pick stories, put ours in a similar order? 0 is no match, 1 is identical.	0.65 range 0.54 to 0.73 152 stories, 11 Oct	An AI is not a person. The same test of another news service gave 0.43 (0.32 to 0.54).
Does a high score mean it mattered?	A story we put first should be one other newsrooms also treat as important, not just one we liked.	In each score band, what share of stories were reported by at least four other outlets within a week?	Under 3: 0% (0 to 5) 3 to 4: 3% (1 to 10) 4 to 5: 6% (2 to 14) 5 to 6: 20% (12 to 31) 6 to 7: 16% (9 to 26) 7 and over: 23% (15 to 34) 70 stories per band	Higher bands clearly beat lower ones, to 6, 6 to 7 and over look alike: our top does not yet separate.
Same story, same score	You get a steady order, not a lottery that depends on which outlet's copy we happened to see.	When two outlets report the same story, do we score both alike?	Median gap 0.6 points 9,027 pairs, 7 Oct	6% of pairs by 2 points more.
Wikipedia's events	It shows what we may be missing: events people record as notable that did not reach you.	Of the events Wikipedia's editors list each day, how many reached an edition, and our first 30 stories?	In an edition: 20% range 15 to 27% First 30: 5% range 3 to 10% 173 events, 10 days	Wikipedia lists many local incidents we leave out by design. Politics and election of 24 events (12%), is the figure we most want to improve.
Human readers	Real people decide whether a brief is worth	Which of two stories would a	Preferred the other news service about 4 to 1	A small group over a few days. It points at

CHECK	WHY IT MATTERS TO YOU	WHAT IT ASKS	LATEST RESULT	WHAT IT DOES NOT SHOW
	reading. This keeps us honest about presentation.	person rather read?	11 people, 250 choices, 2 to 6 Oct	presentation more than a which stories pick.

Where we are weaker

- **Depth.** We pick important stories, but sometimes the thinner article about them, and our own scoring tends to under-rate local news and thin outlets today. Other news services' picks run deeper on average.
- **Fewer incident reports than a news agency.** That is deliberate, but some readers will miss them.
- **Politics outside the big countries.** Elections and party news from mid-sized countries are our weakest area in the Wikipedia check.
- **Non-English reporting.** Some of it is not yet scored, so it cannot reach an edition.
- **Small human test.** Eleven people over a few days. We say so rather than round it up.

What we do about it

- Publish this scorecard with its ranges each month, including the numbers that go the wrong way.
- Replay past days with a proposed change before it goes live, and ship only what improves these checks.
- Tell us what is wrong: tap the ? icon on a story, or write from the Tell us page. You get a reply by email.

Method, briefly

The independent reader is an AI model that sees only the article text and our order; it is not used to select stories. Calibration takes 70 randomly chosen stories per score band, collected 29 Sep to 5 Oct,

and counts outlets in our own collection reporting the same story (by headline word overlap) within seven days. It misses paraphrases and translations, so absolute shares are low and only the comparison between bands means anything. The Wikipedia check pools 10 days (2 to 11 Oct) and matches events to headlines with a small model among word-overlap candidates; an audit of 24 unmatched events found two or three arguably covered. The depth figures come from an AI reader scoring each article on reporting, evidence and analysis. All ranges are Wilson 95% intervals. Updated 12 October 2026.