

AN ILLUSTRATED EXPLAINER

TruthForge (Media)

TruthForge is a browser tool that reads a news article and shows you how it was *built* — not whether its facts are true, but whether the true facts were arranged to make you believe something extra.

The big idea: true and misleading are different questions


Most fact-checkers ask one question: **is this true?** TruthForge asks a second, separate question that a fact-checker skips: **is this candid?** An article can pass every single fact-check and still be built to make you believe something the writer never actually said out loud.

These two questions don't collapse into one score. A statement can be completely true and completely misleading at the same time — and TruthForge's whole design rests on keeping those two readings apart.

Think of it like a magic trick

A good magic trick uses **real objects**. The cup is a real cup. The card is a real card. Nothing on the table is fake — and yet you still get fooled, because the trick isn't in the objects, it's in **where your attention gets pointed**.

A misleading article works the same way. Every sentence in it can be independently, checkably true — real quotes, real numbers, real events — and the article can still walk you to a false conclusion, just by choosing which true things to show you, in which order, with which words around them.



Same shape, two domains: nothing false on the table, and yet a false belief in your head.

TruthForge calls the first question **truth** and the second **candor**. It grades them separately, on purpose, for every claim — never merging them into one true/false verdict.

Eight things it's allowed to look for — nothing else

TruthForge doesn't roam free looking for "bias" in some vague sense. It works down a fixed checklist of exactly **eight named tricks** a piece of writing can play. If an article doesn't do a particular trick, that item is simply left off the report — never shown as an empty, checked-off box.

Framing & tone
word choice tilts the mood

Sourcing balance
who gets quoted, who doesn't

Missing context
no baseline to compare against

Factual gaps
claims that can't be verified

Causal insinuation
implying motive without saying it

Omission
what got left on the floor

Rhetorical technique
fallacies dressed as argument

Fabrication
stated things that aren't so

only these eight
a closed list, not open-ended

An article that's clean on a category shows nothing for it — silence, not a green check.

Eight named tricks, checked one at a time. No trick found means no entry — never a hollow "all clear" badge.

Naming the belief it planted — not just the trick it used

For each finding, TruthForge does something a generic "bias detector" doesn't: it writes down the **exact phrase** in the article doing the planting, and then the grounded, sourced correction next to it.

Plants

"officials scrambled
to contain the fallout"

reads as: a crisis
was barely controlled



Actually

the agency followed its
routine 48-hour protocol

grounded in the
agency's own timeline

Not "this article is biased" — the exact planted phrase, paired with what the sourced record actually shows.

This is deliberately the centerpiece of every finding. A reader can't act on "this feels slanted" — they can act on "this specific phrase implies X, and here's the grounded Y."

Confidence, earned finding by finding — never one badge for

the whole article

What a badge tool does

62%

"reliability score"

One number for the whole piece — hiding that some claims in it are rock-solid and others are shaky.

What TruthForge does

Finding 1 — high confidence

Finding 2 — low confidence

Finding 3 — high confidence

Each read calibrated to what was actually checkable for that one claim.

There is no article-level score anywhere in the tool. A single badge invites you to stop thinking; per-finding confidence keeps the honest uncertainty visible exactly where it lives.

 **Widening the lens, only when you ask**

Sometimes every source backing a finding leans the same political direction — not because the other side is wrong, but because nobody went and looked. On demand, TruthForge goes hunting for the missing side's coverage of the same event.

Before: one-sided coverage



3 sources, same direction

→ search

After: both directions checked



real find, or an honest "nothing found"

Coverage, not a scorecard: it never ranks left against right, only checks whether both looked.

Where the analogy stops: this widens *inputs*, never tilts the *verdict*. It's not a left/right meter — it's the honest question "did anyone else cover this?"

Judging a source by its receipts, not its reputation

Instead of asking "is this outlet generally trustworthy," TruthForge pulls 3–6 **specific past claims** that outlet made and checks each one independently.

Past claim	Outcome
"Plant will add 4,000 jobs by 2024"	held up
"Study found no link between X and Y"	revised
"Official denies any involvement" ... a track record, kept separate from the outlet's politics.	retracted

A ledger of specific, checkable claims — not a star rating.

Turning the checklist on itself

The same discipline TruthForge applies to an article, it applies to *itself*. A control in the header always explains what its own eight-category framework is — and where that framework structurally falls short. A tool that scrutinizes framing but hides its own frame would be hypocritical; TruthForge shows you its own instrument.

⚠ What people assume TruthForge is — and isn't

Easy to assume

"It's a fact-checker that tells me what's true." /
"It'll tell me which side is right." / "It breaks
down any claim into a tree of sub-claims." / "It
adjusts to my own beliefs."

What it actually is

A framing critique of *one article*. No verdicts on
contested claims, no persuasion help, no truth-
claim decomposition tree, no reader-
personalized worldview lens.

The decomposition tree and the worldview lens were real parts of the parent project TruthForge (Media) was extracted from — and both were deliberately left behind. This tool keeps only the piece that actually shipped and worked: one article, one honest read of how it's framed. It's free to grow in its own direction from here, without needing to track wherever the parent project goes next.

In one breath

TruthForge reads the article you're looking at and separates two questions that usually get blurred into one — is it true, and is it candid — then, using a fixed eight-item checklist, names the exact phrases that plant beliefs the article never states outright, grounds each correction in real sources, grades its own confidence claim by claim instead of handing out one badge, and — only when you ask — checks whether the other side of the story got covered at all. Nothing here rules on who's right; it just makes the framing visible.