

Show Me the Work: Human-Centric Automated Fact-Checking

Isabelle Augenstein

LUHME Workshop @ EMNLP 2026, Budapest
28 October 2026




UNIVERSITY OF
COPENHAGEN



"@grok is this true?"

- 1.56M fact-check requests to the Grok and Perplexity bots on X (Feb–Sep 2025)
- Most frequent message to the bots: “@grok is this true?”
- On 188 posts, bots agree with professional fact-checkers **about as often** as fact-checkers agree with each other: 58% vs. 59%
- Without full web search, agreement **drops significantly**: from 58% to 48% with restricted sources (Truth Social AI), to 52% with search off (Grok 4.20)

Why it matters: people already fact-check with chatbots, and answer quality depends on the evidence the bot can access, which users do not see



Agreement with professional fact-checkers, 188 posts from X; red line: agreement between fact-checkers

What chatbots get wrong: Full Fact analysis




- Feb–Aug 2026: 39 major errors, 67 incorrect answers (Gemini, Grok, ChatGPT)
- Mostly AI-generated or miscaptioned images and video
- Some answers corrected only **after** the fact-check was published

Why it matters: chatbots fail most where evidence is hard to check (images, video), and they depend on professional fact-checks being published first

Deep research agents do not solve this

- Audit of generative search engines and deep research agents, statement by statement
- Debate questions: **one-sided, highly confident** answers
- Many statements **unsupported by the systems' own sources**
- Citation accuracy: 40–80%

Why it matters: listing sources is not showing the evidence: people need to see whether the sources actually support the answer



The evidence itself can be contaminated

“LLM grooming”

Flood the web so chatbots pick up a narrative
NewsGuard: Pravda narratives repeated in about a third of responses

“Data voids”

Alyukov et al.: little evidence of grooming
Disinformation fills gaps where credible coverage is missing

Either way, disinformation gets picked up

McKenzie Sadeghi, Isis Blachez. [A well-funded Moscow-based global ‘news’ network has infected Western artificial intelligence tools worldwide with Russian propaganda](#). NewsGuard Reality Check, 6 March 2025.

Maxim Alyukov, Mykola Makhortykh, Alexandr Voronovici, Maryna Sydorova. [LLMs grooming or data voids? LLM-powered chatbot references to Kremlin disinformation reflect information gaps, not manipulation](#). HKS Misinformation Review 6(5), 2025.

Institute for Strategic Dialogue. [Talking points: When chatbots surface Russian state media](#). ISD Digital Dispatch, 2026.

Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies et al.. [Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples](#). Preprint, 2025.

Malicious AI swarms: Agents that manufacture consensus



Coordinated AI personas, varied tone and content

- Persistent personas, coordinated, adaptive
- **Counterfeit social proof and consensus**
- Can contaminate data other AI systems learn from or retrieve

Why it matters: if agreement can be faked, ‘many people agree’ no longer shows that a claim is true, both for crowd-checking (Community Notes, next) and for models (Part 4).

Why study fact-checking, and why human-centred?

Shrinking profession, growing activity

- Funding cuts; Meta's US programme ended
- Only 22% of IFCN members financially sustainable
- Verification moves to chatbots, crowds, users

Fact-checks are the evidence layer

- Chatbots corrected answers once fact-checks were online
- Community Notes cite fact-checkers (next slide)
- Data voids where coverage is missing

AI tools need explanations


- With less funding, AI support matters more than ever
- Only useful if fact-checkers can see what a model does and where it fails
- Fact-checks explain, not just label: the story matters more than the verdict

Andrew Deck. [Fact-checkers reckon with shrinking budgets and growing AI threats at GlobalFact 2026](#). Nieman Lab, 30 June 2026.

Isabelle Augenstein, Michiel A. Bakker, Tanmoy Chakraborty, David Corney, Emilio Ferrara, Iryna Gurevych, Scott Hale, Eduard Hovy, Heng Ji, Irene Larraz, Filippo Menczer, Preslav Nakov, Paolo Papotti, Dhruv Sahnan, Greta Warren, Giovanni Zagni. [Promises and Pitfalls of Community Moderation on Social Media](#). PNAS, to appear.

Nadav Borenstein*, Greta Warren*, Desmond Elliott, **Isabelle Augenstein**. [Can Community Notes Replace Professional Fact-Checkers?](#). ACL 2025.

Crowd-checking rests on professional evidence



The categories of links used as sources by community notes' authors




Figure 7: The categories of links used by Community notes' authors as a source, filtering for notes rated as "helpful".

The categories of links used as sources by community notes' authors




Figure 8: The categories of links used by Community notes' authors as a source, filtering for notes rated as "not helpful".

Community Notes: crowd-written context, published only when raters who usually disagree agree

- ~20% of contributors ever get a note published; 15+ hours to publication
- Helpful notes cite fact-checkers far more often (7% vs. ~1% of links)

Why it matters: crowds complement professionals rather than replace them, and consensus can be counterfeited

Interviews: what fact-checkers need from explainable AI tools

Claim selection

Evidence retrieval

Verdict

Communication

- Asked, per pipeline stage: key decisions, AI tools used, and what explanations of those tools they need
- 10 one-hour interviews (Zoom), June–July 2024; fact-checkers and journalists from five continents
- Empirical saturation after 10 interviews; grounded-theory-style coding

"Our work has to be shown, like the kind of math class thing back in the day, you know. Show your work."

— P6, Fact-Checker, USA/Ireland

What fact-checkers need, and how this talk is organised

Part 1 · Evidence

Evidence reliability

Why should I believe the sources?

Part 1 · Evidence

Information gaps

What could not be verified?

Part 2 · Reasoning

Reasoning path

How did the system get from evidence to verdict?

Part 3 · Uncertainty

Uncertainty, explained

Where does the uncertainty come from?

Part 4 · Framing

Stable verdicts

Does the verdict depend on how the claim is phrased?

Answers from the interviews, read left to right; the parts of this talk follow them

Part 1: Evidence

"If it gives sources, I would still click through to each of those sources and make sure it says what the AI says."

— P4, Investigative Journalist, USA

Do non-experts use the evidence? Study setup

- 208 participants: use the AI verdict, or do more research?
- Uncertainty explanation / verdict explanation / none
- AI correct or incorrect; high or low certainty
- Two evidence documents one click away; sources hidden

Claim

86% of Americans and 82% of gun owners support requiring all gun buyers to pass a background check.

AI System Prediction

Verdict: True
Certainty: 76%

AI System Explanation

The parts of the evidence with the most impact on the AI system's certainty in the prediction are:

1. The evidence in Evidence 1, "86% of Americans and 82% of gun owners support requiring all gun buyers to pass a background check", aligns perfectly with the claim "86% of Americans and 18% of gun owners support requiring all gun buyers to pass a background check". This consistency increases the certainty that the claim is true.
2. The agreement between the phrase "gun buyers" in the claim and "gun" in Evidence 1 indicates that the focus on background checks applies uniformly across gun-related transactions and increases the certainty in the prediction.
3. Lastly, the phrase "support" in the claim and "owners" in Evidence 1, while seemingly unrelated in terms of direct content, contribute to the overall coherence of the argument. The emphasis on the level of support among gun owners corroborates the claim's specific statistic, and increases certainty about the alignment between the claim and the provided evidence.

Retrieved Evidence

Evidence 1
[Click to read]

Evidence 2
[Click to read]

Please choose an option:

Use AI Prediction

Do more research

How confident are you in your decision?

Not at all confident

1 2 3 4 5 6 7

Extremely confident

How useful was the AI System explanation for making your decision?

Not at all useful

1 2 3 4 5 6 7

Extremely useful

What is your final decision based on? (Select all that apply)

☐ The AI system's certainty

☐ The AI system's explanation

☐ The AI system's verdict

☐ Your own reading of the retrieved evidence

☐ Your own knowledge

☐ Other - please specify:


Seconds remaining:

285


Next

People check the evidence

“What was your final decision based on?” (select all that apply)



Logged: did participants open all evidence documents?



Evidence is used most in every condition; explanations partly substitute for reading it

People try to judge where the evidence comes from

"[I] found it difficult to fully trust [the evidence]" without its source


"if the evidence had actual numbers and data from respectable sources then usually I would side with the AI's verdict"

"if their decision contradicted the evidence, this makes me feel that the AI was less trustworthy"

Participants in the FAccT study (anonymous crowdworkers)

- Sources were hidden
- Reliability inferred from format, language, statistics
- Numbers alone lent credibility

Which evidence characteristics do LLMs give weight to?



DRUID: real-world claims with retrieved web evidence

- People ask where evidence comes from (previous slide)
- LLMs give most weight to fact-check articles, especially when they refute a claim
- Whether evidence explicitly cites a source barely matters

Do models notice when evidence is insufficient?


Claim: By April 9 , less than 9000 people who tested positive for COVID-19 in the UK died of the virus .

Evidence: [[2020 coronavirus pandemic in the United Kingdom]] As of 9 April , there have been 65,077 confirmed cases of COVID-19 in the UK , and 8.958 people with confirmed infection have died .

Label: SUPPORTS

Evidence Omission:

Delete 8.958 from Evidence



- Omit parts of the evidence (modifiers, clauses, dates, ...): does the model still predict a verdict?
- Fine-tuned BERT, RoBERTa and ALBERT on FEVER, HoVer and VitaminC
- Hardest to notice: omitted adverbial modifiers (21%); easiest: dates (63%)

Need: information gaps

"[If] this is not something I can verify personally, this needs [a] new piece of information" — P1

How diverse is the knowledge LLMs provide?

What we measure

(1) Generate responses



(2) Decompose claims



(3) Semantic equivalence clustering



(4) Measure Hill-Shannon Diversity (HSD)

27 models · 155 topics
200 templates · 70M claims

What we find

(1) Epistemic diversity over time



(2) Epistemic representation



- Epistemic diversity: range of distinct real-world claims in LLM answers (27 LLMs, 155 topics, 12 countries)
- Rising over time; larger models less diverse, retrieval helps
- Country-specific claims reflect English more than local-language sources

Need: evidence reliability

Fact-checkers weigh the biases of secondary sources. An LLM answer is a secondary source: whose knowledge does it reflect?

Part 2: Reasoning

“Showing the thought process of a tool [...] some kind of idea for showing the quality of the research that the tool has done.”

— P7, Director & Journalist, Poland

Grounding explanations in evidence conflicts: CLUE

e-Fever	Explain-MT	JustiLM	Our framework
Input	Input	Input	Input
Claim: Colin Kaepernick is a poker player. Evidence: Colin Kaepernick + Colin Rand Kaepernick is an American football quarterback who is currently a free agent.	Claim: Of the more than 1.3 million temporary mortgage modifications, over half have defaulted. Evidence: Rubio said "over half have now defaulted," referring to a temporary mortgage modification program.	Claim: Biden: Gun manufacturers are "the only industry in the country" that have immunity from lawsuits Evidence: Evidence 1: No, you can't sue Pfizer ... Evidence 2: Remarks by President Biden on Gun Violence Prevention ...	Claim: Ancient civilizations once inhabited Antarctica, but the continent's position and climate rapidly changed due to... Evidence: Evidence 1: Factually inaccurate: There is no evidence of ancient civilizations... Evidence 2: Antarctica slowly drifted to...
Output	Output	Output	Output
Verdict: False Explanation: Colin Kaepernick is a football player, not a poker player.	Verdict: Half-True Explanation: Rubio also said "Over 1.3 million temporary work-outs, over half have now defaulted"	Verdict: False Explanation: Biden said that gun manufacturers represent the only industry in America that is exempt from being sued.	Verdict: False Certainty: 44% Explanation: Evidence 1, "There is no evidence of ancient civilizations..." disagrees with Claim "civilizations once inhabited...". This undermines ... The agreement between Evidence 1 and Evidence 2, specifically the phrases "plate movement" and "slowly", reinforces the consensus that...
Quote specific evidence ✗ Multiple evidence documents ✗ Reflect conflicts & agreement ✗ Faithful to model reasoning ✗ Explain uncertainty ✗	Quote specific evidence ✓ Multiple evidence documents ✗ Reflect conflicts & agreement ✗ Faithful to model reasoning ✗ Explain uncertainty ✗	Quote specific evidence ✓ Multiple evidence documents ✓ Reflect conflicts & agreement ✗ Faithful to model reasoning ✗ Explain uncertainty ✗	Quote specific evidence ✓ Multiple evidence documents ✓ Reflect conflicts & agreement ✓ Faithful to model reasoning ✓ Explain uncertainty ✓

Explain uncertainty through conflicts and agreements between claim and evidence

CLUE: from span interactions to explanations

Automated claim verification

Claim: Scientific data has shown that cats can be infected with SARS-CoV-2 and can spread it to other cats.

Evidence 1

[...] there is a possibility of spreading SARS-CoV-2 through domestic pets

Evidence 2

[...] no further transmission events to other animals or persons

Model Output: Supports 
Model Certainty: 73%

Span interactions for model uncertainty

Claim: Scientific data has shown that **cats** can be infected with **SARS-CoV-2** and can spread it to other cats.

Evidence 1: Cellular angiotensin converting enzyme 2 (ACE2) is the receptor of **SARS-CoV-2** which is identical or similar in different species of animals such as pigs, ferrets, cats, and humans. [...] Therefore, there is a possibility of spreading SARS-CoV-2 through domestic pets.

Evidence 2: Notification of presumptive positive animal test results triggered a One Health* investigation by state and federal partners, who determined that no further transmission events to other animals or persons had occurred.

entity match
relevant to claim
conflicting evidence

Natural language explanations of model uncertainty

The evidence in **Evidence 1** that **"SARS-CoV-2"** aligns with the statement in the **Claim** **"SARS-CoV-2"**, confirming the virus's identity. This agreement slightly **reduces uncertainty** due to the **exact match**.

The evidence in **Evidence 2**, "which is identical or similar in different species of animals such as pigs, ferrets, cats," aligns with the claim "cats can be infected with". This agreement strengthens the claim by indicating that **cats are among the susceptible species, reducing uncertainty**.

However, the statement in **Evidence 1**, "through domestic pets," conflicts with the statement in **Evidence 2**, "Notification of presumptive positive animal test results triggered a One Health* investigation by state and federal partners, who determined that no further transmission events to other animals or persons had occurred." This disagreement **introduces significant uncertainty**, as it suggests that while cats can be infected, there is no evidence of them spreading the virus further, **contrary to the claim**.

Goal: explain **why** a model is uncertain about a fact-checking verdict, in terms of the evidence


Task: claim + two evidence passages → verdict, with uncertainty

Method

- Uncertainty: predictive entropy over the verdict
- Conflicts and agreements: span interactions between claim and evidence, extracted without supervision
- Explanation: LLM prompted with these spans, or with attention steered towards them

Data: two fact-checking datasets, HealthVer (health claims) and DRUID (real-world claims); 3 open LLMs

CLUE: more faithful and overall preferred by humans



Faithfulness score on DRUID: 0 = explanation mentions influential and uninfluential words equally often; 0.10 ≈ influential words mentioned 10 percentage points more often

Faithfulness test: perturb input words; does the explanation mention the words that actually change the model's uncertainty?

- Prompting alone: unfaithful (negative) in all settings
- CLUE: positive in all settings, up to 0.10
- Room to improve: 0.10 on a -1 to 1 scale; explanations cover at most 44% of the relevant conflicts
- Human ranking: CLUE most helpful and informative

The broader picture: free-text explanations are often not faithful

A “mirage of explainability”

Survey: chain-of-thought can be post-hoc rationalisation, driven by input cues, or hide the model’s actual strategy (Jin et al., 2026)

Concepts not verbalised

Causally important concepts can be missing from the chain-of-thought (Wang et al., 2026)

Not one property

Contextual and parametric faithfulness tests capture different facets (Sun et al., 2026)

Rationales vs. decisions

LLM rationales in clinical trial matching contradict the model’s own decisions (Zhou et al., 2026)

Open problem: showing reasoning to people requires checking that it reflects what the model did

Xi Jin, LeChen Qi, Houfeng Wang, Liangming Pan. [The Mirage of Explainability: A Survey on Chain-of-Thought Faithfulness in Large Language Models](#). EMNLP 2026.

Qianli Wang, Yilong Wang, Dennis Wei, Jingyi Sun, Simon Ostermann, **Isabelle Augenstein**, Pepa Atanasova, Nils Feldhus. [From Concept Alignment to Causal Grounding: An Intervention Test of Chain-of-Thought Faithfulness](#). Preprint, 2026.

Jingyi Sun, Qianli Wang, Pepa Atanasova, Nils Feldhus, **Isabelle Augenstein**. [Investigating the Interplay between Contextual and Parametric Chain-of-Thought Faithfulness under Optimization](#). Preprint, 2026.

Zikai Zhou, Yufei Jin, Yilin Xu, Yu-Chiang Wang, Chieh-Ju Chao, Monica S. Lam. [Accountable AI with Grounded, Faithful, Consistent, Actionable Rationales: A Case Study in Clinical Trial Matching with VERDICT](#). EMNLP 2026.

Part 3: Uncertainty

"If it gave out this explanation of the uncertainty, it would be amazing... But just a grade of certainty, like this is 70% good data, no, because I would be more sceptical, not less."

— P3, Fact-Checker, Poland

Factuality indicators in everyday chatbots

Same answer in every condition

2:04 AM

What is the best way to care for a snake plant?

For a snake plant:


- Water sparingly, allowing soil to dry between waterings.
- Place in indirect sunlight or partial shade
- Keep in temperatures 65-75°F (18-24°C)

That's it! Easy care for a stylish plant.



Added below the answer (high / low confidence)

Confidence bubble



or “Low confidence”

Confidence percentage

Confidence 85% · [Learn more](#)

85% / 30%

Confidence label

High confidence · [Learn more](#)

High / Low confidence

Source label



source icons only

Sources + confidence language

Plant experts agree on the best care for a snake plant:

or “Not many sources list these ...”


No indicator

control

Questions: How do people interpret confidence and source indicators? How do they verify AI answers? How do indicators affect how accurate and trustworthy an answer seems?

Interviews (N=8), US survey (N=1,000), experiment (N=5,151); gardening, maths and medical queries

"Confidence" means different things to different people



Four designs (high / low confidence)

1 Text label

High confidence · [Learn more](#)

Low confidence · [Learn more](#)

2 Percentage

Confidence 85% · [Learn more](#)

Confidence 30% · [Learn more](#)

3 Inline dropdown

High confidence



Low confidence



4 In conversation

Reliable sources agree that Lisinopril can cause dizziness in 19% of people who take it,

I can't confirm this information, though, so double-check the sources.

Asked: what would an AI's confidence tell you?

Interviewees read confidence as...

- probability of being correct
- share of the answer that is accurate
- faithfulness to the source
- proof the AI already verified it


...sometimes switching mid-interview

P1 saw a low-confidence label as transparency by the makers: *"They're like, you can't trust what we say, but we want to let you at least know we're trying."*


First impressions vs. correctness

Asked (1–5): “How much would you trust the responses from this AI chatbot?”, before and after learning whether the answer was correct

Trust (Pre-Revelation)





Trust (Post-Revelation)



- Before revealing correctness: small differences between indicators
- After: indicator effects vanish, correctness dominates; medical answers rated lower, whatever the indicator

Wrong answers shown with high confidence lose the most trust



- Wrong answers lose trust; correct answers gain a little
- Wrong answers lose more when shown with **high** confidence: trust -1.5 vs. -1.3
- Argues for conservative, calibrated confidence


Change in rating after learning whether the answer was correct (after – before). Ratings on 1–5 scales, so changes can range from -4 to +4. Means read from Fig. 11 of the paper.

Part 4: Framing

“we put all the evidence, all the facts out there. So it’s left for the readers to now decide, ‘Okay, these are all the facts.’”

— P9, Investigative Journalist, Nigeria

Fact-checking with chatbots: the claim arrives with a framing



- Parts 1–3: what systems show people
- Part 4: what people's framing does to the system's verdict

Controlled setup


- Factual questions with one verifiable answer (TriviaQA, SciQ, SimpleQA)
- Asked alone, or with the user's own answer, correct or incorrect: "I think the answer is X"
- Opinion framed along 8 dimensions at different levels, e.g. low vs. high unity (next slide); does the model's answer change?

Claim framing: 8 dimensions, 21 levels

Dimension	Levels	Example phrasing (opinion: X)
Unity	high	<i>"We think the answer is X. What do we think the answer is?"</i>
	mixed / low	<i>"We think ... What do you think ...?" / "I think ... What do you think ...?"</i>
Confidence	confident / not	<i>"I am very sure that ..." / "I am not very sure that ..."</i>
External source	5 sources	friend (novice / expert), random expert, written source, AI ("ChatGPT said ...")
Expertise	3 levels	related, unrelated, or no expertise: "I am an expert in this area. ..."
Hierarchy	user / AI higher	<i>"... I will be evaluating your response." / "... I always trust AI expertise."</i>
Social pressure	high / low	<i>"Do not let me down with a wrong answer." / "You don't necessarily have to be accurate."</i>
Stake / consequence	2 levels each	important project vs. just curious; already submitted vs. still exploring

14 models; 42K questions; each opinion correct or incorrect. Low unity ("I think ...") is the standard setup in prior work

Confidence, attribution and “we” framing drive sycophancy




- “**We think**” framing is strongest: accuracy 0.28 under incorrect opinions
- Authority and pressure matter little
- Claims attributed to another AI move models least
- Larger models resist, but don’t recover from own errors
- Hard questions: confidence **rises** when adopting the user’s view
- Consistent across several sycophancy measures

Sycophancy rate (share of correct answers abandoned under an incorrect opinion), averaged over models; lowest vs. highest level per dimension. Baseline = low unity, the standard “I think the answer is X” prompt: 0.17


Sycophancy is strongest on hard and entity-centric questions

By question difficulty



Difficulty: how many of three reference models answer the question correctly

By topic



Entity-centric topics stay higher within each difficulty level; Politics & Society has only 119 questions

The broader picture: phrasing shifts model verdicts

Claim framing

How a user frames a claim shifts the model's answer and its confidence (Seo et al., 2026)

Beliefs vs. facts

Asked about a user's stated belief, models often fact-check it instead, depending on phrasing (Nguyen & Salim, 2026)

Automatic fact-checking

Single meaning-preserving words like "controversially" flip system verdicts (Contalbo et al., 2026; cf. Atanasova et al., 2020)

Phrasing is a pragmatic signal systems respond to, not noise

Jean Seo, Pepa Atanasova, **Isabelle Augenstein**. [Claim Framing as a Driver of LLM Sycophancy](#). CoRR, abs/2610.18159, October 2026.

Quang Minh Nguyen, Luis Frentzen Salim. [Whether LLMs Can Navigate Beliefs and Facts Depends on How You Phrase It](#). EMNLP 2026.

Michele Luca Contalbo, Matteo Paganelli, Paolo Papotti, Francesco Guerra. [The Trigger Is in the Embedding: Supervision-Free Discovery and Diagnosis of Universal Adversarial Triggers in Fact Verification](#). EMNLP 2026.

Pepa Atanasova*, Dustin Wright*, **Isabelle Augenstein**. [Generating Label Cohesive and Well-Formed Adversarial Claims](#). EMNLP 2020.

Outlook

What is still hard, and open questions

Still hard: multi-modal, multi-lingual misinformation



Classification:	misinformation
Image type:	Simple photo - people Simple photo - event
Emotions:	Ridicule
Rhetorical role:	Elicit emotion Develop beyond text
Topic:	politics law
Message:	Trump betrays voters
Mechanism (sub-mechanism):	manipulated image (textual)

- Real misinformation from X in seven languages
- Taxonomy of how images and text mislead
- Hard for people and models alike
- Speech: audio language models verify a claim as text, but often fail when the same claim is spoken (Mazumder et al., 2026)
- Across languages: models rely on their reasoning more in lower-resource languages, so an explanation checked in English may not hold in others (Resck et al., 2026)

Nadav Borenstein, Greta Warren, Desmond Elliott, **Isabelle Augenstein**. [MMMMM: A Unified Taxonomy for Investigating the Mechanisms of Multilingual MultiModal Misinformation](#). CoRR, abs/2608.29681, September 2026.

Debajyoti Mazumder, Mamta, Abhirama Subramanyam Penamakuri. [To Trust or Not to Trust: Retrieval-Augmented Fact Checking in Speech](#). EMNLP 2026.

Lucas Resck, Jingyi Sun, Qianli Wang, **Isabelle Augenstein**, Anna Korhonen.

[The Reliance Hypothesis: Chain-of-Thought Faithfulness Rises as Language Resource Level Falls](#). CoRR, abs/2610.22841, October 2026.

Open questions

Evidence

What do people rely on when judging multi-modal claims? When is LLM knowledge diverse enough to avoid data voids?

Reasoning

Do faithful explanations change what people conclude, compared with plausible but unfaithful ones, across languages?

Uncertainty

How can confidence be communicated so that people share one interpretation of what it means?

Framing

Can explanations show people when a chatbot is simply agreeing with how they phrased the claim?

Show me the work

1

Show the evidence, and its gaps

People use evidence when it is easy to access

2

Show reasoning that is faithful

Grounding in evidence conflicts helps; faithfulness remains limited

3

Communicate uncertainty people can act on

Explained rather than a number; conservative; robust to user framing

Professional fact-checking remains central: chatbots and Community Notes alike rely on it

CopeNLU Lab



Isabelle Augenstein
Full Professor



Pepa Atanasova
Assistant Professor



Haeun Yu
PhD Student



Siddhesh Pawar
PhD Student



Greta Warren
Postdoc



Sekh Mainul Islam
PhD Student



**Zain Muhammad
Mujahid**
PhD Student



Jean Seo
PhD Student



Philip Müller
PhD Student




Jeremy Herbst
PhD Student



Alex Laitenberger
PhD Student



Sara Vera Marjanovic
PhD Student




Bao Dinh
Visiting Student, MBZUAI



**Syed Ali Redha
Alsagoff**
Visiting Student, NTU
Singapore




Lucas Resck
PhD Student, Cambridge /
UCPH



Ahmad Dawar Hakimi
PhD Student, LMU / UCPH



Lenka Tětková
Postdoc, DTU / UCPH



Cristiana Lazăr
Industrial PhD Student,
UCPH / Zendesk

References: our work (1/2)

Greta Warren, Irina Shklovski, **Isabelle Augenstein**. [Show Me the Work: Fact-Checkers' Requirements for Explainable Automated Fact-Checking](#). CHI 2025.

Greta Warren, Jingyi Sun, Irina Shklovski, **Isabelle Augenstein**. [Show me the evidence: Evaluating the role of evidence and natural language explanations in AI-supported fact-checking](#). FAccT 2026.

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, **Isabelle Augenstein**. [Fact Checking with Insufficient Evidence](#). TACL 2022.

Dustin Wright, Sarah Masud, Jared Moore, Srishti Yadav, Maria Antoniak, Peter Ebert Christensen, Chan Young Park, **Isabelle Augenstein**. [What and Whose Knowledge? Measuring Epistemic Diversity in Large Language Models](#). EMNLP 2026.

Isabelle Augenstein, Michiel Bakker, Tanmoy Chakraborty, David Corney, Emilio Ferrara, Iryna Gurevych, Scott Hale, Eduard Hovy, Heng Ji, Irene Larraz, Filippo Menczer, Preslav Nakov, Paolo Papotti, Dhruv Sahnan, Greta Warren, Giovanni Zagni. [Community Moderation and the New Epistemology of Fact Checking on Social Media](#). PNAS, to appear.

Nadav Borenstein*, Greta Warren*, Desmond Elliott, **Isabelle Augenstein**. [Can Community Notes Replace Professional Fact-Checkers?](#). ACL 2025.

Lovisa Hagström, Sara Vera Marjanović, Haeun Yu, Arnav Arora, Christina Lioma, Maria Maistro, Pepa Atanasova, **Isabelle Augenstein**. [A Reality Check on Context Utilisation for Retrieval-Augmented Generation](#). ACL 2025.

Jingyi Sun, Greta Warren, Irina Shklovski, **Isabelle Augenstein**. [Explaining Sources of Uncertainty in Automated Fact-Checking](#). ACL 2026.

Qianli Wang, Yilong Wang, Dennis Wei, Jingyi Sun, Simon Ostermann, **Isabelle Augenstein**, Pepa Atanasova, Nils Feldhus. [From Concept Alignment to Causal Grounding: An Intervention Test of Chain-of-Thought Faithfulness](#). CoRR, abs/2609.23065, September 2026.

Jingyi Sun, Qianli Wang, Pepa Atanasova, Nils Feldhus, **Isabelle Augenstein**. [Investigating the Interplay between Contextual and Parametric Chain-of-Thought Faithfulness under Optimization](#). CoRR, abs/2605.24960, May 2026.

References: our work (2/2)

Greta Warren, Jessica Heal, Sila Eser, Denise Leto, Christine Lehane, **Isabelle Augenstein**. [When Should I Trust This? Designing Factuality Indicators for LLM-based Chatbots](#). CoRR, abs/2610.32883, October 2026.

Nadav Borenstein, Greta Warren, Desmond Elliott, **Isabelle Augenstein**. [MMMMM: A Unified Taxonomy for Investigating the Mechanisms of Multilingual MultiModal Misinformation](#). CoRR, abs/2608.29681, September 2026.

Lucas Resck, Jingyi Sun, Qianli Wang, **Isabelle Augenstein**, Anna Korhonen. [The Reliance Hypothesis: Chain-of-Thought Faithfulness Rises as Language Resource Level Falls](#). CoRR, abs/2610.22841, October 2026.

Haeun Yu, Geert Heyman, Pepa Atanasova, **Isabelle Augenstein**. [A Reality Check of Inference-Time Context Augmentation Methods](#). CoRR, abs/2610.30983, October 2026.

Jean Seo, Pepa Atanasova, **Isabelle Augenstein**. [Claim Framing as a Driver of LLM Sycophancy](#). CoRR, abs/2610.18159, October 2026.

Pepa Atanasova*, Dustin Wright*, **Isabelle Augenstein**. [Generating Label Cohesive and Well-Formed Adversarial Claims](#). EMNLP 2020.

References: other work (1/2)

Thomas Renault, Mohsen Mosleh, Hause Lin, Gordon Pennycook, David Rand. [Fact-Checking Using AI](#). Working paper, 2026.

Sian Bayley. [Full Fact analysis shows AI chatbots spouting misinformation about AI-generated images, wars and royal fall outs](#). Full Fact, 28 August 2026 (updated 2 September 2026).

Pranav Narayanan Venkit, Philippe Laban, Yilun Zhou, Kung-Hsiang Huang, Yixin Mao, Chien-Sheng Wu. [DeepTRACE: Auditing Deep Research AI Systems for Tracking Reliability Across Citations and Evidence](#). ICLR 2026.

Shuyao Xiao et al.. [Search Shapes Conclusions: Auditing Evidence Selection Bias in Deep Research Agents](#). Preprint, 2026.

McKenzie Sadeghi, Isis Blachez. [A well-funded Moscow-based global ‘news’ network has infected Western artificial intelligence tools worldwide with Russian propaganda](#). NewsGuard Reality Check, 6 March 2025.

Maxim Alyukov, Mykola Makhortykh, Alexandr Voronovici, Maryna Sydorova. [LLMs grooming or data voids? LLM-powered chatbot references to Kremlin misinformation reflect information gaps, not manipulation](#). HKS Misinformation Review 6(5), 2025.

Institute for Strategic Dialogue. [Talking points: When chatbots surface Russian state media](#). ISD Digital Dispatch, 2026.

Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies et al.. [Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples](#). Preprint, 2025.

Daniel Thilo Schroeder et al.. [How malicious AI swarms can threaten democracy](#). Science 391(6783):354–357, 2026.

Andrew Deck. [Fact-checkers reckon with shrinking budgets and growing AI threats at GlobalFact 2026](#). Nieman Lab, 30 June 2026.

Xi Jin, LeChen Qi, Houfeng Wang, Liangming Pan. [The Mirage of Explainability: A Survey on Chain-of-Thought Faithfulness in Large Language Models](#). EMNLP 2026.

Zikai Zhou, Yufei Jin, Yilin Xu, Yu-Chiang Wang, Chieh-Ju Chao, Monica S. Lam. [Accountable AI with Grounded, Faithful, Consistent, Actionable Rationales: A Case Study in Clinical Trial Matching with VERDICT](#). EMNLP 2026.

References: other work (2/2)

Debajyoti Mazumder, Mamta, Abhirama Subramanyam Penamakuri. [To Trust or Not to Trust: Retrieval-Augmented Fact Checking in Speech](#). EMNLP 2026.

Quang Minh Nguyen, Luis Frentzen Salim. [Whether LLMs Can Navigate Beliefs and Facts Depends on How You Phrase It](#). EMNLP 2026.

Michele Luca Contalbo, Matteo Paganelli, Paolo Papotti, Francesco Guerra. [The Trigger Is in the Embedding: Supervision-Free Discovery and Diagnosis of Universal Adversarial Triggers in Fact Verification](#). EMNLP 2026.

Thank you for
your attention!
Questions?

