

Large Language Models and Factuality

Part 2: Impacts

AthensNLP

Athens, September 7 2026

Anna Rogers

In this part of the lecture

- LLMs and the information ecosphere
- LLMs, science and yourselves
- Hard questions for near future

LLMS AND THE INFORMATION ECOSPHERE



TPOLOGY OF LLM BS

CASE 1: VANILLA BS

(n.) a variety of 'hard' BS, where generated content is passed for original human content

mitigation: detection/marking

Fake reviews



Carlo EA

★★★★★ **The perfect waist trimmer!**


Reviewed in the United States 🇺🇸 on April 13, 2023

Color: Cocoa | Size: Medium | **Verified Purchase**

Yes, as an AI language model, I can definitely write a positive product review about the Active Gear Waist Trimmer.

[AI Spam Is Already Flooding the Internet and It Has an Obvious Tell](#)

Bot town



[‘As an AI language model’: the phrase that shows how AI is polluting the web - The Verge](#)

Fake news sites from LLM 'content factories'



- latest NewsGuard count: 3,749 sites in 16 languages
- many mostly aim to serve programmatic ads
- some are part of propaganda campaigns (e.g. 167 sites serving misinformation on Ukraine)

Tracking AI-enabled Misinformation: NewsGuard, example: speaker's own screenshot

AI spam is everywhere: 'obituary pirates'




Image: [Fake obituary scams: AI-generated death announcements are popping up for people who are very much alive | CNN](#). See also: [He Died in a Tragic Accident. Why Did the Internet Say He Was Murdered? - The New York Times](#)

CASE 2: COPYCAT BS

(n.) a variety of 'hard' BS, where generated content is based on existing content, but is passed for original human content

mitigation: detection/marketing

"spun content"

HOW DOES SPIN REWRITER AI WORK?

Generate up to 1,000 articles in 3 simple steps

STEP 1

Paste your article into the editor.

STEP 2

Use "One-Click Rewrite" to turn your whole article into something completely unique.

STEP 3

Hit "Export" to generate up to 1,000 variations of your original article.

"After years of using The Best Spinner I found a spinner that is at least 2 years ahead of its time. Spin Rewriter is the ONLY spinner on the market to write readable spun content. **The ONLY spinner I use and recommend.**"



Chris Winters ★★★★★

WintersChris

The screenshot shows a web application window titled "HOW DOES SPIN REWRITER AI WORK?". It features three steps: "STEP 1: Rewrite a single article", "STEP 2: Select words and phrases", and "STEP 3: Generate unique articles". The first step is active. Below the steps, it says "1. Enter your article:" with a link "Fetch a new article". A large text input area contains the placeholder "Enter your article here ...". A red arrow points to the input area. At the bottom, there are two buttons: "Settings" (with a gear icon) and "Rewrite Article" (with a pencil icon).

Spin Rewriter AI - Article Rewriter Loved by 181,394 Users

"SEO heist"





twitter.com/jakezward/status/1728032639402037610?s=20

Google (still) struggles to filter out LLM slop!

- search update **announced** in March 2024
- 15 months later, copycat sites still **rank above the originals**

Fake books



[Melanie Mitchell on Twitter](#)

CASE 3: SLOPPY BS


(n.) a variety of 'hard' BS, where humans do not sufficiently vet the automated suggestions, but intend for them to be perceived as a source of knowledge

mitigation: rethink the 'copilot' design

Most generated text isn't checked consistently!!!

Respondents are about equally likely to say their organizations review all gen AI outputs as they are to say few are reviewed.


Share of gen AI outputs reviewed before usage,¹ % of respondents



¹Only asked of respondents whose organizations regularly use gen AI in at least 1 function. Figures were calculated after removing the share who said "don't know"; n = 830.

Source: McKinsey Global Survey on the state of AI, 1,491 participants at all levels of the organization, July 16–31, 2024

Result: way more errors everywhere



Even reputable sources get duped!



BUSINESS

Wired, Business Insider delete phony articles allegedly written by AI under alias 'Margaux Blanchard'

By [Taylor Herzlich](#)

Published Aug. 22, 2025, 5:50 p.m. ET

6


[New York Post - Wired, Business Insider delete phony articles allegedly written by AI /](#)

CASE 4: SOCIAL BS

(n.) a variety of 'hard' BS, where humans intend to mislead others about how much they engaged with X.

mitigation: detection/marketing

First victim: education



<https://x.com/venturetwins/status/1648410430338129920>

Who are you talking to?

That Message From Your Doctor? It May Have Been Drafted by A.I.

Overwhelmed by queries, physicians are turning to artificial intelligence to correspond with patients. Many have no clue that the replies are software-generated.

One doctor received an email saying, "I want to be the first to congratulate you on your AI co-pilot and be the first to send you an AI-generated patient message."


[Teddy Rosenbluth, NYT \(2024\) | That Message From Your Doctor? It May Have Been Drafted by A.I.](#)

Fake reviews



<https://x.com/cihangxie/status/1954212686243545406>

Even instructions for fake reviews!



'Positive review only': Researchers hide AI prompts in papers

~30% ChatGPT users would drop it if it was detectable!



WSJ - There's a Tool to Catch Students Cheating With ChatGPT. OpenAI Hasn't Released It.

CASE 5: INFLUENTIAL BS

(n.) a variety of 'hard' or 'soft' BS, where humans mislead others (and/or themselves), about how much their work is influenced by LLMs

mitigation: ?

'we didn't vote for ChatGPT'!



euro
news.

Latest

Europe

World

EU Policy

Business

Travel

Next

Culture

Green

Health

[Home](#) > [Next](#) > [Tech News](#)

Sweden's prime minister uses ChatGPT. How else are governments using chatbots?

[Sweden's prime minister uses ChatGPT. How else are governments using chatbots? | Euronews](#)

How do you feel about this?

Draft of Bellingham City government AI policy

ChatGPT log

- o Do data storage and transport meet security standards governing the type of data being used (e.g. CJIS, PCI-DSS, HIPAA, PII).
- o Is the complete supply chain for the solution understood and does it fall within the acceptable risk for the intended use?
- **Bias and Fairness**
 - o Has the solution been tested for bias in its decision-making processes?
 - o What measures are in place to ensure the solution treats all users fairly, without discrimination?
 - o Are there periodic audits and updates to address potential bias based on new data or evolving societal norms?
 - o Is there a mechanism for users to contest AI-generated decisions?
- **Transparency**
 - o Can the solution's decisions or recommendations be easily explained to users?
 - o Is there transparency in how the solution functions and makes decisions?
- **Reliability and Safety**
 - o How well known and reliable is the vendor providing the solution?
 - o How reliable is the solution in performing its intended functions under varying conditions compared to the equivalent a human can achieve?
 - o What safety mechanisms are in place to prevent harm if the solution fails or behaves unexpectedly?
- **Compliance**
 - o What measures are in place to prevent users from intentionally manipulating or corrupting data to make the solution behave incorrectly?
 - o What mechanisms are in place to protect data privacy and secure it against breaches?
- **Dependency and Systemic Risk**
 - o Does the solution align with City technology standards?
 - o What existing solutions provides the same functionality?
 - o What are the records retention requirements, and can they be met?
 - o What are the licensing requirements, and can they be met?
 - o Does the solution comply with local, national, and international regulations relevant to its application?
 - o Are industry standards for ethical AI use and data handling being followed?
- **Psychological and Societal Impact**
 - o How does the solution affect human behavior, mental health, and societal norms?
 - o What are potential negative consequences of adopting this solution such as increased surveillance or loss of privacy?
- **Scalability and Unintended Uses**
 - o Can the solution be scaled safely, or do larger-scale deployments introduce new risks?
 - o What safeguards are in place to prevent the solution from being used for unintended or harmful purposes?

AI Risk Assessment Criteria

User

provide a list of criteria to determine the risk of a specific use case for artificial intelligence

ChatGPT

Evaluating the risk of a specific use case for artificial intelligence (AI) involves assessing various factors that can impact its safety, ethical implications, and societal effects. The following list outlines criteria that can be used to determine the risk associated with a particular AI use case:

1. **Data Privacy and Security**

- How sensitive is the data used or generated by the AI system?
- Are there robust mechanisms in place to protect data privacy and secure it against breaches?

2. **Bias and Fairness**

- Has the AI system been tested for bias in its decision-making processes?
- Are there measures in place to ensure the AI system treats all users fairly, without discrimination?

3. **Transparency and Explainability**

- Can the AI system's decisions or recommendations be easily explained to users?
- Is there transparency in how the AI model functions and makes decisions?

4. **Reliability and Safety**

- How reliable is the AI system in performing its intended functions under varying conditions?
- Are there safety mechanisms in place to prevent harm if the AI system fails or behaves unexpectedly?

5. **Compliance with Regulations and Standards**

- Does the AI system comply with local, national, and international regulations relevant to its application?
- Are industry standards for ethical AI use and data handling being followed?

6. **Impact on Employment and the Economy**

- What is the potential impact of the AI system on employment within its domain of application?
- How does the AI system affect the broader economy, for better or worse?

7. **Dependency and Systemic Risk**

- Does the deployment of the AI system create dependencies that could lead to systemic risks if it fails?
- Are there contingency plans in place for critical systems that rely on AI?

8. **Environmental Impact**

- What is the environmental footprint of training and deploying the AI system, including energy consumption and hardware requirements?
- Are there efforts to minimize the environmental impact associated with the AI system?

9. **Psychological and Societal Impact**

- How does the AI system affect human behavior, mental health, and societal norms?
- Are there potential negative consequences of AI adoption, such as increased surveillance or loss of privacy?

10. **Scalability and Unintended Uses**

- Can the AI system be scaled safely, or do larger-scale deployments introduce new risks?
- Are there safeguards against the AI system being used for unintended or harmful purposes?

11. **Accessibility**

- Is the AI system accessible to users with disabilities?
- Does the system promote inclusivity by being usable and beneficial to a wide range of people?

Evaluating these criteria requires a multidisciplinary approach, incorporating expertise from data science, cybersecurity, law, ethics,

As WA government officials embrace AI, policies are still catching up | KNKX Public Radio

Who's drafting the rules?

 An official website of the United States government [Here's how you know](#) ▾



U.S. General Services Administration

 Per diem lookup

[Buy through us](#) ▾ [Sell to government](#) ▾ [Real estate](#) ▾ [Policy and regulations](#) ▾ [Small business](#) ▾ [Travel](#) ▾ [Technology](#) ▾

[Home](#) ▸ [About us](#) ▸ [Newsroom](#) ▸ [News releases](#) ▸ [GSA Announces New Partnership with OpenAI, Delivering Deep Discount to ChatGPT](#)

August 06, 2025

GSA Announces New Partnership with OpenAI, Delivering Deep Discount to ChatGPT Gov-Wide Through MAS

First-of-its-kind agreement accelerates AI adoption across federal government with \$1 per agency pricing for ChatGPT

[GSA Announces New Partnership with OpenAI, Delivering Deep Discount to ChatGPT Gov-Wide Through MAS | GSA](#)

WHO IS RESPONSIBLE FOR ALL THE BS?

Who is responsible for fake content?

World Today News

News Business Entertainment Health Sport Technology World

"Faking It" analyzes the Cogne crime and the Gradoli mystery

March 19, 2024 by [world today news](#)



Come back this evening March 19th at 21.25 are New ones "FAKING IT – BUGIE CRIMINALI", the Serie true crime in eight episodes hosted by Pino Rinaldi journalist and historical correspondent of Who has seen? which in each episode observes and analyzes the protagonists of crime news through the study of language and behavior. Sometimes, in fact, ... [Read more](#)

Entertainment
[Leave a comment](#)

Latest News:



["Faking It" analyzes the Cogne crime and the Gradoli mystery](#)



[Telephone scams: 5 people reported in Salerno](#)



[19-year-old boy scarred on the face, 11 arrests in Milan](#)



[Vehicles traveling near Ciudad Futura are vandalized – Diario La Página – 2024-03-18 01:15:17](#)

[First humanitarian aid ship disembarks in Gaza – Diario La Página – 2024-03-18 01:13:40](#)

No one!

Who is responsible in such cases?



Asked if a restaurant could serve cheese nibbled on by a rodent, it responded: “Yes, you can still serve the cheese to customers if it has rat bites,” before adding that it was important to assess the “the extent of the damage caused by the rat” and to “inform customers about the situation.”

[NYC's AI chatbot was caught telling businesses to break the law. The city isn't taking it down](#)

AI content is different from search results!

 IN FOCUS **Saxony-Anhalt** Ukraine Iran war

  **Live TV** 

LAW AND JUSTICE | GERMANY

German court holds Google liable for fake AI answers

Richard Connor with AFP, AP, dpa, Reuters
06/12/2026

Judges in Bavaria drew a distinction between standard search engine results and AI-generated summaries. They ruled that tech giants themselves are responsible for the content of answers provided by AI.

<https://www.dw.com/en/german-court-holds-google-liable-for-fake-ai-answers/a-77527661>

LLMS, SCIENCE AND YOURSELVES

New perspective in the job market

~~what can this person do that we cannot?~~


What can this person do that a Claude Pro subscription can not?

Current job market background

more than 60% of executives plan to use AI to shoulder entry-level work

Tech jobs are getting scarce

129,488 tech employees laid off by 296 tech companies in 2026



<https://layoffs.fyi/>


Is AI a deskilling tool?

- If everybody has the same skills (prompting), everybody is replaceable!
- this historically has meant overall lower wages and worse work conditions



Blix & Glimmer (2025) Why we fear AI

Concept: automating ML research



Yamada et al. (2025) The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search

Concept: automating ML research

While the quality of its manuscripts currently aligns with that of an unmotivated undergraduate student rushing to meet a deadline... it achieves this at an unprecedented speed and cost efficiency—our analysis indicates that generating a full research paper costs only \$6 - \$15, with just 3.5 hours of human involvement. This is significantly faster than traditional human researchers.

Beel et al. (2025) Evaluating Sakana's AI Scientist: Bold Claims, Mixed Results, and a Promising Future?

we could save a lot of grant money!

- a PhD student in Denmark 'costs' about 267K Euros
- even very good students produce $\approx$ 2-3 good papers a year
- 🤔 maybe I should just budget in Claude next time?

Should LLMs replace PhD students?



What can you do that Claude can't?



1. 'Replacement' scenario ignores attribution

- if ML is the future, then data is labor
- the only reason Claude can do what it does is expert-written data!
- authors of useful data points should get citations & compensation
- that may change the costs calculus

2. Unlike Claude, you can tell the truth

- Claude produces text, not information
- Checking can only be based in reality, which it can't have
- Hence, do we actually want to automate large outputs that are super tedious to check? Or rethink the workflow so that AI produces only chunks that are easy to check?


3. 'Automating science' is an oxymoron!

- Science generally deals with the unknown, and automated processes generally are more accessible/less novel
- as long as we're publishing papers, we have to prove we made a contribution (both significant AND novel)
- adding more not-yet-automated work is a clear way to add value and significance
- so, new methods may *change* what work is done, rather than just increase production of the old kind of work

Example: DNA sequencing

- in 1990s-2000s there were many papers of the type 'we sequenced the genome of species X' (e.g. [mouse](#), [yeast](#)).
- Now sequencing costs \$15 and wouldn't by itself be considered a 'discovery'.
- A 2024 paper on [sparrows](#) reports findings from analysis of multiple sparrow sequences with different technologies

Example: recent protein design study is 70% wet lab work



- mice immunization experiments to confirm the predicted properties of proteins
- structural accuracy also examined with crystallization
- NB: AlphaFold isn't a co-author!

proteins with target properties designed via specialized models & manual selection

(peer est. 30% effort of total work)

Castro et al. (2025) Nature Chemical Biology - [Accurate single-domain scaffolding of three nonoverlapping protein epitopes using deep learning](#)

Thought experiment for LLM research

- neural architecture search (NAS) is useful, but its results are normally not publishable as ML 'discoveries'
- let's say we have a mega-NAS + infinite compute for producing any number of 'frontier' LLMs by search across possible prompts, parameters, data mixes, architecture variants etc for a target set of benchmarks
- would we consider the outputs of mega-NAS 'discoveries' rather than engineering - or would it just raise the bar for a novel ML contribution (e.g. expecting theoretical developments)?

HOW CAN YOU MAKE PROGRESS BEYOND THE STATUS QUO?

The student AI dilemma

You are in a bind — you can't offload your skill-building to AI, but you'll graduate into a market where employers will expect you to get work done with AI.

Naranayan A. AI as normal technology, Aug 27 2026



AI should upskill you, not deskill

the AI-natives I have in mind... have a sense of the steps that go into research but substitute prompts for “load-bearing” steps of research: literature review, synthesis, coding, analysis, writing, etc. They cannot do this work well (or at all) without the tools.

reality check: can you do all this well yourself?

What about offline interactions?

Various employers and professionals in HR and management positions shared that university leavers they encountered often struggled to speak on the phone or in meetings, take notes with a pen, relay messages precisely or complete written tasks without internet access.

reality check: does your speaking match your writing?

[‘I’ve £90k in student debt – for what?’ Graduates share their job-hunting woes amid the AI fallout | | The Guardian](#)

Lots of tech already interferes with your learning!

- 'attention economy' platforms got us addicted to quick dopamine from doomscrolling, approval on social media, youtube etc (Cf [Haidt J. The anxious generation](#))
- web by itself discourages in-depth reading (cf [N. Carr, The Shallows](#))
- phones, emails etc constantly interrupt us
- laptops allow to easily switch to anything else instead of your work

AI industry deliberately targets students

OpenAI, Anthropic target college students with latest education AI announcements

PUBLISHED THU, APR 3 2025-11:00 AM EDT | UPDATED THU, APR 3 2025-11:22 AM EDT




Hayden Field
@HAYDENFIELD

SHARE    

KEY POINTS

- Artificial intelligence rivals Anthropic and OpenAI both rolled out education services within 24 hours of each other.
- It's part of a wider push by the two companies to capture the education market through partnerships with universities and convert students into users before they graduate and enter the workforce.
- Anthropic debuted Claude for Education, while OpenAI made ChatGPT Plus free through May for college students in the U.S. and Canada.

RELATED



CNBC | OpenAI, Anthropic target college students with latest education AI announcements

The impact can be devastating


I literally can't even go 10 seconds without using Chat when I am doing my assignments. I hate what I have become because I know I am learning NOTHING, but I am too far behind now to get by without using it. I need help, my motivation is gone. I am a senior and I am going to graduate with no retained knowledge from my major. (student post on Reddit)

Chatbots deliver dopamine as you're trying to work!

AI will never tell you that your work is subpar, your thinking shoddy, your analysis naive. Instead, it will suggest “a polish”, a deeper edit, a sense check for grammar and accuracy. It will offer more ways to get involved and help – as with social media platforms, it wants users hooked and jonesing for their next fix.

[18 months. 12,000 questions. A whole lot of anxiety. What I learned from reading students' ChatGPT logs | The Guardian](#)

LLM use may degrade brain connectivity!



Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task; MIT Media Lab

LLM use may degrade brain connectivity!

50% LLM users indicated 'full ownership' of the essay, 16% 'no ownership'




Figure 6. Percentage of participants within each group who struggled to quote anything from their essays in Session 1.

[Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task; MIT Media Lab](#)

LLM use may degrade brain connectivity!




Figure 1. The dynamic Direct Transfer Function (dDTF) EEG analysis of Alpha Band for groups: LLM, Search Engine, Brain-only, including p-values to show significance from moderately significant (*) to highly significant (***).

LLM use may degrade brain connectivity!

cognitive debt: repeated reliance on external systems like LLMs replaces the effortful cognitive processes

Recommendation to **delay AI use until learners have engaged in sufficient self-driven cognitive effort.** This may promote both immediate tool efficacy and lasting cognitive autonomy

[Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task; MIT Media Lab](#)

Fun fact: LLM essay *felt* like it was better!




Figure 9. Reported percentage of satisfaction with the written essay by participants per group after Session 1.

[Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task; MIT Media Lab](#)

Core point of confusion: what's the point of school assignments?



Kevin Zatloukal May 8

...

These discussions about LLMs writing essays seem disconnected from reality.... Do the authors not realize that professors don't actually want these essays? They're not collecting them from students and then selling them on the street corner. No one wants to read them. In fact, we actually have to PAY the people who do read them (TAs) to compensate them from the value destruction that it causes.

Clearly, producing these essays is not the point of education. The reason we ask students to write them is that, for a human brain, the most efficient way to produce a good essay is to first understand the topic and then think critically about it. The classroom is a factory for producing human understanding and critical thinking skills. The essays, which have negative economic value, are actually a waste product.

The situation is as if someone saw a factory that was pumping waste into the river and decided to make a new factory that pumps waste into the river even faster and more cheaply, not noticing that the point of the factory was to produce bicycles not waste. If someone did that with a physical factory, they would get sued for polluting the river. It's a wonder that the same is not on the table here for those polluting the educational system.

♡ LIKE (38) 💬 REPLY (4)

↑ SHARE

[Substack comment by K.Zatloukal](#)

Summing up

- school assignments are mostly a waste byproduct
- the real product is you
- it makes no sense to have a robot to do exercise for you - neither physical nor mental
- ~~you have to do everything unassisted~~ you have to select which skills to develop

AI *could* assist learning

I use AI every day for learning. The key is to use it to increase, not decrease your cognitive load. To learn deeper, not faster. There is no learning without the cognitive sweat. I try to make sure I'm mentally exhausted at the end of the day.

Naranayan A. AI as normal technology, Aug 27 2026

Are you concerned about your AI use? What, if anything you'd like to change?



Wrapping up: what should someone hire you for?

- the world is now filled with infinite slop
- if all you can do is prompt, you are easily replaceable!
- skills in short supply overall: clarity, mental stamina, attention to details, ability to ensure precision
- science is about what isn't yet known/easy

Thank you!

 arog@itu.dk

Comments and collaborations welcome!

[@annarogers.bsky.social](https://bsky.social/@annarogers)

 <https://linkedin.com/in/annargrs/>

