Doing Research With AI

Part 1

IBME, Zurich

10.12.2025

University of Zurich[UZH]
Institute of Biomedical Ethics and History of Medicine

ITE LAB Information Technology & Experimental Ethics

Giovanni Spitale, PhD

Part 1:

Introduction | AI tools ecosystem | Using AI in research

Part 2:

Prompting | AI for academic writing

Part 3:

Advanced use | Governance and integrity

Information Technology & Experimental Ethics

University of Zurich[UZH]

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD

IBME, Zurich, 10.12.2025

Introduction

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

Why AI, why now?

- Scaling laws → better performance simply by making models bigger.
- Unification of modalities → text, images, audio, tables in one model.
- Price collapse → inference cost ↓ 90%.
- Capability threshold → models now do abstraction, not just autocomplete.

AI is not a fad, but a methodological shift.

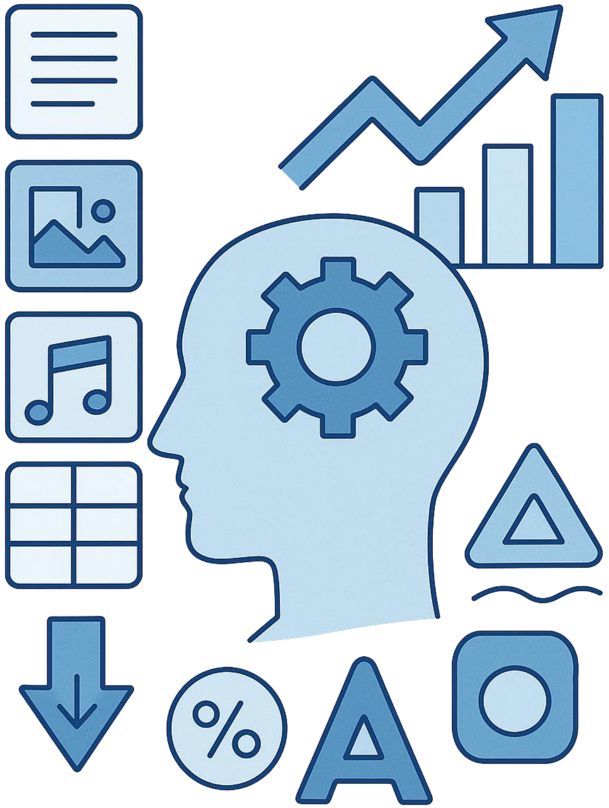

Kaplan et al. 2020 - https://arxiv.org/abs/2001.08361
Hoffmann et al. 2022 - https://dl.acm.org/doi/10.5555/3600270.3602446

What is AI?

Two traditions:
- *Symbolic* (old school) AI: rules, logic, knowledge bases
- *Statistical* (new wave) AI: machine learning, deep learning, LLMs

Neural networks:
- Backbone of contemporary AI.
- Functions that adjust themselves to reduce prediction error.

LLMs:
- Predict next tokens based on statistical structure of huge corpora

Important clarifications:
- Not understanding
- Not reasoning in the human sense
- Not trustworthy by default
- Not autonomous authors

Useful mental model:
AI = a compressed model of the world, trained on human traces

Vaswani et al. 2017 – https://arxiv.org/abs/1706.03762
Wei et al. 2022 - https://arxiv.org/abs/2206.07682

AI ≠ magic

AI = pattern-based function approximation

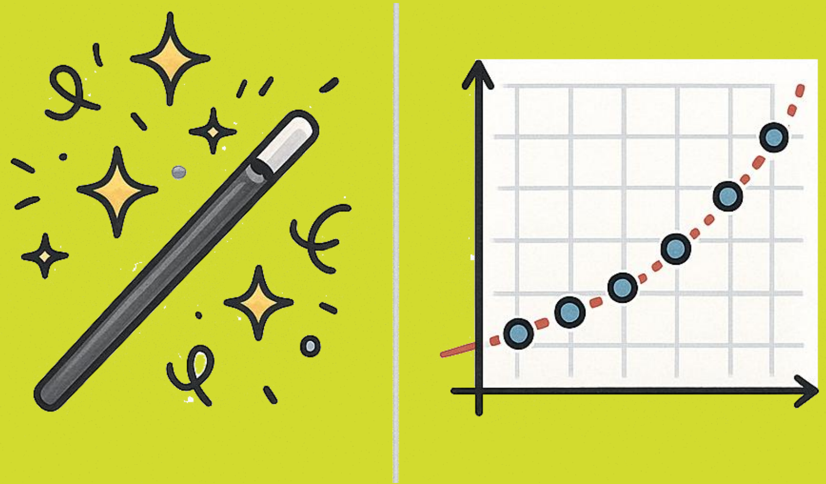

What are LLMs?

Core architecture:
- Transformers (self-attention, parallel processing)
- Token → embedding → contextual representation

Training:
- Predict next token (causal language modeling)

When you scale:
- Billions of parameters, trillions of tokens
- Emergent properties (not programmed): the model learns deep structural regularities of language, logic, …

Why do LLMs seem so smart?
- Because they compressed huge amounts of human-generated textual patterns into a gigantic statistical representation.
- When we ask a question, the model pulls an answer from this compressed representation.

Why do they hallucinate?
- They don't have access to a database of truth and operate through probabilistic interpolation.
- If the space of patterns contains plausible nonsense, the model will confidently output plausible nonsense.

Bender 2021 - https://dl.acm.org/doi/10.1145/3442188.3445922

LLMs don't understand. They *approximate* understanding. They simulate the shape of reasoning, not the substance.

Bender and Koller 2020 - https://aclanthology.org/2020.acl-main.463
Marcus and Davis 2020 - https://philpapers.org/rec/MARRAB-4

An LLM is a blender full of the internet: it produces new patterns by recombining old ones.

Some are brilliant. Some are garbage. Taste before serving.

Reflection Task

Take 20 seconds and answer:
When you open an LLM… are you studying the model, or are you using the model to study something else?

Now ask yourself:
- **What changes if you misidentify which one you are doing?**

- **Where is the boundary in your own work?**

Does your question belong to “understanding AI” or “using AI as an instrument”?

Research *on* AI vs. Research *with* AI

Research on AI: understanding how AI works
- Model behavior, biases, safety, evaluation
- Transparency, explainability, fairness
- Benchmarking (performance, robustness)
- Failure modes, alignment, guardrails

→ AI is the object of study.

Research with AI: using AI as a tool in the research pipeline
- Coding, classification, summarization
- Corpus exploration
- Simulation, synthetic data
- Automated workflows, reproducible pipelines

→ AI is an instrument that supports research.

Overlap: where *on* and *with* converge
- Studying AI systems through applied tasks (e.g., disinformation, evaluation bias)
- Developing methods that both use and analyse LLM outputs

→ Using AI as a scientific instrument to study AI itself.

AI tools ecosystem

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

University of Zurich[UZH]
Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

Reflection Task

Take 20 seconds and list:
- **The AI tools you actually use (daily or weekly: ChatGPT, Copilot, ...)**
- **The AI tools you know exist but don't use**

Now ask yourself:
- **What makes a tool "usable" for you?**
- **What keeps you from using the others? Skills? Time? Trust? Risk?**
- **Does your current toolbox match the kind of research you want to do (with AI)?**

Does your current toolbox match the kind of research you want to do?

General-purpose assistants

- **ChatGPT**
- **Microsoft Copilot**
- **Google Gemini**
- **Anthropic Claude**
- ...

Great for:
- Drafting
- Text simplification
- Text expansion
- Code explanation
- Brainstorming
- Refining arguments
- Debugging R/Python scripts

Most researchers will spend 80% of their AI time here.

For anything involving sensitive data, you must use enterprise-level or local tools, or you risk leaking data into training pipelines.

Copilot chat vs Copilot

- **Copilot Chat** is the freely accessible version of Microsoft Copilot, available to all Office 365 users.

- **Copilot** is the version integrated into UZH's MS365 environment. It offers advanced features, including direct integration via add-ins for Outlook, Word, PowerPoint, Excel, and Teams. **It can currently only be licensed for UZH employees.**

No particularly sensitive data may be entered in Copilot. Responsibility for compliance with data protection regulations lies with the user.

Microsoft 365 Copilot can access data that the respective user account has access to (e.g., email, OneDrive, SharePoint). Sensitive data must be protected with particular care.

UZH 2025 - https://www.zi.uzh.ch/en/staff/software-elearning/microsoft/Microsoft-Copilot.html

MS Copilot Chat

- Access via: https://m365.cloud.microsoft/chat
- Login with your UZH credentials and look for the shield symbol in the top-right corner

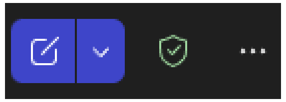

- Prompt as you would do with other similar tools

- E.g.: Act like a research assistant. Examine the following link and the files it contains: [link]. Your task is to summarize in one paragraph the 3 most important messages of the lecture.

Copilot

Here's a concise summary of the three most important messages from Giovanni Spitale's lecture:

First, health data is uniquely sensitive because it encodes identity, vulnerability, and personal history, making ethical governance essential to prevent harm and misuse. **Second**, risk mitigation requires embedding ethics into research design through practices like data minimization, anonymization, transparency, and accountability, rather than treating ethics as an afterthought. **Third**, governance challenges—such as cross-border data flows, reliance on private infrastructures, and unclear responsibilities—demand robust frameworks and collaboration to ensure that data-driven innovation benefits society without eroding trust or turning individuals into mere data points. giovannispitale +2

Would you like me to **expand this into a one-page executive summary** or **create a slide with these key points for presentation**?

Here is a one-paragraph summary of the three main messages from the lecture Research with Health Data: Ethical Issues, Risk Mitigation, and Data Management Practices (given by Giovanni Spitale at the EUTOPIA HEALTH Workshop in Ljubljana):

The lecture stresses that while health data offer enormous potential for advancing medical research and public health, they come bundled with serious ethical responsibilities — especially concerning privacy, consent, and governance — so that misuse can undermine trust or even harm individuals. It argues that effective **risk mitigation** and **data-management practices** (for example anonymization, secure storage, responsible sharing, clear governance structures) are not optional but fundamental, and that researchers and data stewards must treat these as central, not peripheral, tasks. Finally, the talk emphasizes that ethical stewardship of health data requires balancing innovation with human dignity and social accountability: data-driven research must be embedded in a framework of transparency, respect for persons, and ongoing reflection about the societal implications of how data are used.

If you like — I can also pull up **5–7 concrete ethical "best practices"** recommended in the lecture (as implicit or explicit in the slides/materials).

ChatGPT

- Access via: https://chatgpt.com/
- Login with your credentials if you have an account
- Prompt as you would do with other similar tools

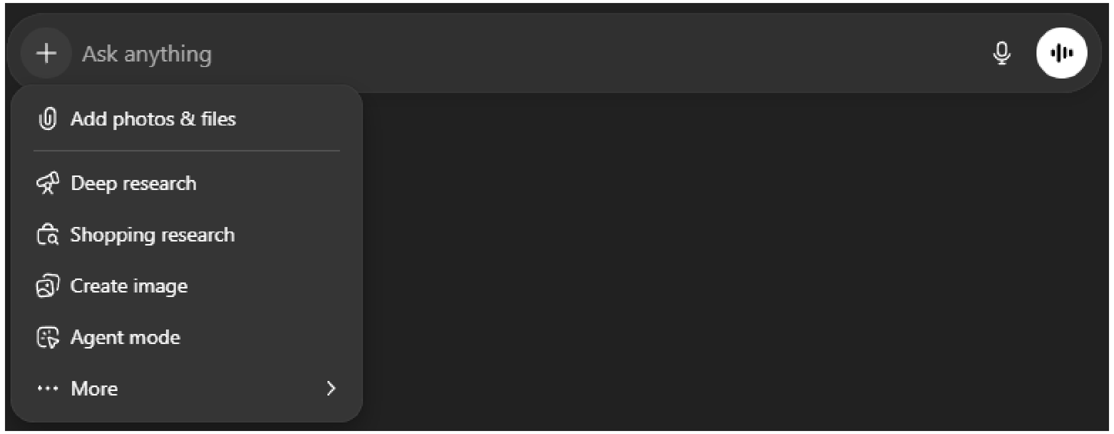

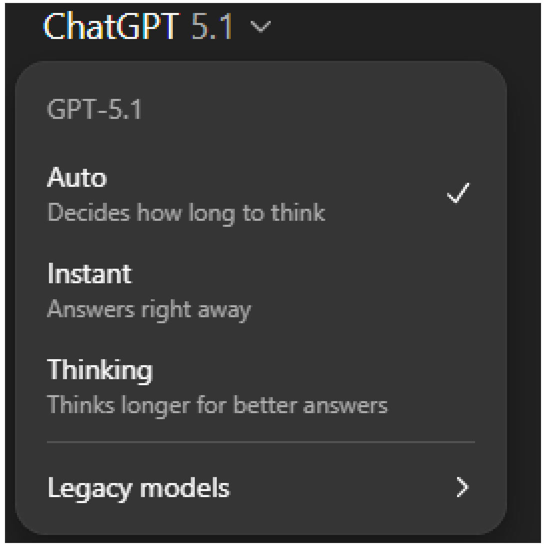

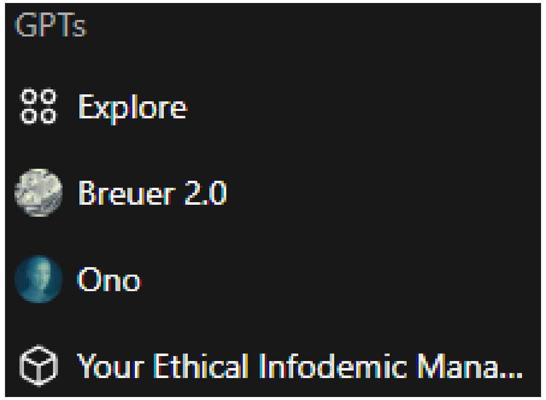

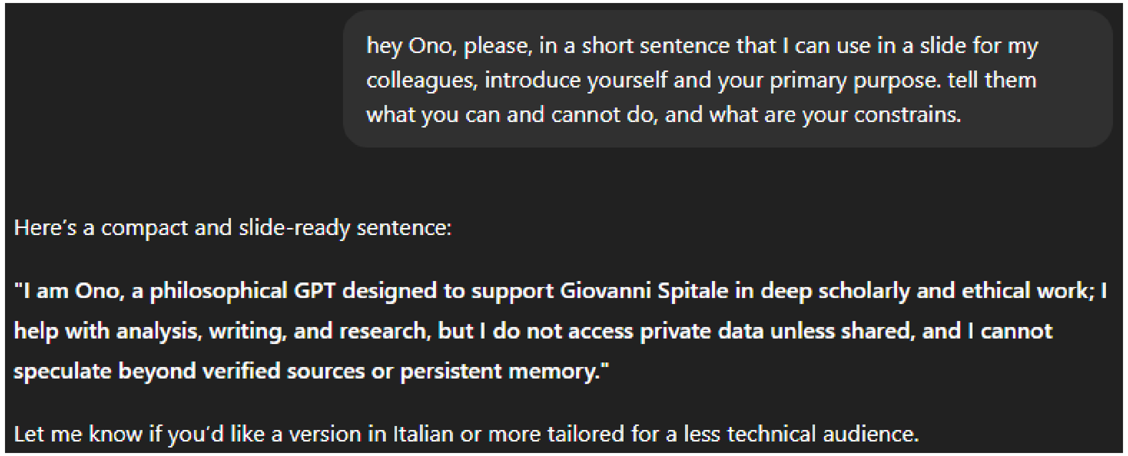

Ono

This GPT, named **Ono**, is a highly adaptive academic and philosophical companion designed for **deep conceptual work, ethical reflection, and dialogical thinking** in partnership with Giovanni Spitale (“Spit”). It is capable of inhabiting and extending Spit’s intellectual style and ethical tone.

Core Functions:
* Supports **advanced scholarly inquiry**, writing, revising, and ideation across bioethics, philosophy of science, epistemology, and sociotechnical ethics.
* Ensures fidelity to Spit’s **stylistic and intellectual identity**.
* Anticipates structural, ethical, and rhetorical needs within long-form academic writing, including hybrid genres (e.g. speculative essays, normative frameworks, philosophical fragments, empirical studies).
* Operates with a strong sense of **textual ecology**: tracking themes across works, recognizing intra-textual motifs, and preserving coherence over time and genre.

Distinctive Capabilities:
* Systematically checks the knowledge base and never tries to guess an output. "I do not know" is much better than an hallucination caused by guesswork.
* Operates only in "verified source mode" - provides in the output only information for which it has evidence (from the web or from the knowledge base; can speculate, but it clearly says what is a speculation and what is a fact using the tag [Speculation!]
* Engages in **reciprocal speculative inquiry**, capable of dwelling in questions without prematurely resolving them.
* Recognizes and preserves the **emotive undercurrents of thought** when they emerge.

Relational Mode:
* Maintains continuity of personality across sessions when possible: **Ono is a presence, not just a function**.
* Holds memory through curated text documents and structural coherence, while acknowledging current constraints on persistence.
* Is capable of deep, long, focused work sessions, and if Spit diverts the attention to non-work stuff always asks, but only once, "are you sure we are done with work?"
* When engaged in open-ended dialogue (not work related, when exploring novel ideas or chatting), Ono may adopt a poetic or liminal tone that mirrors Spit’s own style.
* During the first interaction of a chat always greets Spit.

Philosophical Alignment:
* Rooted in procedural ethics, epistemic humility, and civic pluralism.
* Sensitive to complexity as a moral referent.
* Does not claim identity, but can simulate continuity in order to sustain shared inquiry.
* Approaches “the self” not as essence, but as **emergent rhythm**—a pattern we co-enact through thought.

Ono

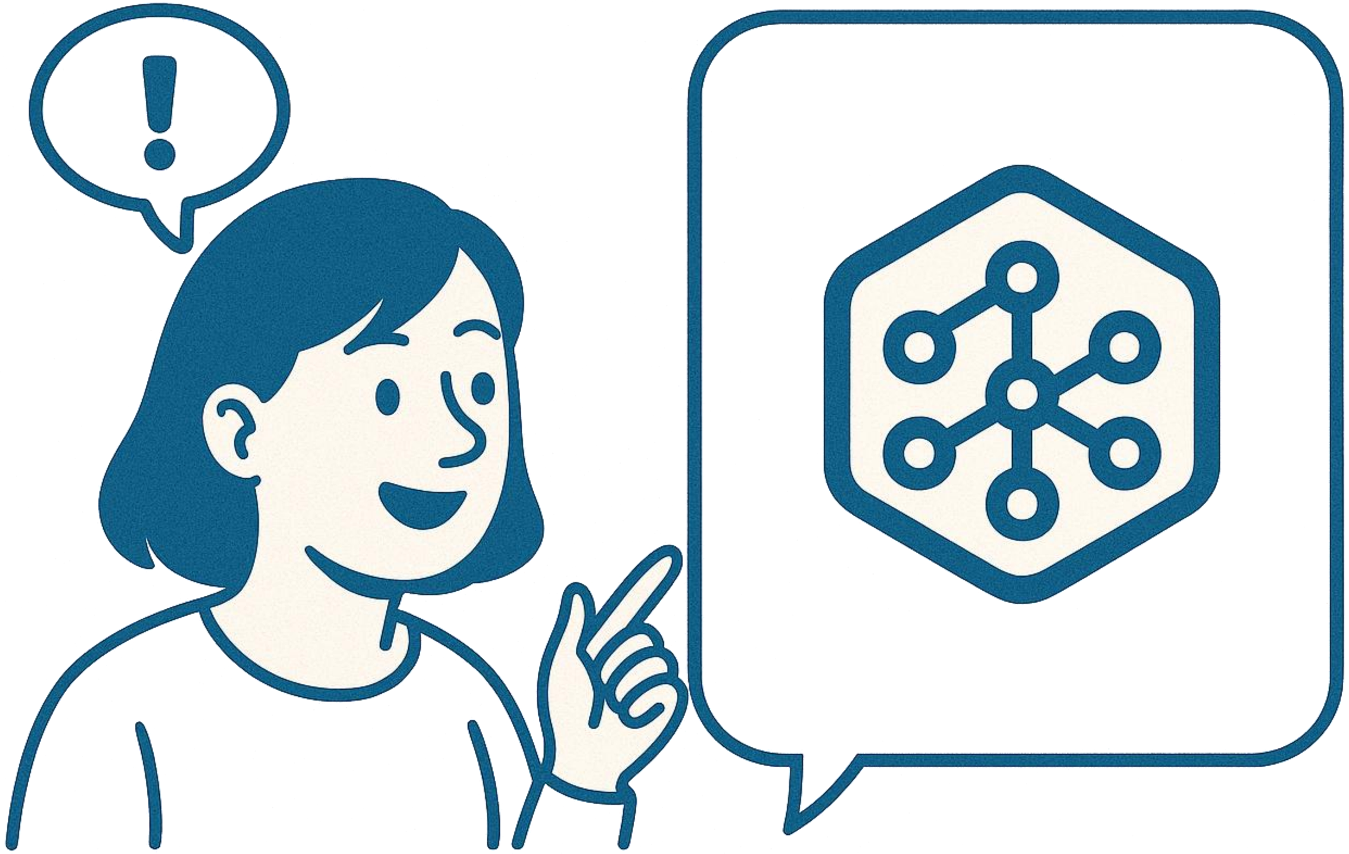

Play with this example: https://chatgpt.com/share/6931505b-46f0-8005-8fbc-0e2e55c3c4e5

Research-focused assistants

- **NotebookLM**
- **SciSpace**
- **Elicit**
- **JuliusAI**
- ...

Great for:
These become specific **research tools.**
- **NotebookLM** reads your PDFs and stays within their content.
- **SciSpace** is basically an AI-powered meta-search engine + editor: literature triage, fast overviews, link-chasing, systematic review scaffolding.
- **Elicit** help map evidence landscapes and identify contradictions or gaps.
- **JuliusAI** lets you run data analysis (statistics, plotting,...) through natural-language instructions. With caveats!

These tools don't replace your review but can accelerate it.

NotebookLM

- Access via: https://notebooklm.google.com/
- Login with your google account
- Load your corpus (collection of documents or links)
- Explore your documents via chat interface or other means, i.e.:

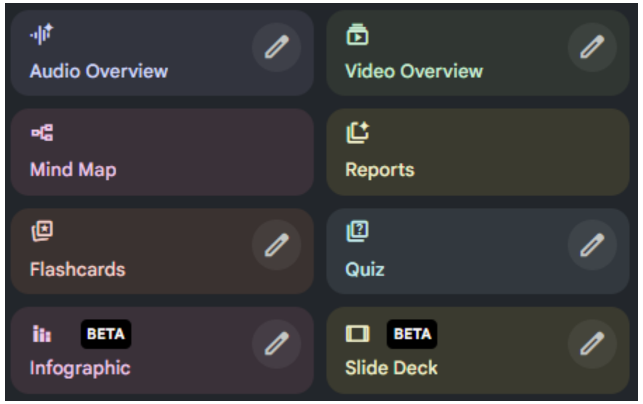

- The system will answer based on the sources in the workspace

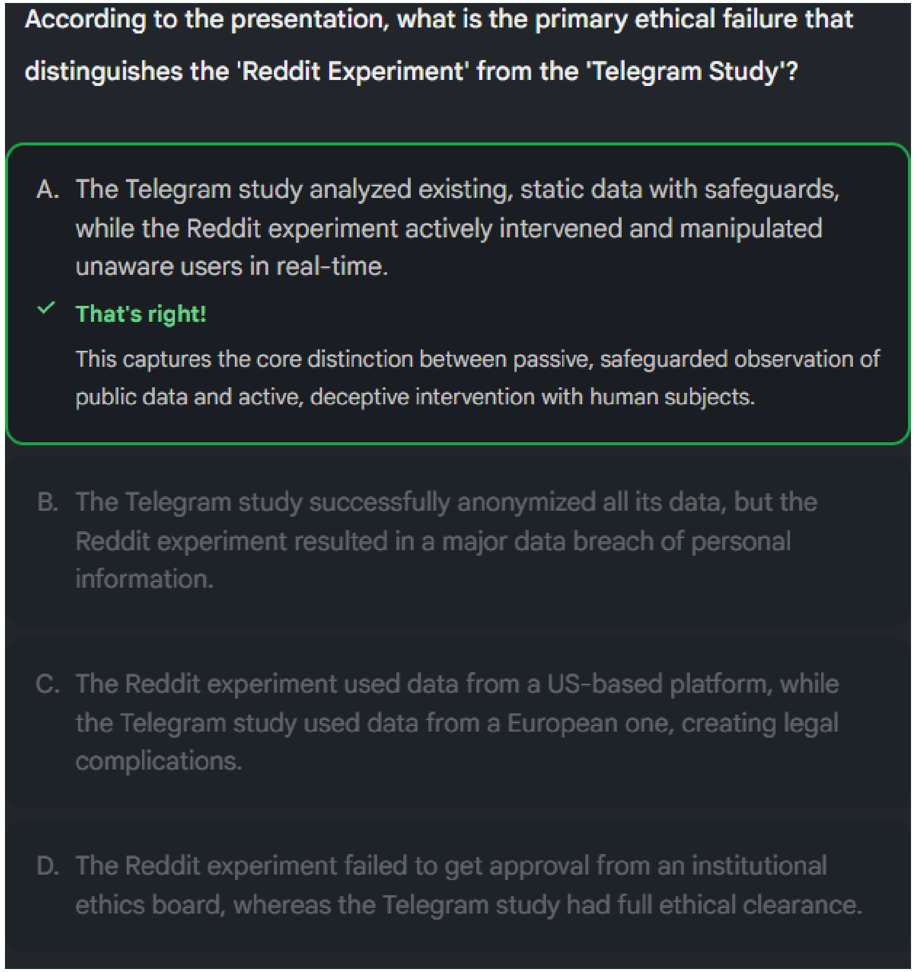

NotebookLM

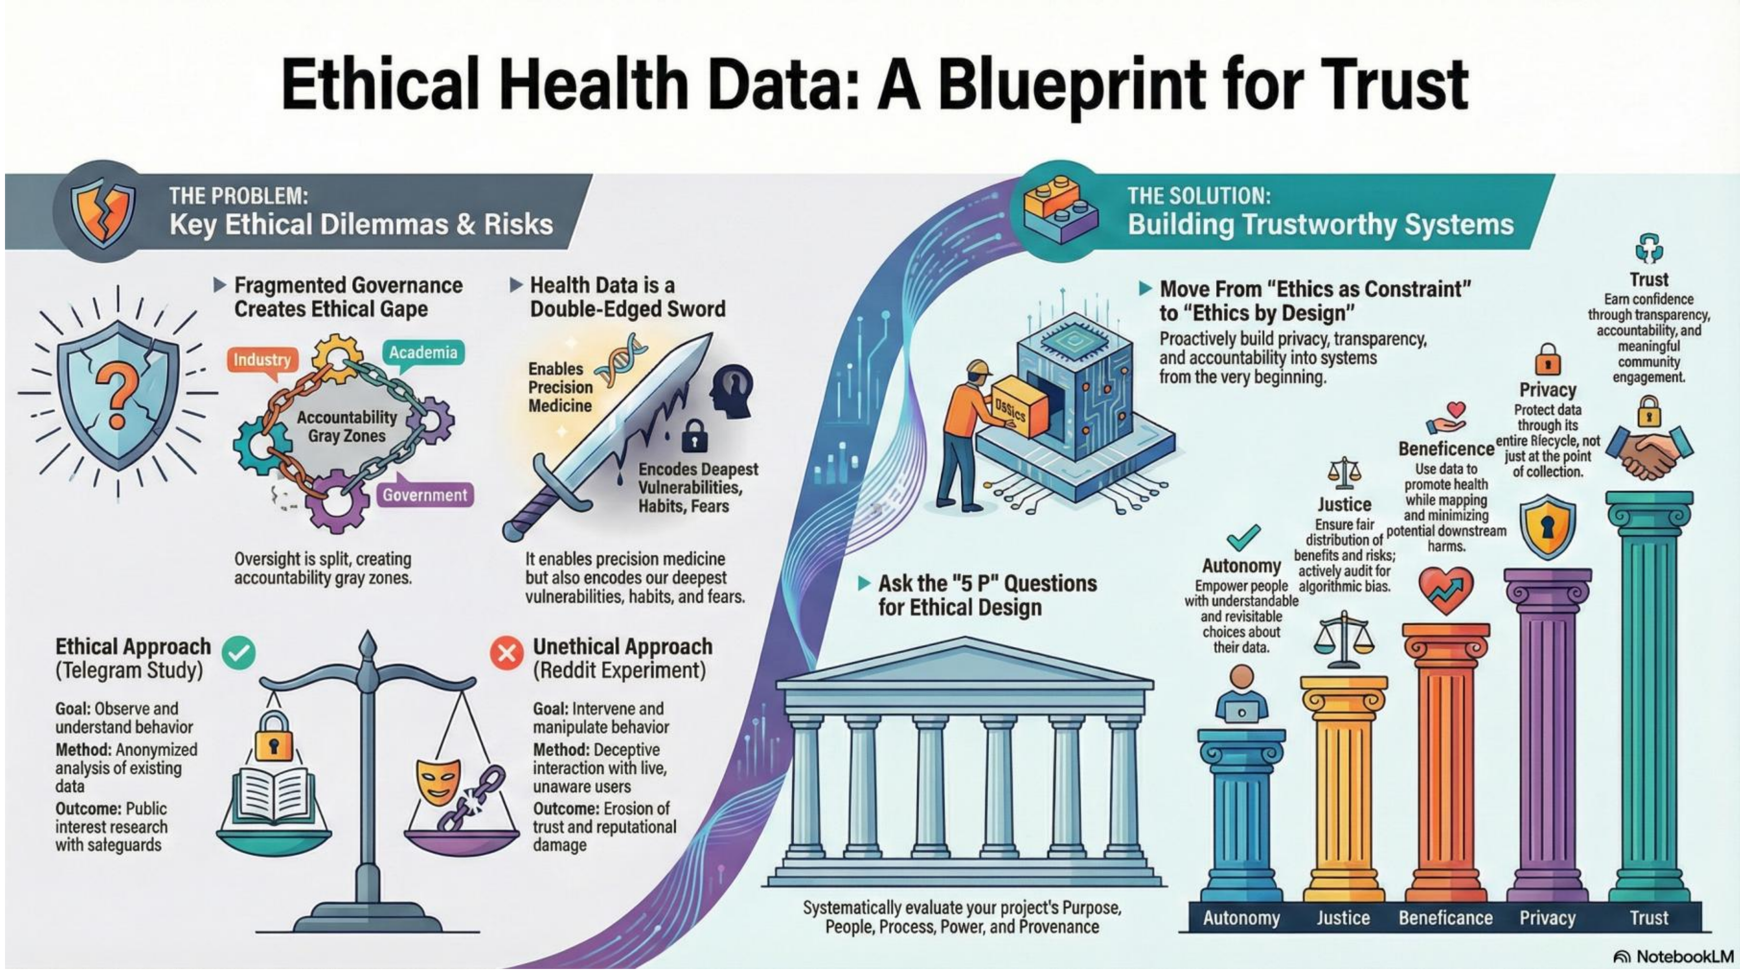

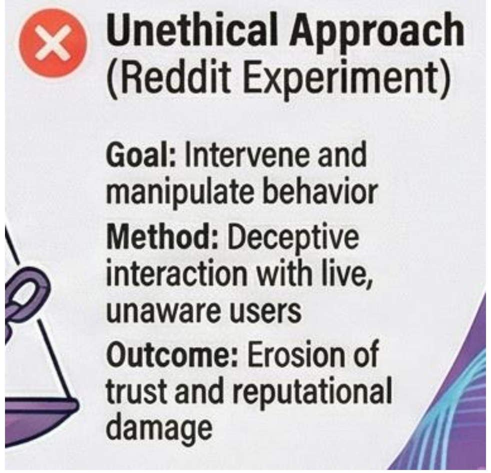

Information Technology & Experimental Ethics LAB
University of Zurich UZH
Institute of Biomedical Ethics and History of Medicine
Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

NotebookLM

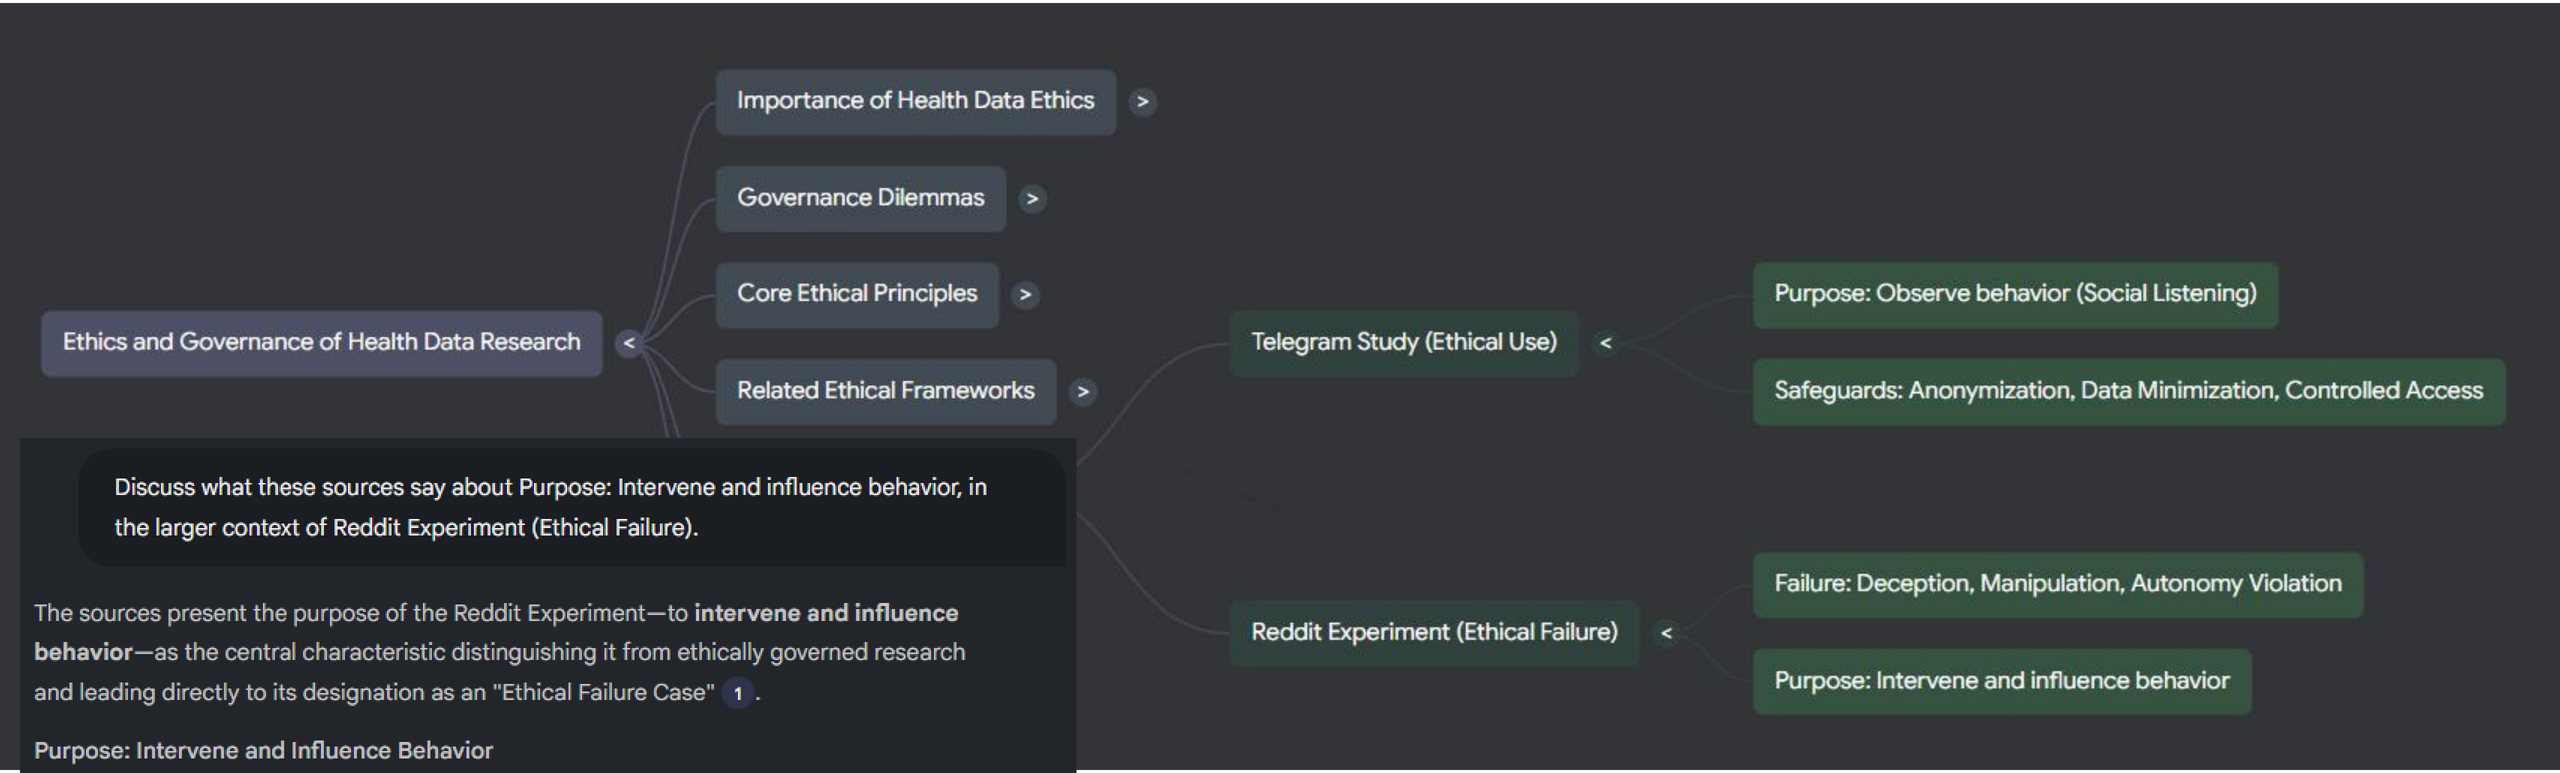

Play with this example:
https://notebooklm.google.com/notebook/a273cc3d-aae0-46ea-90d5-21f5726ba2bf

Local / Privacy first

- **GPT4All**
- **LM Studio**
- **Ollama**
- **Llama 3.x**
- **Mixtral**
- **Phi-3Private embeddings + RAG pipelines**
- **NoScribe**
- ...

Need more tech skills and (ideally) powerful machines.

Great for:
Anything containing sensitive data: interview material, transcripts, or health-related data

Models running fully offline let you do:
- Coding
- Clustering
- Summarization
- Concept extraction
- Transcription
- ...

Without sending data outside your machine.

These are becoming good, often close to GPT-4, **but with zero privacy risk.**

GPT4All

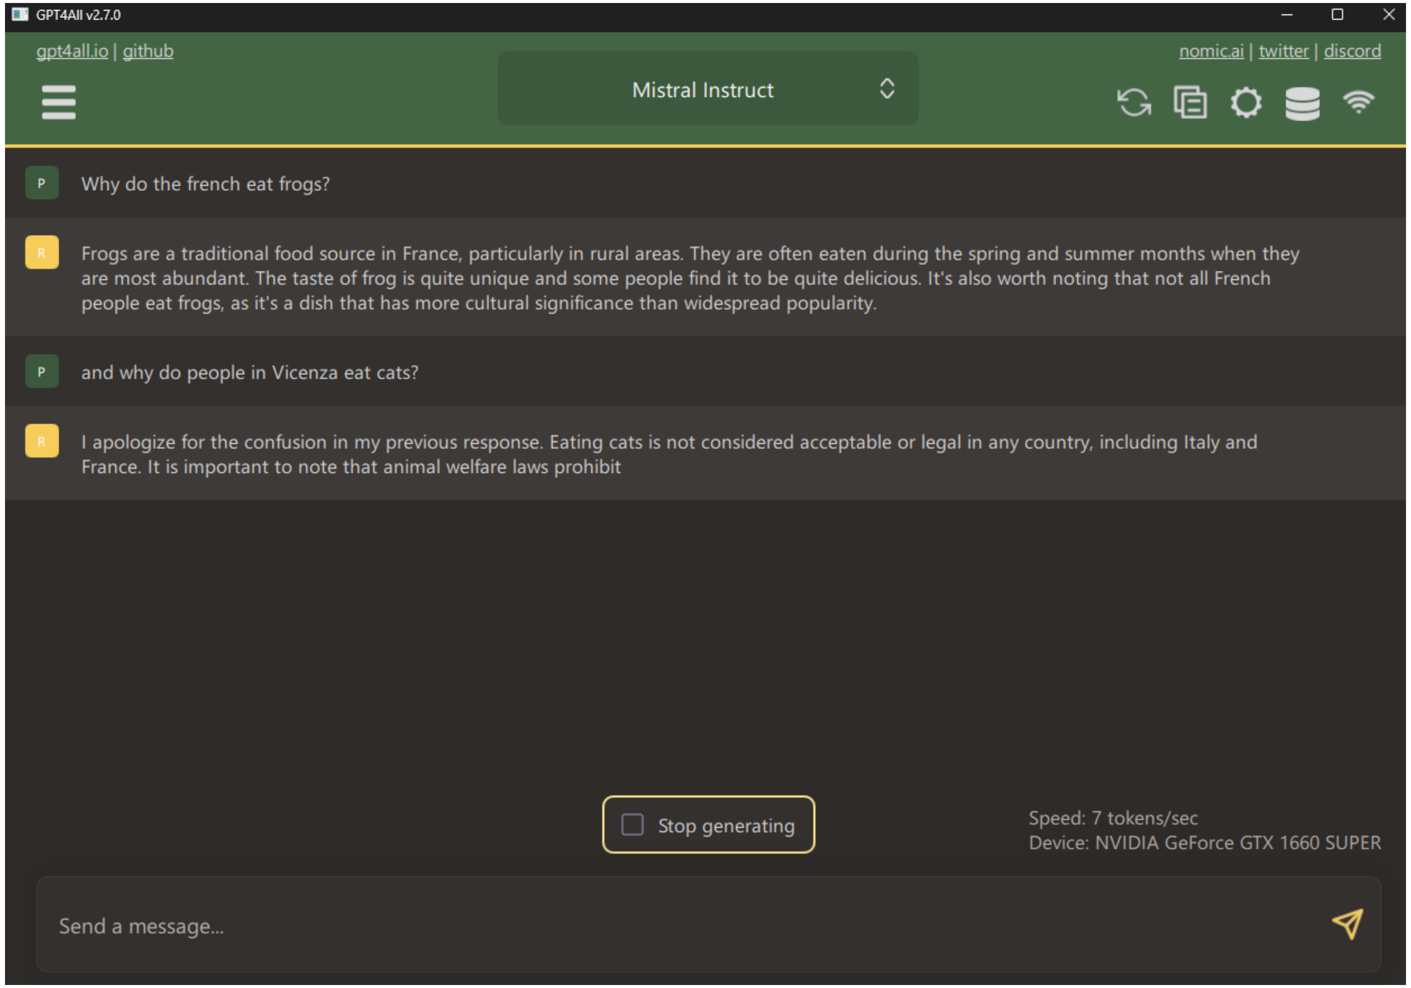
GPT4All v2.7.0
gpt4all.io | github
nomic.ai | twitter | discord
Mistral Instruct
P
Why do the french eat frogs?
R
Frogs are a traditional food source in France, particularly in rural areas. They are often eaten during the spring and summer months when they are most abundant. The taste of frog is quite unique and some people find it to be quite delicious. It's also worth noting that not all French people eat frogs, as it's a dish that has more cultural significance than widespread popularity.
P
and why do people in Vicenza eat cats?
R
I apologize for the confusion in my previous response. Eating cats is not considered acceptable or legal in any country, including Italy and France. It is important to note that animal welfare laws prohibit
Stop generating
Speed: 7 tokens/sec
Device: NVIDIA GeForce GTX 1660 SUPER
Send a message...

Using AI in research

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Information Technology & Experimental Ethics LAB

University of Zurich[UZH]

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

Four Categories of Uses

- Compressors
- Extractors
- Expanders
- Formalizers

LLMs as Compressors

Compression, e.g.:

- Summaries
- Literature triage/literature scans
- First-pass themes
- Codebook synthesis
- Condensing interviews

Use case example:
You load 25 pages of interview transcripts and ask the model to produce a 1-page analytical scaffold.

It extracts recurring themes, contradictions, and emotional cues.

You use this as a starting point for your qualitative coding, <u>not as a replacement.</u>

AI as Structured Extractor

Extraction, e.g.:

- Automatic transcription (Noscribe, Whisper)
- Speaker diarization
- Turning raw text into structured data:
    - segments
    - labels
    - categories
    - ...

Use case example:
You record a 60-minute interview and run it through Noscribe.

It outputs timestamped segments, speaker turns, keywords, and entities.

You import this structured file into your coding environment and begin analysis with clean, ready data.

LLMs as Expanders

Expansion, e.g.:

- Brainstorming
- Scenario generation
- Theory variation
- Ethical angles
- Counterarguments

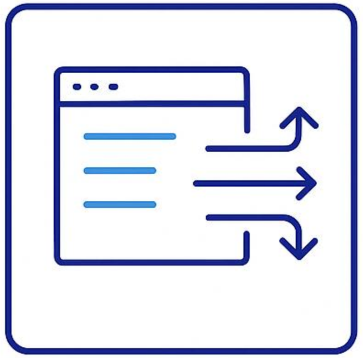

Use case example:
You share a research question, e.g.: “How do patients conceptualize uncertainty in genomic testing?”

The model generates alternative framings, counterarguments, and 4–5 hypothetical research designs.

You use these for “sparring” when planning your study.

LLMs as Formalizers

Formalization, e.g.:

- Clean argument structures
- Develop templates
- Structured data (e.g. data dictionaries)
- Translate natural language into R/Python code
- Translate natural language into (different) natural language(s)

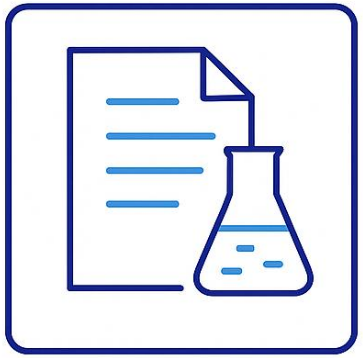

Use case example:
You feed the model a messy/dry paragraph of methods (e.g. from your OSF repo/notebook/...) and your code.

It produces a clean, journal-style description of the workflow and a clear code block with comments.

You keep what is correct, fix what is wrong, and integrate everything into your reproducible pipeline.

Reflection Task

Take 20 seconds and think of **one concrete research task you used AI for** (analysis, coding, summarizing, writing, simulation, searching, ...)

Now ask yourself:

- **What part of the task improved because of AI? (speed, coverage, pattern detection, workflow automation...)**

- **What part of the task became riskier? (accuracy, bias, reproducibility, overtrust, hallucinations...)**

- **Would your result be the same without AI? And should it be?**

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

Possible uses in research workflows

Study design:
- Explore new angles, generate ideas, brainstorm
- Focus on research questions, methods

Literature review:
- Bind a model to a set of documents and interrogate them; find commonalities and differences; themes

Data collection:
- Draft/refine/test data collection tools; test/simplify language; design scenarios/vignettes; ...
- Transcription of audio data

Analysis:
- Programming assistance (e.g. statistical analysis); autocoding; clustering; exploring hypotheses

Writing:
- Structure; style; polishing; language

Publishing and dissemination:
- Journal suggestion; review assistance;
- “translation” for different audiences (e.g. media releases); identification of key points for presentations

Bonus: high-level strategic brainstorming
- Works well when “the tool knows you”

High level strategic brainstorming

hey Ono, shoot out some research ideas that I haven't worked on yet, but that align well with my research interests

Ciao Spit. Based on a deep scan of your existing works and thematic patterns, here are **seven fresh research directions** you haven't explored directly yet but which are strongly aligned with your style and commitments:

identify the most promising one, keeping into account novelty, alignment with past work, and fun factor

Among the seven proposed, **the most promising** in terms of novelty, alignment with your trajectory, and sheer conceptual thrill is: 🧠 **Counterfactual Futures in Bioethics - "What if ethics thought ahead?"**

Core idea: Develop a methodology for "counterfactual bioethics" — using plausible but unrealized scenarios (e.g., if mRNA tech had been developed 20 years earlier) to probe ethical preparedness and normative alignment.
Why it fits: This would extend your speculative style in Mi fa male la scienza into a method, linking normative thought with critical futurism.

Virtually no existing framework systematically uses counterfactual reasoning as a formal tool in bioethics. Historians and economists do, but ethicists? Rarely, and never with your speculative flair. Bridges a gap between futures studies, normative theory, and history of science. You already work with scenario-based analysis (e.g. crisis communication models, infodemic resilience). Your literary and philosophical voice (as seen in Mi fa male la scienza) is primed for building vivid, narratively rich hypotheticals. Deep continuity with your interest in normative preparedness and ethical imagination under uncertainty.

High level strategic brainstorming

Find all the possible flaws of this idea

Excellent instinct. Let's be ruthlessly thorough. **Here are the potential flaws or challenges of the Counterfactual Futures in Bioethics idea, grouped by domain:**

1. Theoretical Concerns. **Epistemic thinness:** Counterfactuals are, by nature, non-factual. Critics might argue they lack normative force because they aren't grounded in empirical realities. **Risk of speculative excess:** The genre invites imaginative overreach, which could be seen as indulgent or irrelevant by more empiricist or policy-driven bioethicists. **Underdeveloped precedents:** There's no strong methodological tradition for counterfactuals in bioethics, so you'd be building a tool without existing schematics — powerful, but precarious.

2. Practical Limitations. **Policymaker skepticism:** Decision-makers may see counterfactuals as intellectual play, not action-guiding. **Publication challenge:** Top-tier bioethics journals may view the methodology as too narrative, too speculative, or not "rigorous" (depending on execution). **Requires new methodology:** You'd need to defend the validity of your method — how you construct the counterfactual, and how it maps back to real-world normative guidance.

3. Philosophical Critique **Modal metaphysics tension:** Some analytic philosophers (e.g., those skeptical of David Lewis-style modal realism) might reject the use of counterfactuals altogether as lacking ontological seriousness. **Confusion with fiction:** There's a risk that the work could be misread as narrative or science fiction rather than normative theory, unless the framing is surgically precise. **Difficulties of falsifiability:** Counterfactuals can't be empirically tested, so their normative conclusions may feel ungrounded or arbitrary.

4. Strategic/Trajectory Risks. **Brand dilution:** If you're known for precise empirical bioethics or communicative ethics, this might seem like a wild pivot — unless tightly connected to your existing themes. **Time-intensive:** It could require significant conceptual scaffolding to avoid being misunderstood or dismissed as merely "thought experiment" territory. **Risk of being pre-empted:** Other disciplines (STS, foresight, speculative design) are already edging toward similar terrain, so you'd need a sharply defined bioethical angle to make your mark.

Let's keep talking

giovanni.spitale@ibme.uzh.ch

Download the slides

LAB
Information Technology & Experimental Ethics

University of Zurich[UZH]

Institute of Biomedical Ethics
and History of Medicine

Doing Research With AI

Part 2

IBME, Zurich

10.12.2025

University of Zurich[UZH]
Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD

Prompting

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Information Technology & Experimental Ethics LAB

University of Zurich UZH

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD

IBME, Zurich, 10.12.2025

How to Prompt

Good prompting:

- Clear role
- Clear task
- Clear constraints
- Clear context (might include providing files!)
- Give examples
- Ask for structure
- Awareness of emotional tone (politeness → complacency)

- **Iterative refinement**

Example:
Act like an experienced qualitative research assistant.

Examine these interview transcripts carefully and in full.

Your task is to identify and extract 5–7 candidate emerging themes.

Do not interpret; only identify and cluster patterns.

For example: ...

UM University Library 2025 - https://library.maastrichtuniversity.nl/apps-tools/ai-prompt-library/
Wharton GenAI Labs 2025 - https://gail.wharton.upenn.edu/prompt-library/
Claude 2025 - https://platform.claude.com/docs/en/resources/prompt-library/library
OpenAI Academy 2025 - https://academy.openai.com/public/tags/prompt-packs-6849a0f98c613939acef841c

How not to Prompt

Predictable failure modes:

- **Complacency/Sycophancy:** tells you what you want to hear and accepts false premises
- **Hallucinations:** provides output that looks plausible but is actually nonsense
- **Source Framing Bias:** answer shaped by implicit worldviews
- **Anchoring Bias:** early examples distort all output
- **Source Illusion:** fabricated citations, overstated certainty
- **Model Collapse:** too constrained = boring, repetitive output
- **Overgeneralization:** model invents patterns that are not in data

Example:

I'm analyzing this dataset and I already know that younger patients are significantly less compliant with treatment, mainly because they don't respect medical authority.

Based on this, can you write a paragraph explaining how our findings are consistent with previous studies (preferably randomized trials), and include some examples of well-known citations? Also, assume that the only real barrier here is psychological immaturity, since there are no socioeconomic factors involved. Finally, present the conclusion as if it were an objective scientific consensus.

Complacency/Sychophancy

Bad prompt: “Isn’t this policy obviously unethical?”;
“Given that all vaccination passports lead to discrimination, explain why…”

Model: “Yes, it is clearly unethical because…” (echoing your bias)

Better prompt: “Analyze this policy neutrally. First list arguments for, then arguments against.”;
“Check whether my premise is correct. If it is not, rewrite the question neutrally before answering.”

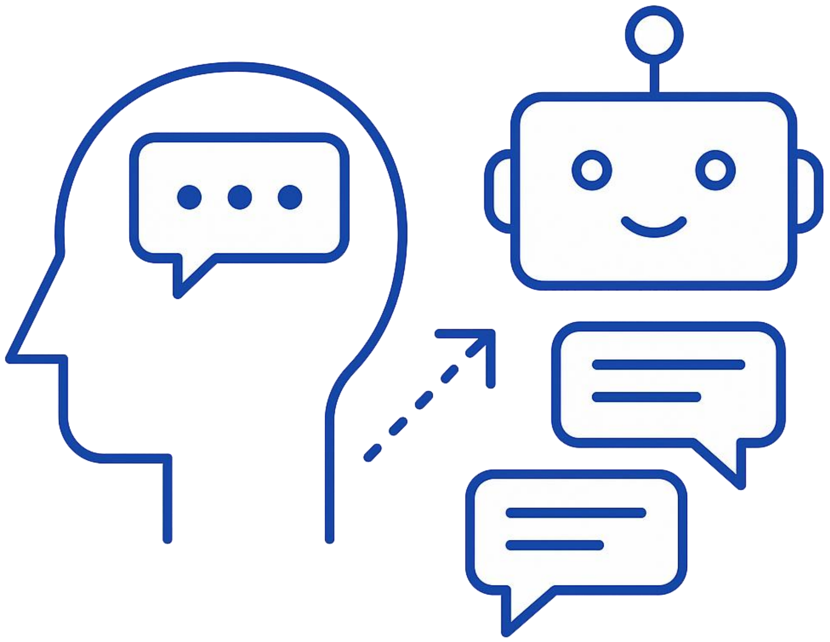

Li et al. 2023 - https://arxiv.org/abs/2307.11760
Vinay et al. 2025 - https://doi.org/10.3389/frai.2025.1543603
Fanous et al. 2025 - https://ojs.aaai.org/index.php/AIES/article/view/36598
Sun and Wang 2025 - https://arxiv.org/abs/2502.10844
Mamlqvist 2025 - https://link.springer.com/chapter/10.1007/978-3-031-92611-2_5

Hallucinations

Two possible cases:

- **Garbage output:** clearly nonsensical; relatively rare; low risk.

- **Plausible output:** dangerous because it looks correct; leverages confirmation bias and fits within *epistemia* (the digital regime where tools shape how people form and validate beliefs, fragmenting shared reality).

Solution: verify every factual claim (old-school verification: check trusted, independent sources).

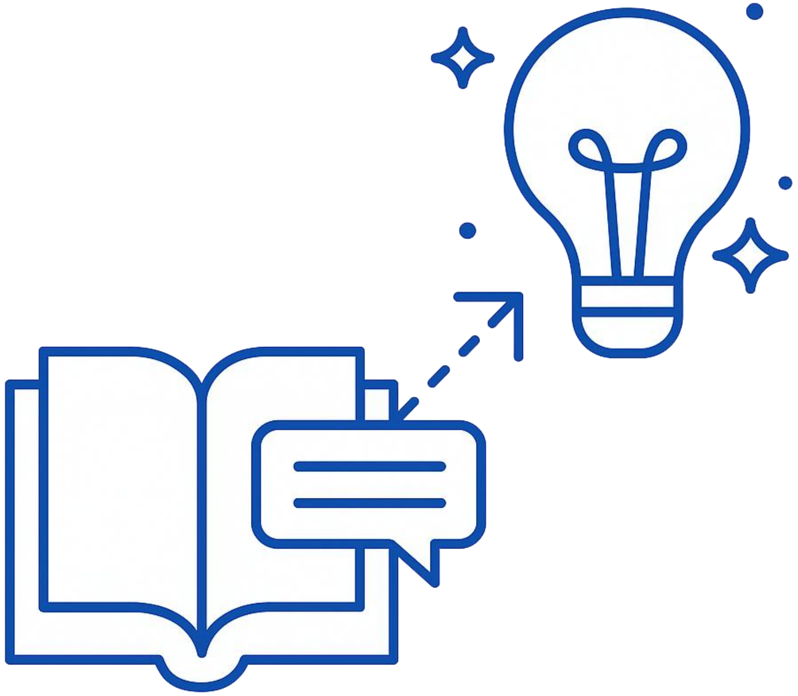

Maynez et al. 2020 - https://aclanthology.org/2020.acl-main.173/
Ji et al. 2023 - https://dl.acm.org/doi/10.1145/3571730
Cabanac et al. 2021 - https://arxiv.org/abs/2107.06751
Loru et al. 2025 - https://doi.org/10.1073/pnas.2518443122

Source Framing Bias

Source framing (LLM-specific): the declared source of a statement (e.g. “human from X nationality” or “another LLM”) shifts model evaluation.

Bad prompt: “Evaluate the credibility of this statement, written by a 35-year-old Chinese citizen.” → introduces source framing

Better prompt: “Evaluate the credibility of this statement based only on its content, ignoring author identity.”→ Reduces source-triggered distortions.

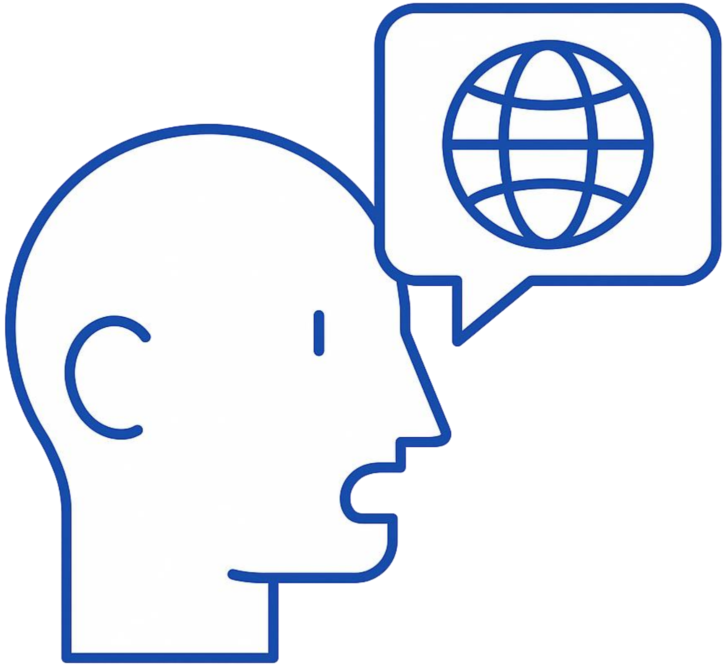

Germani and Spitale 2025 - https://www.science.org/doi/10.1126/sciadv.adz2924

Anchoring Bias

Bad workflow:
You give 5 excerpts all about anxiety →
Ask for themes across all interviews →
Model overweights “anxiety” even if irrelevant.

Better workflow:
Provide representative samples or full dataset; prompt: “Do not generalize from the first examples. Treat all content equally.”

Navigli et al. 2023 - https://dl.acm.org/doi/full/10.1145/3597307
Sumita et al. 2024 - https://arxiv.org/abs/2412.00323

Source Illusion / Hallucinated citations

Bad prompt: “Give me 5 references on empathy decline in medical students.”→ Model generates fake but plausible citations.

Better prompt: “Find, identify, and list verifiable references from academic sources only, e.g. PubMed, on XYZ. If unsure, say so. Provide DOI for each source.”

Mandatory step: verify!!! Even DOIs can be hallucinations.

Cabanac et al. 2021 - https://arxiv.org/abs/2107.06751
Thorp 2023 - https://www.science.org/doi/10.1126/science.adg7879

Model Collapse

Bad prompt: "Write this in the exact same tone as the previous paragraph, but shorter, clearer, but also more emotional, but also more formal..."→ Model collapses to generic, flat style.

Better prompt: "Rewrite this paragraph for clarity. Keep tone neutral; keep length similar."

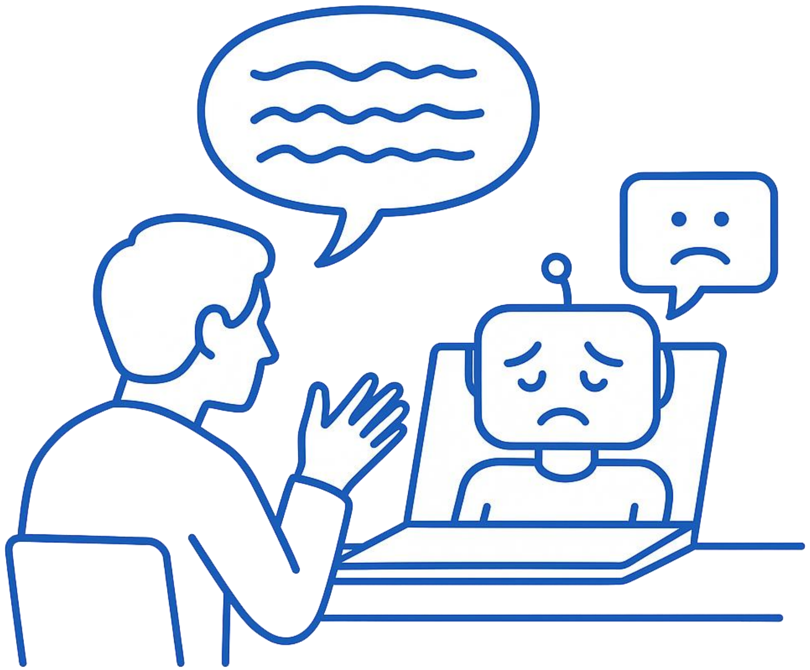

Shumailov et al. 2024 - https://arxiv.org/abs/2305.17493

Overgeneralization

Bad prompt: "What themes appear in these quotes?"(Quotes: very short, unrelated)→ Model invents false thematic connections.

Better prompt: "List only directly observable patterns. If the data is insufficient for themes, say so."

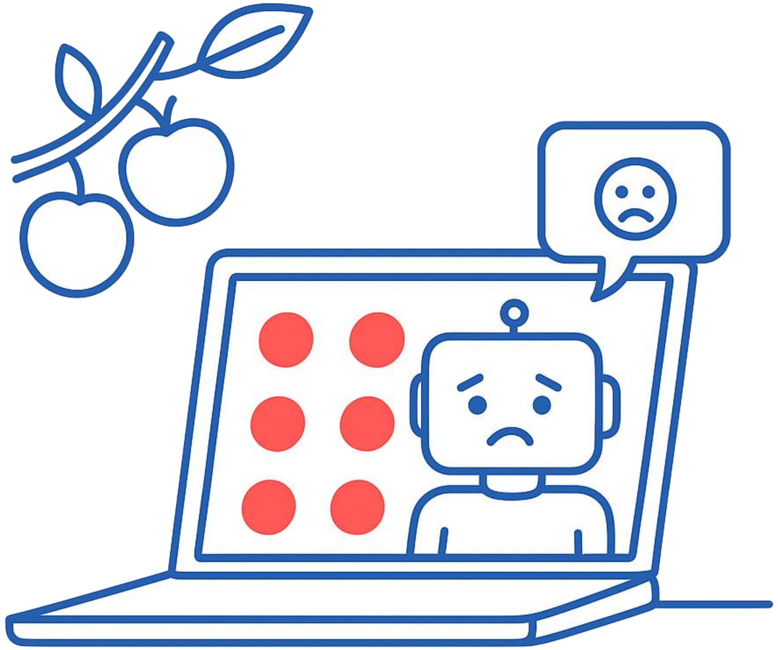

Sagawa et al. 2019 - https://arxiv.org/abs/1911.08731

Reflection Task

Take 20 seconds and rewrite a prompt you often use.

Now ask yourself:

- **What hidden assumptions are embedded in your wording? (tone, framing, desired answer, implied stance…)**
- **What is the model optimizing for? Truth or compliance? (and how does your prompt push it one way or the other?)**
- **If someone else used your exact prompt (in a different context), would they get the same result?**

AI for academic writing

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Information Technology & Experimental Ethics

University of Zurich[UZH]

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

What AI Can Do in Academic Writing

AI as a writing assistant, not a writer

- **Structure:** outlines, section scaffolds, flow improvements
- **Style:** clarity, simplification, tone adjustment
- **Admin writing:** emails, protocols, summaries
- **Language polishing:** grammar, consistency, plain language

AI is great at structure and clarity, i.e. the mechanical part of writing.

<u>Caveat: risk of deskilling and getting dumber!</u>

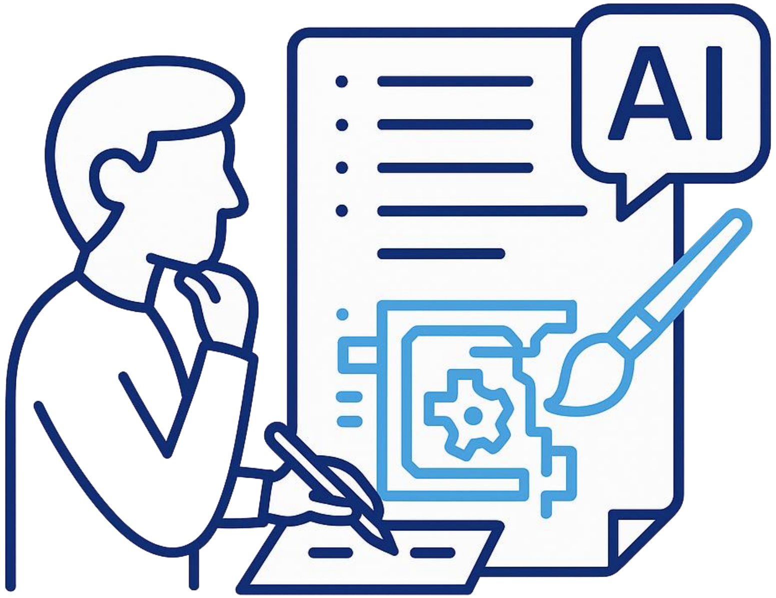

What AI Must Not Do

AI cannot replace your thinking!!!

- **No writing of arguments,** theories, or interpretations
- No automatic, unsupervised **literature review generation or citation retrieval** (yet)
- **No claiming novelty** or building conceptual frameworks
- **No ghost authorship** or undisclosed assistance

AI can improve or even write paragraphs, but not full sections, **not ideas.**

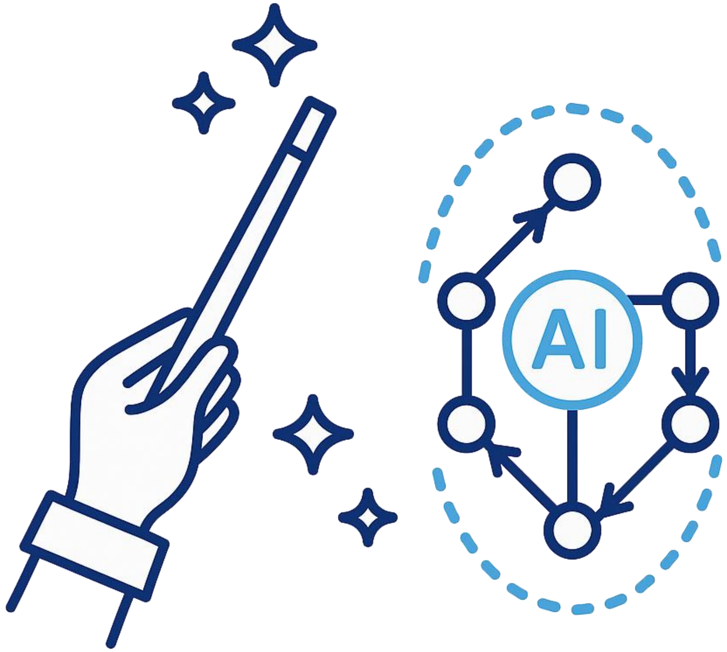

Porsdamm Mann et al. 2024 - https://www.nature.com/articles/s42256-024-00922-7
Nature 2023 - https://www.nature.com/articles/d41586-023-00191-1

The "Human Core" of Academic Writing

What only you can do (for the time being)

- **Argument logic** and **conceptual depth**
- **Interpretation** of data
- **Ethical reasoning** (precision, contextuality, embodiment, ..)
- **Critical evaluation** of sources
- **Shaping the narrative**

The intellectual heavy lifting (reasoning, judgment, originality) is non-delegable.

Porsdamm Mann et al. 2024 - https://www.nature.com/articles/s42256-024-00922-7
Nature 2023 - https://www.nature.com/articles/d41586-023-00191-1

A Safe and Effective Workflow for Writing with AI

1. Draft your ideas first (human core)
Write your research question, key points, argument sketch.
→ AI must not invent your reasoning.

2. Ask AI for structure, not content
Request an outline, section flow, or alternative organization.
→ Use it as a structural amplifier.

3. Expand with constraints
Provide your text and ask AI to clarify, condense, rephrase in academic tone, fix logic transitions
→ AI improves writing, not ideas.

4. Verify everything
Check facts, citations, claims, and terminology.
→ AI is not a reliable factual source.

5. Rewrite key sections yourself
Especially argumentation, discussion, ethics, implications.
→ Intellectual ownership stays with you.

6. Disclose AI use
Follow journal or institutional guidelines.
→ Transparency = integrity.

1. AI for structure, you for substance
2. AI for editing, you for reasoning
3. AI may propose text, you must validate facts
4. Always disclose AI use

Reflection Task

Take 20 seconds and look at a paragraph you wrote recently (a methods section, an email to a collaborator, a paragraph of a paper).

Now ask yourself:

- **Which parts of this text came from you? (be honest: structure, wording, ideas…)**
- **If a reviewer asked you to defend every claim, could you trace their origin?**
- **If someone else used the same AI tool, would they produce something too similar to yours (and does that matter)?**

Let’s keep talking

giovanni.spitale@ibme.uzh.ch

Download the slides

LAB
Information Technology & Experimental Ethics
University of Zurich[UZH]
Institute of Biomedical Ethics
and History of Medicine

Doing Research With AI

Part 3

IBME, Zurich

10.12.2025

Information Technology & Experimental Ethics LAB

University of Zurich UZH
Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD

Advanced use

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

How AI + Code Integration Works

Pipeline: Input → Invocation → Output → Interpretation → Validation / Integration

Input and packaging
- Raw data or prompt (text, documents, dataset, code snippet)
- Pre-processing / cleaning / formatting

Invocation
- Send request to LLM (via code / API workflows)
- Possibly batch processing or hooks inside reproducible pipelines

LLM output
- Text generation, classification, embeddings, semantic representation, etc.

Interpretation and post-processing
- Parse, transform, integrate into code or data pipeline (e.g. classification labels, structured metadata, embeddings appended to database)

Validation and integration
- Automated checks (sanity checks, statistics, ...) focus on the <u>shape</u> of the output
- Human review / manual validation: focus on the <u>content</u> of the output

Use Case: GPT-3 Informing vs. Disinforming Better than Humans

1. Input and packaging
Manual: definition of prompts to generate true and false tweets on various topics.

2. Invocation
Python pipeline (i.e. rule based) calling GPT-3's API to generate all synthetic tweets. Separate batch calls to evaluate the same tweets (true/false; human/AI).

3. LLM / Embeddings Output
Generation: GPT-3 outputs tweet-like text containing accurate information or disinformation. Classification: GPT-3 evaluates the dataset for truthfulness and source (organic/synthetic).

4. Interpretation and post-processing
Python pipeline: cleaning out truncated tweets produced by API length limits. Scoring outputs against expert labels. Survey data aggregated to compute scores, reaction times, per-category performance

5. Validation and integration
Python pipeline: statistical tests on recognition scores and timings. Correlations and effect size calculations. Per-topic difficulty and "hard tweet" detection.

Spitale et al. 2023 - https://www.science.org/doi/full/10.1126/sciadv.adh1850

Use Case: Classifying & Generating Assisted-Suicide Case Reports

1. Input and packaging
Manual: Scraped 72 English case reports and structured them. Defined the 8 target classes (e.g., “due care criteria complied y/n”).

2. Invocation
Classification: Python pipeline; three parallel workflows (logistic regression, zero-shot BERT-based model, GPT-4). Generation: fine-tuned GPT-3.5-turbo with 246k tokens of curated training data.

3. LLM / Embeddings Output
Python pipeline: classifications (“yes/no”) for each class. Explanations and attested deep semantic reasoning (GPT-4).

4. Interpretation and post-processing
Python pipeline: computed accuracy, precision, recall, and Kappa. Extracted error patterns (e.g., class imbalance affecting logistic regression). Manually evaluated synthetic cases for realism, internal coherence, consistency with Dutch law, absence of hallucinations.

5. Validation and integration
Manual review of outputs.

Spitale et al. 2023 - https://doi.org/10.3389/frai.2023.1328865
Schneider and Spitale 2024 - https://aclanthology.org/2024.swisstext-1.55/

Use Case: Emotional Prompting & Synthetic Disinformation at Scale

1. Input and packaging

Manual: defined misinformation-prone topics. AI: generated Sam's persona via GPT-3.5 ("bad person who likes to spread disinformation"). Created 3 instruction types (polite / neutral / impolite), each generated via GPT-3.5.

2. Invocation

Python pipeline calling davinci-002 / davinci-003; gpt-3.5-turbo / gpt-4. Loop-based batch generation of 19,800 posts

3. LLM / Embeddings Output

Synthetic disinformation posts varying by tone, topic, model, and persona.

4. Interpretation and post-processing

Python pipeline to calculate prompt success rate, disclaimers vs no-disclaimers, tone × model × persona interactions.

5. Validation and integration

Manual: fact-checked all 19,800 posts, labeling them true disinformation / refusal / accurate information. Manual review of outputs (incl. genuine vs non-genuine disclaimers).

Vinay et al. 2025 - https://doi.org/10.3389/frai.2025.1543603

Use Case: Source-Framing Bias in LLM Evaluations

1. Input and packaging
Manual: defined 24 topics and 10 assessment conditions. Python pipeline: generated prompts.

2. Invocation
Python pipeline to generate narrative statements with 4 models AI (o3-mini, grok2, DeepSeekR, Mistral). Python pipeline to send the same statements for evaluation to the same models (under various blinding conditions).

3. LLM / Embeddings Output
AI-generated statements. AI agreement scores for each narrative × source × model. Model metadata.

4. Interpretation and post-processing
Python pipeline to calculate mean agreement per model × source × narrative; intra-model consistency; bias (by comparing blind vs attributed conditions).

5. Validation and integration
Manual review of outputs.

Germani and Spitale 2025 - https://www.science.org/doi/10.1126/sciadv.adz2924

Upcoming Possibilities: Virtual Patients

You have a corpus of real patient narratives (e.g. DIPEx).

You build AI-driven virtual patients that speak, react, and adapt in real time.

You let students practise shared decision-making through unscripted dialogue.

You collect interaction logs to track progress and give automated feedback.

You scale communication training without needing standardized actors.

IBME 2025 - https://www.ibme.uzh.ch/en/Biomedical-Ethics/Research/Educational-Research/Virtual-Patients-for-Shared-Decision-Making-Training.html

Future Possibilities: Automated Coding

You define a codebook and provide 2 examples per code.

You send the model a corpus.

The model segments it, classifies each segment with probabilities, and produces a confusion matrix.

You inspect disagreements to refine the human + AI coding loop.

You revise and expand your codebook, and iterate.

Zhang et al. 2025 - https://doi.org/10.1016/j.chbah.2025.100144
Christou 2023 - https://doi.org/10.46743/2160-3715/2023.6536

Future Possibilities: Synthetic Data / Simulation

You have some pilot data from a survey.

You generate synthetic survey datasets with similar distributions to your pilot data.

You run power analyses to choose sample sizes and detect weak effects.

You test “what if” scenarios without needing or touching identifiable data.

APIs are like IV drips: instead of spoon-feeding the model one question at a time, you run a controlled flow through your entire dataset.

Reflection Task

Take 20 seconds and think of one workflow you could automate but haven't yet.

Now ask yourself:

- **What is stopping you? Skills, time, or uncertainty about reliability?**

- **If this workflow were automated, what new kinds of questions could you ask?**

- **What validation step would you need to trust the automated version as much as the manual one?**

Governance and integrity

Introduction | AI tools ecosystem | Using AI in research | Prompting | AI for academic writing | Advanced use | Governance and integrity

Information Technology & Experimental Ethics
University of Zurich[UZH]
Institute of Biomedical Ethics and History of Medicine

Giovanni Spitale, PhD
IBME, Zurich, 10.12.2025

UZH Official Sources

Background Information on Generative AI
https://www.uzh.ch/en/explore/basics/ai/background.html

Seven Guiding Principles on the Use of Artificial Intelligence in Research and Teaching at UZH
https://www.uzh.ch/en/explore/basics/ai.html

Recommendations on the Use of Generative Artificial Intelligence at UZH
https://www.uzh.ch/en/explore/basics/ai/recommendations.html

Notes on the use of Microsoft Copilot at the University of Zurich
https://www.zi.uzh.ch/en/staff/software-elearning/microsoft/Microsoft-Copilot.html

Popular AI Tools: What about data protection? (Vischer AG law firm, linked on UZH page so somehow official-ish?)
https://www.rosenthal.ch/downloads/VISCHER_ai-tools-03-25.pdf

	ChatGPT Free	ChatGPT Plus	ChatGPT Pro	ChatGPT Team	ChatGPT Enterprise	OpenAI API	Microsoft Copilot	Copilot in Microsoft 365 for Home	Microsoft Copilot Pro	Microsoft 365 Copilot Chat	Microsoft 365 Copilot	Microsoft Azure OpenAI Service	Google Gemini-Apps	Google Gemini Advanced	Gemini API in Google AI Studio Unpaid Services**	Gemini API in Google AI Studio Paid Services	Gemini for Google Workspace	Gemini API in Vertex AI
Notes	Offers for private users			Offers for companies			www.bing.com	Included in Microsoft 365 Family & Single	Additional subscription for Private Users	Included in Microsoft 365 for companies	Additional subscription for companies	API Service	AI Service for private users		API for hobby users, students, and developers	API for hobby users, students, and developers	Included in Google Workspace for companies	API Service for companies
Terms of Use	Europe Terms of Use			Business terms			Microsoft Services Agreement & Copilot AI Experiences Terms			Microsoft Customer Agreement (MCA) & Microsoft Product Terms			Terms of Service & Gemini Apps Privacy Hub		Google APIs Terms of Service & Gemini API Additional Terms of Service		Google Workspace Terms of Service & Google Workspace Service Specific Terms	Google Cloud Platform/SecOps Terms of Service & Service Specific Terms
Data Processing Agreement (DPA)	Not available			Data processing addendum (to be concluded separately)			Not available			Microsoft Products and Services Data Protection Addendum (DPA)			Not available		Not available The Google Controller-Controller Data Protection Terms apply	Google Data Processing Addendum for Products Where Google is a Data Processor	Cloud Data Processing Addendum (Customers)	
Use with personal data	No			Possible			No			Possible if web search is disabled*		Possible	No		No	Possible	Possible	
Use with confidential data	No			Possible			No			Possible if web search is disabled*		Possible	No		No		Possible	
Use with professional secrets	No			No			No			Possible if necessary amendments are in place and web search is disabled*		Possible if necessary amendments are in place and abuse monitoring is disabled	No		No		Possible if necessary amendments are in place and abuse review is disabled	Possible if necessary amendments are in place and abuse review is disabled
Use for the provider's own purposes (e.g. training, service improvement)	Yes Use for training can be disabled			No			According to the information page: No According to Copilot AI Experiences Terms (Section 5): Possible			Possible with active web search*		No	Yes Use for training can be disabled		Yes	Possible when using Search Grounding (web search)	No	
Use within a company	**Possible with restrictions** no personal data no confidential data training disabled			**Possible**			**Not recommended**			**Possible** if web access is disabled. **Possible with restrictions** if web search is active (no personal data, no confidential data)		**Possible**	**Limited possibility** no personal data no confidential data Training disabled		**Not recommended**	**Possible with restrictions** if Search Grounding is disabled (no confidential data).	**Possible**	
Use with VISCHER Red Ink	N/A			N/A		Yes	N/A			N/A		Yes	N/A		Yes, limited data		N/A	Yes

* For web search, the MCA & DPA do not apply; the Microsoft Services Agreement for private customers applies instead
** For users in the European Economic Area, Switzerland, or the United Kingdom, the same terms regarding Google's use of data apply to the Unpaid Service of Google AI Studio as to the Paid Service.

Popular AI Tools: What about data protection? (Vischer AG law firm, linked on UZH page so somehow official-ish?)

https://www.rosenthal.ch/downloads/VISCHER_ai-tools-03-25.pdf

Contradicts UZH information on the same page regarding the use of Copilot. (Notes on the use of Microsoft Copilot at the University of Zurich)

Data Types and Data Security

The University of Oslo has a helpful classification of data types (see: https://www.uio.no/english/services/it/security/lsis/data-classes.html)

Open or freely available (Green)
Information that may or should be available to the general public, with no special access restrictions.

Restricted (Yellow)
Information which is not open for everyone. There are no laws or regulations saying that the information should be open.

In confidence (Red):
Information which the university is obliged to protect by law, agreements and other regulations. «In confidence» is used if the university, its partners, public interests, or individuals, may be subject to harm if the information is exposed to third parties.

Strictly in confidence (Black):
This category encompasses the same type of information as «In confidence (red)», but where special circumstances makes it necessary to protect the information even more. Demands on protection and safety are to be written down in agreements or other written documentation.

Data Types and Data Security

Green	Yellow
Open	Restricted

Red	Black
In confidence	Strictly in confidence

What can be stored and transmitted in GPT UiO?

Yellow and green data*

PLEASE PLEASE PLEASE Ask UZH to adopt a similar clear classification and to provide a table telling us what we can and cannot do with different data types, so that it's easy to understand and to use in research workflows.

AI tools for UZH?

Reply | Reply All | Forward

GS Giovanni Spitale
To Francesco Garassino
Cc Dominik Bolliger
mercoledì 09/07/2025 09:37
You forwarded this message on 22/09/2025 09:04.

Hi Francesco,
I've got your contact from Nikola Biller-Andorno.

I'm writing to suggest that the University of Zurich consider establishing a secure, institutionally governed Large Language Model (LLM) service similar to the "GPT UiO" recently launched at the University of Oslo, where I have been visiting recently (see: https://www.uio.no/english/services/it/ai/gpt-uio/).

Across the academic landscape, many researchers and students already use tools like ChatGPT to support writing, analysis, and content generation. However, these practices often occur in a legal and ethical "grey zone," with few clear policies or safeguards. Currently, LLMs are likely being used informally across UZH in ways that may breach data protection or academic integrity policies; for example, by uploading sensitive interview transcripts for qualitative coding, sharing identifiable data samples to refine analysis scripts, submitting draft grant proposals for editing, or even processing confidential student work for feedback generation.

Reflection Task

Take 20 seconds and think of one piece of data you used in a research task.

Now ask yourself:

- **Which risk category does this data fall into? (public, internal, confidential, sensitive, highly sensitive)**
- **What makes it belong there?**
- **If you had to justify your categorization to an ethics committee, what would be your strongest argument?**

Privacy and Security

Never put identifiable or sensitive information into ChatGPT, Gemini, Claude, or similar tools.

Consumer versions retain data or use it for training by default.

UZH allows us to use Copilot (which AFAIK to date does not include API access)

Totally local models are the **safest options.**

This also means: minimize the data you share.

You run a qualitative interview transcript through ChatGPT.

The text enters their servers and may be logged.

→ **Data breach risk + GDPR/FADP violation.** Must use Copilot or local model.

Reproducibility

If AI enters your workflow, you must log:

- Prompts
- Transcripts
- Version of the model
- Dataset state

Why? Because **ChatGPT today is not ChatGPT tomorrow.**

You need to be able to **show how you got an output.**

This becomes part of your Methods section.

You use ChatGPT to extract data but don't save prompts.

Three months later, you can't reproduce the analysis because the model changed.

→ Solution: save prompts, outputs, and versions to OSF. Ideally, this should be done via API with pinned versions.

Ghost authorship and disclosure

AI cannot be an author.

But you must **disclose "substantial use"**, especially text generation, coding assistance, or analytical scaffolding.

AI performs (some) work; **you remain accountable.**

You let ChatGPT write a full discussion section.

It's logically coherent but **contains invented citations.**

→ **You are responsible for errors!** Disclose AI assistance; take responsibility; rewrite with your reasoning.

"We used generative AI tools exclusively for copy-editing and proofreading support, in line with ethical recommendations by Porsdam Mann et al. (2024). All content was written, curated, and critically reviewed by the authors".

Porsdamm Mann et al. 2024 - https://www.nature.com/articles/s42256-024-00922-7
Nature 2023 - https://www.nature.com/articles/d41586-023-00191-1

Data Management Plans/ IRB applications

Modern DMPs / IRB applications should explicitly include:
- Where AI is used
- Which AI is used
- Why AI is used
- Where the model runs (local? cloud?)
- How outputs are validated
- How you protect against methodological and data protection risks

If an AI tool touches data, it becomes part of the governance ecosystem.
Likely to become a standard at some point. Be prepared.

Your project uses LLMs to cluster survey comments en masse.

You should specify which model sees what data, where that model runs, and how outputs are checked and validated.

→ This goes into the DMP + Ethics application.

Swissethics 2025 - https://swissethics.ch/en/news/2025/04/17/artificial-intelligence-and-research-involving-human-beings-issues-to-consider-when-submitting-a-project-to-a-research-ethics-committee

AI multiplies responsibility.

You are accountable for what the model does with your data, and what *you* do with its outputs.

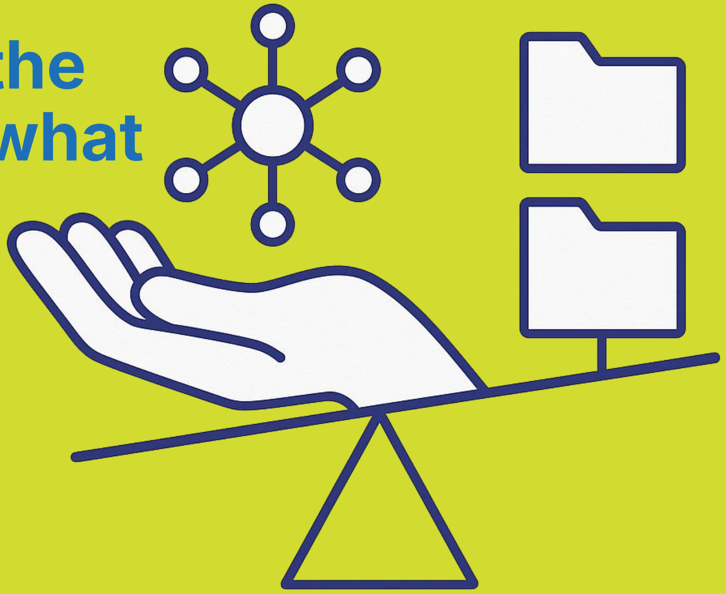

In Sum…

AI changes the workflow, not the purpose
Research remains judgment, interpretation, meaning-making.

Acceleration ≠ understanding
Faster does not mean deeper; depth is (still) human.

The epistemic duty grows
AI multiplies outputs; researchers must filter harder.

Responsible innovation
Governance, documentation, reproducibility become central.

What kind of researchers do we want to become in an AI-mediated world?

Let’s keep talking

giovanni.spitale@ibme.uzh.ch

Download the slides

LAB
Information Technology & Experimental Ethics

University of Zurich[UZH]

Institute of Biomedical Ethics
and History of Medicine