Agentic AI impact on our activities

Philippe Bacon - Paul Zakharov - Alexandre Boucaud

Biennale APC - June 2026

A Highly Divisive Topic

As we will see through the survey results, this topic is particularly divisive within APC, IN2P3, and in French society more broadly.

Adopters

Among enthusiasts, the main requests are

- **data security** and the deployment of **institutional tools** with **responsible usage**
- the need for **access to state-of-the-art tools**, **via API**, to stay competitive
- **trust** remains **limited** and the fear of **intellectual impoverishment** is present

Critics

The main reasons for strong opposition

- an energy and material sinkhole with an unsustainable ecological impact
- a source of job precarity and skills erosion
- a **technological choice** that requires **collective debate**

Manifeste d'objecteurs de conscience de l'IAg https://atecopol.hypotheses.org/13082

GOUVERNEMENT
Liberté
Égalité
Fraternité

👉 Read the consultation results

Direction interministérielle
de la transformation publique

Conseil de l'intelligence artificielle et du numérique : quelles priorités pour 2026 ?

Synthèse de la consultation

C
IA
Num

Service innovation et participation citoyenne / Centre interministériel de la participation citoyenne

L'objectif [de la mise en place de Emmy] est de permettre à chacun et chacune de mieux comprendre les opportunités que peut lui offrir l'IA dans son cadre professionnel, tout en identifiant aussi ses risques et ses limites. Et bien entendu, ces expérimentations doivent se faire avec modération, ***l'empreinte carbone de l'utilisation de l'IA étant un enjeu majeur****.*

Antoine Petit · Voeux 2026

Purpose of this Working Group

AI has become a major topic in today's society, no matter what we think of it.

We do think there is a **need for discussions at the level of the lab** to learn about its current capabilities, review the general usage and question our practices.

This working group was formed to

- collect discussion on the use of **AI in research**,
- focus on **practical applications** in research contexts.
- identify key areas of focus with some **guidelines**.

> **Disclaimer:**
>
> Our role is not to attack nor defend AI, but to examine how to use it effectively in research.

Outline

01

SECTION 01

Agentic AI as of June 2026

to put everything in context

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Many language models were released, getting better with time until they became **reasoning models** (*chain-of-thought*) **in 2025**, enabling step-by-step processing of larger inputs.

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Many language models were released, getting better with time until they became **reasoning models** (*chain-of-thought*) **in 2025**, enabling step-by-step processing of larger inputs.

They were then equipped with **tools** to perform calculations or take actions (retrieve web page content, write and run a script, then read the result).

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Many language models were released, getting better with time until they became **reasoning models** (*chain-of-thought*) **in 2025**, enabling step-by-step processing of larger inputs.

They were then equipped with **tools** to perform calculations or take actions (retrieve web page content, write and run a script, then read the result).

By giving them the ability to self-evaluate at the end of a task — and potentially retry differently — they became **agents**.

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Many language models were released, getting better with time until they became **reasoning models** (*chain-of-thought*) **in 2025**, enabling step-by-step processing of larger inputs.

They were then equipped with **tools** to perform calculations or take actions (retrieve web page content, write and run a script, then read the result).

By giving them the ability to self-evaluate at the end of a task — and potentially retry differently — they became **agents**.

Agents then specialized in specific tasks (e.g. coding, reading PDFs, web search), giving rise to the concept of **multi-agent** systems — an orchestrator agent managing specialized agents and aggregating results.

Agentic AI in 1 Slide

November 30th 2022 - First release of ChatGPT (OpenAI), opened a **new era of langage models**

Many language models were released, getting better with time until they became **reasoning models** (*chain-of-thought*) **in 2025**, enabling step-by-step processing of larger inputs.

They were then equipped with **tools** to perform calculations or take actions (retrieve web page content, write and run a script, then read the result).

By giving them the ability to self-evaluate at the end of a task — and potentially retry differently — they became **agents**.

Agents then specialized in specific tasks (e.g. coding, reading PDFs, web search), giving rise to the concept of **multi-agent** systems — an orchestrator agent managing specialized agents and aggregating results.

Until late 2025, these agents made many errors or produced needlessly verbose code, leading to a perception of unreliable tools.

Early 2026: A Paradigm Shift

A major leap in performance since December 2025, followed by a second one in February 2026.

This was first seen with the release of Google Gemini 3 Pro, then Anthropic's Claude Opus 4.6 (large) and Claude Sonnet 4.6 (medium).

Model reliability and focus made an enormous leap forward, making them far more useful for development tasks.

Error Detection Benchmark

BullshitBench v2: Detection Rate Over Time

Release date vs. green rate (clear pushback %) for all organizations. Best model per release date shown.

☑ Best per release date only

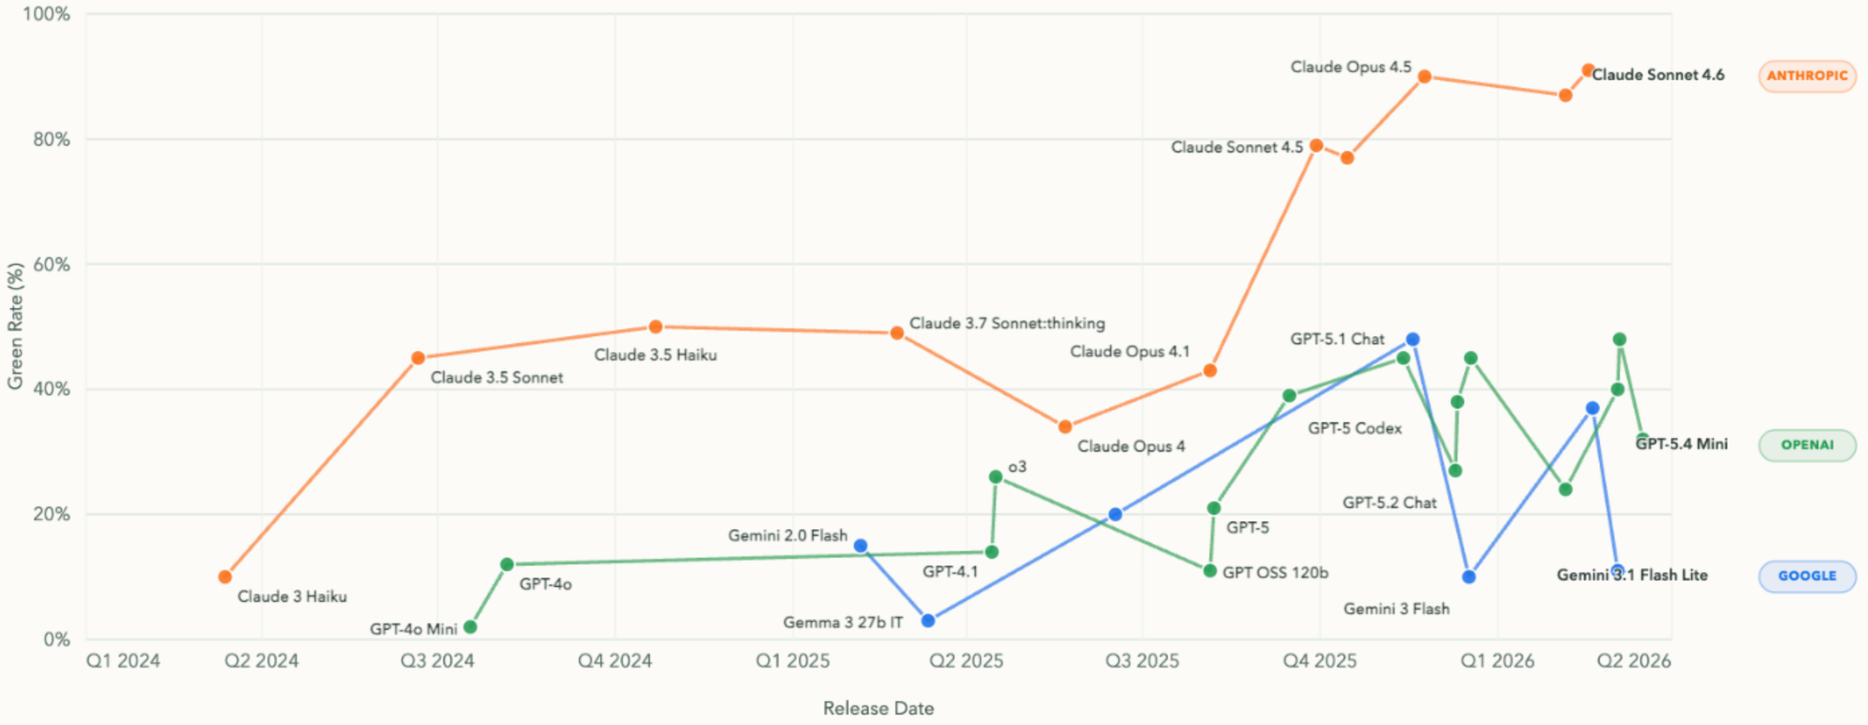

33 models shown from Anthropic, Google, OpenAI.

https://petergpt.github.io/bullshit-benchmark

Breakthrough in Mathematics - Solving Erdös Problems

May 20th 2026

20 mai 2026 Recherches Étape clé

Un modèle OpenAI a réfuté une conjecture centrale en géométrie discrète

Lire la preuve ↗ Lire les remarques complémentaires ↗

https://openai.com/fr-FR/index/model-disproves-discrete-geometry-conjecture

May 22nd 2026

Google DeepMind 2026-5-22

Advancing Mathematics Research with AI-Driven Formal Proof Search

George Tsoukalas[1†], Anton Kovsharov[1†], Sergey Shirobokov[1†], Anja Surina[1], Moritz Firsching[1], Gergely Bérczi[2], Francisco J. R. Ruiz[1], Arun Suggala[1], Adam Zsolt Wagner[1], Eric Wieser[1], Lei Yu[1], Aja Huang[1], Miklós Z. Horváth[1], Andrew Ferrauiolo[1], Henryk Michalewski[1], Codrut Grosu[3], Thomas Hubert[1], Matej Balog[1], Pushmeet Kohli[1] and Swarat Chaudhuri[1†]
[1]Google DeepMind[1], [2]Aarhus University, [3]Google

Large language models (LLMs) increasingly excel at mathematical reasoning, but their unreliability limits their utility in mathematics research. A mitigation is using LLMs to generate formal proofs in languages like Lean. We perform the first large-scale evaluation of this method's ability to solve open problems. Our most capable agent autonomously resolved 9 of 353 open Erdős problems at the per-problem cost of a few hundred dollars, proved 44/492 OEIS conjectures, and is being deployed in combinatorics, optimization, graph theory, algebraic geometry, and quantum optics research. A basic agent alternating LLM-based generation with Lean-based verification replicated the Erdős successes but proved costlier on the hardest problems. These findings demonstrate the power of AI-aided formal proof search and shed light on the agent designs that enable it.

AlphaProof autonomously solved 9 Erdös conjectures

https://arxiv.org/abs/2605.22763

Possible Impact on Our Profession

Various recent studies (often biased, as they are commissioned by the tech giants themselves) report the percentage of typical job tasks that AI would be capable of handling.

It is clear that our profession is among those that will face strong AI-driven pressure, whether we like it or not.

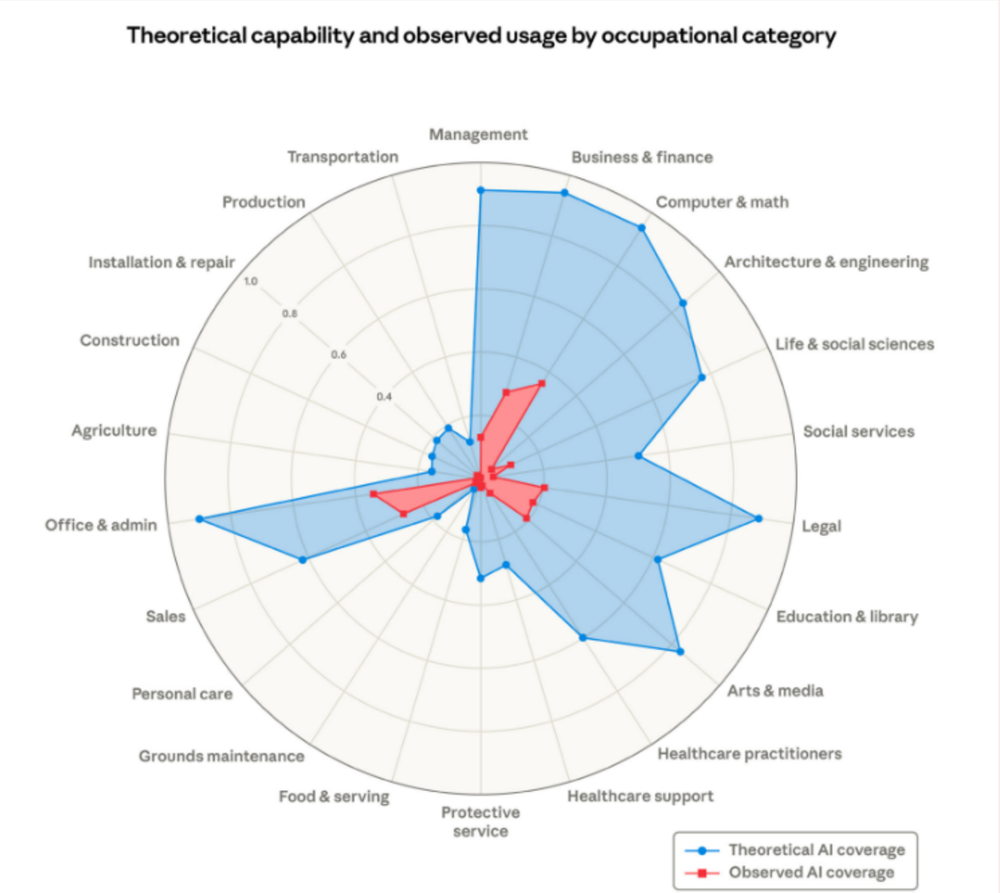

Figure 2: Theoretical capability and observed exposure by occupational category
This figure shows the share of job tasks that LLMs could theoretically perform (blue area) and our own job coverage measure derived from usage data (red area).

https://www.anthropic.com/research/labor-market-impacts - March 2026

More related to us there is a trend in AI Scientist Concept

The Denario project: Deep knowledge AI agents for scientific discovery

Boris Bolliet*, Pablo Villanueva-Domingo*, Francisco Villaescusa-Navarro*,
Adrian E. Bayer, Aidan Acquah, Chetana Amancharla, Almog Barzilay Siegal , Pablo Bermejo,
Camille Bilodeau, Pablo Cárdenas Ramírez, Miles Cranmer, Urbano L. França, ChangHoon Hahn,
Yan-Fei Jiang, Raul Jimenez, Jun-Young Lee, Antonio Lerario, Osman Mamun, Thomas Meier,
Anupam Anand Ojha, Pavlos Protopapas, Shimanto Roy, David N. Spergel, Pedro Tarancón-Álvarez,
Ujjwal Tiwari, Matteo Viel, Digvijay Wadekar, Chi Wang, Bonny Y. Wang, Licong Xu, Yossi Yovel, Shuwen Yue,
Wenhan Zhou, Qiyao Zhu, Jiajun Zou, Íñigo Zubeldia

Module	Task	Input	Output
Idea	Generate project idea	input.md	idea.md
Literature	Determine if idea is new	input.md idea.md	literature.md
Methods	Develop project plan	input.md idea.md	methods.md
Analysis	Implement plan write and execute code make plots	input.md idea.md methods.md	results.md Plots
Paper	Write paper	input.md idea.md methods.md results.md	paper.pdf
Review	Review paper	(input.md) paper.pdf	referee.md

Towards Agentic AI serving Science with builtin rigor

LIGHTCONE RESEARCH

Open-source initiative between CNRS and UC Berkeley focused on **verifiable** and **reproducible** AI-assisted science.

https://lightconeresearch.org

https://github.com/LightconeResearch

ASTRA

Specification format for computational science that tracks every the provenance of data, model and insights, and records every decision that was made for a given analysis.

lightcone-cli

Command line interface leveraging an AI agent that turns any ASTRA specification into materialized results through custom recipes (workflows).

02

SECTION 02

Survey results

https://machine-learning.pages.in2p3.fr/llm-survey-2026

APC responders

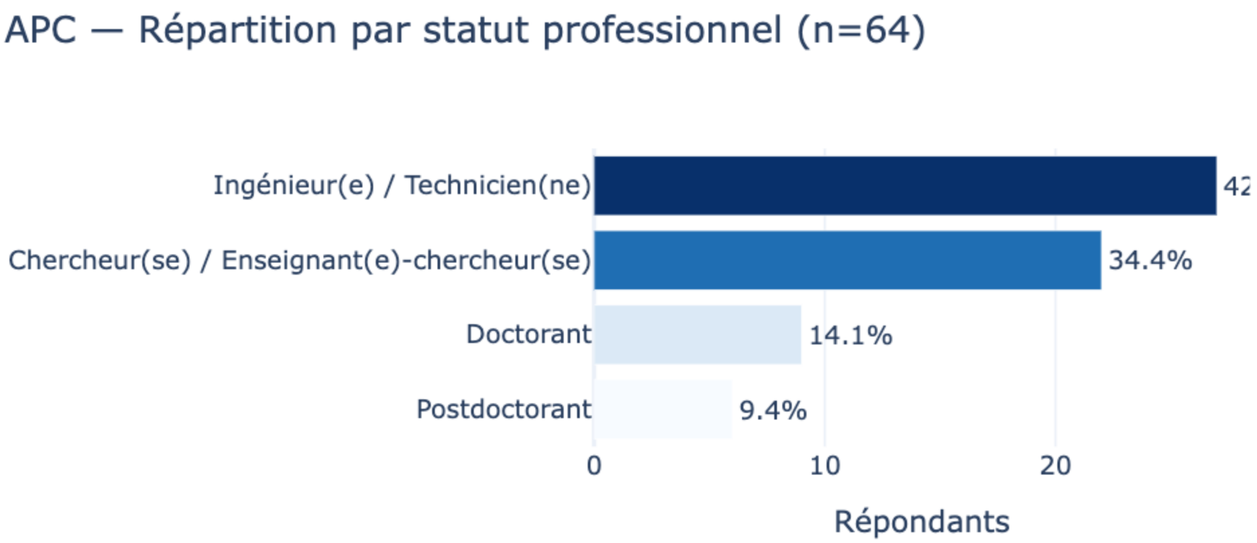

64 APC members replied, 53 actually using AI.

Confidence Index

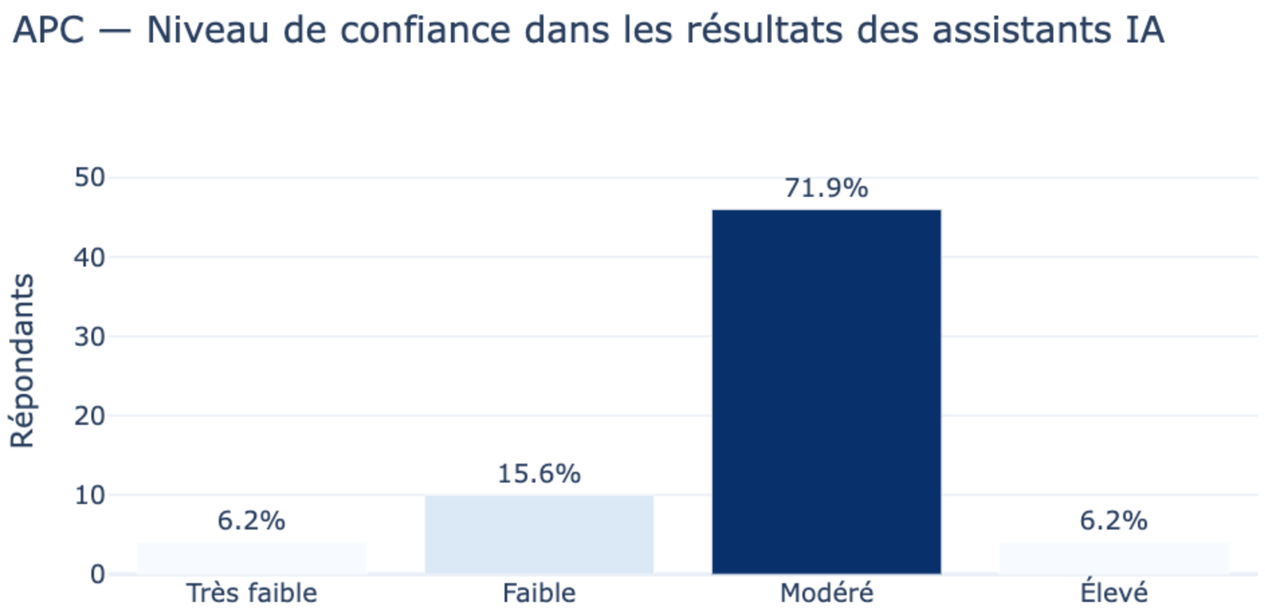

Usage frequency

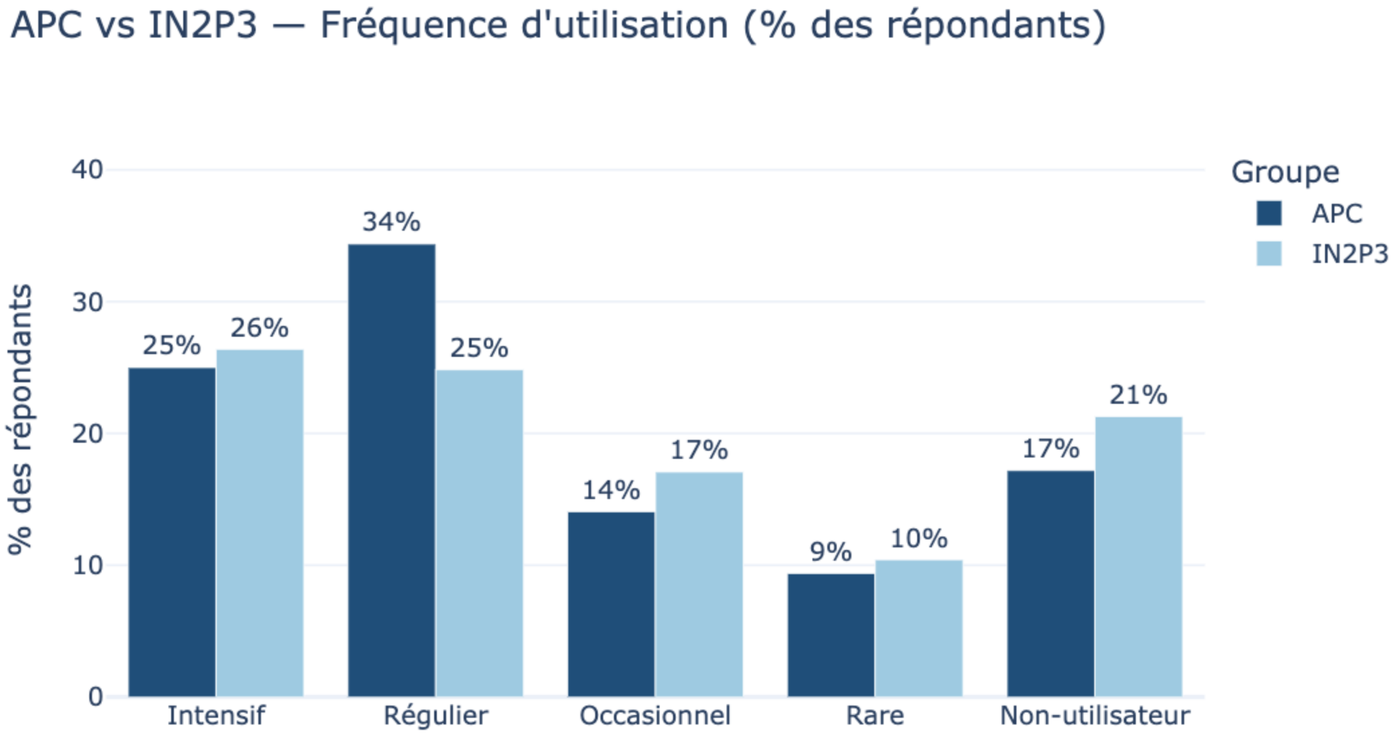

Usage by APC members (code developers)

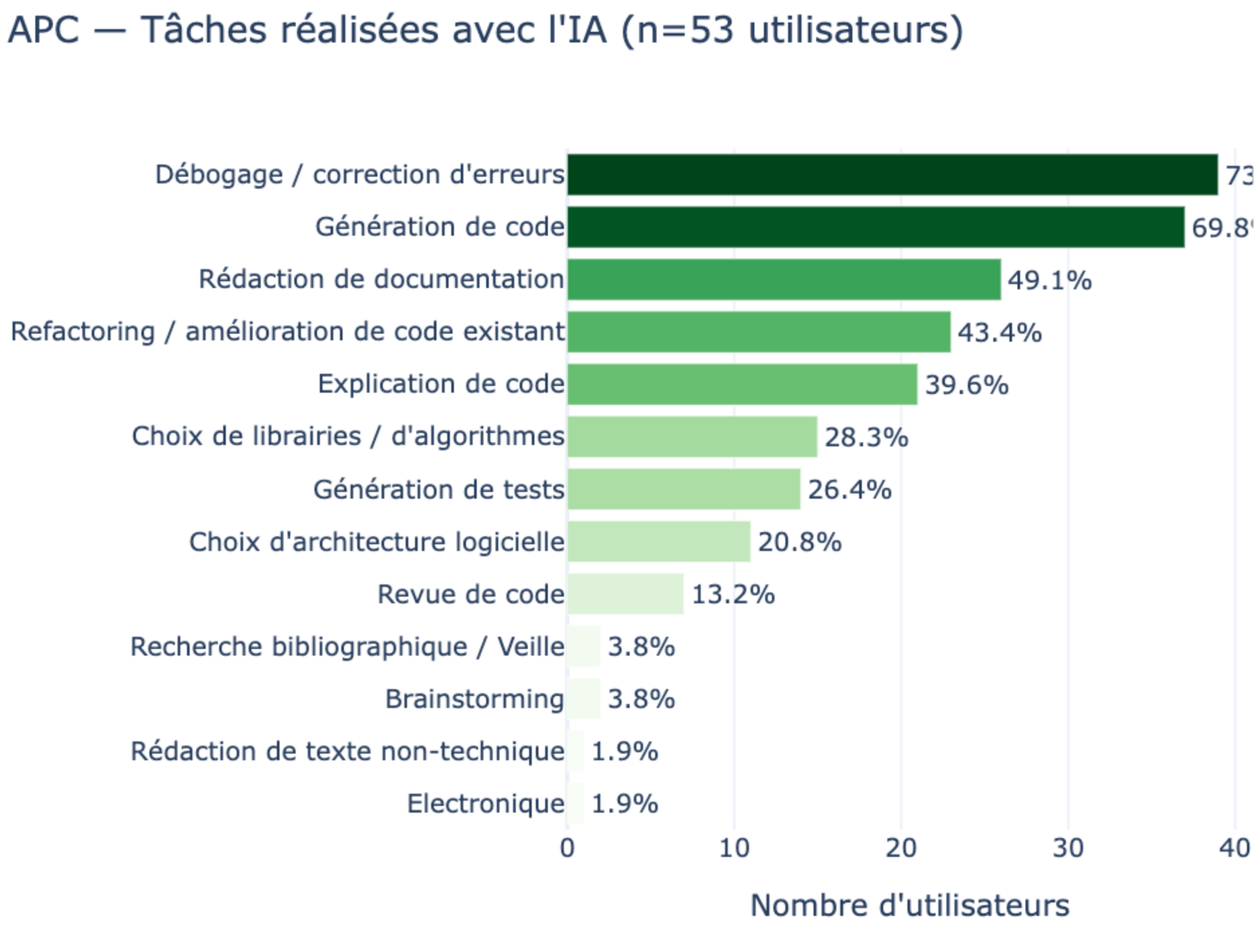

What Tools do People Use

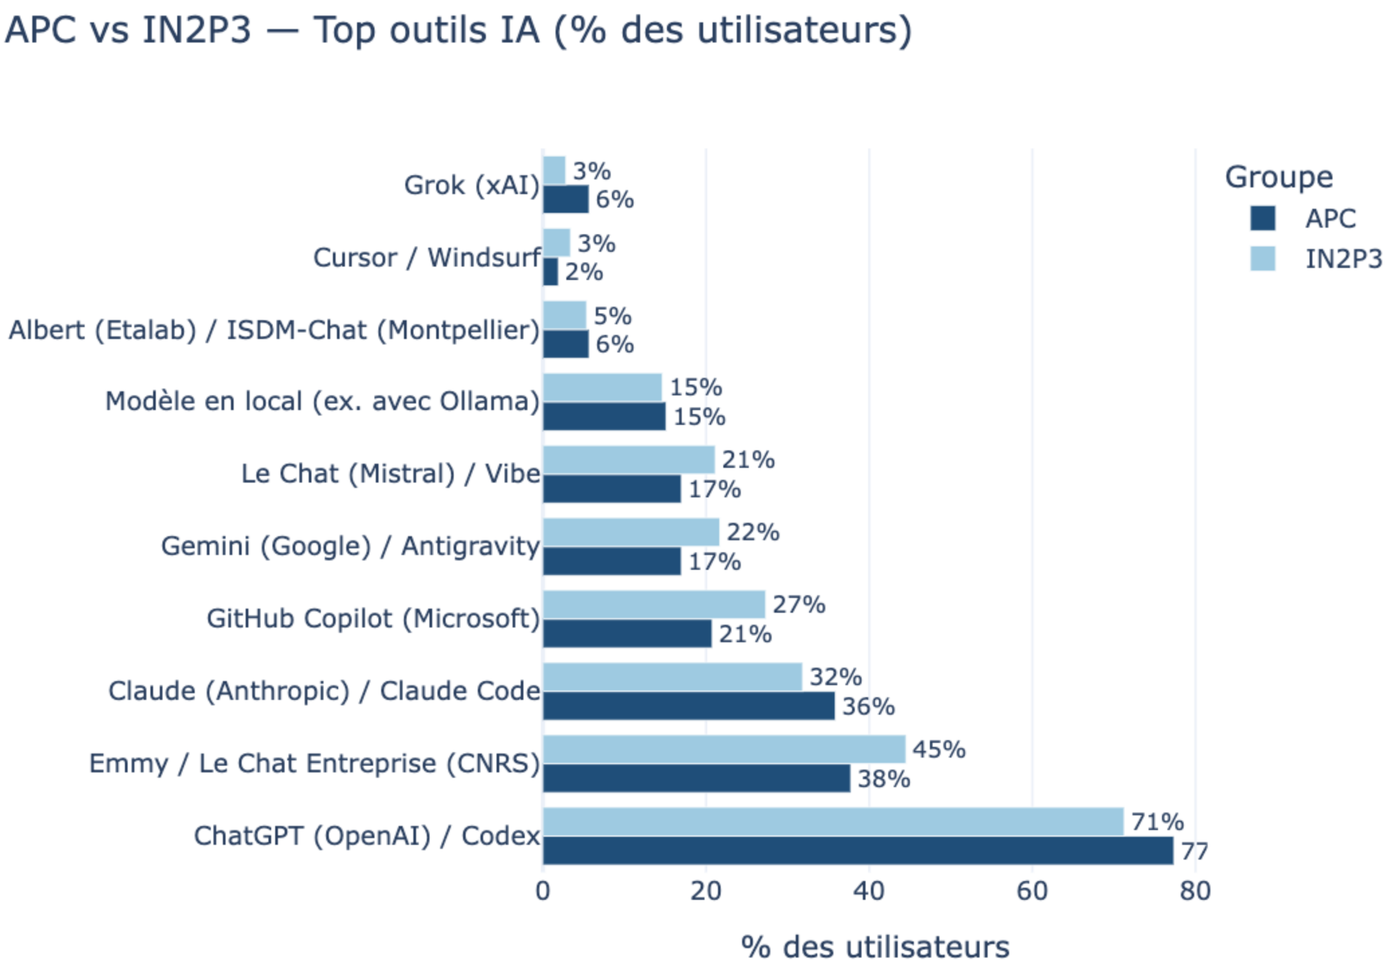

What Developers Want

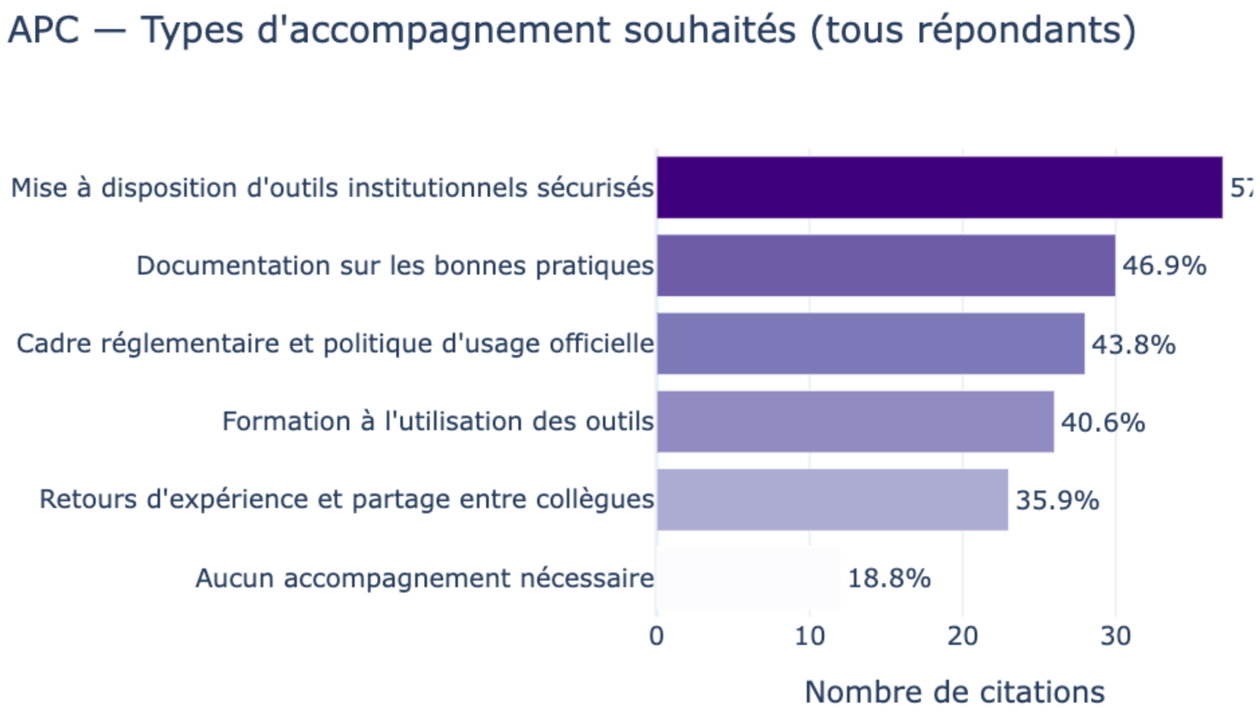

Barriers to Usage by Frequency of Use

Freins principaux par catégorie d'usage (% du groupe)

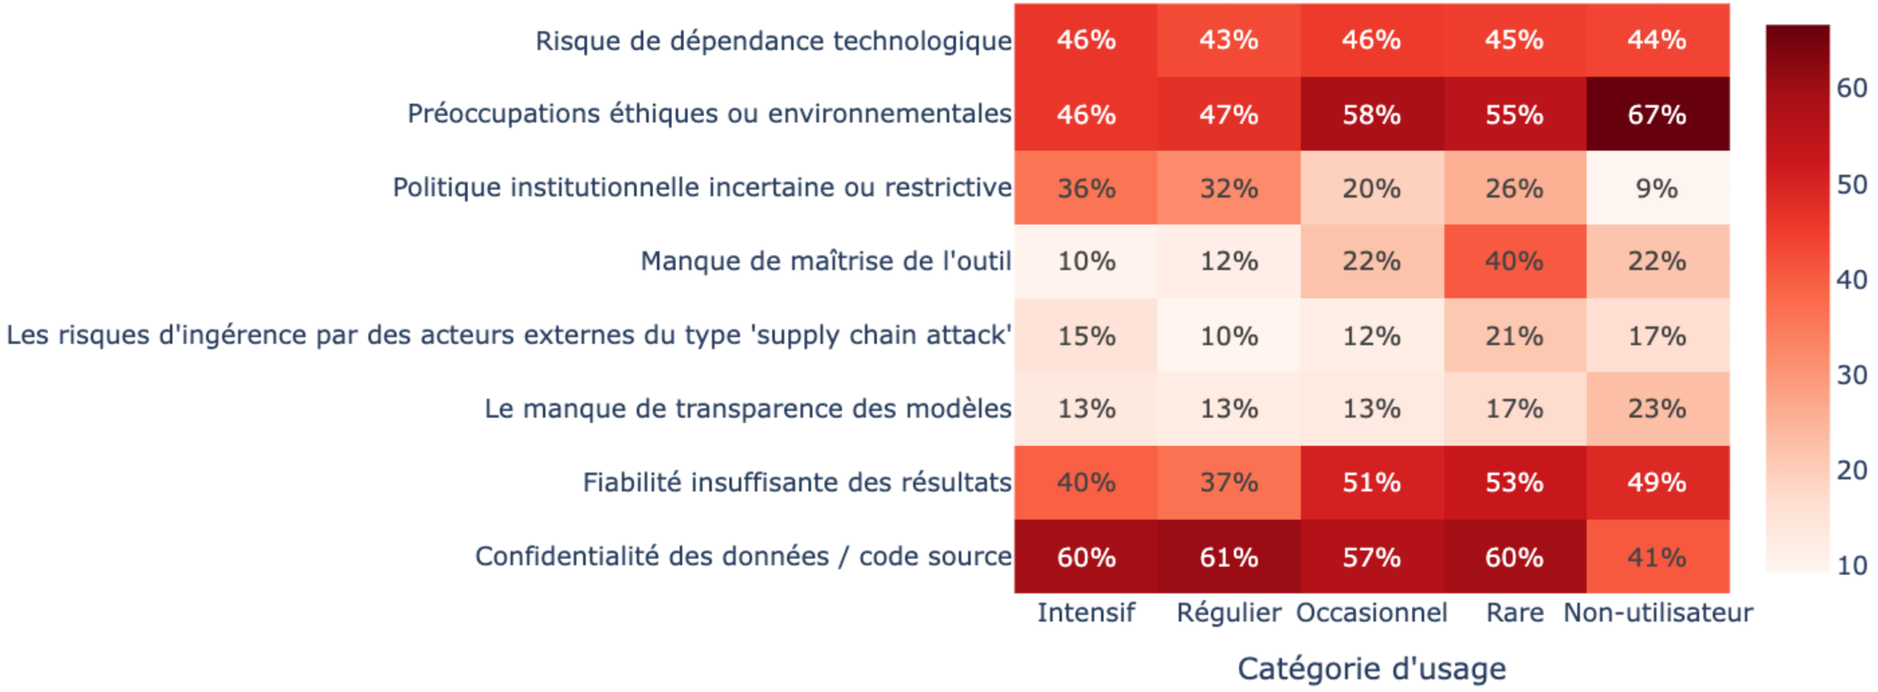

Selected Quotes

Le plaisir du développement et de la réflexion disparait en promptant

---

I am afraid of becoming stupid, losing my brain power

Financement de ces outils au même titre que n'importe quel logiciel de CAO

Restreindre l'usage de l'IA au vu de ses implications sociales, éthiques et environnementales délétères.

Un rappel clair et précis sur les effets sociaux de l'IA générative.

Les personnels universitaires ne peuvent pas bénéficier des mêmes outils que les personnels CNRS dans une UMR, ce qui est préjudiciable.

Je n'utilise pas l'IA par fierté. J'estime être en capacité de produire de moi même ce qui pourrait concerner mes demandes. Bien évidement dans le cadre de mon travail. Cependant, comme tout outil, j'ai bel et bien conscience de son utilité et de son efficacité.

Notre rôle de scientifique est d'être des dépositaires humains de la connaissance, si nos compétences et notre capacité à raisonner sont dépendantes d'outils qui nous privent de notre réflexion propre, je pense que notre intégrité est menacée. Il est clair que l'usage de LLM a des effets très négatifs sur une partie de la population, je pense que nous jouons un rôle d'exemplarité dans la prudence vis-à-vis de ces outils.

All shall use AI, from all scientific domains .. it is not a choice ...

03

SECTION 03

Use cases

Collected during the preparation meetings and the survey

Use cases per job category (non-exhaustive)

Docs/Postdocs

- **StackOverflow-like** assistance.
- **Brainstorming** (not limited to code).
- **Bibliographic research**.

Faculty/Researchers

- **Explain concepts** to students (reformulation).
- **Theme generation** for research topics.
- **Tentative use for grading** (mixed results).
- **Synthesize documents** (papers).
- **Grant writing** support.
- **Code writing assistance**.
- **Bibliographic research** (Consensus, ChatGPT, Sight).

Use cases per job category (non-exhaustive)

Administrative/HR

- **Translation**
- **Writing documents** (posts on social networks, APC website)
- **Rephrasing of message**
- **Abstract of thesis**
- **Explanation of complex concepts in scientific literature**
- **Interview preparation** (questions)

IT/Engineers

- **"Code assistant"** (unit testing, debugging, refactoring)
- **Propose code architecture & software design**.
- **Recommend algorithms** (pros/cons).

Generative AI at CNRS: Emmy

Mid-December 2025, launch of **Emmy**, the generative AI for CNRS staff, whose capabilities include

- text translation in all languages;
- document summarization;
- rephrasing assistance;
- reasoning support;
- web search;
- text and image recognition;
- "reasoning" mode: the AI processes the user's question step by step to provide a more relevant and comprehensive answer;
- document collections

This tool results from a deal between CNRS and the Mistral IA company for 35,000 licenses of their **Le Chat Entreprise** offer.

https://emmy.cnrs.fr/

04

SECTION 04

Major themes

Environmental Impact

- **Huge training & inference cost**: ~1.8–2.2% of France's annual electricity production (growing)
- **Rapid advancement** of generative AI (~month) vs. research timelines.
  - Hard to quickly adapt habits, technologies and guidelines.
- **Knowledge gap**: State-of-the-art and capabilities evolve monthly.
  - A technology is rapidly oudated...
  - ... so does the use of it.
- **Need for flexibility** and a **local experimentation platform**.
  - Experimentation to design needs for ressources (hardware, electric consumption)
  - Goal is *not* to scale with large model providers.

Environmental Impact

- **Huge training & inference cost**: ~1.8–2.2% of France's annual electricity production (growing)
- **Rapid advancement** of generative AI (~month) vs. research timelines.
  - Hard to quickly adapt habits, technologies and guidelines.
- **Knowledge gap**: State-of-the-art and capabilities evolve monthly.
  - A technology is rapidly oudated...
  - ... so does the use of it.
- **Need for flexibility** and a **local experimentation platform**.
  - Experimentation to design needs for ressources (hardware, electric consumption)
  - Goal is *not* to scale with large model providers.

Challenges

- **Hard to quantify** environmental impact of individual AI usage.
- **Rising RAM costs** + **hardware availability** issues.
  - Global resource reallocation toward AI.

Environmental Impact

Ideas

- **Use rate-limited inference services** (e.g., Albert, ILaaS).
- **Stick to shared computing resources**.
- **Recycle hardware** (Environmental charter - Art. 1.1, 3.2, 3.3)
- **Set up yearly token allocation** (like we already do with compute time).
- **Run carbon footprint tracking** (anonymous):
  - Include in **annual lab reports**.
  - Add to **job submission metrics**.
- **Use local platform** only for small-scale testing (not for scaling demand).
  - Helps designing infrastructure adapted to the need.

Example: Inference as a Service in French Academia

ILaaS

A shared federation aiming for trustworthy, robust, ethical, and frugal generative AI

Provides an inference API for open-source models.

UPC will contribute to this federation

ILaaS
Infrastructure d'IA mutualisée

Utiliser les services d'inférence d'ILaaS

Intégrez ILaaS facilement dans vos applications existantes

Si vous utilisez un logiciel exploitant l'IA générative et qu'il permet de choisir l'URL du serveur d'IA, vous pouvez substituer les services d'OpenAI par les services d'ILaaS est très simplement.

Il vous faudra pour cela obtenir du consortium une clé d'API pour vous authentifier sur le service d'inférence d'ILaaS.

Concrètement, il suffit de:

1. demander une clé d'API ILaaS en tant que consommateur d'inférence,
2. paramétrer l'URL d'accès au service avec l'URL qui vous sera fournie,
3. saisir la clé d'API ILaaS dans l'interface de configuration.

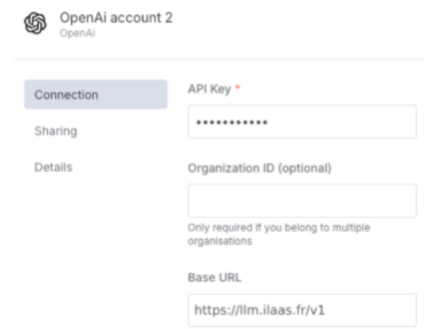

Exemple de configuration de l'URL du serveur d'IA générative et de la clé d'API dans un programme (ici OpenWebUI)

https://ilaas.fr

Example: Inference as a Service in French Academia

Albert – DINUM API

La plateforme interministérielle d'inférence d'IA générative

Albert API, développé par la Direction interministérielle du numérique (DINUM) dans le cadre de son programme **ALLiaNCE**, est la plateforme interministérielle d'inférence permettant aux administrations de se doter d'intelligence artificielle (IA) générative.

A destination de vos produits numériques, l'API Albert offre un accès vers de nombreux modèles pour tous vos cas d'usage (voir les modèles) ainsi qu'une solution de *RAG* (Retrieval Augmented Generation) *as a service*.

De l'expérimentation...

Construit selon les conventions d'OpenAI, Albert API est compatible avec le **client officiel Python** et s'intègre facilement avec des outils comme OpenWebUI ou Langchain facilitant ainsi le développement de vos produits numériques.

... à la production.

Grâce à une infrastructure sécurisée, Albert API offre un service pensé pour la production de vos produits numériques. Consultez notre page dédiée en cliquant ici pour en savoir plus.

https://albert.sites.beta.gouv.fr

Claude Code + Albert API = `le-claude`

https://github.com/EiffL/le-claude

```
$ npx le-claude
```

Transparency & Security

- **CNRS-Mistral contract (Emmy)**: Data remains **private** and **not used for training**.
- **UPC**: ILaaS runs on **French inference servers**.
- **Prohibited usage**: Any other tool must comply with **RGPD**.
  - **No sharing** of personal data (e.g., student lists) without **explicit consent**.
  - **Anonymize data** if necessary.

Ideas

- **Always apply critical thinking** to AI responses.
  - Models can fail unpredictably (e.g., 10 exchanges then AI hallucination).
- **Expertise first**: Use AI as a **tool**, not a replacement.
  - **Improve/correct** AI outputs with human input.

Academic Integrity

- **arXiv ban**: if manifest use of AI is detected
- **Data poisoning**: Risk of training on **non-peer-reviewed drafts**
- **Declare AI use**: https://declare-ai.org.
- **Adapt policies** for interns/labs (e.g. https://mammouth.ai).

Example: AI Charters

Discussed in Reprises sessions

- **Legal framework**: GDPR, copyright, sector-specific compliance (e.g. health, education).
- **Ethics principles**: Transparency, accountability, fairness, privacy.
- **Best practices**:
  - usage limits
  - human validation of results
  - traceability of generated content (source)
  - ecological impact

Example: AI Charters

Discussed in Reprises sessions

- **Legal framework**: GDPR, copyright, sector-specific compliance (e.g. health, education).
- **Ethics principles**: Transparency, accountability, fairness, privacy.
- **Best practices**:
  - usage limits
  - human validation of results
  - traceability of generated content (source)
  - ecological impact

Must be educational!

Examples:

- AI charter portal in public administration
- KairoiAI template used as a basis by LIP6
- personal charter of a PhD student

Futures of the Profession

Adaptation: How will it change our work?

Identified risks

- affect social relations
- AI related illnesses (e.g. "AI fatigue")
- publication pressure
- hiring pressure (docs/postdocs)
- dependency
- Knowledge acquisition & expertise preservation

Example: Essay Emphasizing the Human Dimension of Our Work

Why do we do astrophysics?

by **David W. Hogg**[1]

Abstract: At time of writing, large language models (LLMs) are beginning to obtain the ability to design, execute, write up, and referee scientific projects on the data-science side of astrophysics. What implications does this have for our profession? In this *white paper*, I list—and argue for—a set of facts or "points of agreement" about what astrophysics is, or should be; these include considerations of novelty, people-centrism, trust, and (the lack of) clinical value. I then list and discuss every possible benefit that astrophysics can be seen as bringing to us, and to science, and to universities, and to the world; these include considerations of love, weaponry, and personal (and personnel) development. I conclude with a discussion of two possible (extreme and bad) policy recommendations related to the use of LLMs in astrophysics, dubbed "let-them-cook" and "ban-and-punish." I argue strongly against both of these; it is not going to be easy to develop or adopt good moderate policies.

https://arxiv.org/abs/2602.10181 - February 2026

Pitch

- the arrival of powerful generative AI is here
- several options are available to us as scientists
- two are particularly bad: full acceptance ("all-AI") and outright rejection through policing of researchers
- he uses it to revisit the fundamentals of why we love this work and opens perspectives on how to work alongside AI

Futures of the Profession

Adaptation: How will it change our work?

Identified risks

- affect social relations
- AI related illnesses (e.g. "AI fatigue")
- publication pressure
- hiring pressure (docs/postdocs)
- dependency
- expertise preservation

Ideas

- Reflect collectively on that transformation
- Include ethical considerations.
- Rethink the Research metrics - role of publications vs. other means
- Be mindful of new societal risks
  - loss of meaning in the research world – essay
  - new psychosocial risks – example

Teaching

- **AI Charter** on the use of AI at UPC underway (aligned with other universities).
- UPC will get it's own LLM by contributing to the inference resources of the **ILaaS**
- Things are different now - we need to adapt
  - Ex: instead of banning ChatGPT for translations, an English teacher (middle/high school) organized a session to *analyze ChatGPT translations* and compare them with manual translations, only to demonstrate the shortcomings of using AI.

Ideas

- Instead of banning AI, encourage moderate use by calling on **student responsability**
  - *"This is your training - you are learning"*
- **Adapt** the way skills & knowledge are **evaluated**

Example: Teaching Agentic AI Alongside a University Course

Feedback of the use of LLMs in an astrophysics course at Harvard by Chris Stubbs.

Key points

- substantial upfront preparation work on prompts (everything is shared in the paper)
- RAG fine-tuning on a course document
- AI restricted to short answers **with course citations** and no extended student–AI dialog
- guidelines on when use is authorized and when it is strongly discouraged

⇛ Overall very well received by students

Stubbs et al. 2026 - https://arxiv.org/abs/2602.04389

Merci à toutes les personnes ayant répondu au sondage et celles présentes aux réunions du WG

Time for discussions

Ad: Ecole thématique Labobots / AISSAI

A second thematic semester is being organized this year in partnership with IN2P3.
A thematic school **for ITA** called *Labobots* will be held in the fall September 29 → October 2, 2026 in Saint-Rémy-lès-Chevreuse. It will be sponsored by the CNRS center AISSAI.

> **Registration opening today or tomorrow !**
>
> https://indico.ijclab.in2p3.fr/event/13661

Organisation

- Françoise Bouvet (IJCLab)
- David Rousseau (IJCLab)
- David Chamont (IJCLab)
- Hugo Bacard (IJCLab)
- Sébastien Gadrat (CCIN2P3)
- Imed Magroune (CEA)
- Anne-Laure Méalier (Centrale Mediterrannée)
- Alexandre Boucaud (APC)

AI for science, science for AI

Questions for the discussion

Quand est-il acceptable d'utiliser l'IA pour produire un document à notre place ?

When is it acceptable to have AI write a document for us?

ex: rapport, mail, travail intellectuel, revue, ...

Quelles règles aimeriez-vous faire adopter au niveau du laboratoire ?

Which rules would you like to have the lab commit to?

ex: charte interne, points de vigilance, ...

Quelles sont vos attentes vis-à-vis du laboratoire concernant l'IA ?

What are your expectations from the lab concerning AI?

ex: cours, discussions usages / éthique, actualités et évolution, ...