

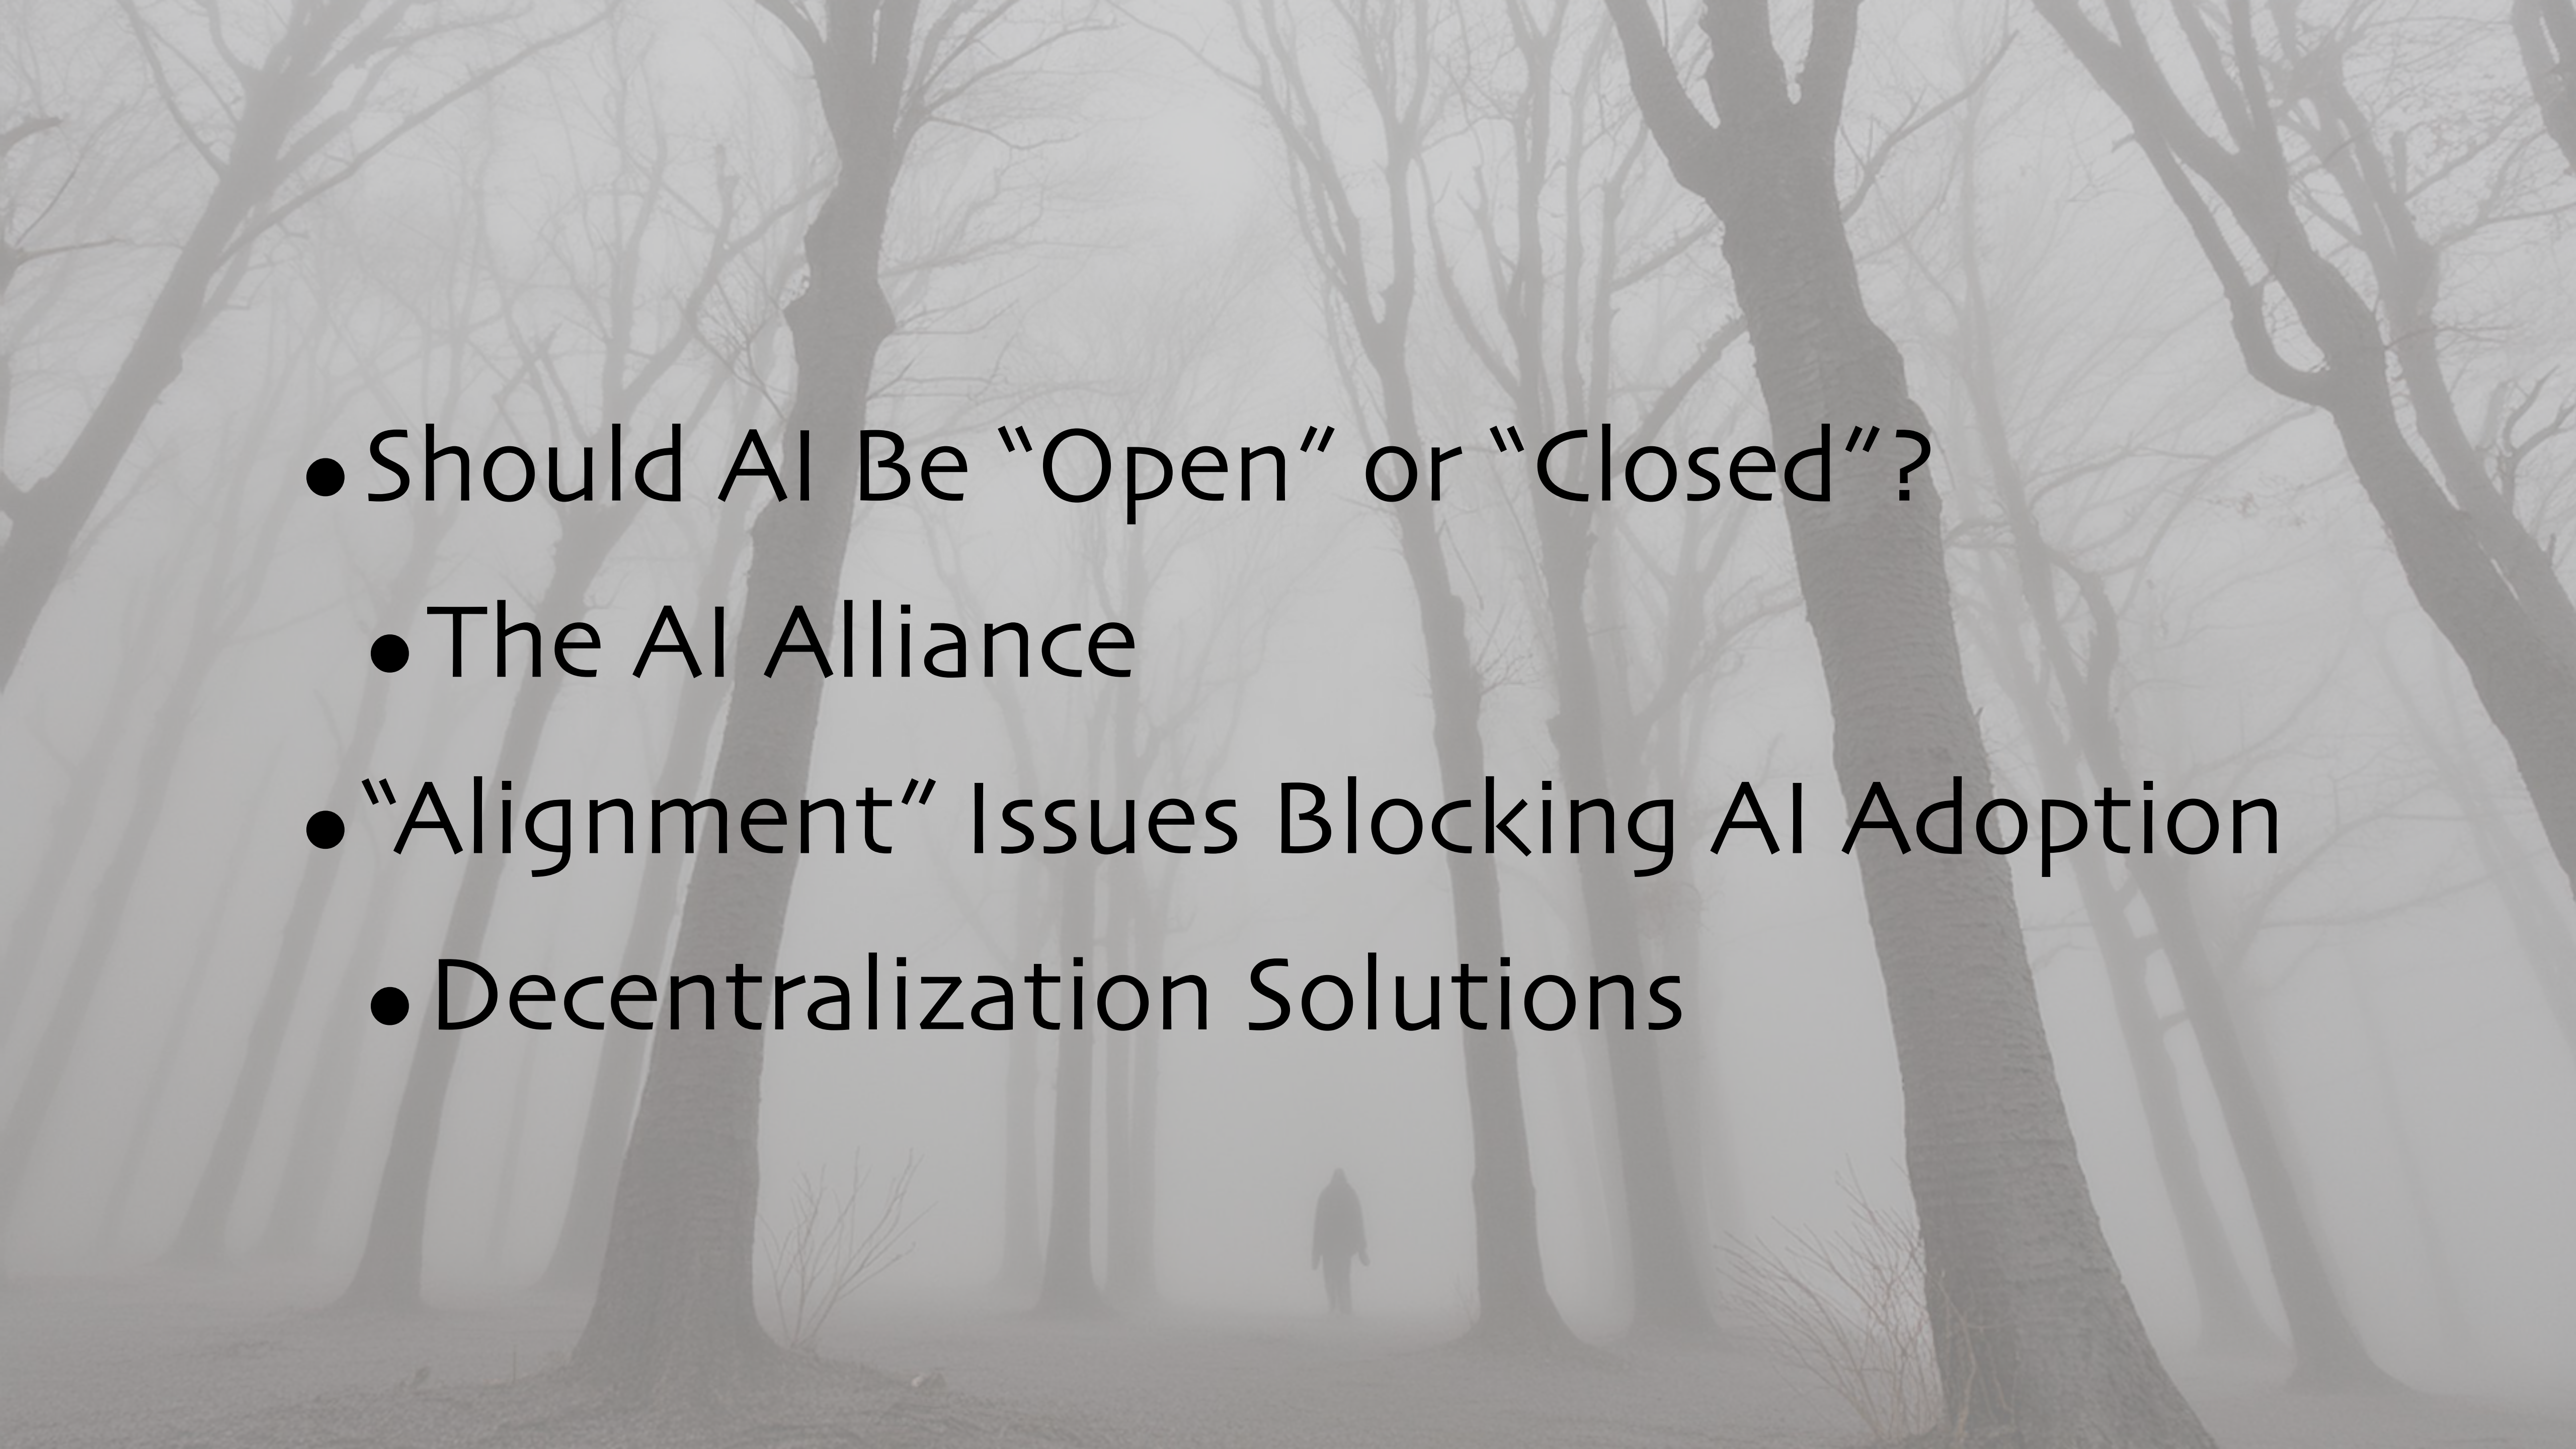
Issues Blocking AI Adoption: A Decentralization Perspective

Summit on Responsible Decentralized Intelligence
- Future of Decentralization and AI,
NYC August 6, 2024

Dean Wampler, Ph.D.
The AI Alliance
IBM Research
thealliance.ai
dwampler@thealliance.ai

deanwampler.com/talks

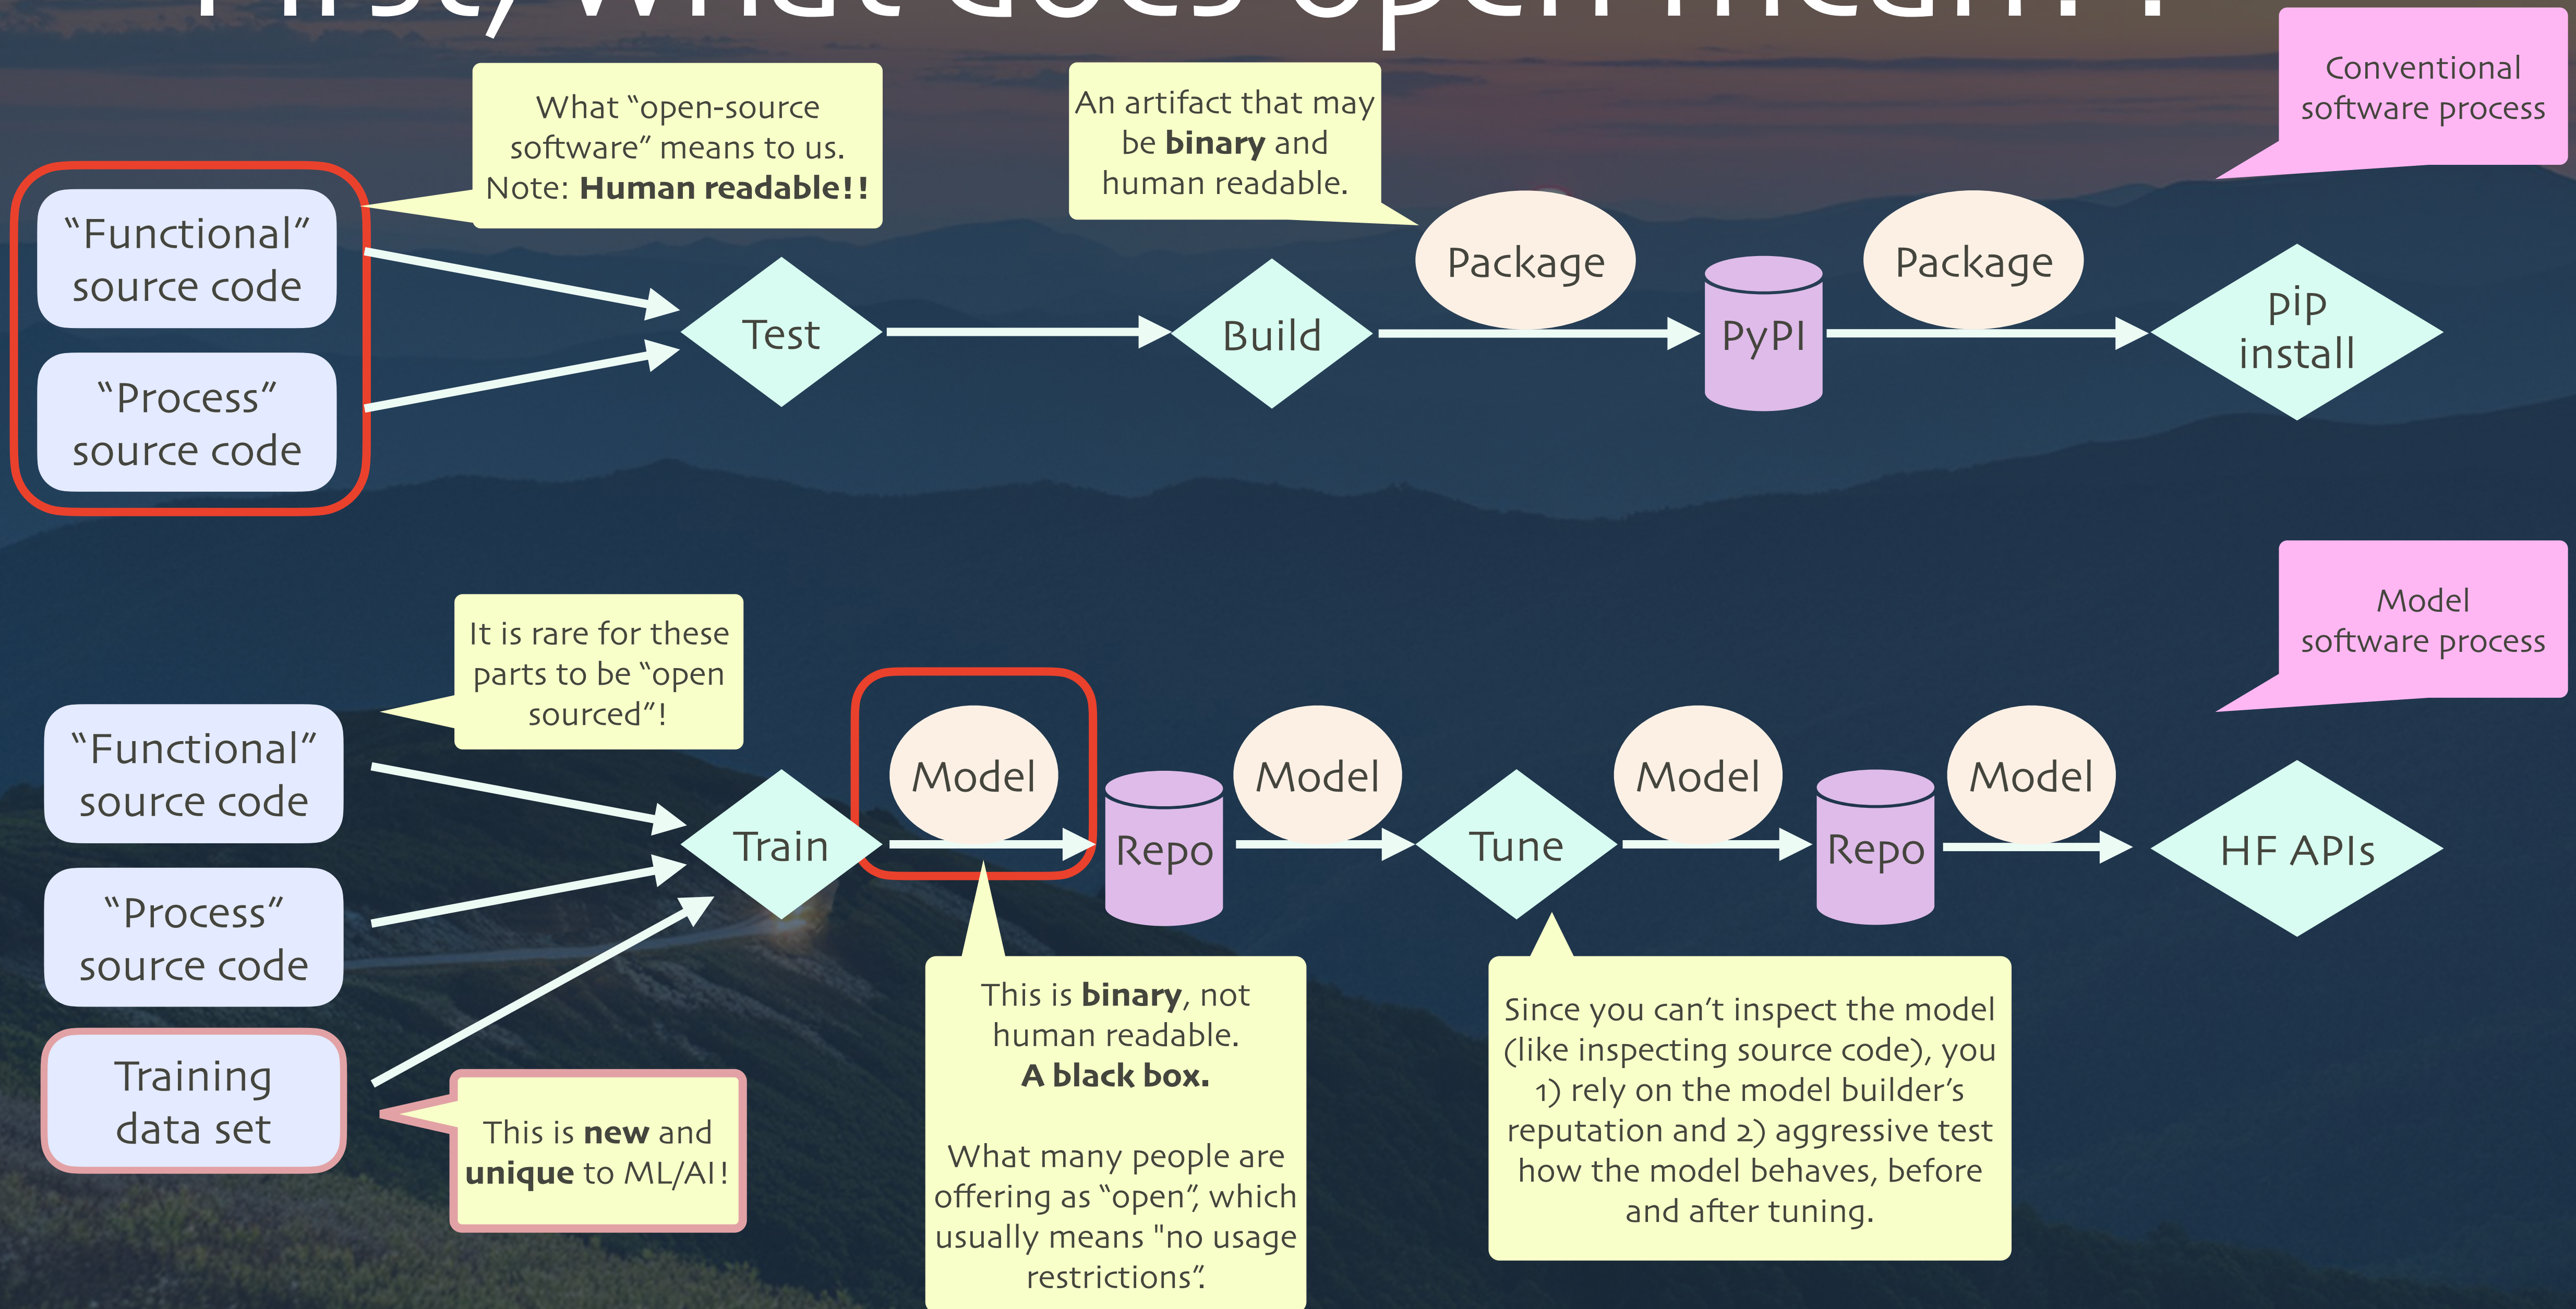


- 
- A misty forest with bare trees and a lone figure in the distance. The scene is dimly lit, with a soft, greyish light filtering through the branches. The trees are tall and thin, their silhouettes standing out against the fog. In the center of the path, a small, dark figure of a person is walking away from the viewer, adding a sense of solitude and mystery to the scene.
- Should AI Be “Open” or “Closed”?
 - The AI Alliance
 - “Alignment” Issues Blocking AI Adoption
 - Decentralization Solutions

A scenic landscape at sunset. The sky is a gradient of orange and red, with a thin layer of clouds. In the distance, a range of mountains is visible, with the sun setting behind one of the peaks. In the foreground, a winding road curves across a green hillside. A car is driving on the road, its headlights on, illuminating the path ahead. The overall mood is serene and contemplative.

Should AI be
“open” or “closed”?

First, what does open mean??



Why Open (as Much as Possible)?

- “The benefits can be summarized as autonomy, transparency, frictionless reuse, and collaborative improvement.”
- Although it is common to go from proprietary to open:
 - E.g., Windows/Unix/macOS → Linux

Do we want AI technology controlled by a few entities or more widely available?

The AI Alliance

A community of technology creators, developers and adopters collaborating to advance safe, responsible AI rooted in open innovation.

Diagram as of February.
>110 Now



The AI Alliance

Six Focus Areas:

1. Education, skills building, and research
2. Trust and safety
3. Tools for building models and applications
4. Hardware portability
5. Open models and datasets
6. Policy and regulations



Spreading knowledge, research

Technical initiatives

Maximize access, with safety



“Alignment” Issues Blocking AI Adoption

How can
decentralization help?

Alignment



Alignment

Alignment - Assuring that the model or AI application works as intended, i.e., that the results satisfy requirements for:

- Usefulness for user goals
- Secure
- Free of bias
- Free of objectionable speech and concepts
- Free of copyrighted material
- Factually correct, i.e., free of hallucinations

Alignment

Alignment - Assuring application works as results satisfy requirements

- Usefulness for users
- Secure
- Free of bias
- Free of objection
- Free of copyright
- Factually correct

Alignment is the hardest problem blocking broader adoption of Gen AI.

No successful production system today is just a chat interface and a "bare" model.

Decentralization

- Federated Learning:
 - Align models to local “requirements” (e.g., your face) while preserving privacy
 - Distribute the compute overhead
 - <https://arxiv.org/abs/2405.10853>

The screenshot shows the arXiv website interface. At the top, there is a search bar with the text 'Search...', a dropdown menu set to 'All fields', and a 'Search' button. Below the search bar, the breadcrumb navigation reads 'arXiv > cs > arXiv:2405.10853'. The main content area is titled 'Computer Science > Machine Learning' and features the paper title 'The Future of Large Language Model Pre-training is Federated' in bold. Below the title, the authors are listed: 'Lorenzo Sani, Alex Iacob, Zeyu Cao, Bill Marino, Yan Gao, Tomas Paulik, Wanru Zhao, William F. Shen, Preslav Aleksandrov, Xinchu Qiu, Nicholas D. Lane'. A short abstract begins with 'Generative pre-trained large language models (LLMs) have demonstrated impressive performance over a wide range of tasks, thanks to the unprecedented amount of data they have been trained on. As established scaling laws indicate, LLMs' future performance improvement depends on the amount of computing and...'. On the right side, there is a section titled 'Access Paper:' with links for 'View PDF', 'TeX Source', and 'Other Formats'. Below this, it shows the 'Current browse context: cs.LG' and navigation links like '< prev', 'next >', 'new', 'recent', and '2024-05'. At the bottom right, it says 'Change to browse by:'.

Decentralization


- Distributed Ledgers:
 - Secure the data supply chain for training, tuning, and inference

Decentralization

- Distributed Ledgers:
 - Secure the data supply chain for training, tuning, and inference

Governance of the supply chain:
While most AI users currently accept “open” models with opaque training data lineage, that will likely change.

What Actually Works?



Dare Obasanjo

@carnage4life@mas.to

There is a big difference between tech as augmentation versus automation. Augmentation (think Excel and accountants) benefits workers while automation (think traffic lights versus traffic wardens) benefits capital.

LLMs are controversial because the tech is best at augmentation but is being sold by lots of vendors as automation.

Jun 10, 2024, 10:31 · 🌐 · Ivory for iOS · ↻ 109 · ★ 188



<https://mas.to/@carnage4life/112593042823322764>

Emphasize Augmentation

- Keep humans in the loop, but improve productivity:
 - Distilling information more quickly
 - Translating between human and “domain languages”
 - SQL, Python, but also domain jargon (medical, finance, science, ...)
 - Use complementary tools
 - Use deterministic tools for factual accuracy, logical reasoning, and planning
 - What Agent frameworks enable

Emphasize Augmentation

***A.I. Needs Copper. It Just Helped
to Find Millions of Tons of It.***

An exploration site run by KoBold Metals in Chililabombwe, Zambia, in June. Zinyange Aunto
York Times

Train model with geological and mining data to predict where the copper is likely to be

The deposit, in Zambia, could make billions for Silicon Valley, provide minerals for the energy transition and help the United States in its rivalry with China.

<https://www.nytimes.com/2024/07/11/climate/kobold-zambia-copper-ai-mining.html>

Thank You!

- Visit thealliance.ai
- Let me know what you think!
 - dwampler@thealliance.ai
 - Mastodon and Bluesky: @deanwampler
 - Other talks: deanwampler.com/talks: use this →

deanwampler.com/talks

