The Evolution of AI in a Changing Climate

From Linear Regression to Foundation Models

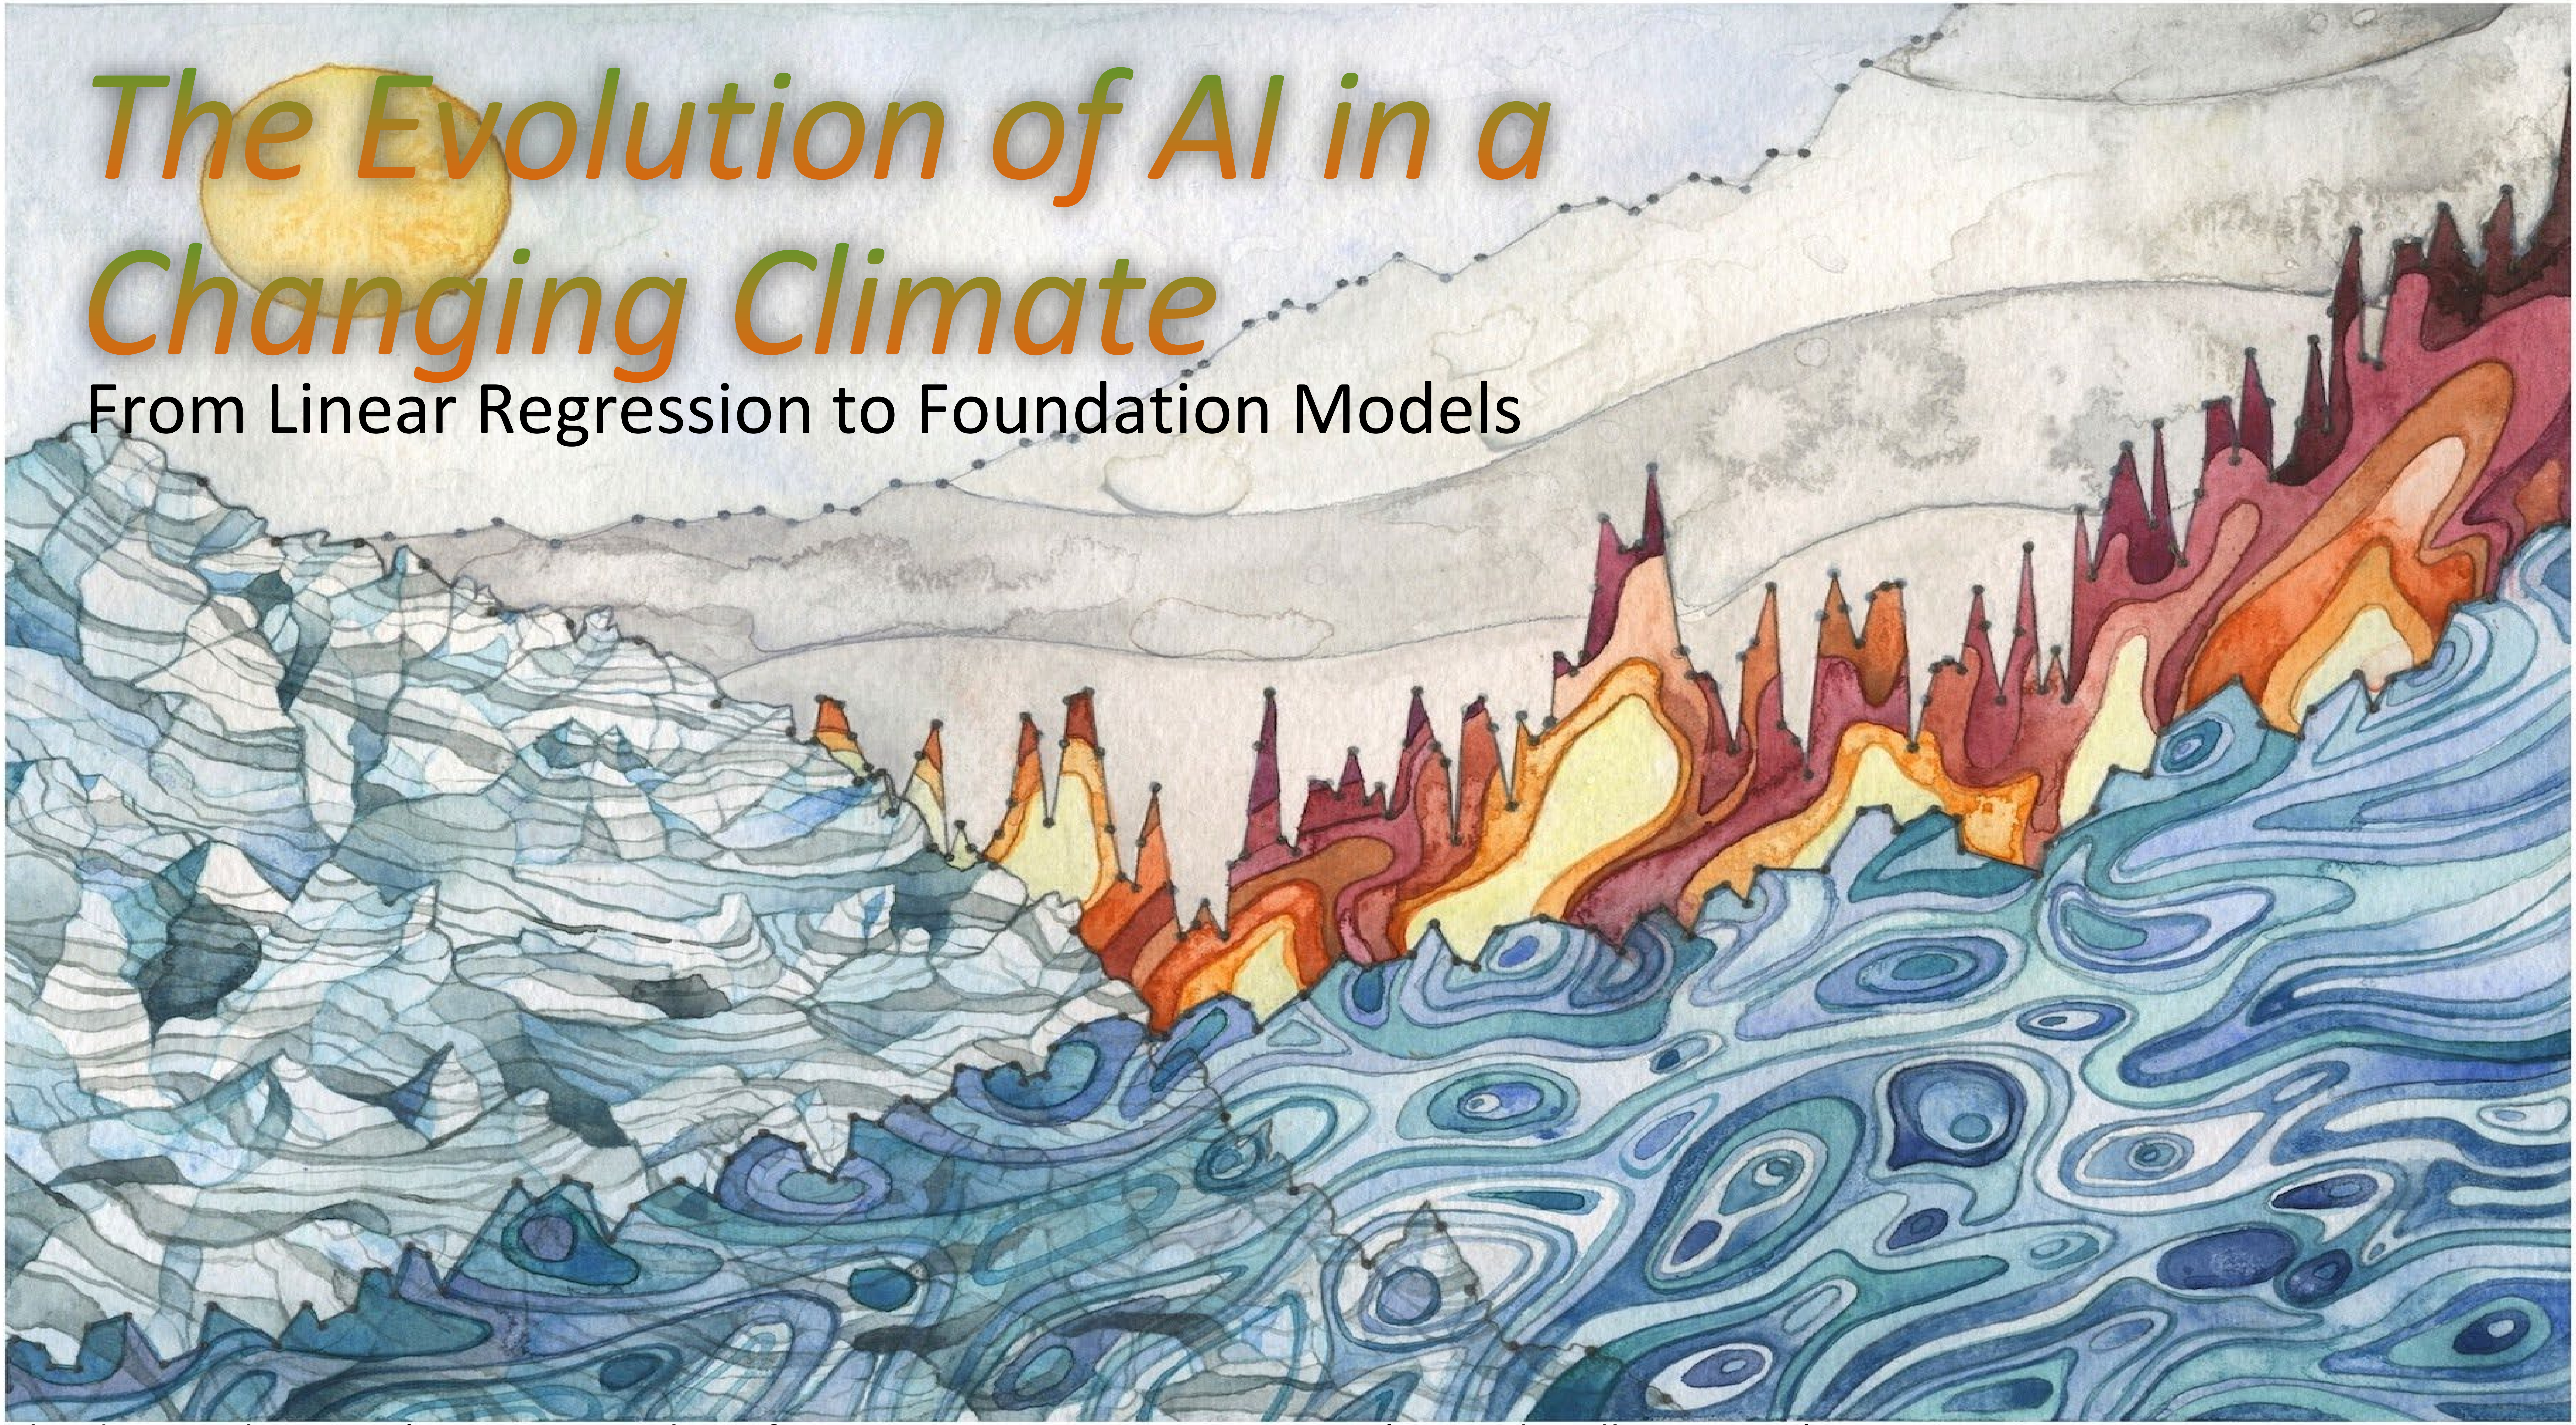

Claudio Battiloro, HEI’s 2025 Annual Conference, Austin, TX, May 4, 2025 (image by Jill Perot Art)

The Potential of AI+Climate+Health

Really?!?!?!

- Climate change brings more extreme weather, wildfires, and shifting disease patterns
  - Understanding and mitigating health impacts is *complicated* – e.g. heatwaves affecting vulnerable people, wildfire smoke causing respiratory illnesses
- **AI's Promise**: AI can analyze ***unprecedentedly massive multimodal*** data to find ***generalizable*** patterns and make predictions ***more accurately*** than traditional methods
  - This can inform early warnings and adaptive responses (e.g. alerting hospitals of an incoming heat-related patient surge)

Outline

1. AI's evolution in health
2. AI's evolution in climate science
3. Do we even need AI for climate adaptation?
4. The convergence – new AI models to jointly tackle climate *and* health challenges
    - **AI-Informed Healthy Climate Adaptation**
5. The dark side of AI and how we should deal with it
6. Reckoning with the future of technology, society, and climate

AI in Health

Early Days

- “AI for Health” began with basic statistical models
  - Linear and logistic regression have been used to predict outcomes (e.g. heart disease risk from cholesterol) ➟ **Simple, interpretable, linear, tabular data only, non-generalizable**

> N Engl J Med. 1985 Dec 26;313(26):1629-36. doi: 10.1056/NEJM198512263132604.

Simple linear regression in medical research

K Godfrey

PMID: 3840866 DOI: 10.1056/NEJM198512263132604

Abstract

This article discusses the method of fitting a straight line to data by linear regression and focuses on examples from 36 Original Articles published in the Journal in 1978 and 1979. Medical authors generally use linear regression to summarize the data (as in 12 of 36 articles in my survey) or to calculate the correlation between two variables (21 of 36 articles). Investigators need to become better acquainted with residual plots, which give insight into how well the fitted line models the data, and with confidence bounds for regression lines. Statistical computing packages enable investigators to use these techniques easily.

AI in Health

Middle Age

- “AI for health” moved to less basic statistical models ➟ Machine Learning
    - Random forests etc. have been used to predict to predict patient trajectories ➟ **Simple, interpretable, non-linear, mostly tabular/temporal data, non-generalizable**

› Sci Transl Med. 2010 Sep 8;2(48):48ra65. doi: 10.1126/scitranslmed.3001304.

Integration of early physiological responses predicts later illness severity in preterm infants

Suchi Saria [1], Anand K Rajani, Jeffrey Gould, Daphne Koller, Anna A Penn

Affiliations + expand

PMID: 20826840 PMCID: PMC3564961 DOI: 10.1126/scitranslmed.3001304

Abstract

Physiological data are routinely recorded in intensive care, but their use for rapid assessment of illness severity or long-term morbidity prediction has been limited. We developed a physiological assessment score for preterm newborns, akin to an electronic Apgar score, based on standard signals recorded noninvasively on admission to a neonatal intensive care unit. We were able to accurately and reliably estimate the probability of an individual preterm infant's risk of severe morbidity on the basis of noninvasive measurements. This prediction algorithm was developed with electronically captured physiological time series data from the first 3 hours of life in preterm infants (< or =34 weeks gestation, birth weight < or =2000 g). Extraction and integration of the data with state-of-the-art machine learning methods produced a probability score for illness severity, the PhysiScore. PhysiScore was validated on 138 infants with the leave-one-out method to prospectively identify infants at risk of short- and long-term morbidity. PhysiScore provided higher

AI in Health

Modern Times

nature > npj digital medicine > articles > article

Article | Open access | Published: 08 May 2018

Scalable and accurate deep learning with electronic health records

Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M. Dai, Nissan Hajaj, Michaela Hardt, Peter J. Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, Patrik Sundberg, Hector Yee, Kun Zhang, Yi Zhang, Gerardo …e, Kurt Litsch, Alexander Mossin, Justin Tansuwan, De… …Dean + Show authors

…18) | Cite this article

…n record (EHR) data is anticipated to drive …hcare quality. Constructing predictive statistical …rated predictor variables from normalized EHR data, …e vast majority of information in each patient's …atients' entire raw EHR records based on the Fast …HIR) format. We demonstrate that deep learning … multiple medical … validated our

NeurIPS Proceedings Search

MiME: Multilevel Medical Embedding of Electronic Health Records for Predictive Healthcare

Part of Advances in Neural Information Processing Systems 31 (NeurIPS 2018)

…electronic health records (EHR) data, but these models typically require training data volume that …sed to bridge the data volume constraint, but this approach is often not directly applicable or …ge the inherent multilevel structure of EHR data and, in particular, the encoded relationships …bedding of EHR data while jointly performing auxiliary prediction tasks that rely on this inherent …diction and sequential disease prediction, where MiME outperformed baseline methods in diverse …ilure … …ng its

> BMC Med Res Methodol. 2023 Dec 7;23(1):285. doi: 10.1186/s12874-023-02112-2.

Attention-based neural networks for clinical prediction modelling on electronic health records

Egill A Fridgeirsson [1], David Sontag [2], Peter Rijnbeek [3]

Affiliations + expand

PMID: 38062352 PMCID: PMC10701944 DOI: 10.1186/s12874-023-02112-2

Abstract

Background: Deep learning models have had a lot of success in various fields. However, on structured data they have struggled. Here we apply four state-of-the-art supervised deep learning models using the attention mechanism and compare against logistic regression and XGBoost using discrimination, calibration and clinical utility.

Methods: We develop the models using a general practitioners database. We implement a recurrent neural network, a transformer with and without reverse distillation and a graph neural network. We measure discrimination using the area under the receiver operating characteristic curve (AUC) and the area under the precision recall curve (AUPRC). We assess smooth calibration using restricted cubic splines and clinical utility with decision curve analysis.

Periodicals > Journal of Biomedical Informatics > Vol. 144, No. C > Deep learning prediction models based on EHR trajectories: A systematic review

REVIEW-ARTICLE

Deep learning prediction models based on EHR trajectories: : A systematic review

Authors: Ali Amirahmadi, Mattias Ohlsson, Kobra Etminani Authors Info & Claims

Volume 144, Issue C • https://doi.org/10.1016/j.jbi.2023.104430

Published: 01 August 2023 Publication History

AI in Health

Contemporary Times

HARVARD
MEDICAL SCHOOL

Education & Admissions | Research | News & Events | About HMS

News & Research

News Topics

A New Artificial Intelligence Tool for

ation mo

nt choi

once pa
eases, h
sease m
s in mac
nmet ne
electror
conditi
on individuals, BEHRT shows a striking
e precision scores for different tasks), over the
existing state-of-the-art deep EHR models. In addition to its scalability and superior accuracy,

Stanford University
Human-Centered
Artificial Intelligence

About | Research | Education | Policy | AI Index

RESEARCH

A Multi-Center Study on the Adaptability of a Shared Foundation Model for Electronic Health Records

Article | Open access | Published: 26 December 2022

A large language model for electronic health records

Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E. Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Anthony B. Costa, Mona G. Flores, Ying Zhang, Tanja Magoc, Christopher A. Harle, Gloria Lipori, Duane A. Mitchell, William R. Hogan, Elizabeth A. Shenkman, Jiang Bian & Yonghui Wu

npj Digital Medicine 5, Article number: 194 (2022) | Cite this article

rocess and
wered by
clinical
rained in the
llions of
els with
is study, we
n words of
valuate it on
tion,
answering
up the size

Health AI Developer Foundations Overview

Developing AI for health applications can be challenging for several reasons:

- data availability is limited;
- modelling experimentation requires a lot of compute
- expert labeling is expensive; and
- the number of potential applications to tackle is vast.

Health AI Developer Foundations (HAI-DEF) turbocharges AI development for healthcare applications with open weight models specialized for health domains. These models can be adapted for the use-cases you care about requiring less data and compute than traditional approaches. HAI-DEF includes rich examples of how to use the models in the form of Google Colab Notebooks. It also includes libraries to streamline ingestion of processing of specialized medical data like 3D imaging and digital pathology. HAI-DEF is tightly integrated into Google Cloud so developers can take ideas from research through to productionization on the same platform. HAI-DEF is an evolving effort; we will continue to add new models and resources to cover more healthcare applications.

AI in Climate Science

It's ~~raining~~ burning cats and dogs

AI in Climate

The Classical Approach

- "~~AI for climate~~" is often not AI ➟ Numerical Weather Prediction (NWP)
  - Physics-Based Models ➟ **Reliable, refined, model mismatch, scarce data assimilation**

The quiet revolution of numerical weather prediction

Peter Bauer, Alan Thorpe & Gilbert Brunet

Nature **525**, 47–55 (2015) | Cite this article

Abstract

Advances in numerical weather prediction represent a quiet revolution because they have resulted from a steady accumulation of scientific knowledge and technological advances over many years that, with only a few exceptions, have not been associated with the aura of fundamental physics breakthroughs. Nonetheless, the impact of numerical weather prediction is among the greatest of any area of physical science. As a computational problem, global weather prediction is comparable to the simulation of the human brain and of the evolution of the early Universe, and it is performed every day at major operational centres across the world.

AI in Climate

The GenAI approach

- "AI for climate" is based on foundation models ➟ Generative AI
  - Atmosphere Foundation Models ➟ **Reliable, refined, mitigate model mismatch, huge data ingestion, huge**

Microsoft Research Blog

ducing Aurora: The first large-
foundation model of the
sphere

Google DeepMind About Research Technologies Discover

RESEARCH

GraphCast: AI model for faster and more accurate global weather forecasting

14 NOVEMBER 2023

The potential of AI+Climate+Health

- Climate change brings more extreme weather, wildfires, and shifting disease patterns
  - Understanding and mitigating health impacts is *complicated* – e.g. heatwaves affecting vulnerable people, wildfire smoke causing respiratory illnesses
- **AI's Promise**: AI can analyze ***unprecedentedly massive multimodal*** data to find ***generalizable*** patterns and make predictions ***more accurately*** than traditional methods
  - This can inform early warnings and adaptive responses (e.g. alerting hospitals of an incoming heat-related patient surge)

Where we could go

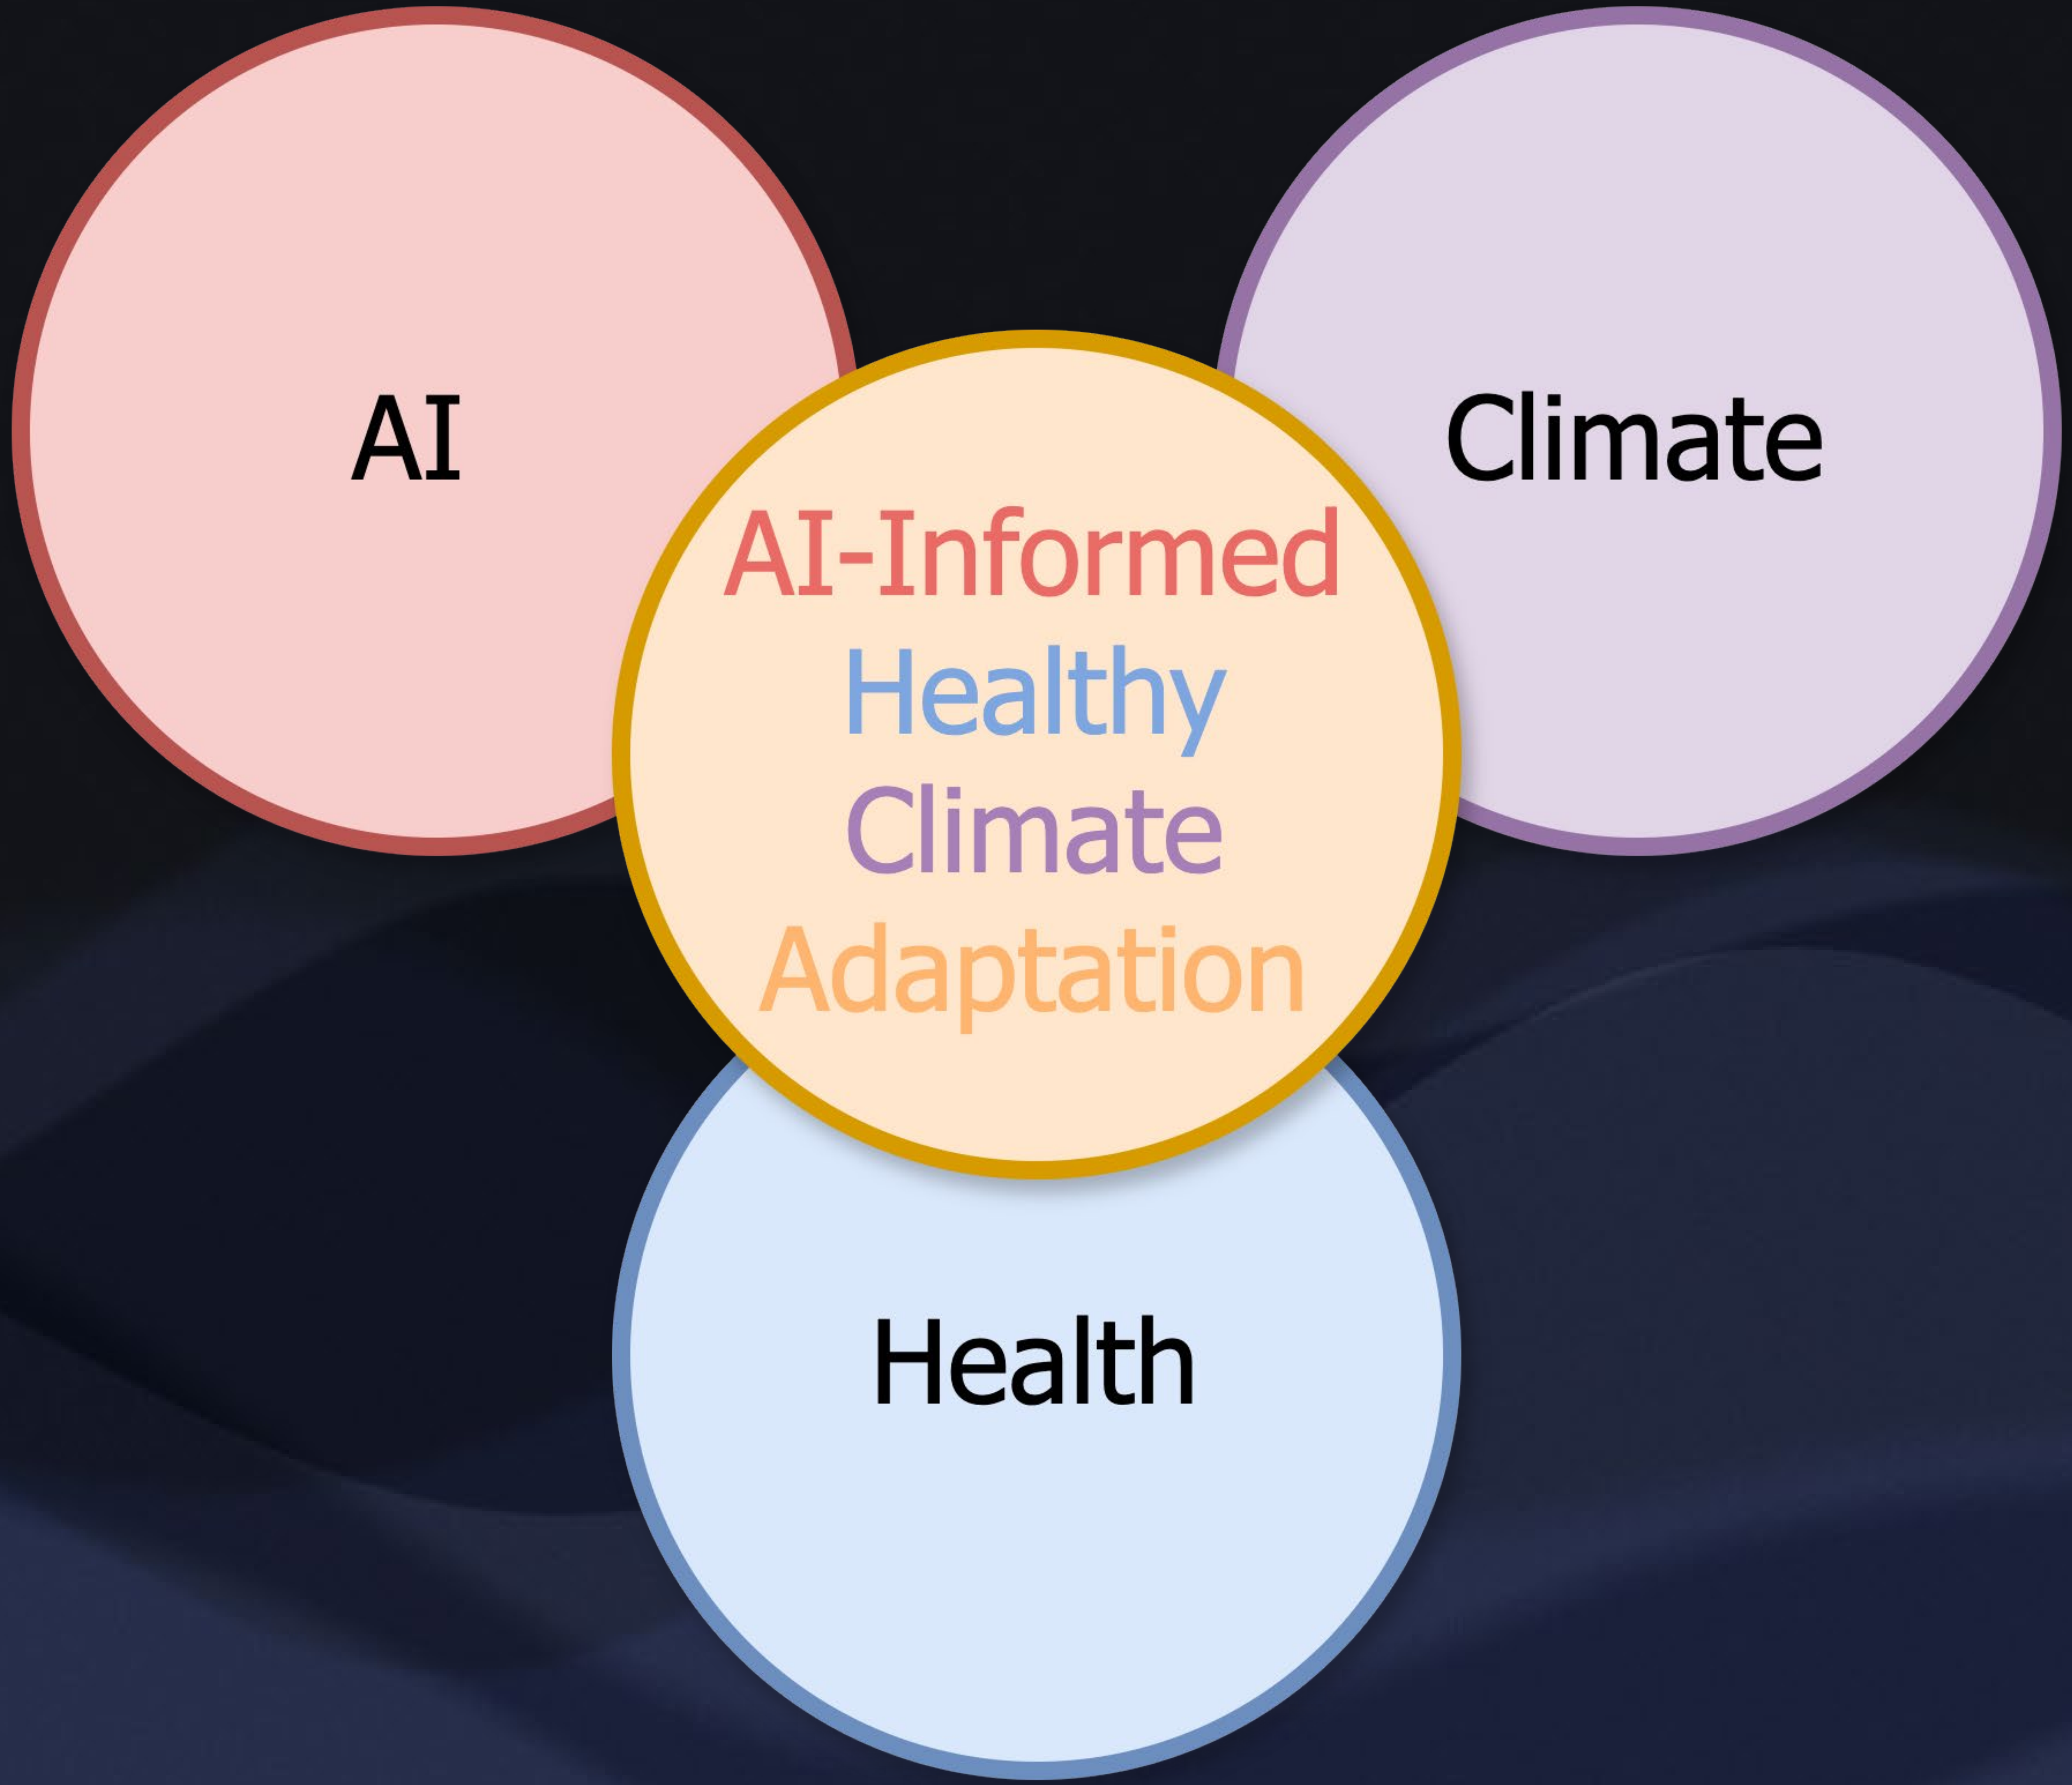

Where we are

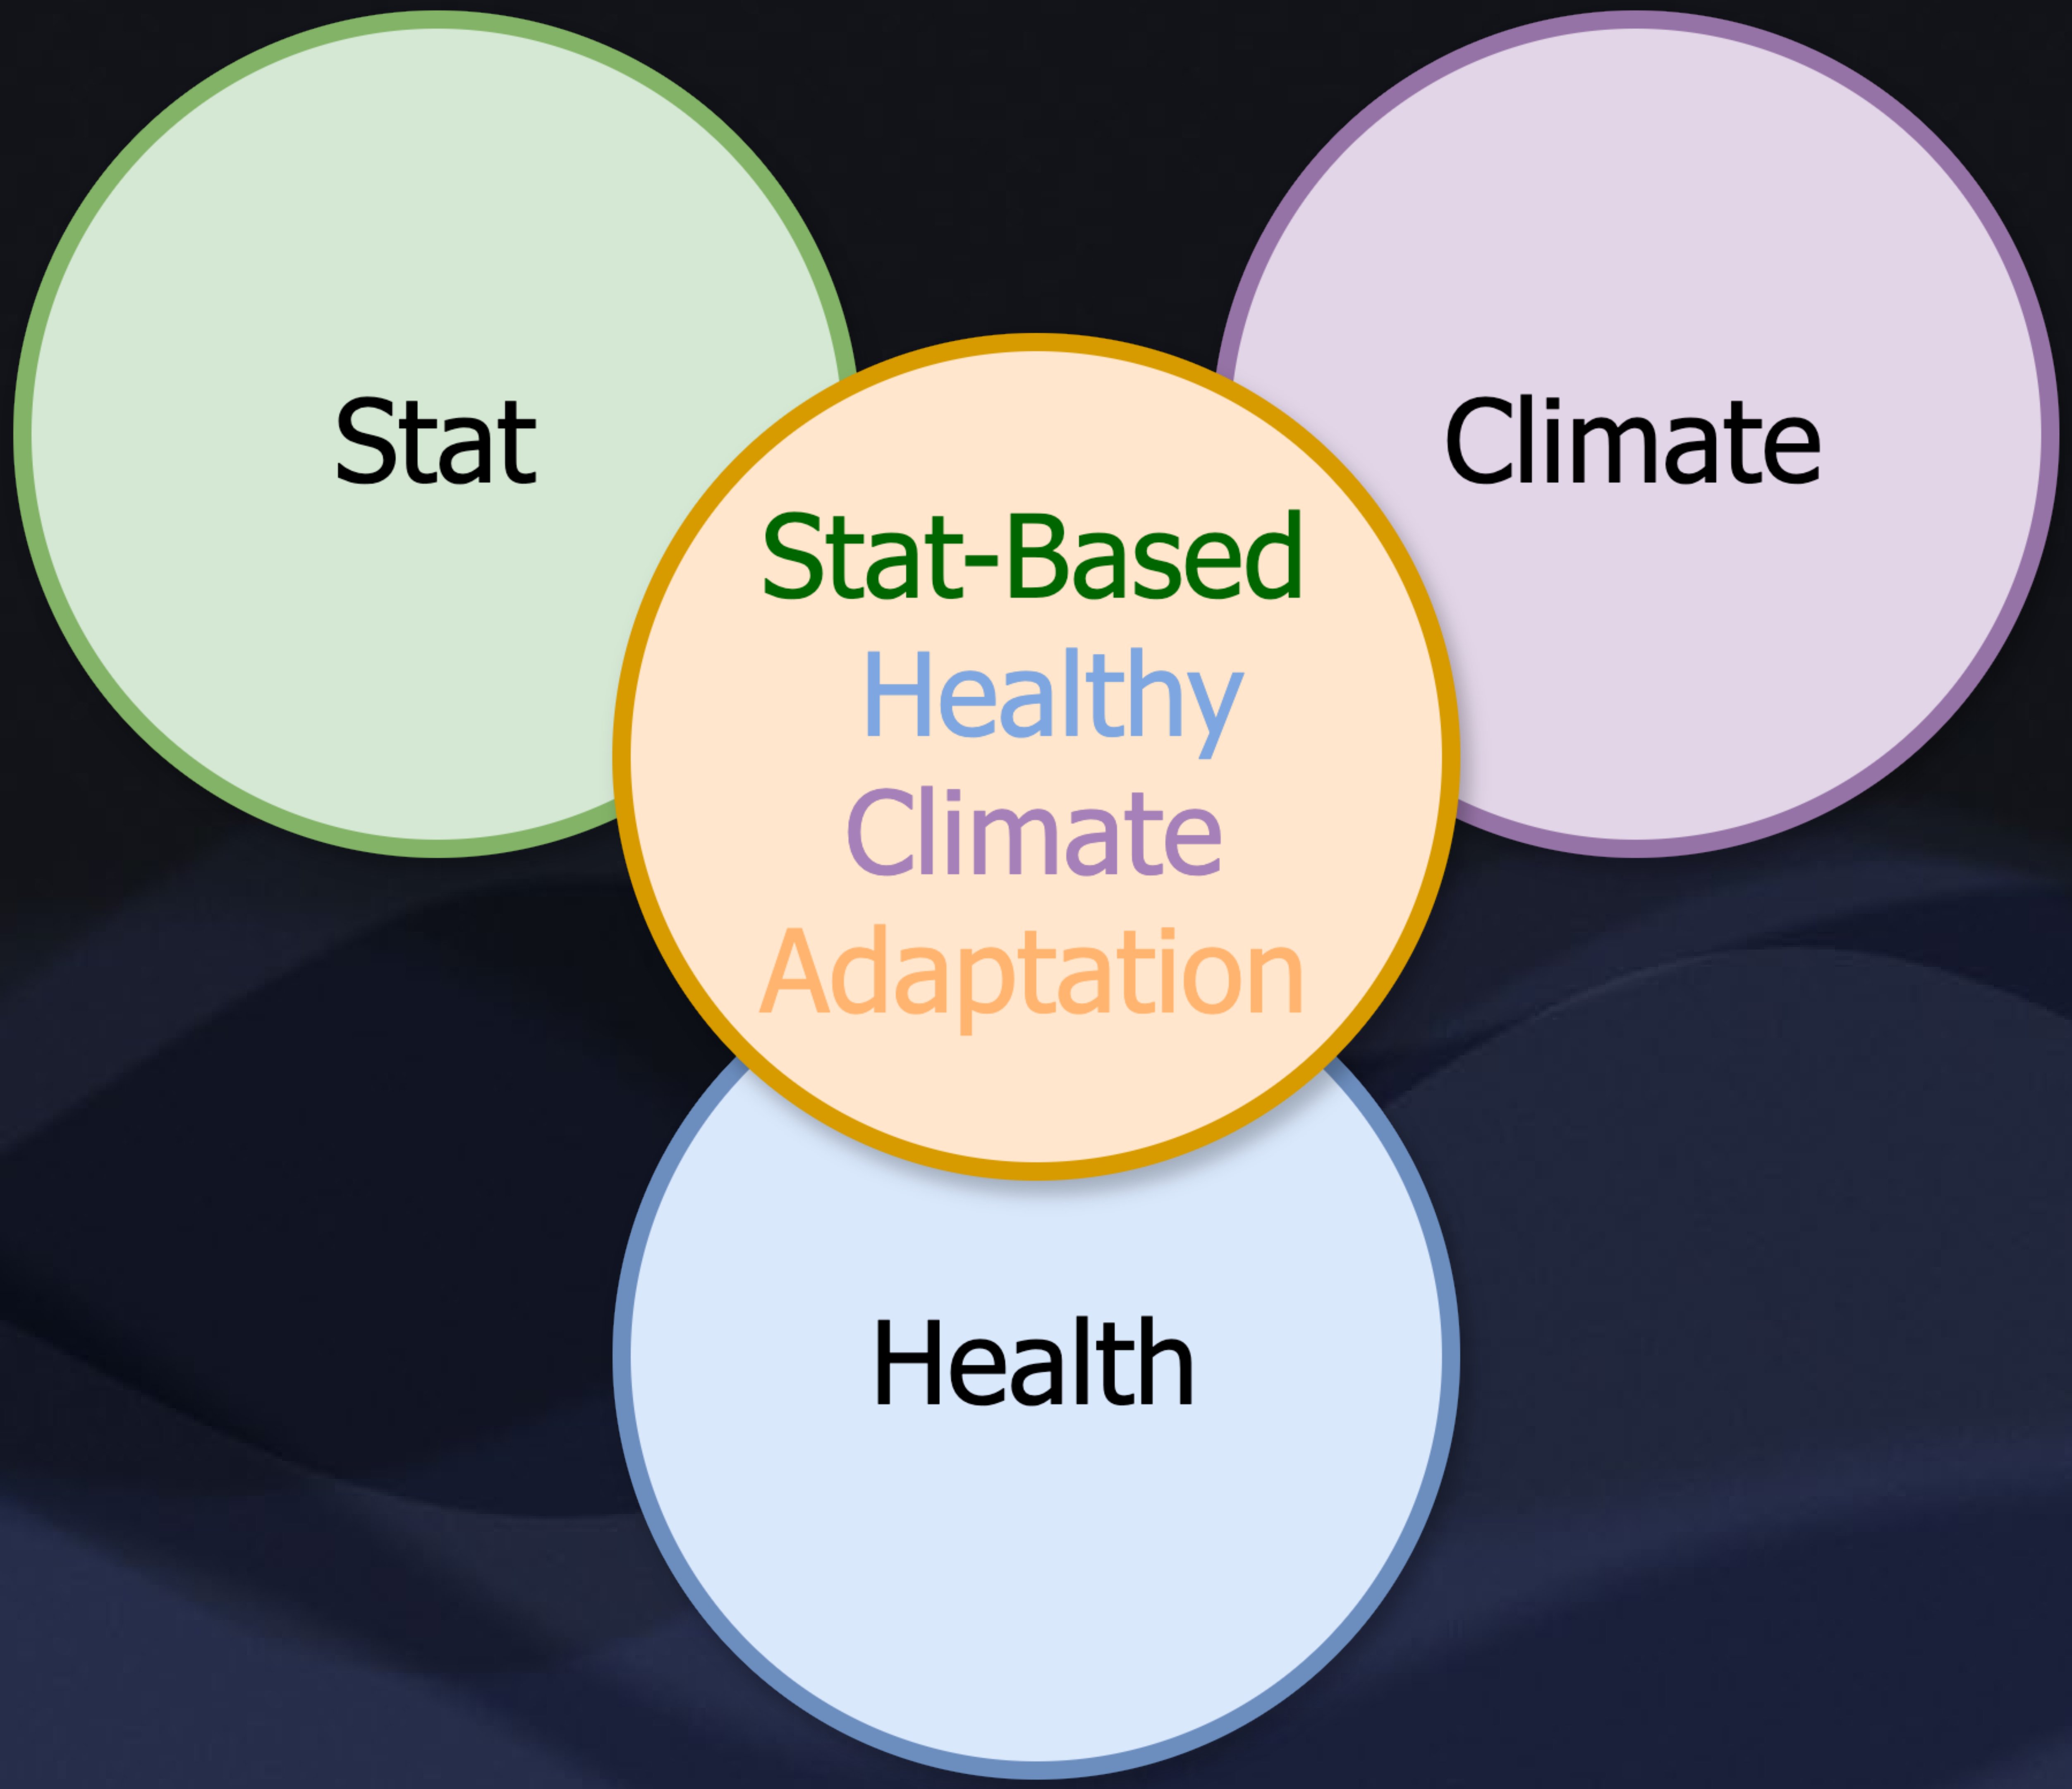
Stat
Climate
Stat-Based
Healthy
Climate
Adaptation
Health

Where we are

It works extremely well!

SCIENCE ADVANCES | RES

HEALTH AND MEDICINE

Evaluating the impa particulate matter c

X. Wu[1]*, D. Braun[1,2]*, J. Schwartz[3],

Many studies link long-term fine particl standards (12 micrograms per cubic me tional statistical approaches does not enrollees—we provide strong evidence

SPECIAL SECTION

REVIEW

Policy-induced Statistical and

Daniel Mork[1], Scott Delane

Air pollution causes premat lower-income groups in the in air pollution exposure an policies, air pollution, and exposures, socioeconomic estimands accounting for health outcomes; and (iii) effects on disparities to id new data science approach exposure and health outco

Biden Administration Moves to Tighten Limits on Deadly Air Pollution

A new rule would, for the first time in a decade, reduce emissions of soot that disproportionately harm communities of color.

Give this article | 95

NGLAND MEDICINE

29, 2017 VOL. 376 NO. 26

in the Medicare Population

, Ph.D., Yun Wang, Ph.D., Petros Koutrakis, Ph.D., nici, Ph.D., and Joel D. Schwartz, Ph.D.

- Is exposure to PM2 … mortality risk?
- …

Optimization in the Artificial ice Era

[Sc*1, Falco J. Bargagli-Stoffi, PhD‡2,1, and iinici, PhD[1]

RESEARCH ARTICLE | ENVIRONMENTAL SCIENCES

Childhood $PM_{2.5}$ exp mobility in the Unite

Sophie-An Kingsbury Lee, Luca Merlo, and France

Edited by Christopher Barrett, Cornell University, Ithaca, NY; re

September 9, 2024 | 121 (38) e2401882121 | https:/

THIS ARTICLE IS DISCUSSED ON A PODCAST

Where we could go

Climate Adaptation through the Lens of Foundation Models (FMs)

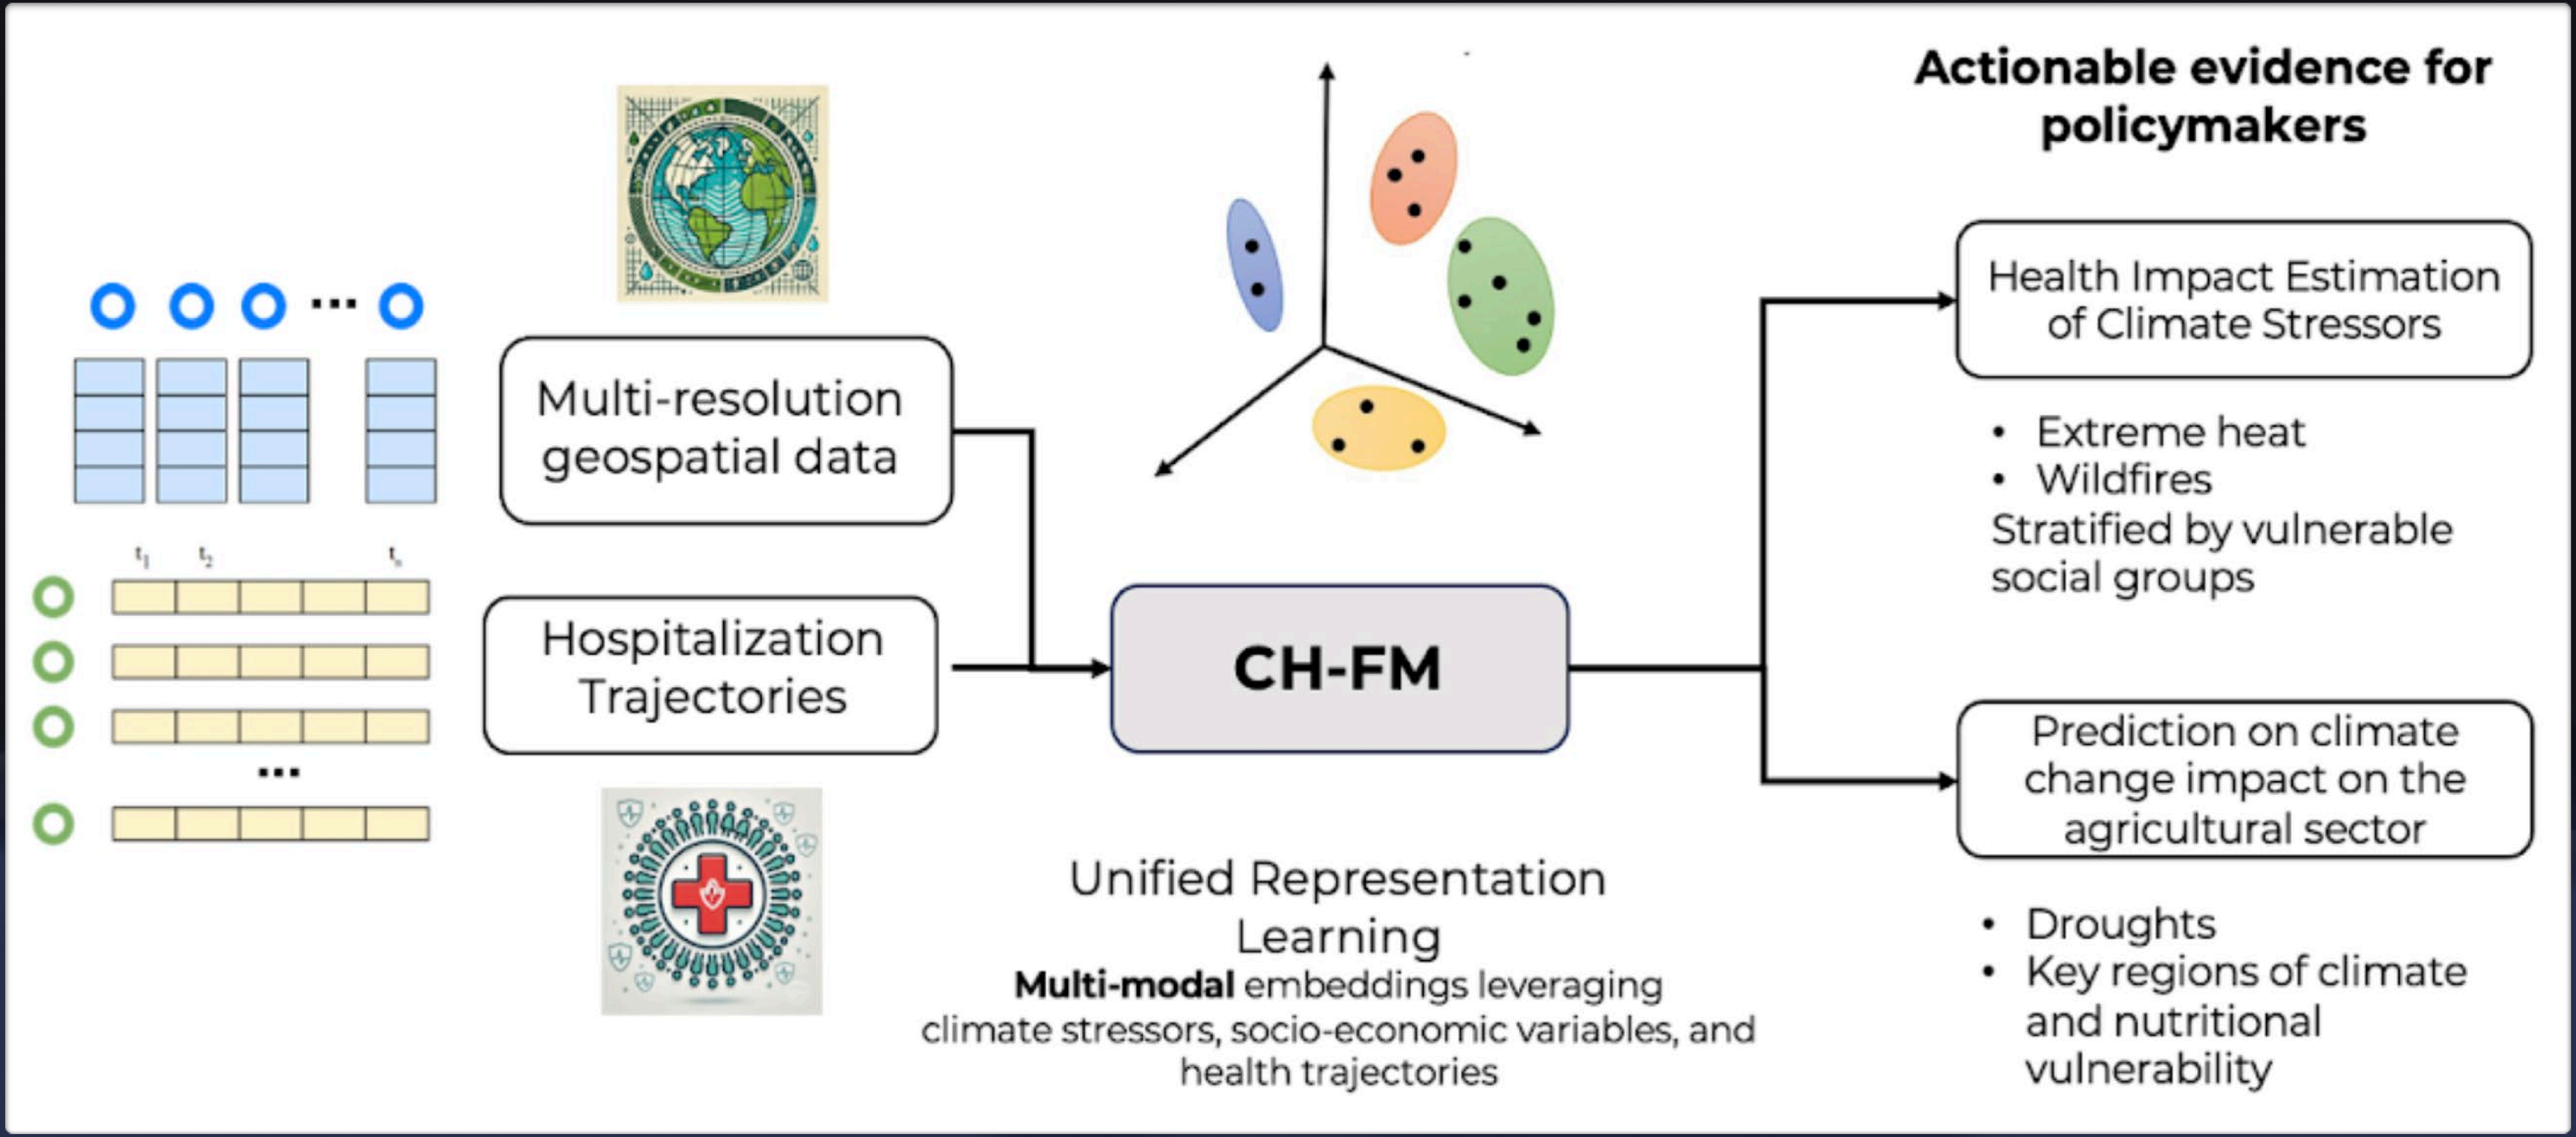

Adaptation without health?

- **All these efforts to suggest healthy adaptation strategies on the base of physical activity and aggregated vulnerability indices?!?!?!?**
  - **We need Climate and Health Foundation Models to really enable AI-Informed Healthy Climate Adaptation!**

The Dark Side of the Moon

Environmental Burden of United States Data Centers in the Artificial Intelligence Era

Gianluca Guidi, Francesca Dominici, Jonathan Gilmour, Kevin Butler, Eric Bell, Scott Delaney, Falco J. Bargagli-Stoffi

The rapid proliferation of data centers in the US – driven partly by the adoption of artificial intelligence – has set off alarm bells about the industry's environmental impact. We compiled detailed information on 2,132 US data centers operating between September 2023 and August 2024 and determined their electricity consumption, electricity sources, and attributable CO_2e emissions. Our findings reveal that data centers accounted for more than 4% of total US electricity consumption – with 56% derived from fossil fuels – generating more than 105 million tons of CO_2e (2.18% of US emissions in 2023). Data centers' carbon intensity – the amount of CO_2e emitted per unit of electricity consumed – exceeded the US average by 48%. Our data pipeline and visualization tools can be used to assess current and future environmental impacts of data centers.

The Unpaid Toll: Quantifying the Public Health Impact of AI

Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, Shaolei Ren

The surging demand for AI has led to a rapid expansion of energy-intensive data centers, impacting the environment through escalating carbon emissions and water consumption. While significant attention has been paid to AI's growing environmental footprint, the public health burden, a hidden toll of AI, has been largely overlooked. Specifically, AI's lifecycle, from chip manufacturing to data center operation, significantly degrades air quality through emissions of criteria air pollutants such as fine particulate matter, substantially impacting public health. This paper introduces a methodology to model pollutant emissions across AI's lifecycle, quantifying the public health impacts. Our findings reveal that training an AI model of the Llama3.1 scale can produce air pollutants equivalent to more than 10,000 round trips by car between Los Angeles and New York City. The total public health burden of U.S. data centers in 2030 is valued at up to more than $20 billion per year, double that of U.S. coal-based steelmaking and comparable to that of on-road emissions of California. Further, the public health costs unevenly impact economically disadvantaged communities, where the per-household health burden could be 200x more than that in less-impacted communities. We recommend adopting a standard reporting protocol for criteria air pollutants and the public health costs of AI, paying attention to all impacted communities, and implementing health-informed AI to mitigate adverse effects while promoting public health equity.

AI models collapse when trained on recursively generated data

Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson & Yarin Gal

Nature **631**, 755–759 (2024) | Cite this article

An Author Correction to this article was published on 21 March 2025

This article has been updated

Abstract

Stable diffusion revolutionized image creation from descriptive text. GPT-2 (ref. [1]), GPT-3(.5) (ref. [2]) and GPT-4 (ref. [3]) demonstrated high performance across a variety of language tasks. ChatGPT introduced such language models to the public. It is now clear that generative artificial intelligence (AI) such as large language models (LLMs) is here to stay and will substantially change the ecosystem of online text and images. Here we consider what may happen to GPT-{n} once LLMs contribute much of the text found online. We find that indiscriminate use of model-generated content in training causes irreversible defects in the resulting models, in which tails of the original content distribution disappear. We refer to this

The shaky foundations of large language models and foundation models for electronic health records

Michael Wornow, Yizhe Xu, Rahul Thapa, Birju Patel, Ethan Steinberg, Scott Fleming, Michael A. Pfeffer, Jason Fries & Nigam H. Shah

npj Digital Medicine **6**, Article number: 135 (2023) | Cite this article

Abstract

The success of foundation models such as ChatGPT and AlphaFold has spurred significant interest in building similar models for electronic medical records (EMRs) to improve patient care and hospital operations. However, recent hype has obscured critical gaps in our understanding of these models' capabilities. In this narrative review, we examine 84 foundation models trained on non-imaging EMR data (i.e., clinical text and/or structured data) and create a taxonomy delineating their architectures, training data, and potential use cases. We find that most models are trained on small, narrowly-scoped clinical datasets (e.g., MIMIC-III) or broad, public biomedical corpora (e.g., PubMed) and are evaluated on tasks that do not provide meaningful insights on their usefulness to health systems. Considering these findings, we propose an improved evaluation framework for measuring the benefits of clinical foundation models that is more closely grounded to metrics that matter in healthcare.

The Dark Side of the Moon

Reconstructing Training Data From Trained Neura

Niv Haim, Gal Vardi, Gilad Yehudai, Ohad Shamir, michal Irani

Published: 31 Oct 2022, Last Modified: 06 Apr 2025 NeurIPS 2022 Accept Readers: Everyone Show Bibtex Show Revisions

Keywords: implicit bias, dataset reconstruction, privacy attacks

TL;DR: We provide a novel scheme for reconstructing large portions of the actual training samples from a trained neural network. Our scher
in training neural networks.

Abstract: Understanding to what extent neural network
In this paper we show that in some cases a significant fra
We propose a novel reconstruction scheme that stems fr
To the best of our knowledge, our results are the first to
This has negative implications on privacy, as it can be us
We demonstrate our method for binary MLP classifiers o

Supplementary Material: pdf

Community Implementations: 1 code implementat

The Exploited Labor Behind Artificial Intelligence

cs and the promise
eral intelligence
orres

lligence (AI) is to develop artificial
ence than anything we have ever seen.
uld be built, researchers are working to
e that, unlike systems with specific
ng principles, undefined systems like
lding AGI often framed as an
he normative framework that motivates
ion of the twentieth century. As a result,
enicists in the past (*e.g.*, racism,
thin the movement to build AGI,
power, while using the language of
"safety" and "benefiting humanity" to evade accountability. We conclude by urging researchers to work on defined tasks for which we can develop safety protocols, rather than attempting to build a presumably all-knowing system such as AGI.

BUSINESS • TECH

Exclusive: OpenAI Used Less Than $2 Per Hour to Toxi

15 MINUTE R

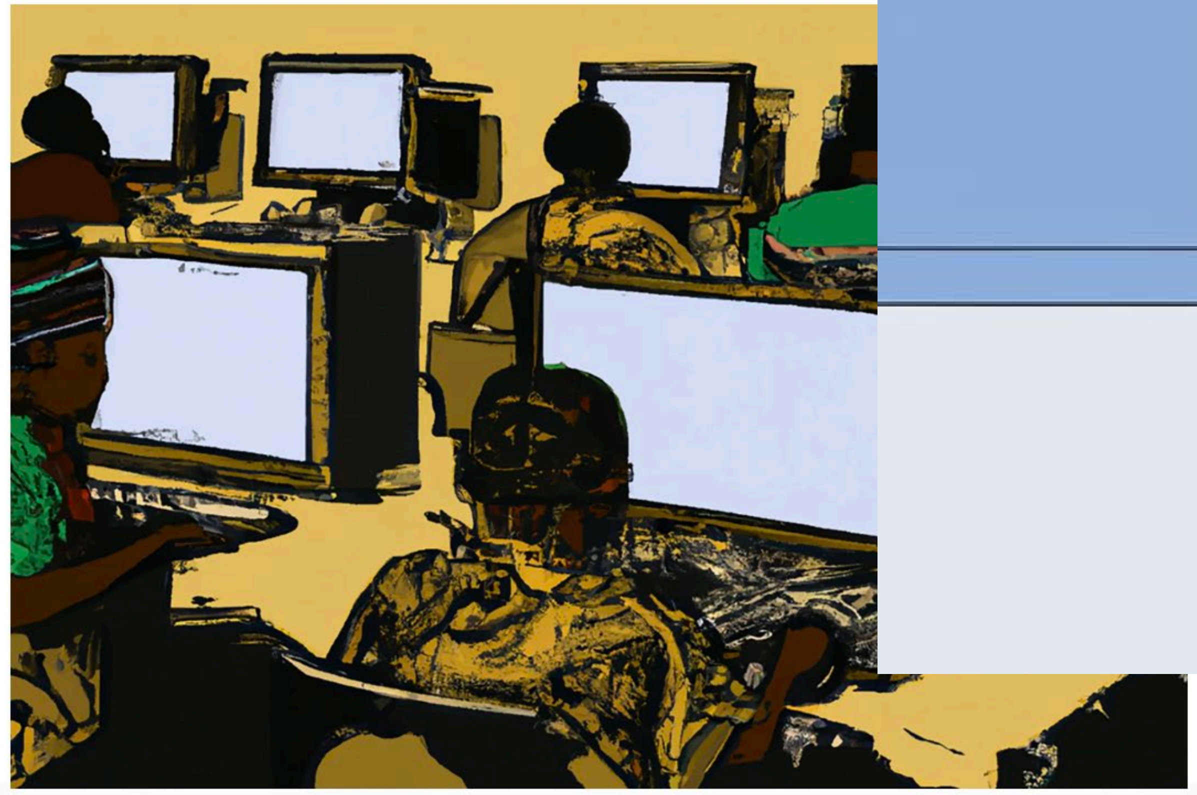

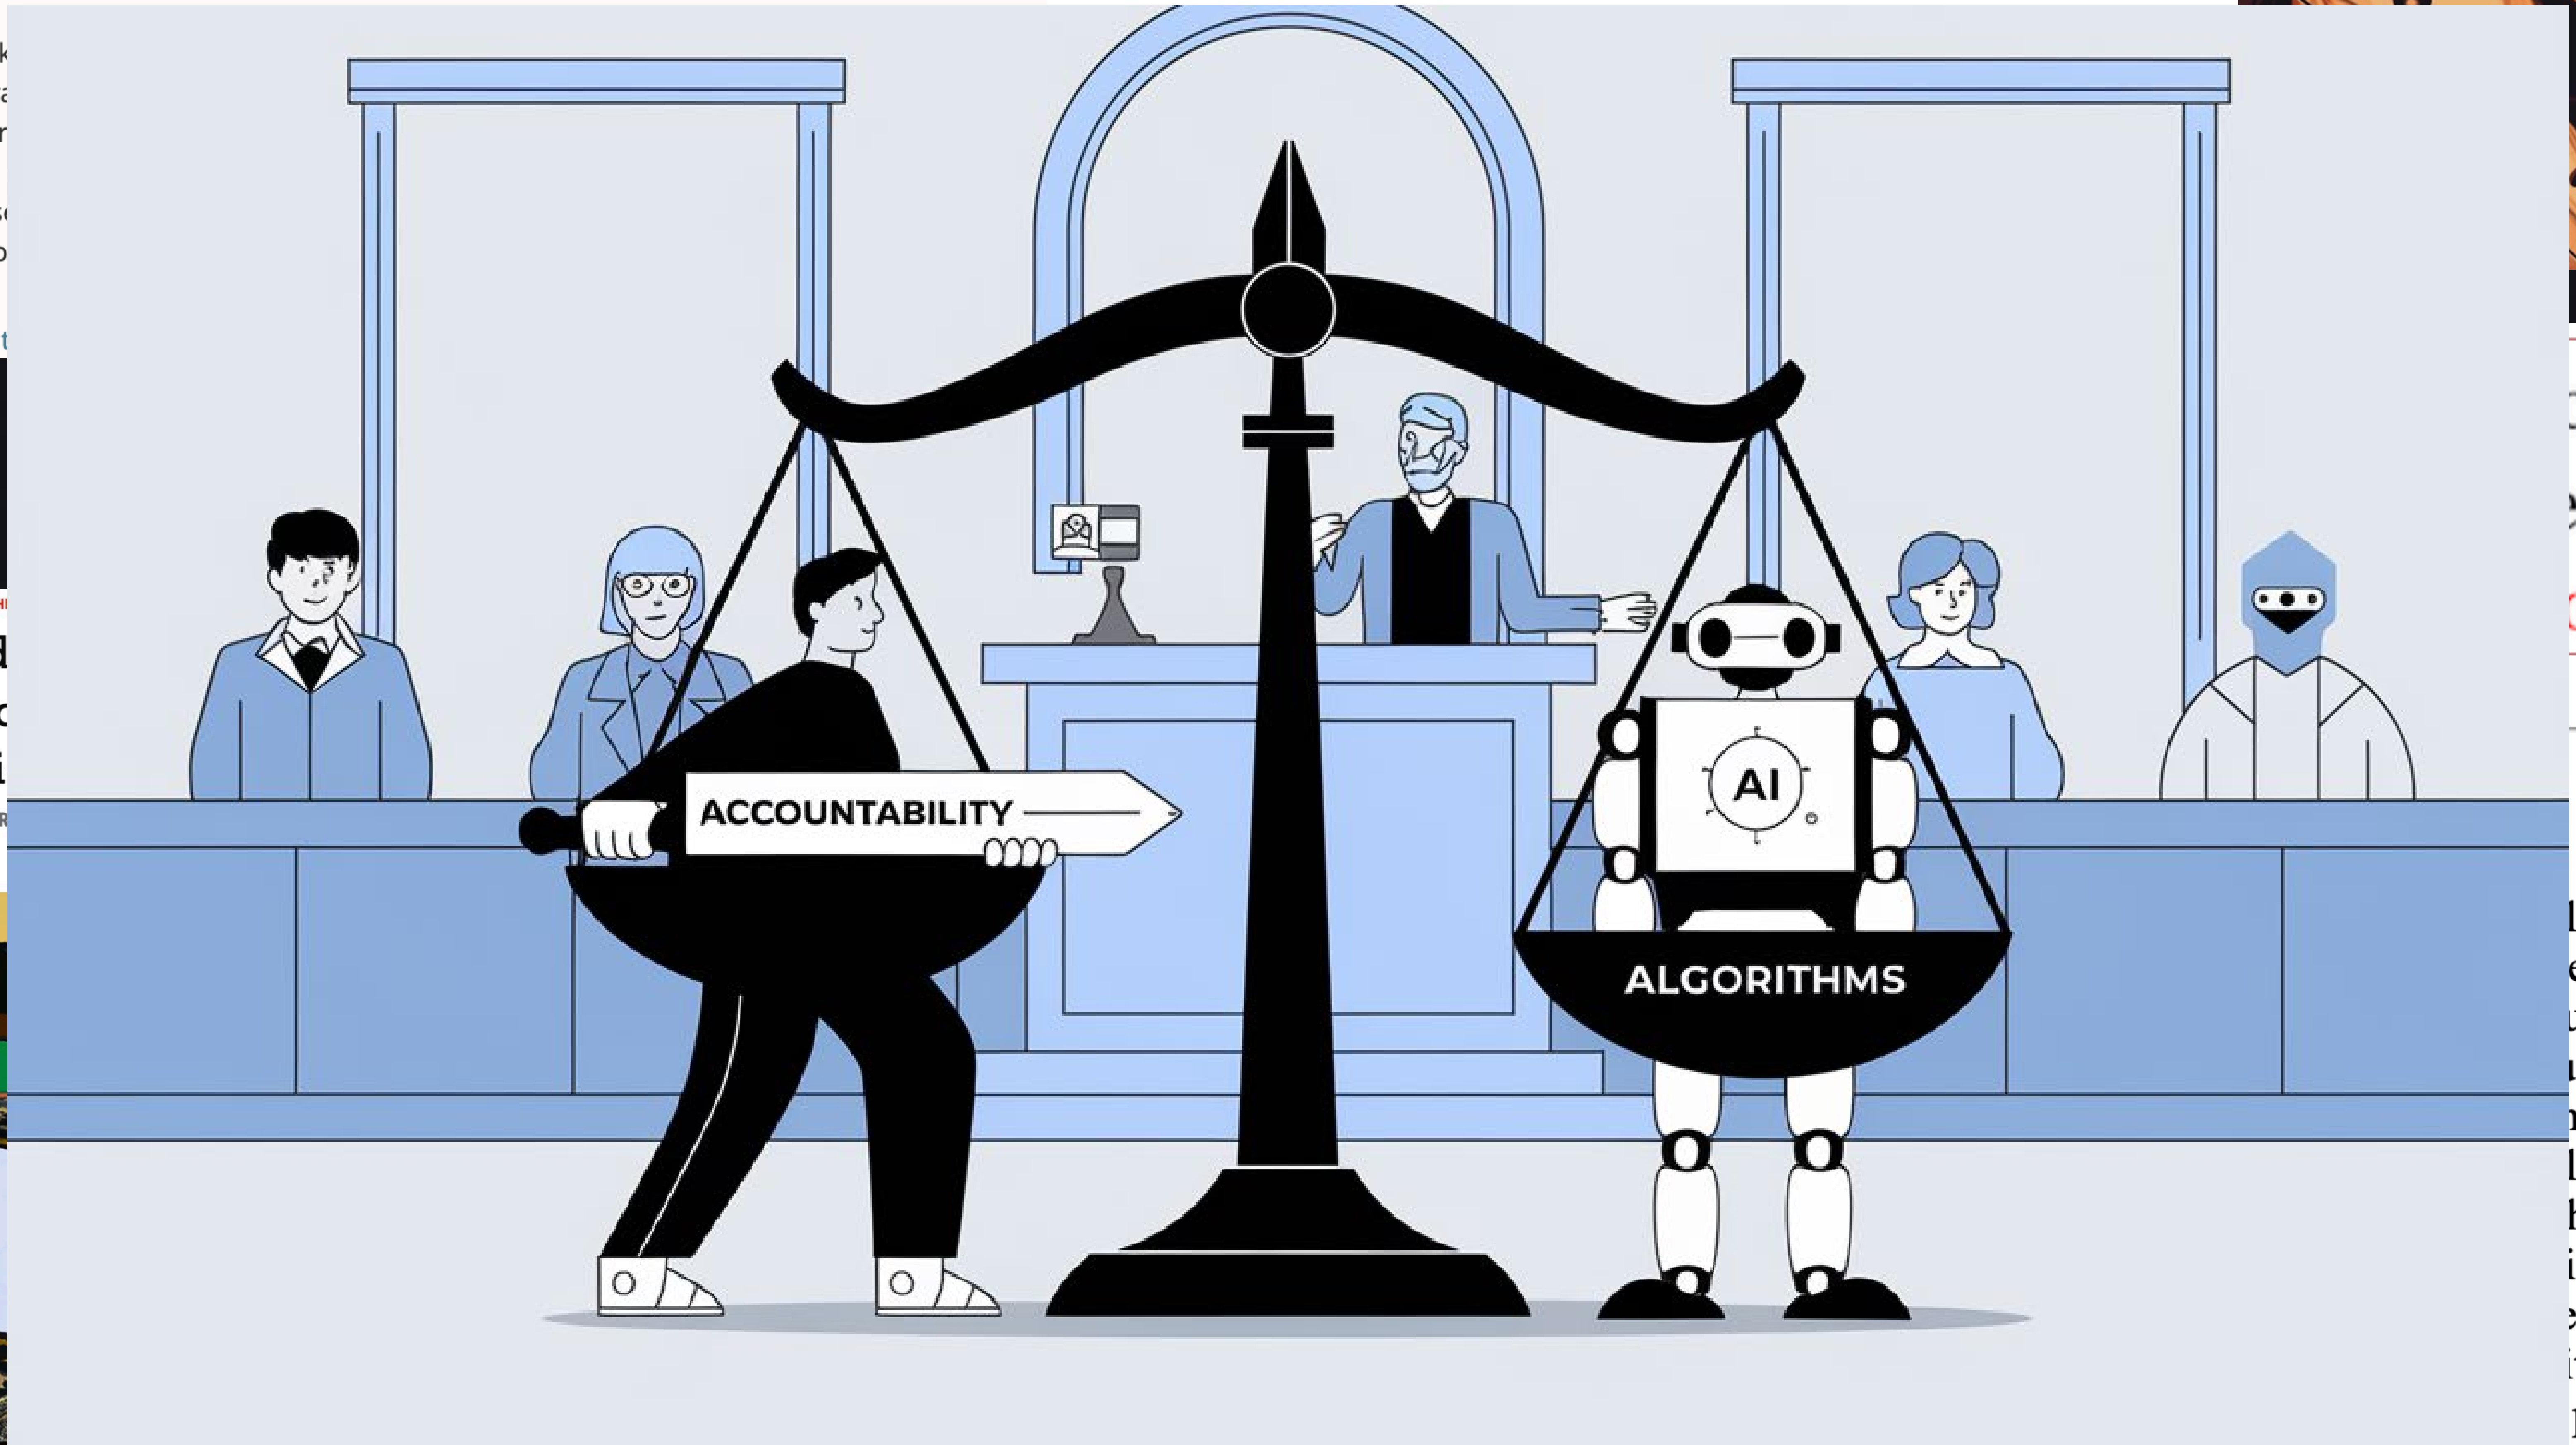

Climate UStopia

Thanks to Ruha Benjamin for inspiring these visuals

Thank You!

Claudio’s X

Claudio’s Linkedin

Claudio’s Website