HDSI | Harvard Data Science Initiative

Uncertainty in AI: Solution or Obstacle for Climate Action?

- Francesca Dominici, PhD
- Professor of Biostatistics, Population Health and Data Science
- Harvard T.H. Chan School of Public Health
- Director of the Harvard Data Science Initiative

Outline

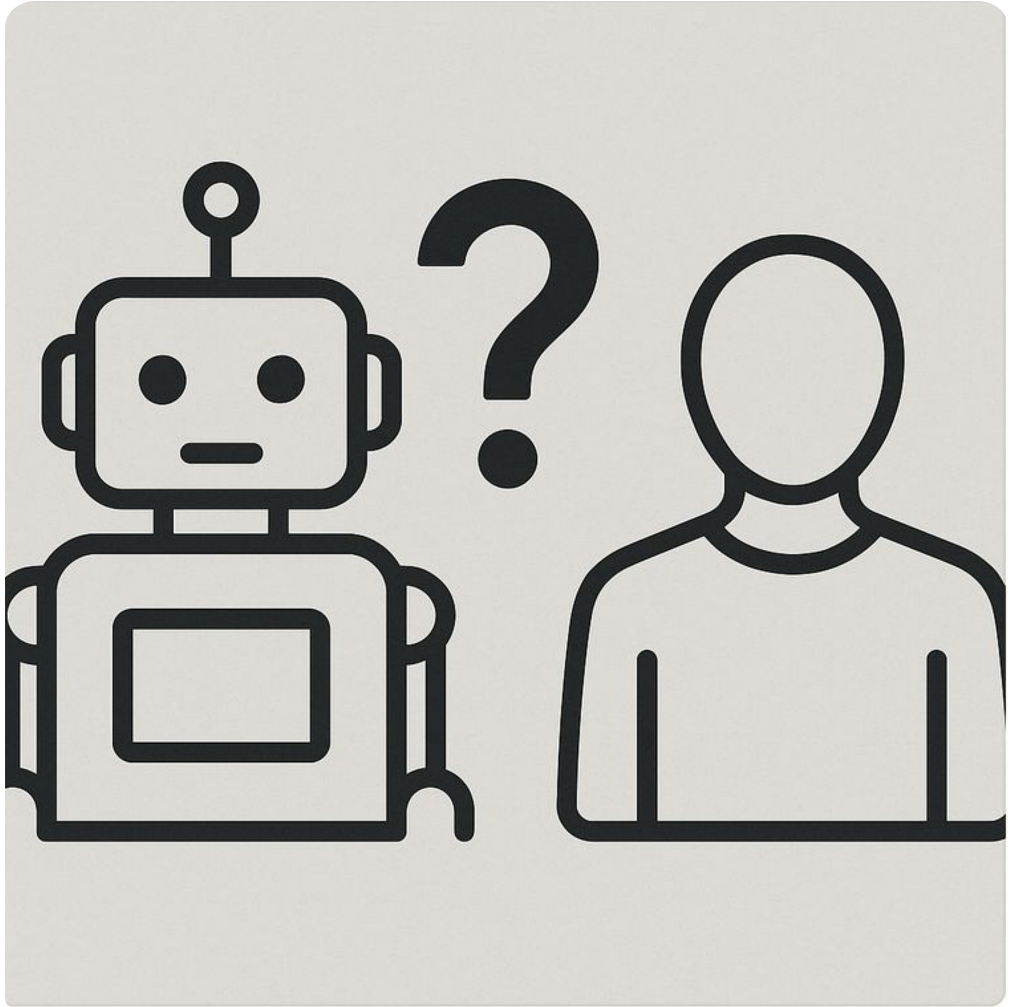

- The trillion-dollar, highly political scientific questions
- AI as a promising solution
- My journey
- Geo-AI for Air Pollution Exposure Estimation
- Causal Inference in Artificial Intelligence
- Responsible AI
- The Environmental Impact of AI
- Conclusions

AI as Eutopia or Dystopia?

The trillion-dollar, highly political scientific questions

- Does exposure to fine particulate matter, even at low levels, cause an increase in hospitalizations?
- Does exposure to wildfire cause cancer?
- Is air pollution from coal-fired power plants more toxic than air pollution from other sources?

1. The Causal Framework: seeing versus doing versus imagining

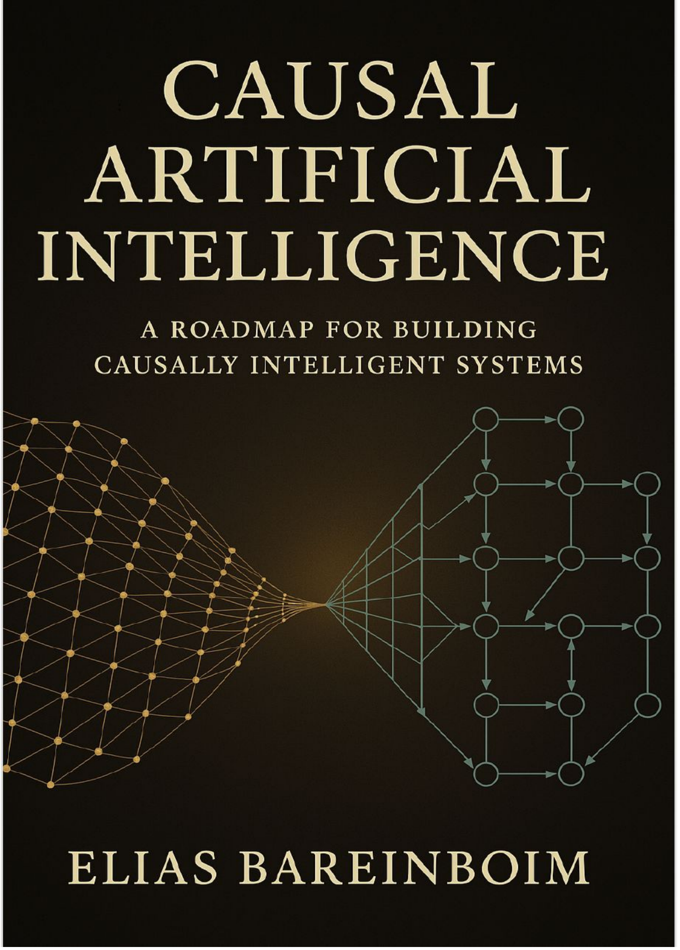

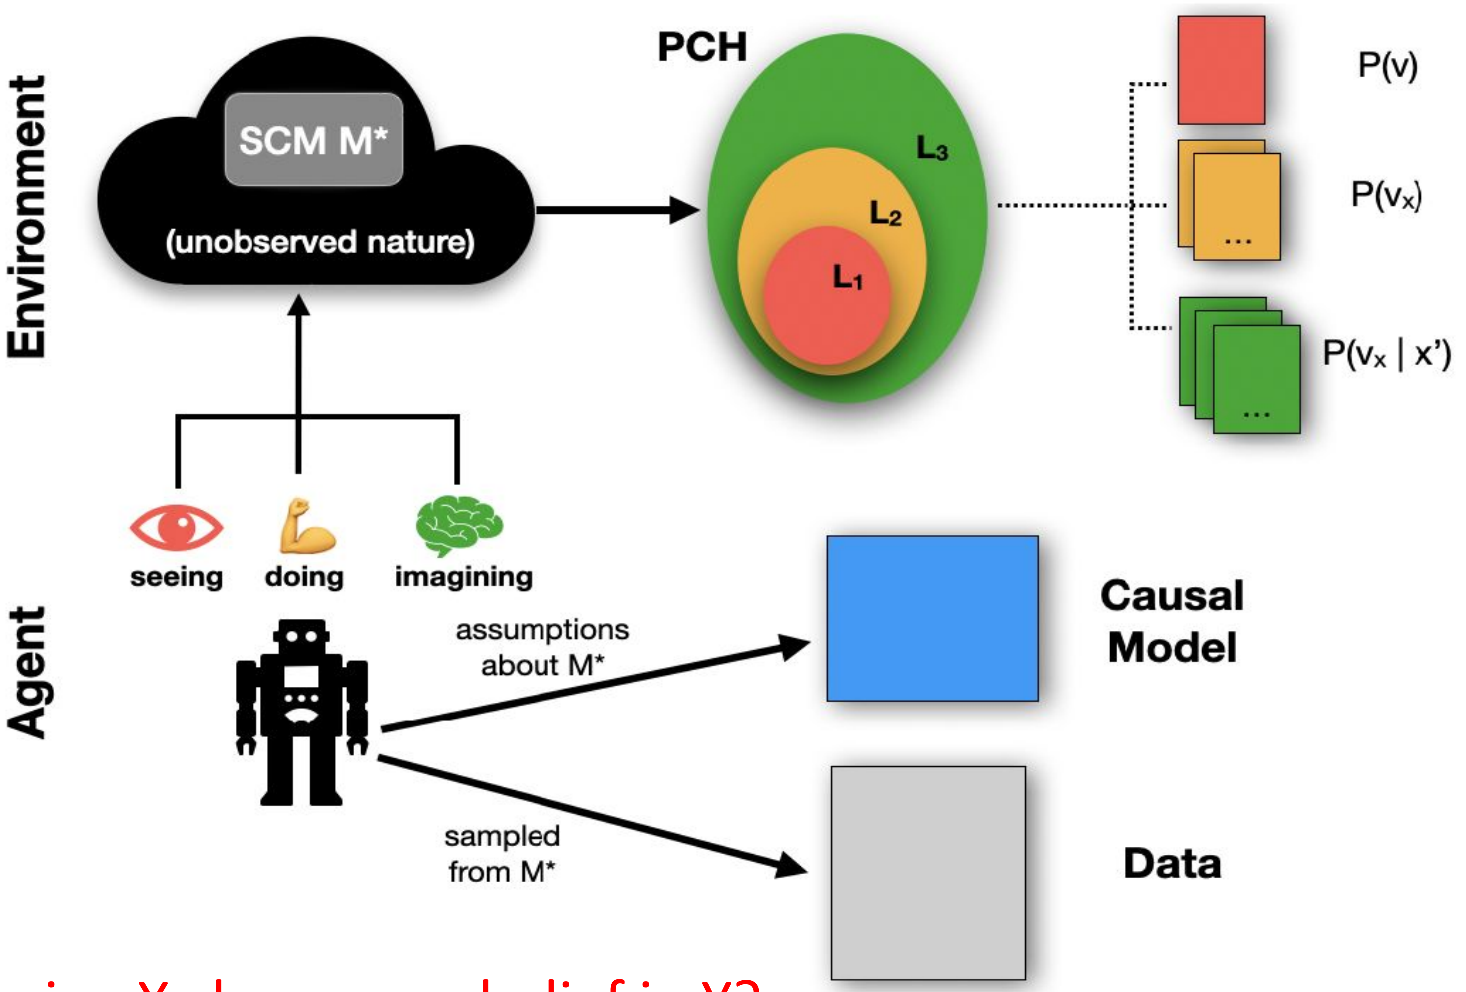

1. **Associational**

Seeing What is: How would be seeing X change my belief in Y?

2. **Interventional**

Doing What if: What if I do X?

3. **Counterfactual imagining**

Imagining -> Why? What if I acted differently?

Point-Counterpoint

Point: Clarifying Policy Evidence With Potential-Outcomes Thinking—Beyond Exposure-Response Estimation in Air Pollution Epidemiology

Corwin Matthew Zigler* and Francesca Dominici

* Correspondence to Dr. Corwin Matthew Zigler, 655 Huntington Avenue, Building 2, 4th Floor, Department of Biostatistics, Harvard School of Public Health, Boston, MA 02115 (e-mail: czigler@hsph.harvard.edu).

Initially submitted January 8, 2014; accepted for publication May 19, 2014.

Seeing: Is exposure to $PM_{2.5}$ below the NAAQS (12 $\mu g/m^3$) associated with an increased mortality risk?

Doing: If I intervene by implementing a new air quality policy, how many lives would it save?

Imagining: Let’s imagine a world where we can adapt to extreme weather events. What will be the positive consequences to society?

DATA

Data integration of over 20 government data repositories

- All Medicare participants (n=67,682,479) in the continental United States from 2000 to 2021
- Outcomes: all-cause mortality and cause specific hospitalization
- Individual level information: date of death, age of entry, year of entry, sex, race, whether eligible for Medicaid (proxy for SES)
- Zip code of residence and other covariates

Causal Reasoning AI for Policy Decisions (9 TB of data)

EXPOSURES AND INTERVENTIONS (E OR I)

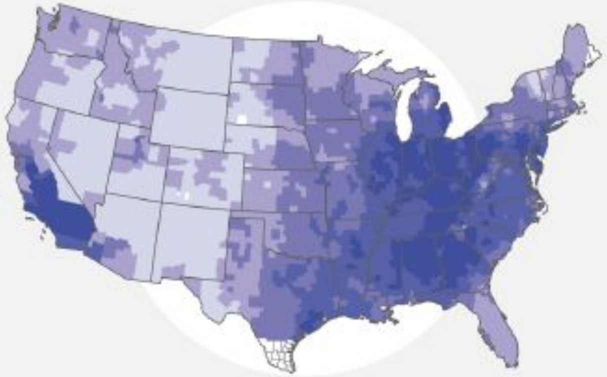

$PM_{2.5}$ exposure levels by county (average 2000-2012)

DATA SOURCES

Criteria air pollutants
EPA AQS daily average of $PM_{2.5}$, ozone, NO_2, 1995-2015;
Daily 1km x 1km predictions of $PM_{2.5}$, ozone, NO_2, 2000-2014

Methane
1km x 1km predictions at 3-day intervals, 2009-present

Weather
NOAA daily estimates (temperature, precipitation, humidity, ...) on a 0.3° grid

Power plants
EPA AMPD daily emissions, 1995-2015

Coal mines
MSHA location and producting pits, 1970-2015

Fracking wells and disposal wells
Drillinginfo database with well location and depth, daily production

Traffic
Annual traffic counts and density from the Department of Transportation

Residential community green space
NASA vegetation index on a 250m² grid

Factories and industrial sites
Geocoded locations of businesses

HEALTH OUTCOMES (Y)

Medicare mortality rate by county (average 2000-2012)

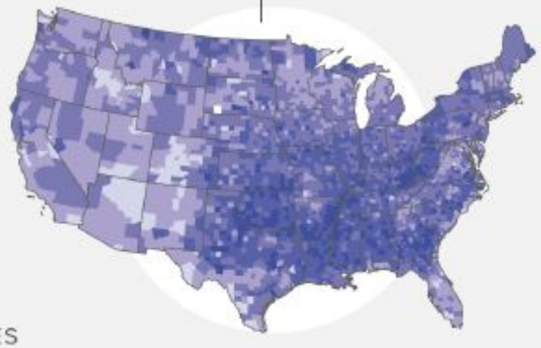

DATA SOURCES

Medicare
28 million per year, 1999-2015

Medicaid
28 million per year, low income, 2010-2011

Aetna
40 million, all ages, above-average income, 2008-2016

CONFOUNDERS (X)

Poverty prevalence by county (average 2000 and 2010)

DATA SOURCES

Individual demographics
Age, sex, race, ZIP code of residence

Individual medical history
Previous diagnoses, medications prescribed

ZIP code level variables
Income, education, demographics, employment, household size

County-level variables
Crime, smoking, BMI

Seeing

Doing

Imaging

What are the unique challenges to answering these questions?

- **Misaligned data**
- **Multiple diseases that naturally interact**
- **Multimodal data**
- **Spatial and temporal**
- Spurious correlations (confounding)
- Massive, noisy data
- Causal effects are hard to detect

- Interpretability
- Reproducibility
- Responsability
- Uncertainty

→ Are not negotiable

Artificial Intelligence in Air Pollution Exposure Estimation

Machine Learning Methods

Tree-Based Models

Neural Networks

Ensemble and Hybrid Models

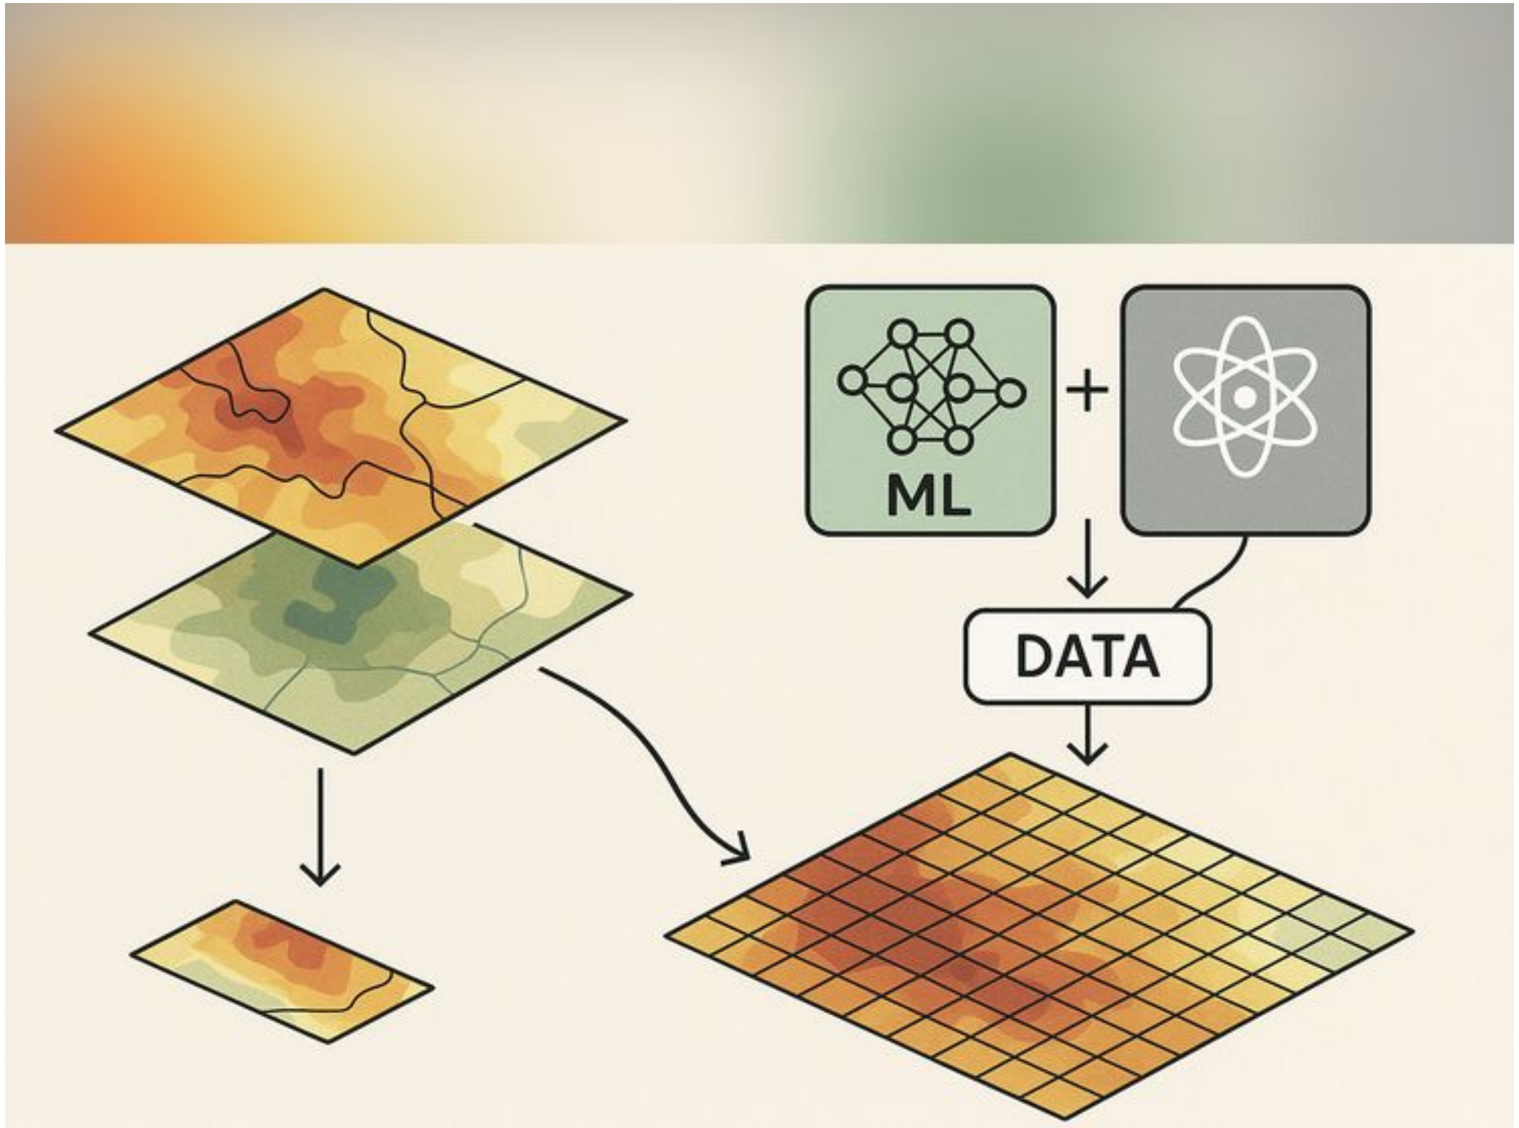

Neurips 2019

Accurate Uncertainty Estimation and Decomposition in Ensemble Learning

Jeremiah Zhe Liu*
Google Research & Harvard University
zhl112@mail.harvard.edu

John Paisley
Columbia University
jpaisley@columbia.edu

Marianthi-Anna Kioumourtzoglou
Columbia University
mk3961@cumc.columbia.edu

Brent A. Coull
Harvard University
bcoull@hsph.harvard.edu

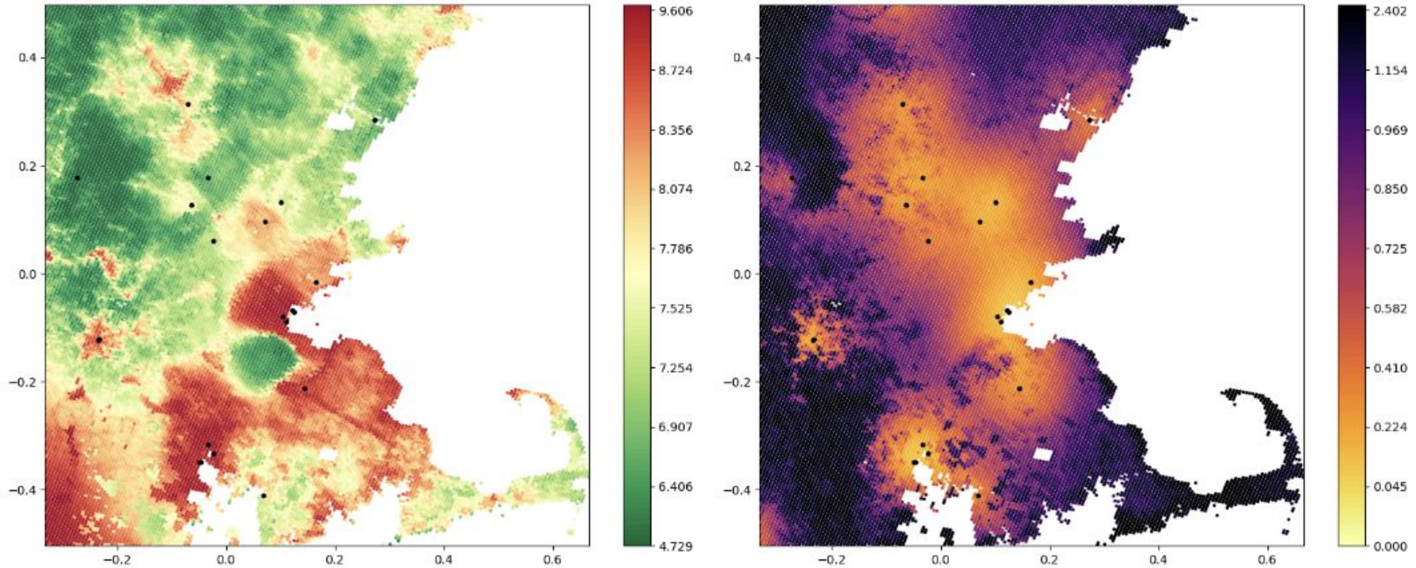

(a) Posterior Mean (b) Overall Uncertainty

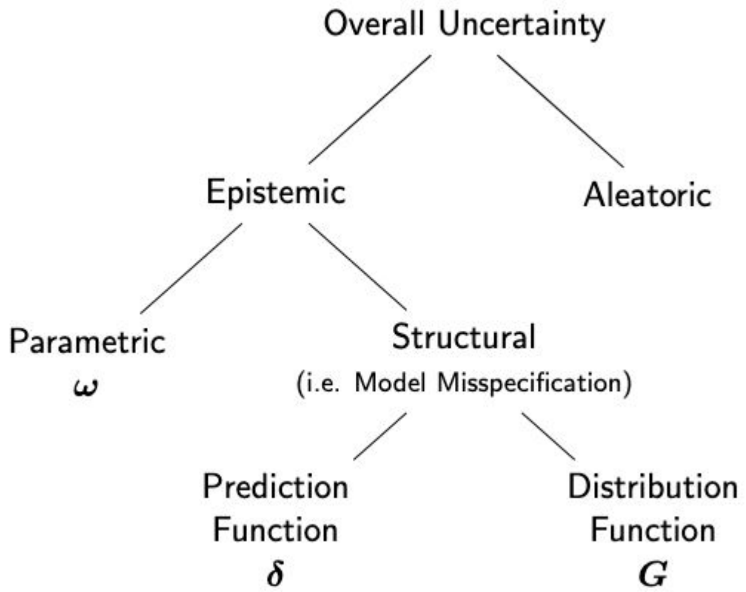

Figure 1: A decomposition of different types of uncertainty by the Bayesian Nonparametric Ensemble (BNE).

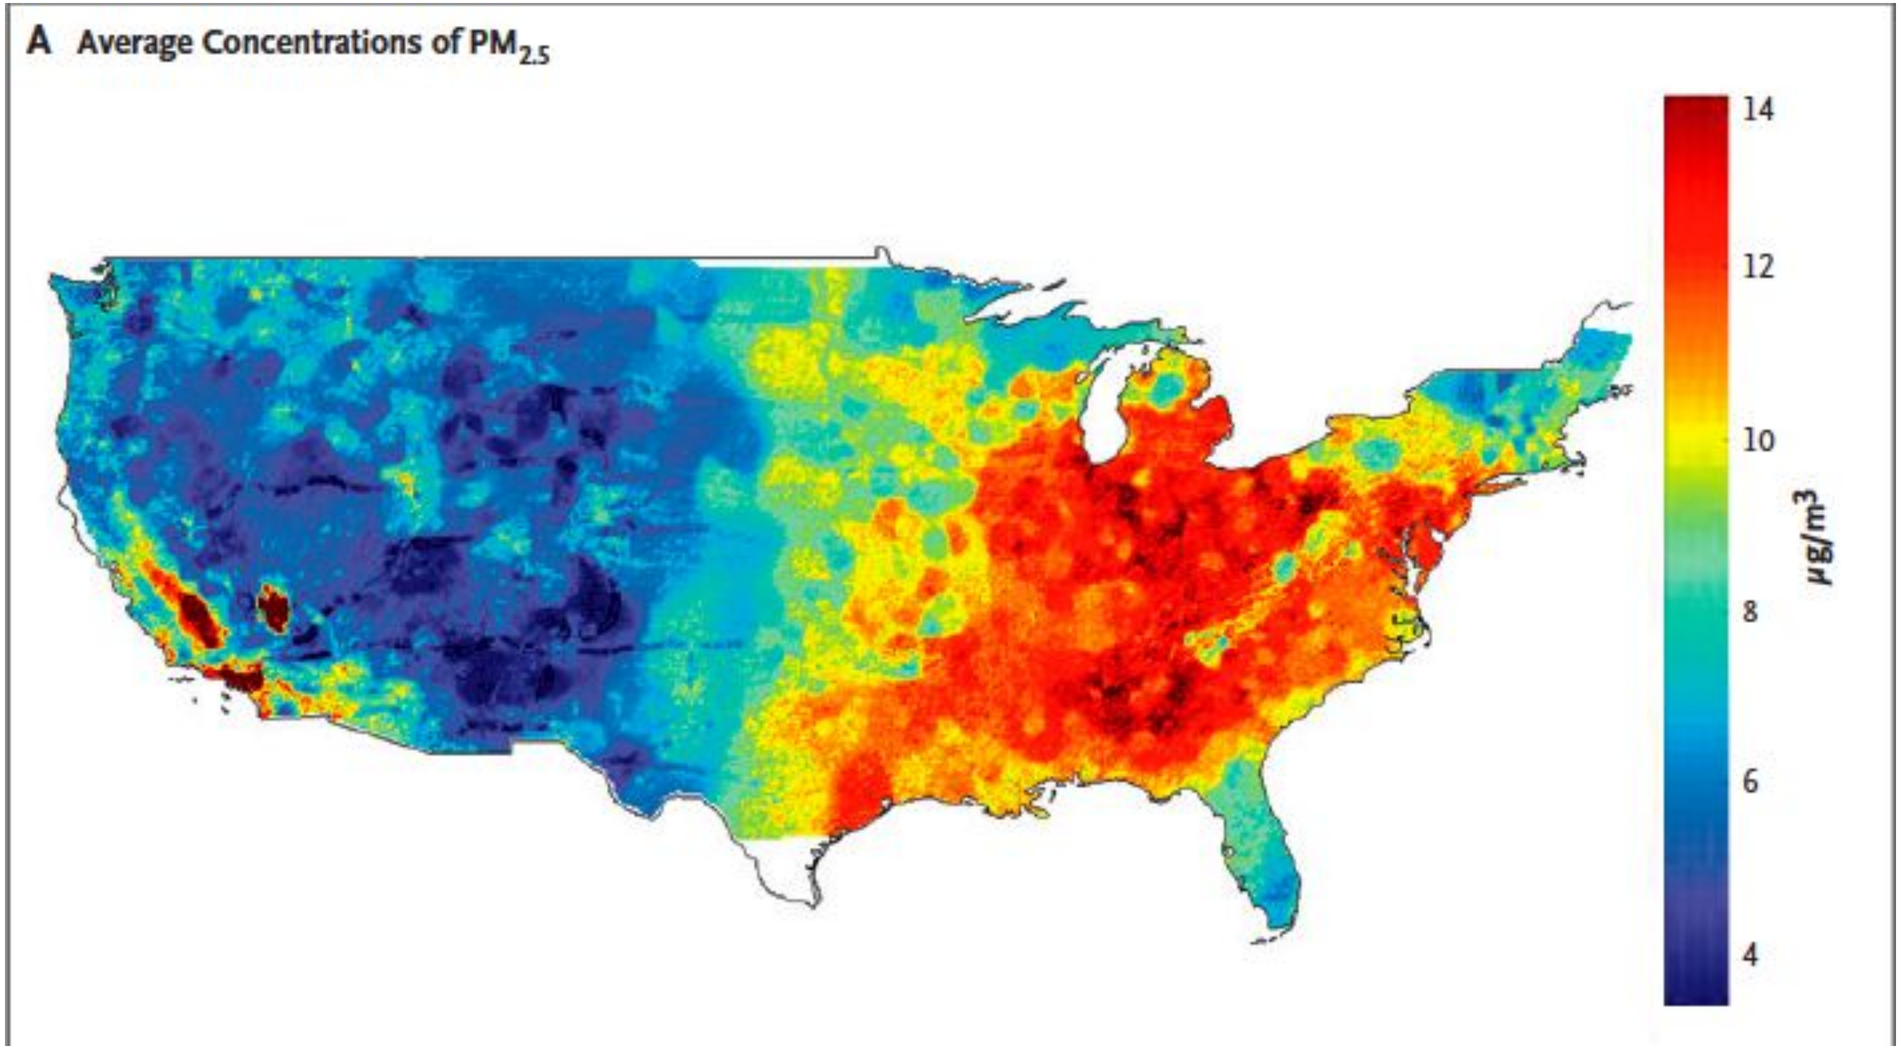

Di Q, Amini H, Shi L, Kloog I, Silvern R, Kelly J, et al. 2019. An ensemble-based model of PM2.5 concentration across the contiguous United States with high spatiotemporal resolution. Environ Int 130:104909, 10.1016/j.envint.2019.104909

2. Causal Inference in Artificial Intelligence:

Exploring Potential Outcomes for Decision-Making

Approaches to Causal Inference

Potential Outcomes Model	VS	Causal Graph Models
• Typically estimates the effect of a single, pre-specified intervention at a time. • Relies on constructing or mimicking randomized experiments using statistical adjustments such as matching or regression. • Focuses on estimating average treatment effects or individual causal effects. • Designed for targeted and specific causal inquiries.		• Can model and analyze multiple interventions and causal pathways simultaneously. • Identify confounders, mediators, and colliders explicitly to guide unbiased causal estimation. • Facilitate discovery of new causal relationships and simulate hypothetical interventions across complex systems. • Ideal for exploring complex and interconnected causal structures.

Timeline of Bayesian, Causal, and AI Developments in Inference

Late 90s: Before Causal Inference

Seeing (with uncertainty)

- Bayesian regression.
- Bayesian Model Averaging

2000s: Bayesian + Causal

Doing (with uncertainty)

Potential outcome framework

Uncertainty in the confounder selections.

2010s: Bayesian + Causal + ML

Doing (and learning)

GPS, Causal Forests,

2023+: AI

Imagining

Foundational models

2025+: Enabling AI/Causal FMs

Imagining with uncertainty

1 2 3 4 5

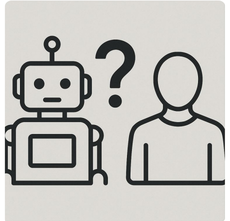

2050+: AGI

Wang C, Parmigiani G, **Dominici F** (2012) Bayesian Effect Estimation Accounting for Adjustment Uncertainty. *Biometrics*, (68)3:681-689.

2000s: Bayesian + Causal Doing (with uncertainty)
Potential outcome framework
Uncertainty in the confounder selections.

$$E\{X_i\} = \sum_{m=1}^{M} \alpha_m^X \delta_m^{\alpha^X} U_{im}\,, \tag{1}$$

$$E\{Y_i|X_i\} = \beta^{\alpha^Y} X_i + \sum_{m=1}^{M} \alpha_m^Y \delta_m^{\alpha^Y} U_{im}\,, \tag{2}$$

$$\boxed{\frac{P\left(\alpha_m^Y = 1 \middle| \alpha_m^X = 1\right)}{P\left(\alpha_m^Y = 0 \middle| \alpha_m^X = 1\right)} = \omega,}$$

Wang C, **Dominici F**, Parmigiani G, Zigler CM (2015) Accounting for Uncertainty in Confounder and Effect Modifier Selection when Estimating Average Causal Effects in Generalized Linear Models. *Biometrics*

Uncertainty in Propensity Score Estimation: Bayesian Methods for Variable Selection and Model Averaged Causal Effects

JASA 2014

Corwin Matthew Zigler [1], Francesca Dominici [1]

$$g_x(E[X_i|U_i]) = \sum_{k=0}^{p} \alpha_k \gamma_k U_{ik},$$

$$g_y(E[Y_i|X_i, U]) = \beta_0 + \beta_X X_i + h(PS(\gamma, \alpha, U_i); \xi) + \sum_{k=1}^{p} \alpha_k \delta_k U_{ik},$$

2000s: Bayesian + Causal
Doing (with uncertainty)
Potential outcome framework
Uncertainty in the confounder selections.

2010-2020s: Bayesian + Causal + ML
Doing (and learning)

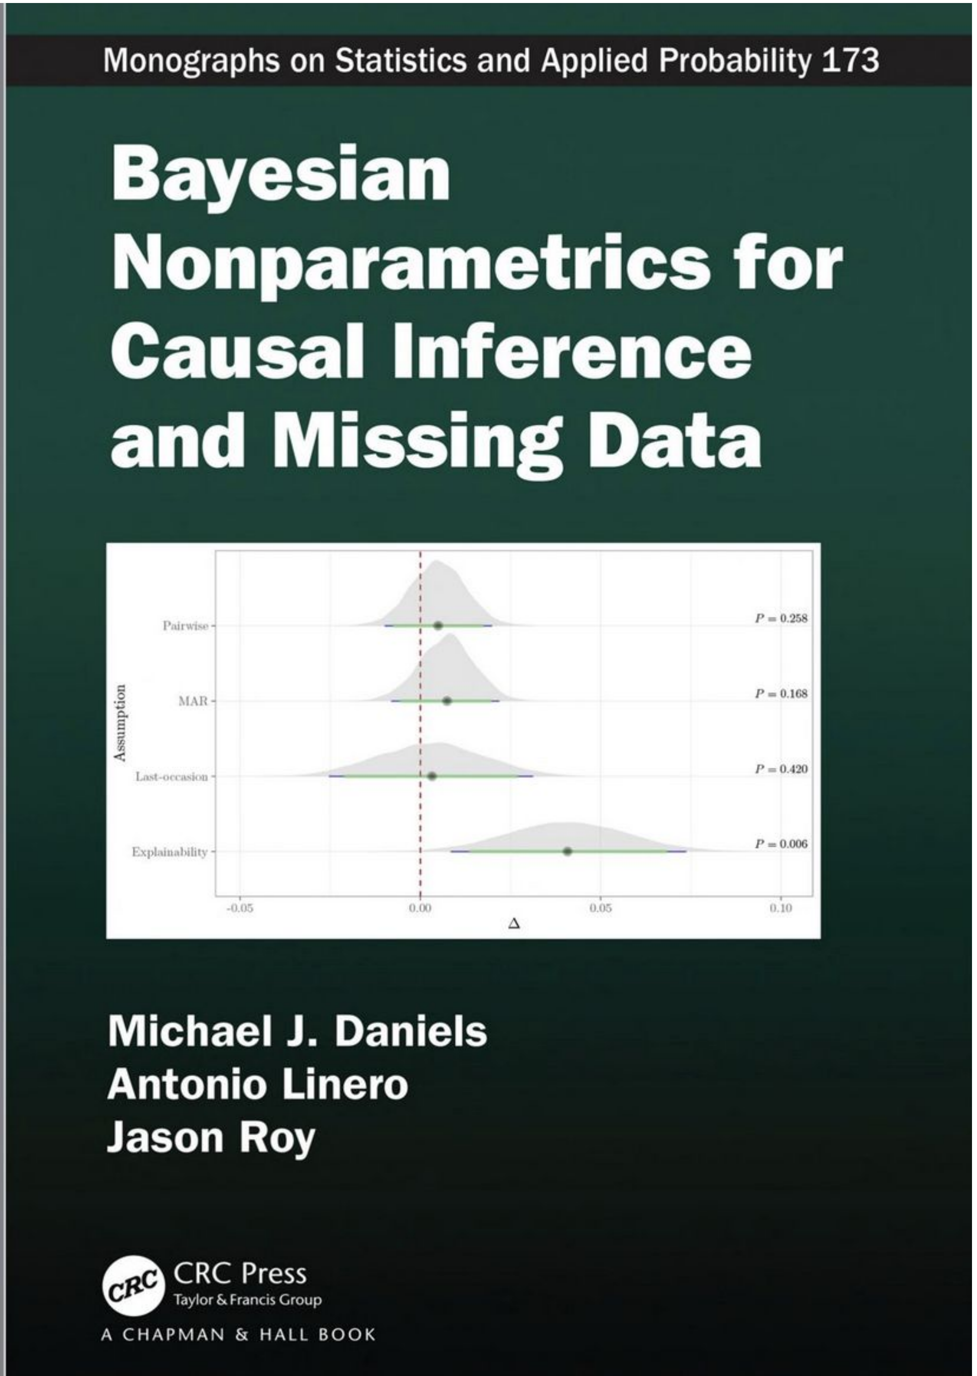

Received: 7 April 2022 | Accepted: 19 January 2023

DOI: 10.1111/biom.13833

Biometrics WILEY

BIOMETRIC METHODOLOGY

Bayesian nonparametric adjustment of confoundi

Chanmin Kim[1] | Mauricio Tec[2] | Corwin Zigler[3]

WIREs COMPUTATIONAL STATISTICS

Overview

The how and why of Bayesian nonparametric causal inference

Antonio R. Linero, Joseph L. Antonelli

First published: 06 April 2022 | https://doi.org/10.1002/wics.1583 | Citations: 1

Edited by: James Gentle, Commissioning Editor and Co-Editor-in-Chief and David Scott, Review Editor and Co-Editor-in-Chief

Read the full text >

PDF TOOLS SHARE

The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Decision Making and Reinforcement learning

Optimizing Heat Alert Issuance with Reinforcement Learning

Ellen M. Considine*[1], Rachel C. Nethery[1], Gregory A. Wellenius[2], Francesca Dominici[1], Mauricio Tec*[1,3]

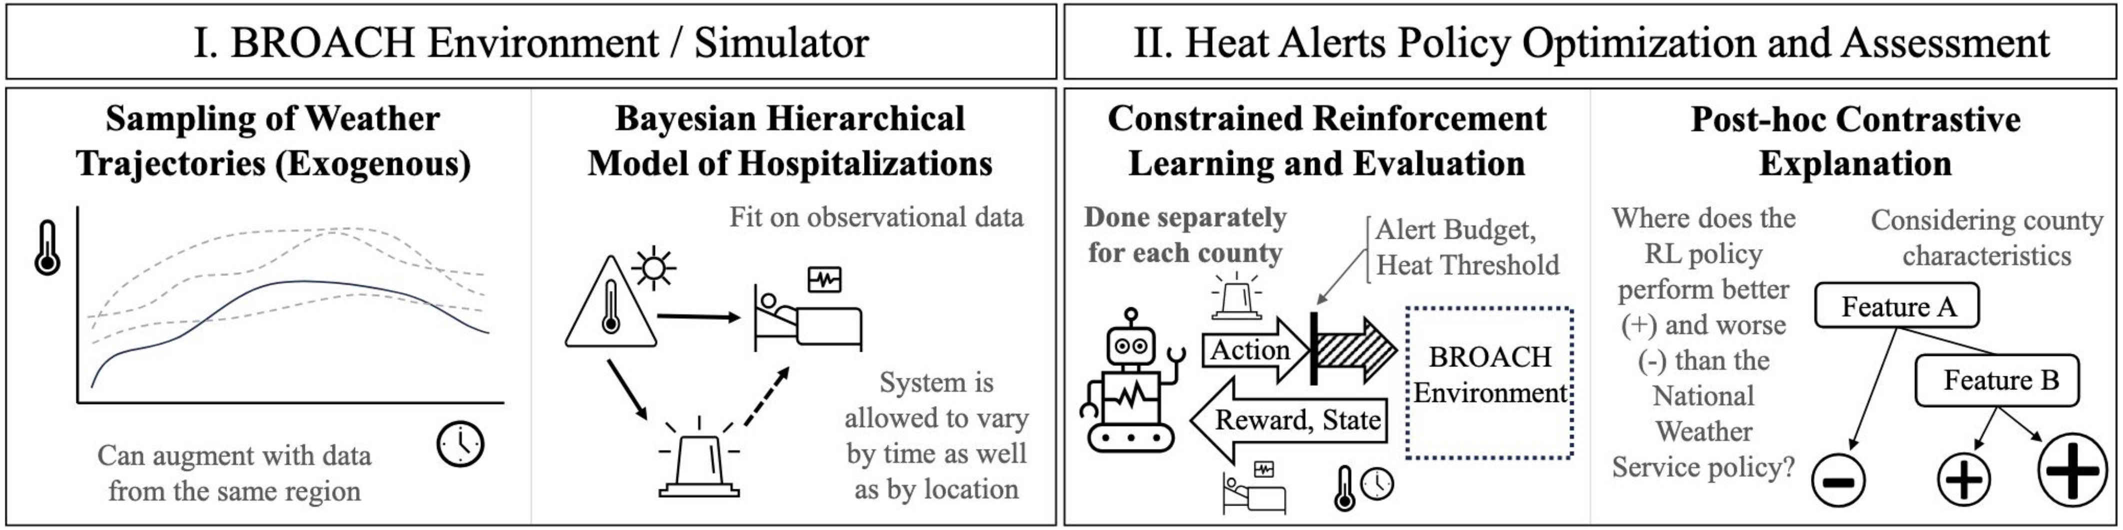

Figure 1: Overview of the heat alerts RL framework.

3. The impact to policy

The NEW ENGLAND JOURNAL of MEDICINE

Air Pollution and Mortality at the Intersection of Race and Social Class

Kevin P. Josey, Ph.D., Scott W. Delaney, Sc.D., J.D., Xiao Wu, Ph.D., Rachel C. Nethery, Ph.D., Priyanka DeSouza, Ph.D., Danielle Braun, Ph.D., and Francesca Dominici, Ph.D.

Table 1. Characteristics of the Medicare Cohort, 2000 through 2016.*

Characteristic	Full Cohort†	Black Persons		White Persons	
		Higher Income‡	Low Income§	Higher Income‡	Low Income§
Persons — no. (% of full cohort)	73,129,782 (100)	4,872,714 (6.7)	1,671,776 (2.3)	56,422,414 (77.2)	4,989,457 (6.8)
Person-yr — no. (% of total person-yr)	623,042,512 (100)	37,862,780 (6.1)	14,886,928 (2.4)	483,479,863 (77.6)	48,247,908 (7.7)
Deaths — no. (% of total deaths)	29,467,648 (100)	1,488,555 (5.1)	1,154,227 (3.9)	20,773,208 (70.5)	4,769,240 (16.2)
Median follow-up time — yr	8.0	7.0	8.0	8.0	8.0
Age at entry — %					
65–74 yr	80.6	86.2	77.4	80.4	72.7
75–84 yr	14.8	10.7	15.6	15.3	17.2
85–94 yr	4.2	2.5	6.2	4.0	9.0
≥95 yr	0.4	0.6	0.8	0.3	1.1
Female sex — %	55.4	54.9	68.1	54.3	68.0
Medicaid eligible — %	11.6	0	100	0	100

Lowering exposure from 12 to 9 unit □ 5 % mortality reduction among Black Americans; 2.5% mortality reduction among White Americans

NEW RESEARCH

Cleaner Air Helps Everyone. It Helps Black Communities a Lot.

A new study quantified the benefits of pollution reduction in terms of race and class.

Share full article

St. James, La., one of several Mississippi River towns dotted by chemical plants and oil refineries. William Widmer for The New York Times

The NEW ENGLAND JOURNAL of MEDICINE

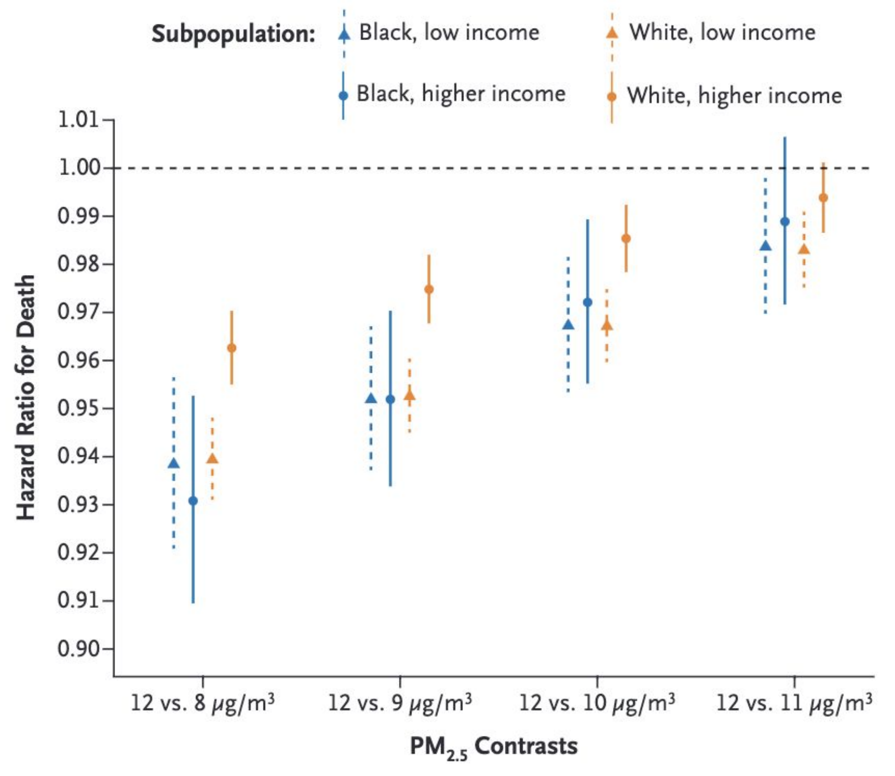

Figure 4. Differences in Mortality with Decreasing $PM_{2.5}$ Exposure among Marginalized Subpopulations.

Shown are point estimates and 95% confidence intervals of the hazard ratio for death comparing different levels of annual average $PM_{2.5}$ exposure (12 μg per cubic meter vs. 11, 10, 9, or 8 μg per cubic meter) on average for subpopulations defined in selected ways. Low income was defined as dual eligibility for both Medicare and Medicaid. Confidence intervals were not adjusted for multiplicity; therefore, they should not be used in place of hypothesis testing.

Doing

How my lab has impacted this decision

Data Science

April 2023

The NEW ENGLAND JOURNAL of MEDICINE

SPECIAL ARTICLE | VOL. 388 NO. 15, APR 13, 2023

Air Pollution and Mortality at the Intersection of Race and Social Class

K.P. Josey and Others | N Engl J Med 2023; 388:1396-1404

In this large study, the mortality benefits of reducing levels of fine particulate matter air pollution were greater for low-income and higher-income Black persons and for low-income White persons than for higher-income White persons.

Nov 2023

Science

RESEARCH ARTICLE

POLLUTION

Mortality risk from United States coal electricity generation

Policy

Dec 2023, COP28 Opening day, Kerry talk about my study: Climate Crisis is a Health Crisis.

A lot of seeing and doing

Impact

February 2024

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

- ✔ Cleaner Air
- ✔ Lives Saved
- ✔ Less GHG

On February 7, 2024, the U.S. Environmental Protection Agency (EPA) announced a final rule to strengthen the nation's National Ambient Air Quality Standards (NAAQS) for fine particle pollution (12 □ 9 mg/m3)

EPA United States Environmental Protection Agency

Estimated Monetized Benefits, Costs, and Net Benefits Associated with the Final Standard Levels in 2032 for the U.S. (2017$)

	9/35 μg/m³
Benefits[a]	\$22 billion to \$46 billion
Costs[b]	\$590 million
Net Benefits	\$22 billion to \$46 billion

Notes: We focus results to provide a snapshot of costs and benefits in 2032, using the best available information to approximate social costs and social benefits recognizing uncertainties and limitations in those estimates.
[a] The benefits are associated with two point estimates from two different epidemiologic studies, and we present the benefits calculated at a real discount rate of 3 percent.
[b] The costs are annualized using a 7 percent interest rate.

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

4. Counterfactual Imagining: A role for AI foundation model

The potential of AI for Climate Adaptation

- 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)

2023+: AI
Imaging
Foundational models

ClimaCare: A Foundation Model for Healthy Climate Adaptation

Claudio Battiloro[*,†], James Kitch[*,†], Bret Nestor[*,†],
Mauricio Tec[†], Michelle Audirac, Danielle Braun,
Francesca Dominici

1. Pre-trained on the entire US health care system x environmental data x societal data
2. It produces unified embeddings that capture the complex spatiotemporal relationships between climate stressors, socioeconomic variables, and health outcomes.
3. We evaluate the model on benchmark downstream tasks, i.e., health outcomes interpolation, extrapolation, downscaling, and forecasting
4. We implement "what-if" scenario forecasting for climate adaptation using synthetic ground-truth data to validate counterfactual predictions when any input exposure is altered.

Towards a One-of-a-Kind geo-AI for Healthy Climate Adaptation

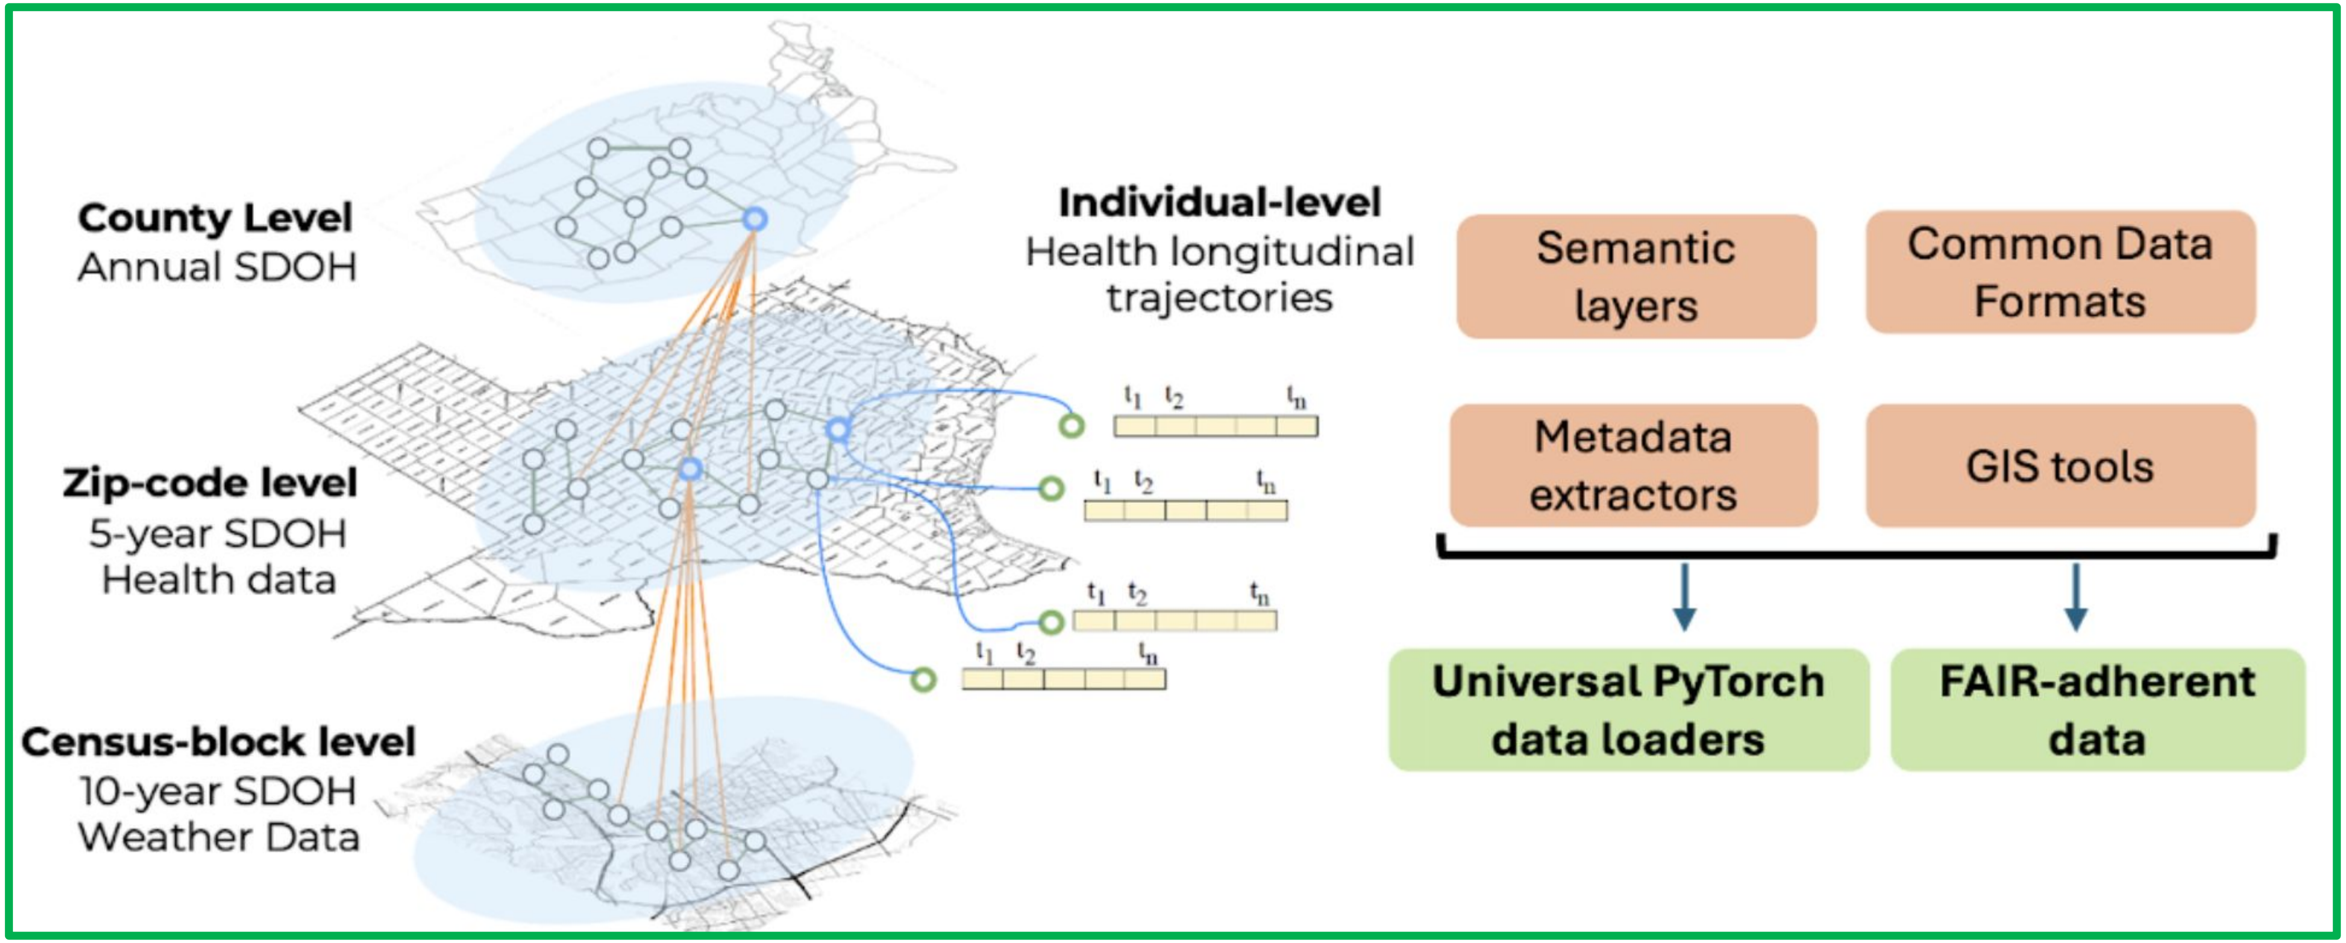

ClimaCare: Downstream Tasks

- **Spatiotemporal Downstream Tasks:**
  - Spatial Interpolation
  - Spatial Extrapolation
  - Forecasting

What-If Downstream Tasks:

ERC Estimation

Enhanced Causal Inference

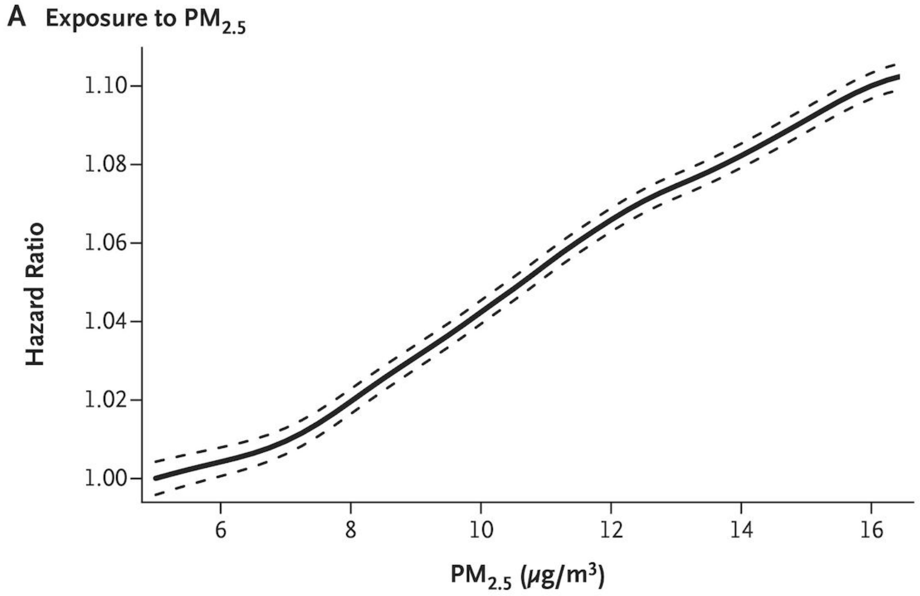

Poster session 2 (Wednesday)

The Relativity of Causal Knowledge

Gabriele D'Acunto[1,2] Claudio Battiloro[3]

[1]Information Engineering, Electronics and Telecommunications Dept., Sapienza University, Rome, Italy
[2]National Inter-University Consortium for Telecommunications (CNIT), Parma, Italy
[3]Biostatistics Dept., Harvard University, Cambridge, MA, USA

Conceptually, the relativity of causal knowledge can drive a paradigm shift in how causality is typically understood in AI/ML. By stripping causality of its oracular and absolute meaning, the relativity of causal knowledge situates it within a different ontological setting, where truth is not monolithic but emerges inevitably and relatively from a set of relationships. This is mathematically formalized by the path-dependent RCK in Equation (10). Our framework paves the way to multiple intriguing areas of reasearch, pivotal to fully characterize RCK and make it applicable.

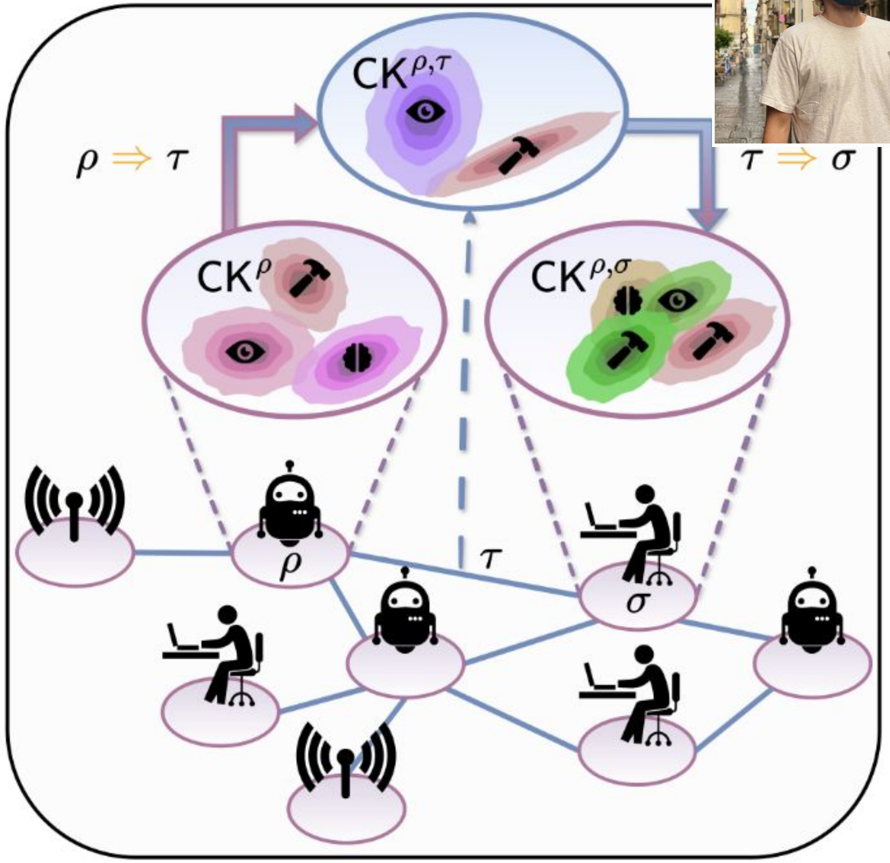

Figure 1: The *relativity of causal knowledge* states that causal knowledge (CK) is subjective and interconnected rather than objective and isolated. Multiple subjects of/in the same system will develop multiple and different instances of CK describing the system. Informally, CK can be seen as a set of probability measures corresponding to seeing, doing, and imagining. The

2025+: Enabling AI/Causal FMs
Imagining with uncertainty
Grand challenges.

- **Grand Challenge #1: We need to scale causal reasoning to the size and complexity of massive data.**
- **Grand Challenge #2: We need to learn how to use real world data to understand the consequences of complex actions and strategies**
- **Grand Challenge #3: We need to learn the precise pathways between actions and their consequences.**

+

Uncertainty

2050: Let’s imagine a world where we have solved causal reasoning

- AI agents are everywhere
- Is this a better future?

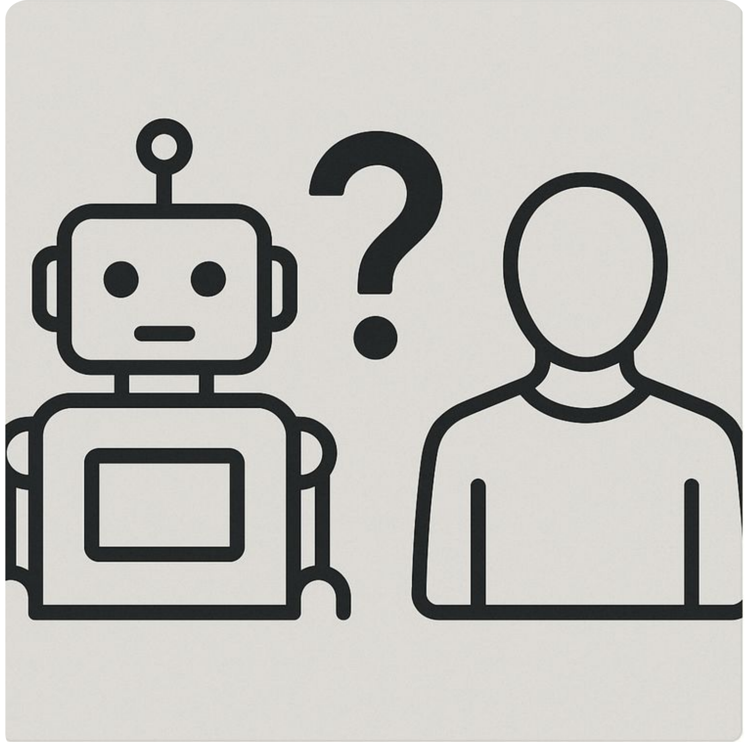

Evaluating the Environmental Impact of Hyperscale Data Centers in the U.S.

Joint work with Gianluca Guidi, Tiziano Squartini, Callaway Sprinkle, Jonathan Gilmour Kevin Butler, Eric Bell, Scott Delaney, Falco J. Bargagli-Stoffi

At Amazon's Biggest Data Center, Everything Is Supersized for A.I.

On 1,200 acres of cornfield in Indiana, Amazon is building one of the largest computers ever for work with Anthropic, an artificial intelligence start-up.

403 Hyperscale data centers and 3318 energy supplier power plants in the US (May 2024 to April 2025)

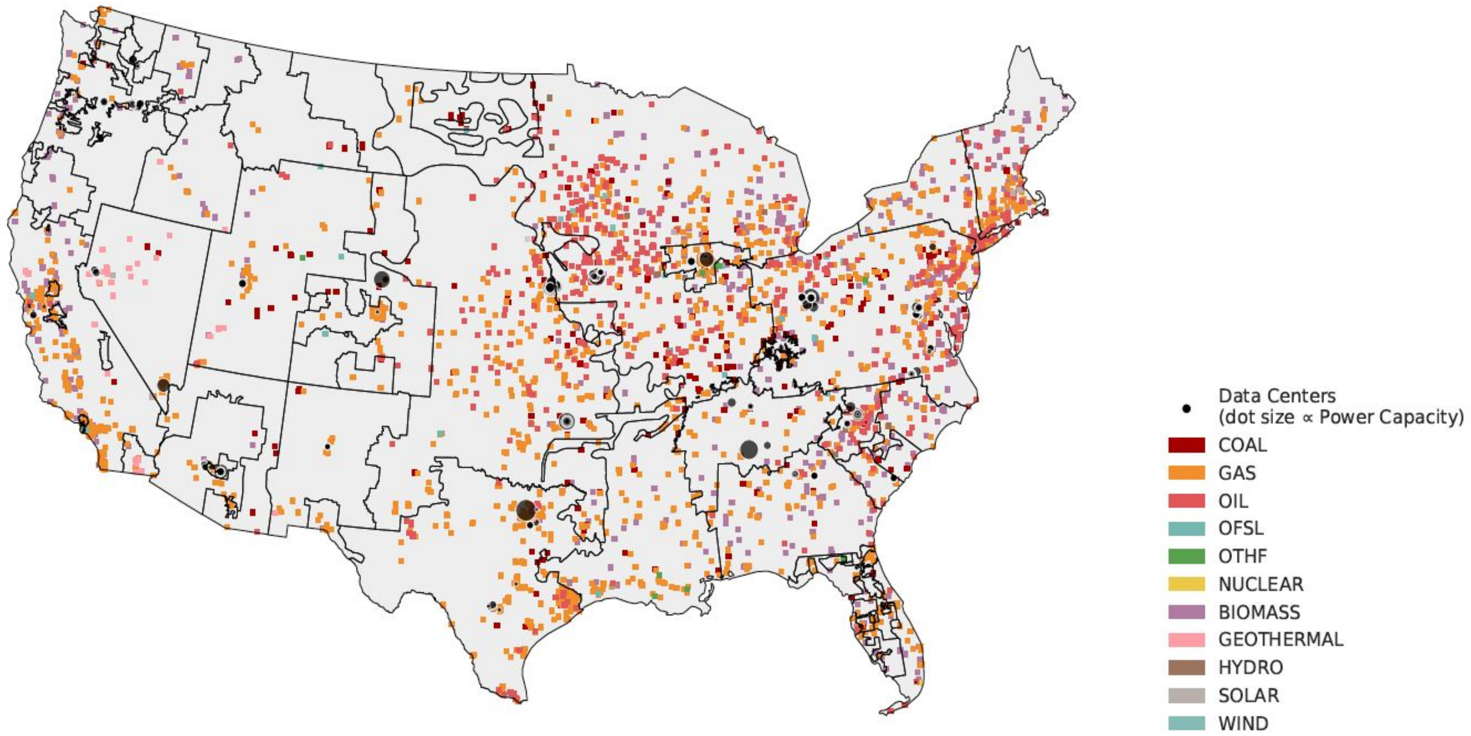

Fig. 1. Geographic distribution of hyperscale data centers and power plants in the contiguous US, overlaid with balancing authority regions. This figure shows the 403 hyperscale data centers and 3,318 operational power plants included in our analysis for the study period from May 2024 to April 2025. The map is displayed at the balancing authority (BA) level, representing regions where electricity supply and demand are managed in real time. The size of each hyperscale data center marker is proportional to its power capacity, while power plants are colored by their primary fuel type.

Scientific questions

1. What are the electricity consumption, sources, and attributable CO_2 emissions of those 403 data centers?

2. What is the fuel mix of the power plants supplying electricity to data centers?

3. Which states have the highest CO_2 emissions attributable to data centers?

Hint: With a data pipeline that can answer those questions, we make informed decisions, such as: Where should I place a data center? Where should I intervene on the power grid? How can we decarbonize this sector?

Materials and Methods

01 Dataset Compilation and Validation

Our analysis began with compiling and validating a dataset of 403 HDCs.

This was achieved using private data providers, web scraping, and satellite imagery.

02 Power Capacity Estimation

Missing power capacities were estimated using a Gradient Boosted Regression Tree model.

03 Electricity Consumption Calculation

Annual electricity consumption was calculated by multiplying the power capacity, hours per year, and a utilization rate of 66.3%.

The utilization rate was determined empirically.

04 CO_2 Emissions Estimation

Each HDC was assigned to its balancing authority and corresponding power plants.

CO_2 emissions were estimated using an energy generation-weighted model and EPA emission factors.

05 Carbon Intensity Computation

Carbon intensity was computed as grams of CO_2 emitted per kilowatt-hour consumed.

This computation was performed at various geographic levels.

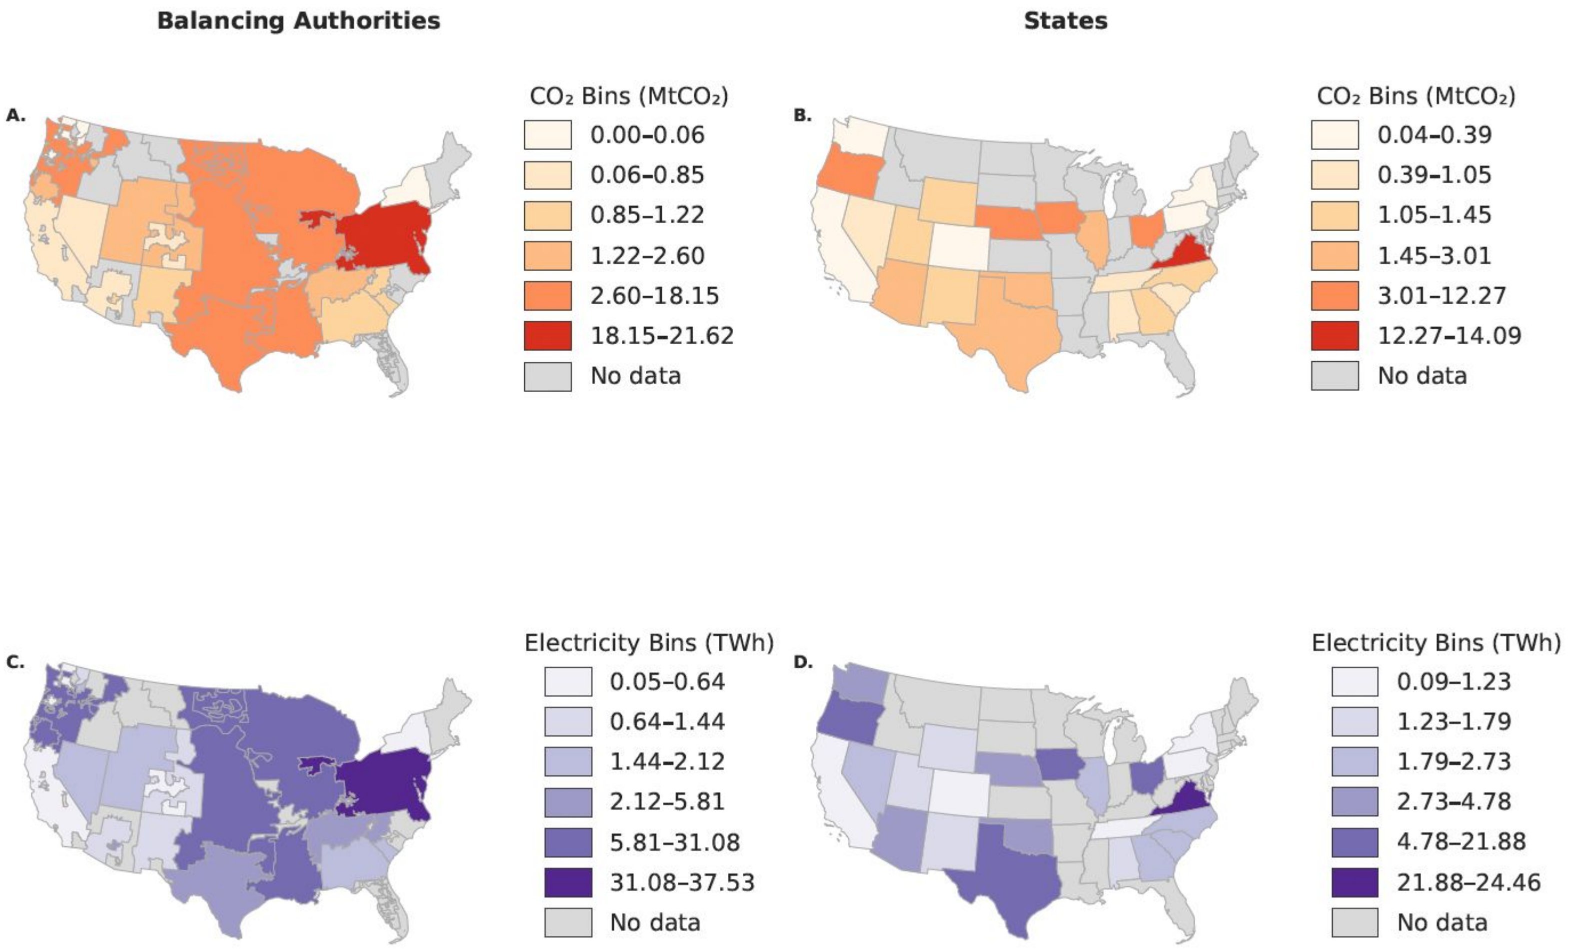

Fig. 2. Hyperscale Data Center electricity consumption and CO_2 emissions. (Left column, A and C) The balancing authority (BA) region in which a hyperscale data center is located determines the mix of power plants that supply its electricity and thus its attributable emissions. See fig.S.4.1 for BA regions and corresponding names. (Right column, B and D) Maps at the state level show electricity consumption and emissions for which the hyperscale data centers within the state are responsible for. Color bins represent percentile-based ranges: 0–20%, 20–40%, 40–60%, 60–80%, 80–99%, and 99–100%.

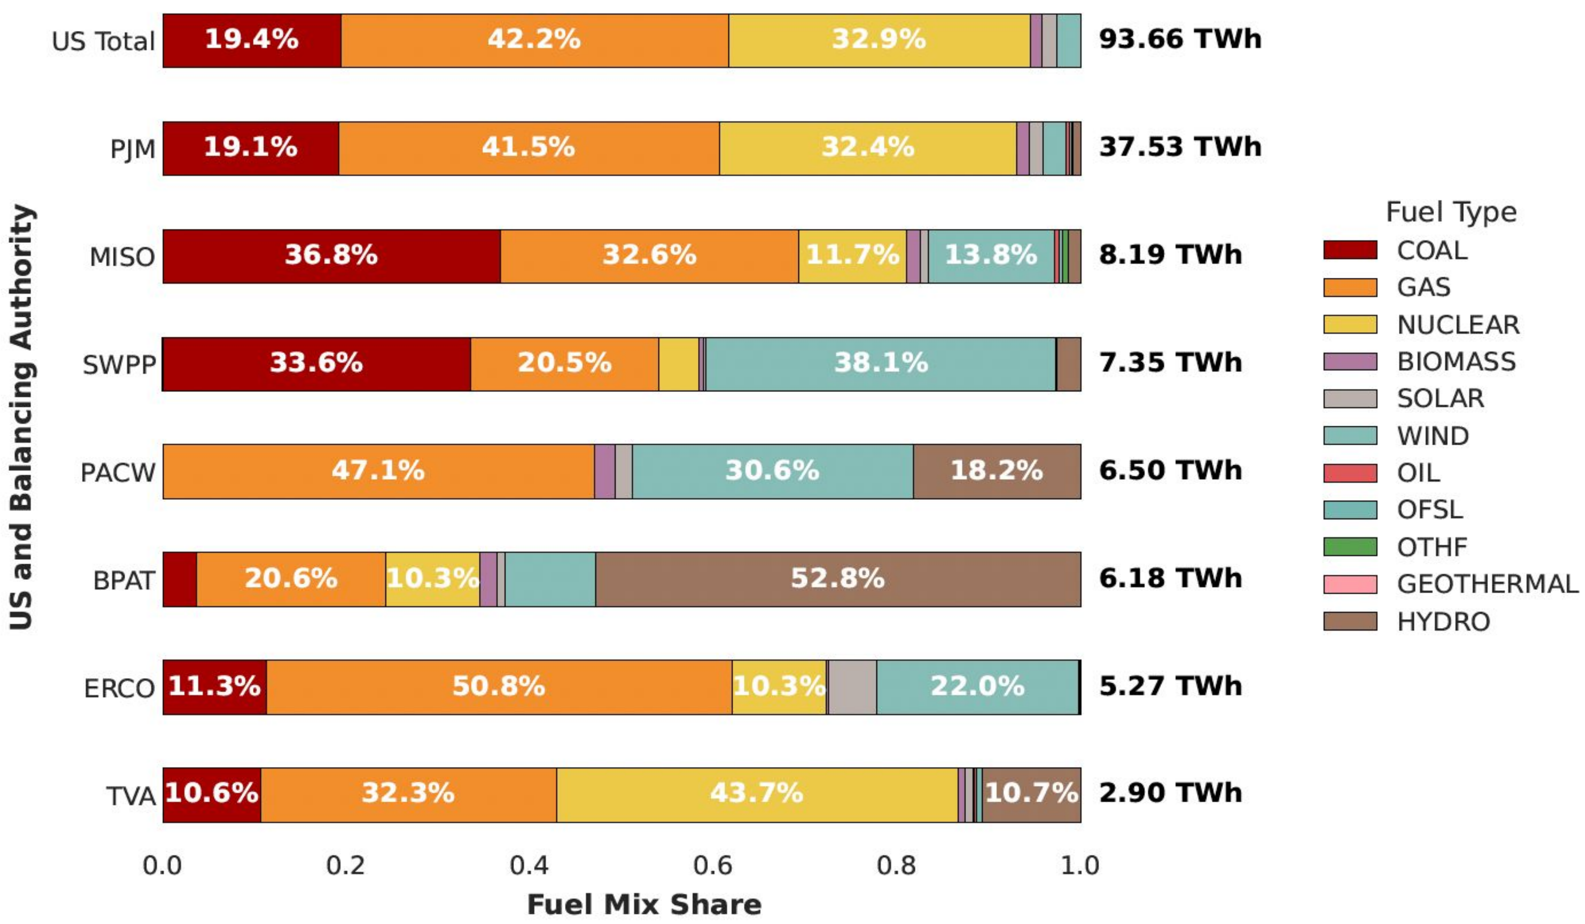

Fig. 4. Fuel mix of power plants supplying electricity for hyperscale US data centers. The top bar represents the distribution of fuel types used by the power plants supplying electricity for hyperscale US data centers in our study. The bottom bars show the largest balancing authorities ranked by aggregated power capacity of hyperscale data centers (shown on the vertical axis), and the amount of electricity produced per fuel type. See fig.S.4.1 for BA regions and corresponding names.

Carbon Emissions Attributable to Hyperscale Data Centers

Total CO2 Emissions from HDCs

52.69M

The total CO2 emissions attributable to the 403 hyperscale data centers (HDCs) amounted to 52.69 million metric tons.

Proportion of US Carbon Emissions

1.10%

This represents approximately 1.10% of the total US carbon emissions from electricity consumption in 2023.

Increase Since 2018

5x

This is more than five times the total emissions reported for HDCs in 2018.

Highest Emissions by State

24.46M

Virginia had the highest CO2 emissions attributable to HDCs, amounting to 24.46 million metric tons.

Significant State Contributions

5.82M

Ohio followed with 5.82 million metric tons of CO2 emissions attributable to HDCs.

- 52.69 M represents the annual CO_2 emissions of a major U.S. city or a sizable portion of the U.S. aviation industry.

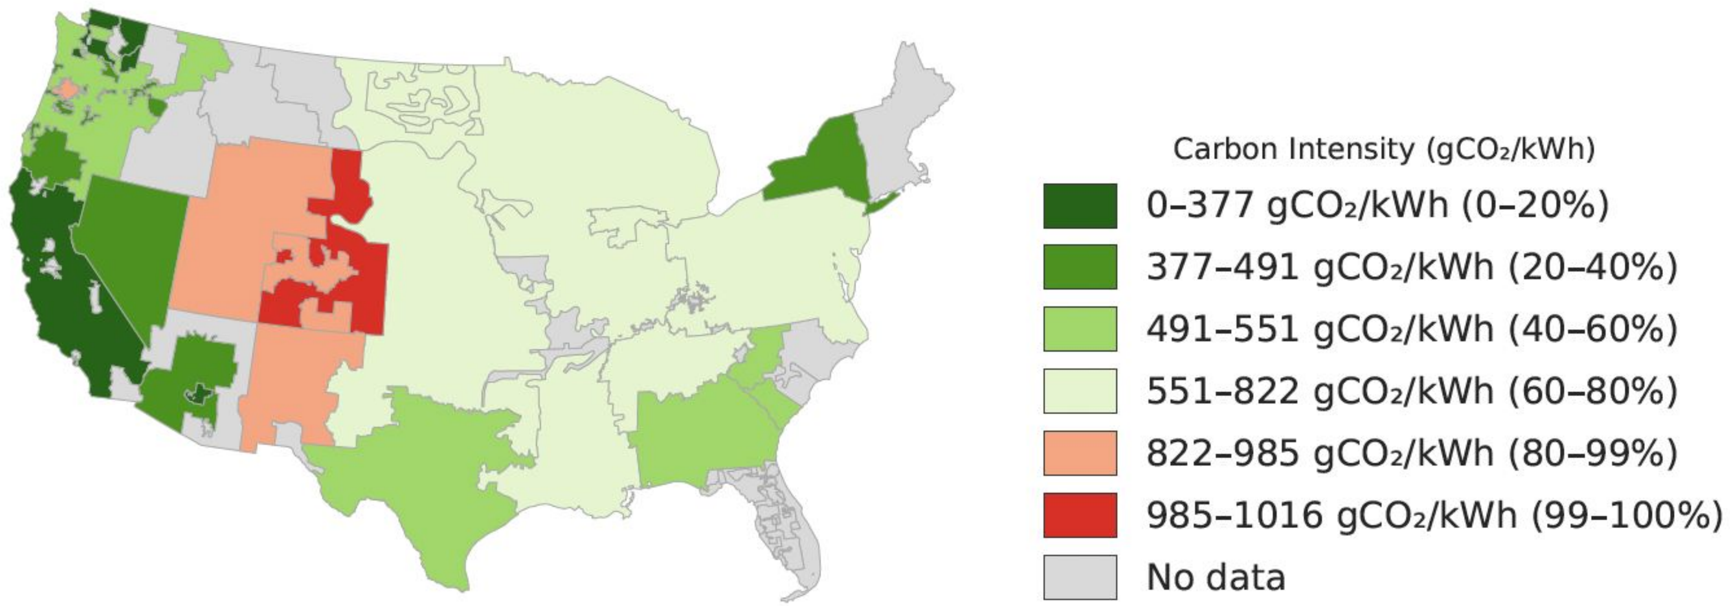

Fig. 3. Carbon intensities of electricity consumption for hyperscale US data centers by balancing authority. Carbon intensity is defined as the amount of carbon dioxide emissions produced per unit of electricity generated, or consumed, and is expressed in units such as grams of CO_2 per kilowatt-hour (gCO_2/kWh) for electricity generation. The figure shows HDCs' carbon intensity for electricity consumption at the balancing authority level, in grams of CO_2 per kWh. Color bins represent percentile-based ranges: 0–20%, 20–40%, 40–60%, 60–80%, 80–99%, and 99–100%.

AI as Eutopia or Dystopia?

The EPA's 2024 regulatory analysis projects that new standards for coal (and some new gas) power plants will cut about 55 million metric tons per year

The total CO_2 emissions attributable to the 403 hyperscale data centers (HDCs) amounted to 52.69 million metric tons.

.

Role of Causal Reasoning AI In Creating A Sustainable Future

1. Huge opportunity for causal reasoning AI to impact policy
2. Still a lot of work to develop AI that can reasonate about causality and with uncertainity
3. Responsible AI is playing an increasing important role

The world is facing enormous challenges use your expertise to solve them!

Livable Cities

Global Prosperity

Human Health

Disaster Response

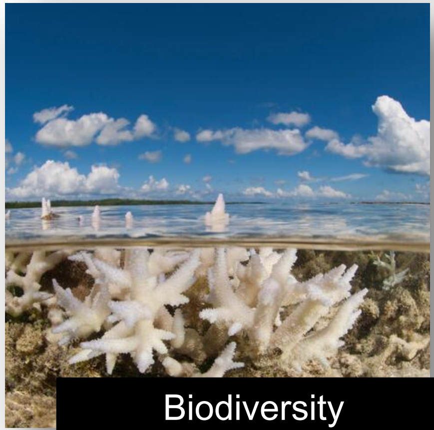

Biodiversity

Clean Energy

- Thank you

- Interpretability
- Reproducibility
- Responsability
- Uncertainty quantification

AI systems require causal intelligence, which I broadly construe as the cognitive ability to learn and exploit causal relationships. It has become increasingly clear that understanding causal intelligence is fundamental to advancing AI and data science.

2010s: Bayesian, Causal, and Machine Learning Synergy

01

Integration of Machine Learning with Bayesian and Causal Methods

Machine learning techniques were integrated with Bayesian and causal methods.

02

Innovations in Causal and Machine Learning Techniques

Innovations included Generalized Propensity Scores (GPS), Causal Forests, Superlearner algorithms, and double robustness approaches.

03

Addressing Complex Data Challenges

Machine learning addressed complex nonlinear patterns and high-dimensional data challenges.

04

Harmonizing Frameworks through Research

Significant technical progress was made to harmonize ML with causal and Bayesian frameworks, as demonstrated by researchers like Xiao, Falco, and Chanmin.

2000s: Integration of Bayesian and Causal Inference

1

Bridging Bayesian and Causal Inference

The potential outcome framework bridged Bayesian statistics with causal inference.

2

Explicit Causal Assumptions

Enabled clearer, more explicit assumptions about causal relationships.

3

Flexible Confounding Adjustment

Confounding adjustment became more flexible through propensity score methods.

4

Incorporating Bayesian Uncertainty

Bayesian uncertainty was incorporated into confounder selection processes.

5

Limitations in Complexity

Despite these advances, applications were mostly limited to relatively simple models.