Causal AI
What is ‘Causal Inference’ and Why is it Key to Machine Learning?

Causal Machine Learning for Healthcare and Precision Medicine

The Causal AI Revolution is Happening Now

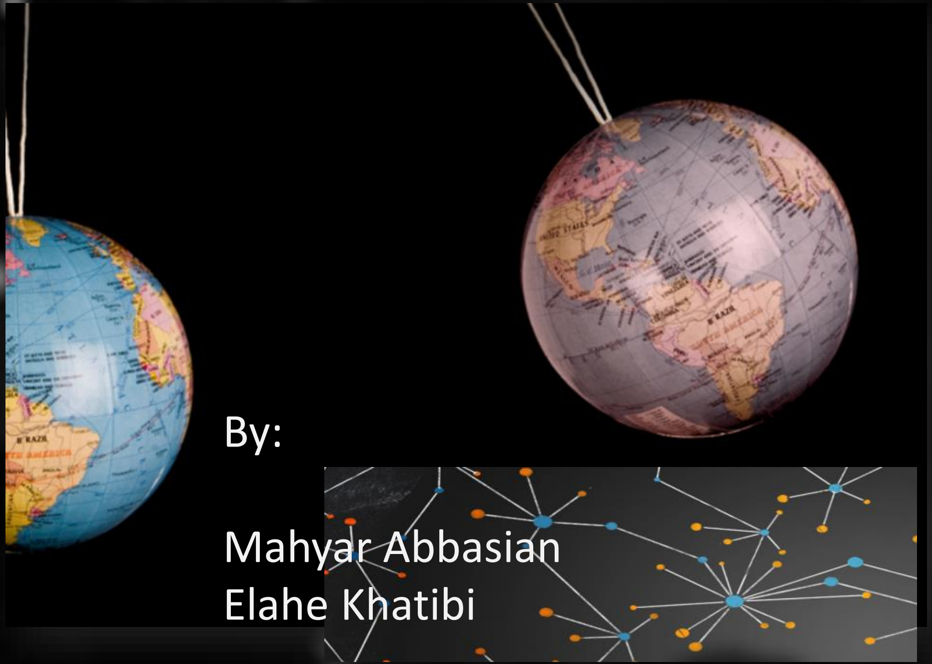

By:

Mahyar Abbasian
Elahe Khatibi

Agenda

- ✓ Background and Motivation
- ✓ Causal AI in Healthcare
- ✓ Alzheimer's disease practical example
- ✓ Our Causal Analysis on Affect Dataset

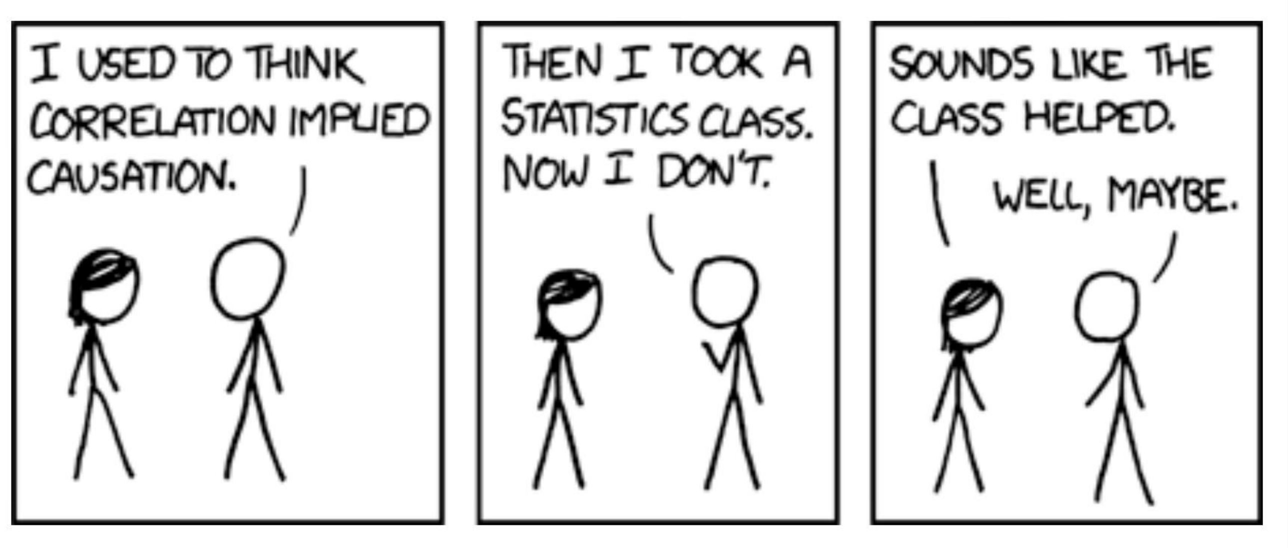

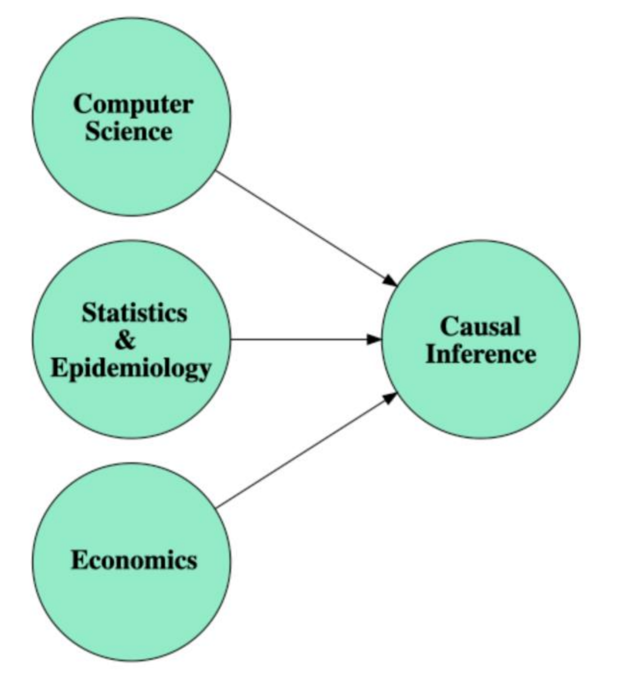

There are three main sources of influence in causal inference: computer science, statistics and epidemiology and econometrics. Active research on causality started in the 80's.

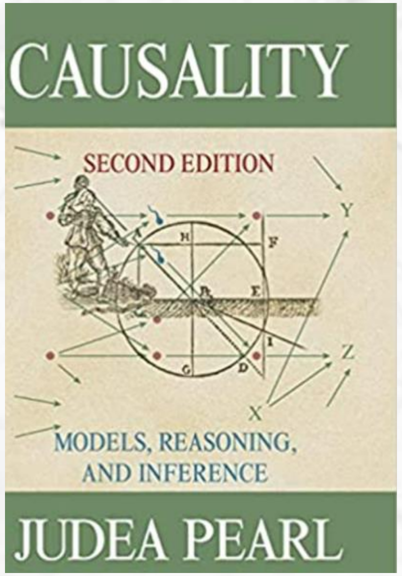

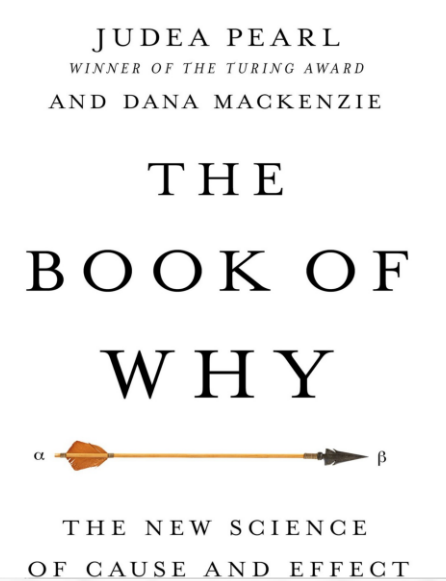

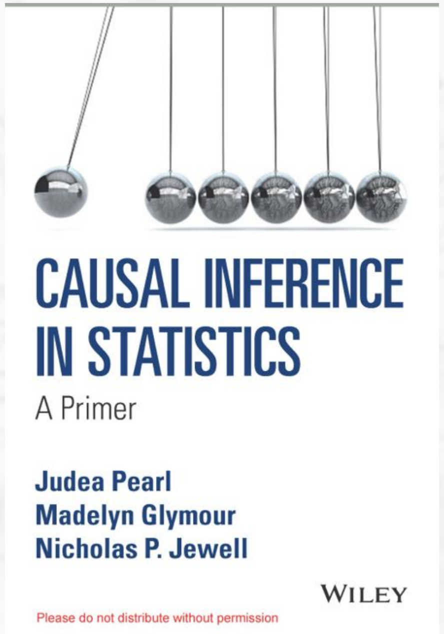

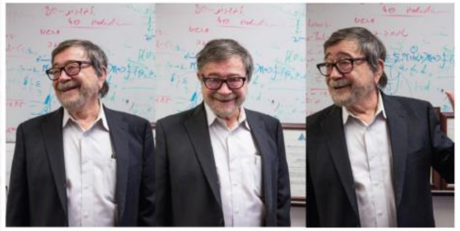

PROFESSOR OF COMPUTER SCIENCE AND DIRECTOR OF THE COGNITIVE SYSTEMS LAB

Professor Judea Pearl is a recipient of the A.M. Turing Award, often called the "Nobel Prize in Computing." The foundations of modern artificial intelligence are built on his breakthrough work, paving the way for AI tech such as driverless cars and voice recognition software.

Why Tech Companies Hire So Many Economists

Traditional ML vs Causal AI

State-of-the-Art AI vs Causal AI

Correlation-based ML	Causal AI
Predictions only	Decision augmentation
Limited explainability	Intrinsic explainability
Spirals out of control in novel situations	Adapts to new conditions
Minimal human-machine interaction	Human-machine partnership
Constrained by historical data	Equipped with machine imagination
No guarantees on fairness	Fair and bias-free
Needs a lot of data	Performs with all datasets, big and small

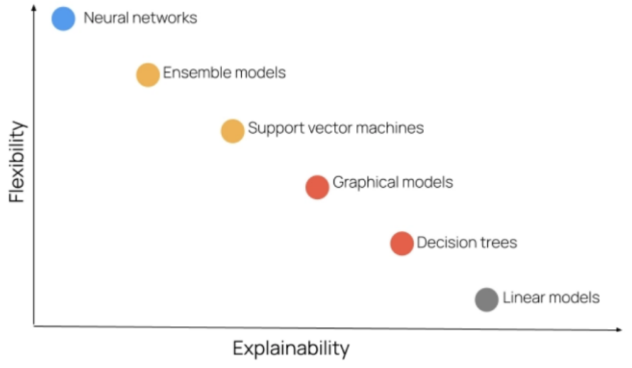

Figure. There is a trade-off between flexibility and explainability in conventional machine learning.

Advantages of Causal AI

Decision-making AI

Causal AI doesn't just predict the future, it shapes it.

Explainable AI

Put the "cause" in "because" with next-generation explainable AI.

Adaptable AI

Causal AI continuously adapts to real-world dynamics.

Human-centric AI

Human-plus-Causal AI partnership allows organizations to harness the benefits of AI.

Imaginative AI

Causal AI can explore hypothetical worlds, uncovering insights that explain why events happened.

Fair AI

AI has a bias problem and Causal AI is the solution.

AI for small data

70% of organizations are shifting their focus from big to small data — Causal AI can help.

Trustworthy AI

Trust is the most important but often-overlooked ingredient in successful AI adoption.

ROYAL SOCIETY OPEN SCIENCE

royalsocietypublishing.org/journal/rsos

Causal machine learning for healthcare and precision medicine

Pedro Sanchez[1], Jeremy P. Voisey[2], Tian Xia[1], Hannah I. Watson[2], Alison Q. O'Neil[1,2] and Sotirios A. Tsaftaris[1]

[1]School of Engineering, University of Edinburgh, Edinburgh, UK
[2]AI Research, Canon Medical Research Europe, Edinburgh, Lothian, UK

PS, 0000-0003-2435-3049

Review

Check for updates

Cite this article: Sanchez P, Voisey JP, Xia T, Watson HI, O'Neil AQ, Tsaftaris SA. 2022 Causal machine learning for healthcare and precision medicine. *R. Soc. Open Sci.* **9**: 220638. https://doi.org/10.1098/rsos.220638

Causal machine learning (CML) has experienced increasing popularity in healthcare. Beyond the inherent capabilities of adding domain knowledge into learning systems, CML

Causal AI Challenges in Healthcare

Modern healthcare data are:

- Insufficient
- Multi-modal (Time-series, Imaging, HER, Annotations)
- High-dimensional
- Often unstructured

Question?

How to acquire the necessary information to causally reason about treatments and outcomes?

ML in Healthcare

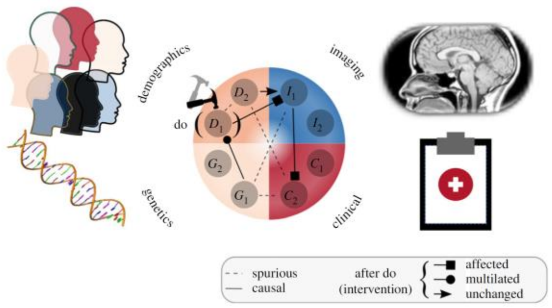

Current ML systems are based on **previous correlations in data**

- Out of distribution data
- bias over dataset
- Insufficient data

Precision medicine (also known as personalized medicine) need to answer complex queries about how individuals would respond to interventions.

Question?

How to increase the accuracy of the ML models for individuals as well as ensure generalization?

Alzheimer's disease practical example

Modelling the data generation process

age range (years)		60–70	70–80	80–90
naive	precision	87.7	91.4	75.5
	recall	92.5	94.2	97.1
counterfactually augmented	precision	88.3	93.6	84.2
	recall	91.5	96.5	95.7

- ✓ Illustration of how a naively trained classifier (a convolutional neural network on ADNI dataset) fails when the data generation process and causal structure are not identified.
  - ➢ Issue: healthy older patients (80–90 years old) are less accurately predicted because ageing itself causes the brain to have Alzheimer's-like patterns.

Results:

After training with counterfactually augmented data, the classifier's precision for the worse performance age group improved.

Alzheimer’s disease practical example

- ✓ Alzheimer’s disease (AD) is a type of cognitive decline that generally appears later in life [6].
- ✓ AD is associated with brain atrophy, i.e. volumetric reduction of grey matter [7].
- ✓ AD causes the symptom of brain morphology change [6-8].
- ✓ It is well established that atrophy also occurs during normal ageing [7].

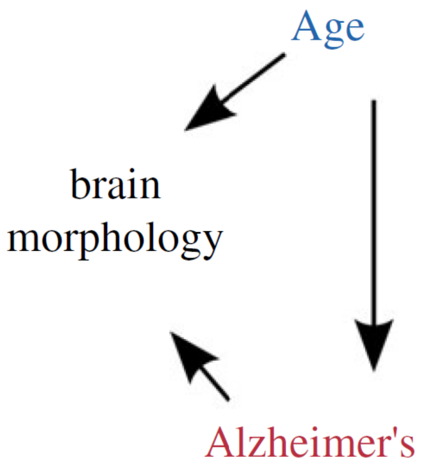

Causal graph

Conclusion:

In this scenario, we can assume that age is a confounder of brain morphology, measured by the MR image, and AD diagnosis.

Alzheimer’s disease practical example

- ✓ To model the effect of having *age* as a *confounder* of brain morphology and AD, conditional generative model has been applied.
- ✓ Synthesize images of a patient at different ages and with different AD status.

Outcome:

- ✓ By observing qualitatively the difference between the baseline and synthesized images, we see that ageing and AD have similar effects on the brain. That is, that both variables change the volume of brain when intervened on independently.
- ✓ This causal knowledge enables the formulation of best strategies for mitigating data bias(es) and improving generalization

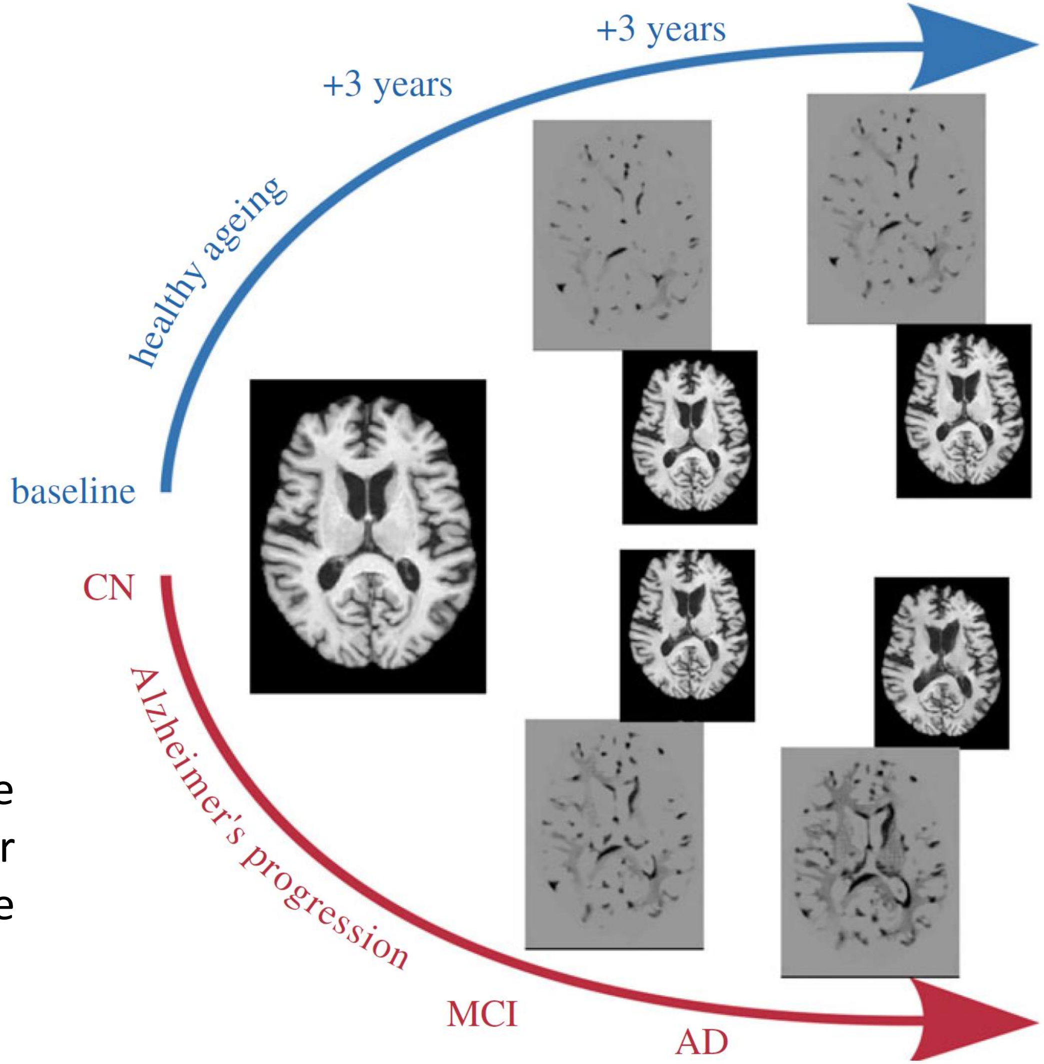

The images with grey background are difference images obtained by subtracting the synthesized image from the baseline.

Proposed Model (Conditional GAN)

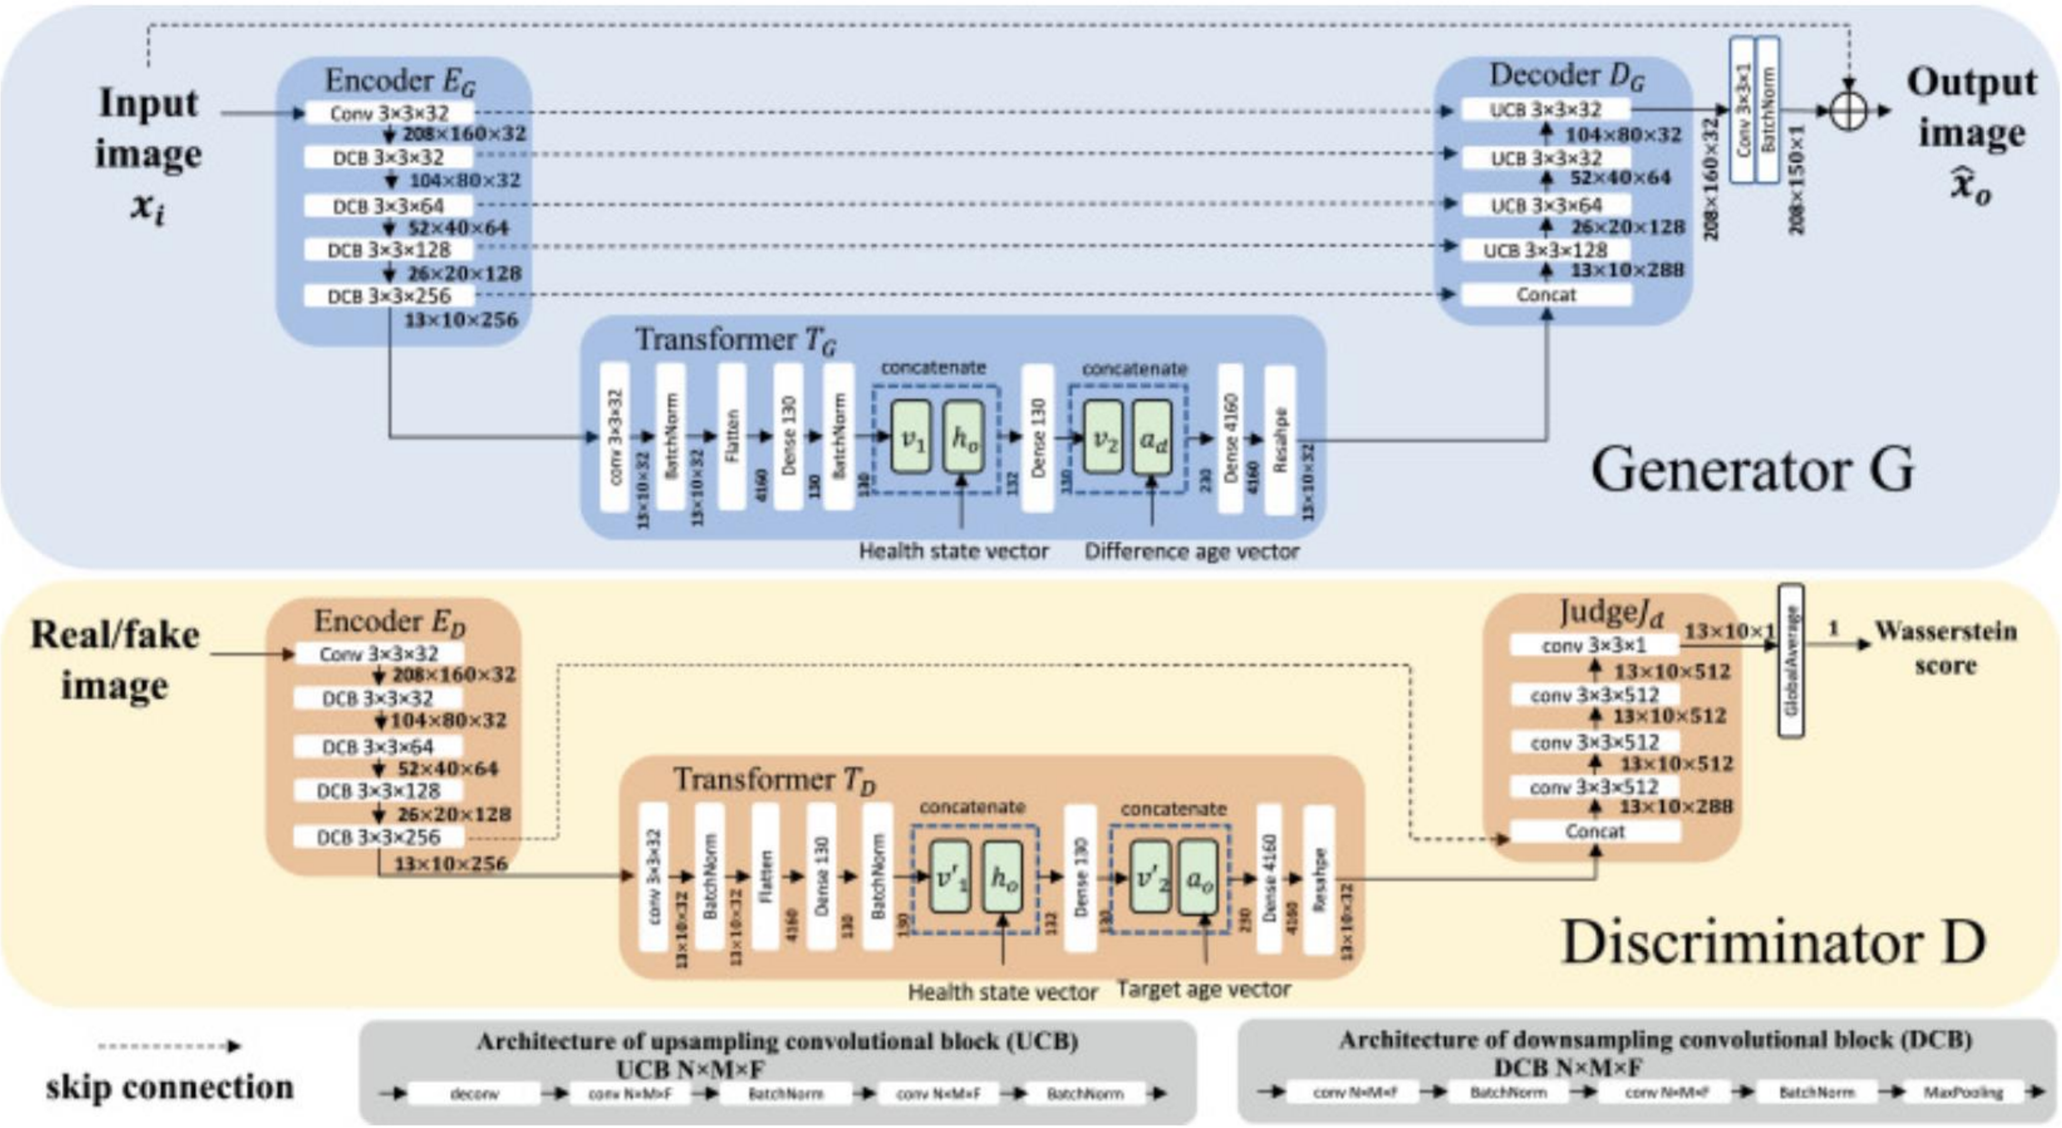

Quantitative results and Summary

- ✓ A method has been presented that learns to simulate subject-specific aged images *without* longitudinal data. It relies on a Generator to generate the images and a Discriminator that captures the joint distribution of brain images and clinical variables, i.e. age and health state (AD status). Also, it offers an embedding mechanism to encode the information of age and health state into our network, and age-modulated and self-reconstruction losses to preserve *subject identity*.
- ✓ Qualitative results showing that this method is able to generate consistent and realistic images conditioned on the target age and health state.

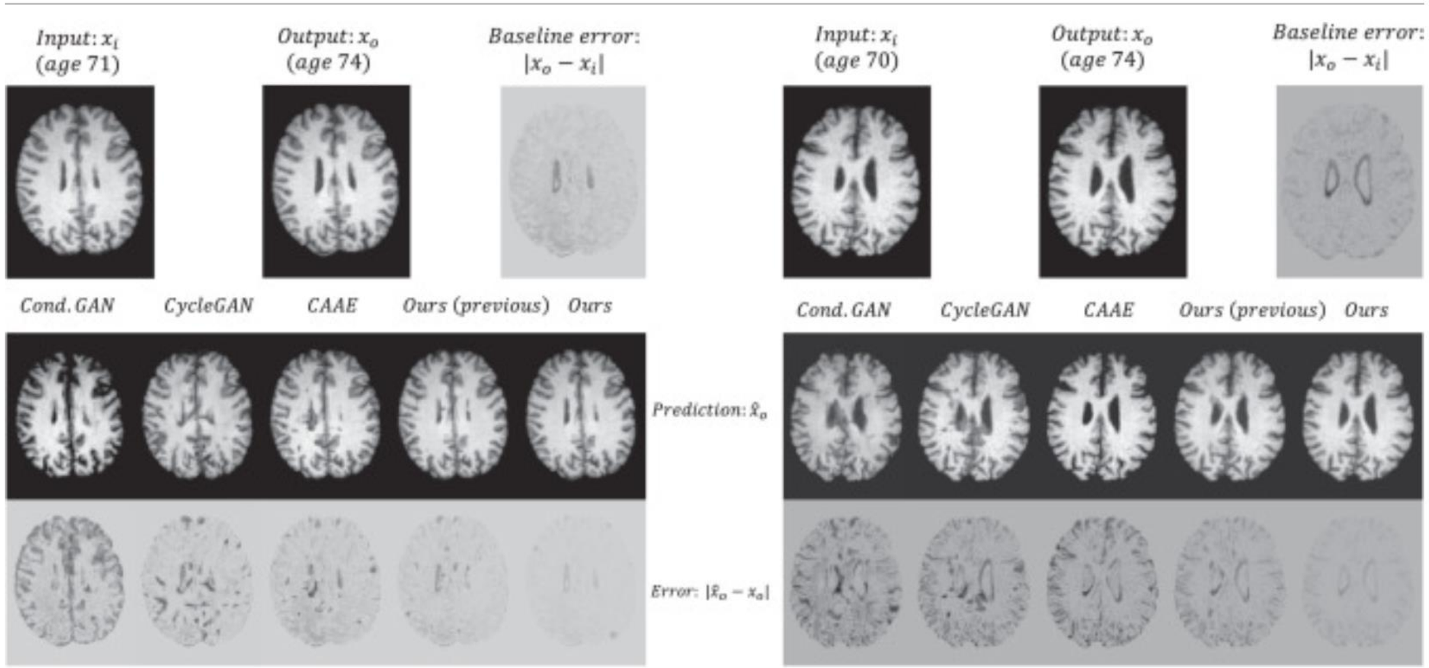

Modelling the data generation process

age range (years)		60–70	70–80	80–90
naive	precision	87.7	91.4	75.5
	recall	92.5	94.2	97.1
counterfactually augmented	precision	88.3	93.6	84.2
	recall	91.5	96.5	95.7

✓ Illustration of how a naively trained classifier (a convolutional neural network on ADNI dataset) fails when the data generation process and causal structure are not identified.

- Issue: healthy older patients (80–90 years old) are less accurately predicted because ageing itself causes the brain to have Alzheimer's-like patterns.

Results:

After training with counterfactually augmented data, the classifier's precision for the worse performance age group improved.

Individualized Treatment Effect

Our goal is to use machine learning to estimate the effect of a treatment on an individual using static or time-series observational data.

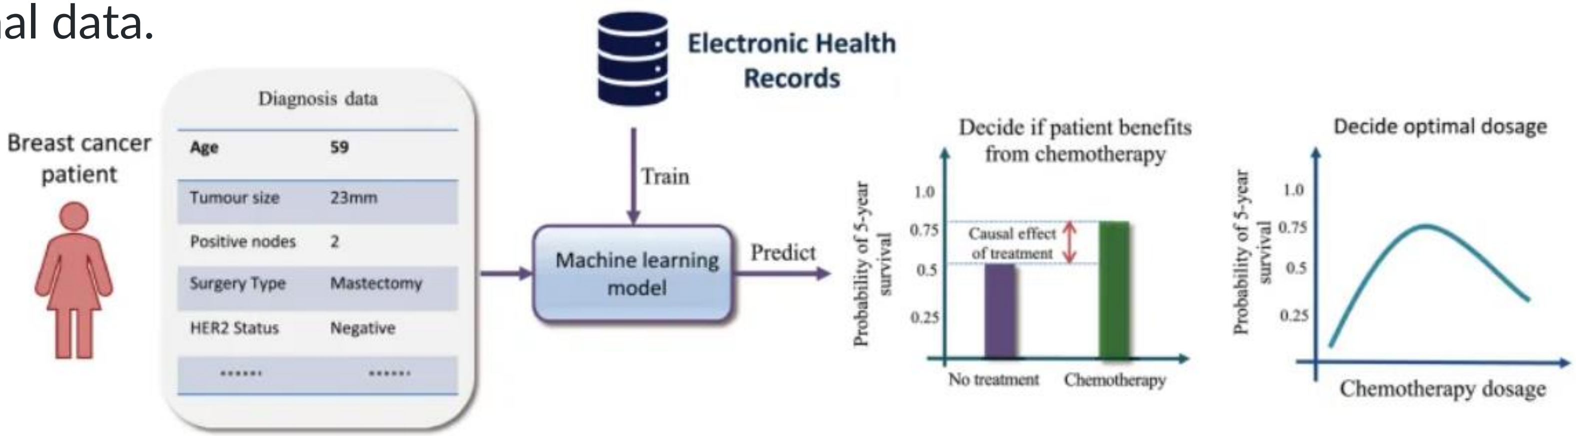

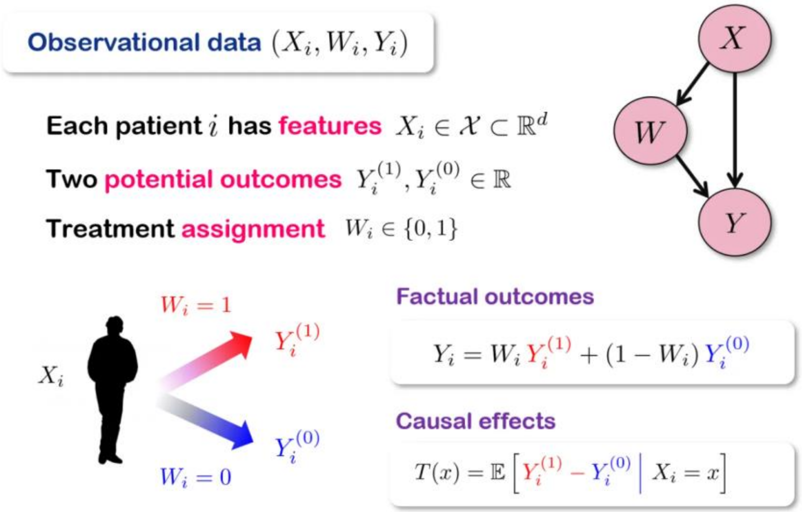

Treatment effect and precision medicine

Difference between individualized and average treatment effect (ITE versus ATE).

$$E[Y^1|X_i = 1] - E[Y^0|X_i = 0] = \mathrm{ATE}$$

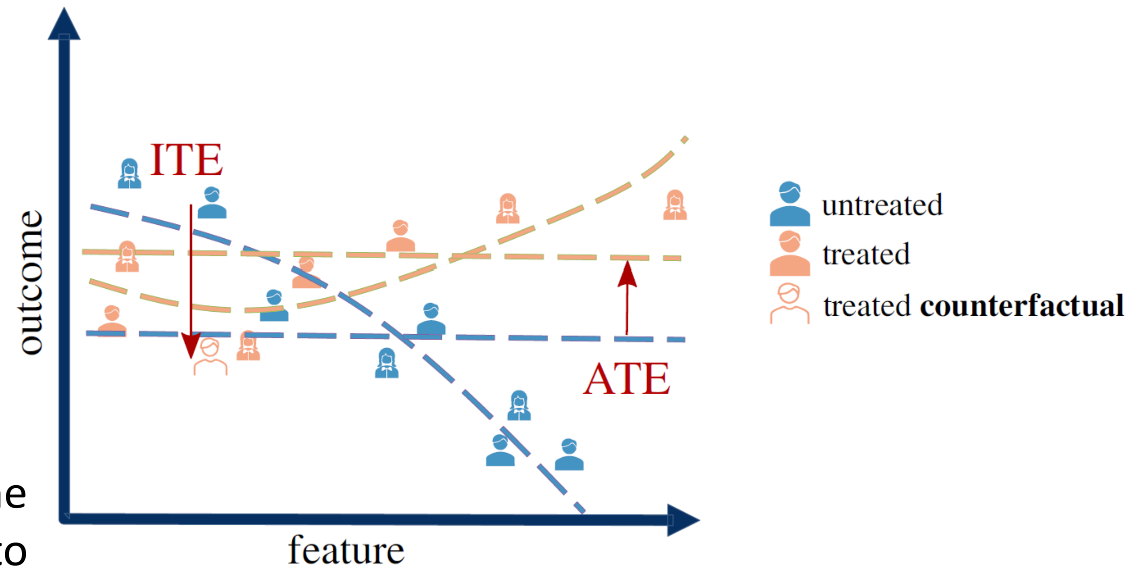

The ITE for each patient is the difference between actual and the counterfactual outcome. We show an example counterfactual to highlight that ITE for some patients might differ from the average (ATE).

Features: (patient characteristics)

Outcome: measure of response to the treatment

Conclusions: ITE and precision medicine

- ✓ The estimation of treatment affect using observational data is subject to *confounding* as patient characteristics affect both the selection of treatment and outcome. Causal inference methods need to mitigate this.

- ✓ Employing causal inference methods to estimate individualized treatment effects, we can understand which patients benefit from certain medication and which patients do not, thus enabling us to make personalized treatment recommendations

Conclusion: Causal machine learning for complex data (Future Research Trends)

- ✓ Multi-modal data
- ✓ Temporal data
- ✓ Out-of-distribution generalization with unstructured and high-dimensional data

Research directions in causal machine learning

- ✓ Causal Representations
- ✓ Causal Discovery
- ✓ Causal Reasoning

Discovering Causal Signals in Images

David Lopez-Paz
Facebook AI Research
dlp@fb.com

Robert Nishihara
UC Berkeley
rkn@eecs.berkeley.edu

Soumith Chintala
Facebook AI Research
soumith@fb.com

Bernhard Schölkopf
MPI for Intelligent Systems
bs@tue.mpg.de

Léon Bottou
Facebook AI Research
leon@bottou.org

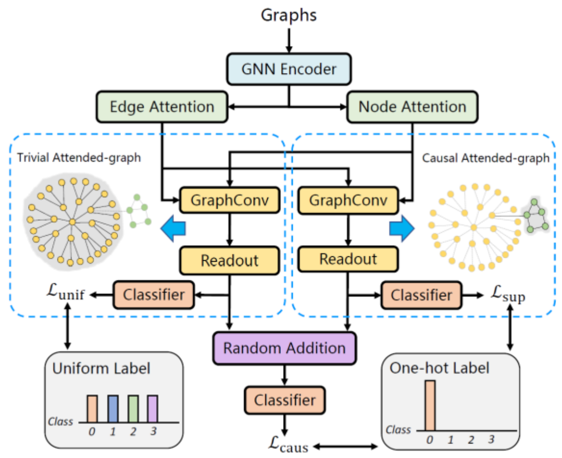

Figure 2: The overview of the proposed Causal Attention Learning (CAL) framework.

RESEARCH-ARTICLE

Causal Attention for Interpretable and Generalizable Graph Classification

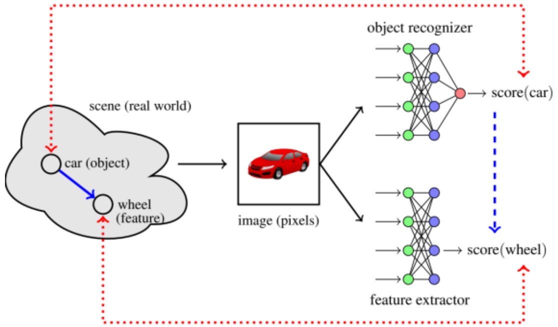

Causal Generative Neural Networks

Olivier Goudet* [1], **Diviyan Kalainathan*** [1], **Philippe Caillou**[1], **Isabelle Guyon**[1], **David Lopez-Paz**[2], **Michèle Sebag**[1]
[1]TAU, CNRS – INRIA – LRI, Univ. Paris-Sud, Univ. Paris-Saclay
[2]Facebook AI Research

Causal Inference in Natural Language Processing: Estimation, Prediction, Interpretation and Beyond

Amir Feder[‡,10], Katherine A. Keith[2], Emaad Manzoor[3], Reid Pryzant[4], Dhanya Sridhar[5], Zach Wood-Doughty[6], Jacob Eisenstein[7], Justin Grimmer[8], Roi Reichart[1], Margaret E. Roberts[9], Brandon M. Stewart[10], Victor Veitch[7,11], and Diyi Yang[12]

Towards Causal Representation Learning

Bernhard Schölkopf †, Francesco Locatello †, Stefan Bauer *, Nan Rosemary Ke *, Nal Kalchbrenner
Anirudh Goyal, Yoshua Bengio

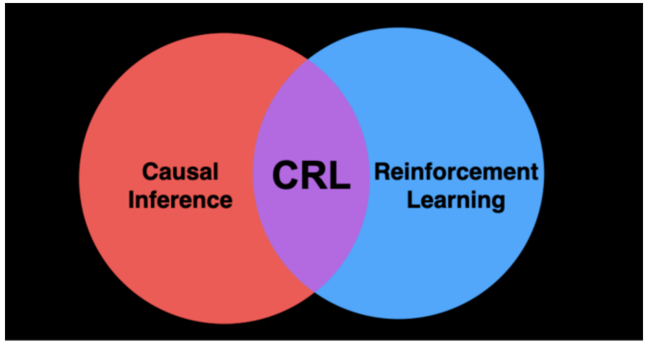

Causal Reinforcement Learning

Causal Autoregressive Flows

Ilyes Khemakhem*
Gatsby Unit, UCL

Ricardo P. Monti*
Gatsby Unit, UCL

Robert Leech
King's College London

Aapo Hyvärinen
University of Helsinki

Our Causal Analysis on Affect Dataset

Causal Discovery

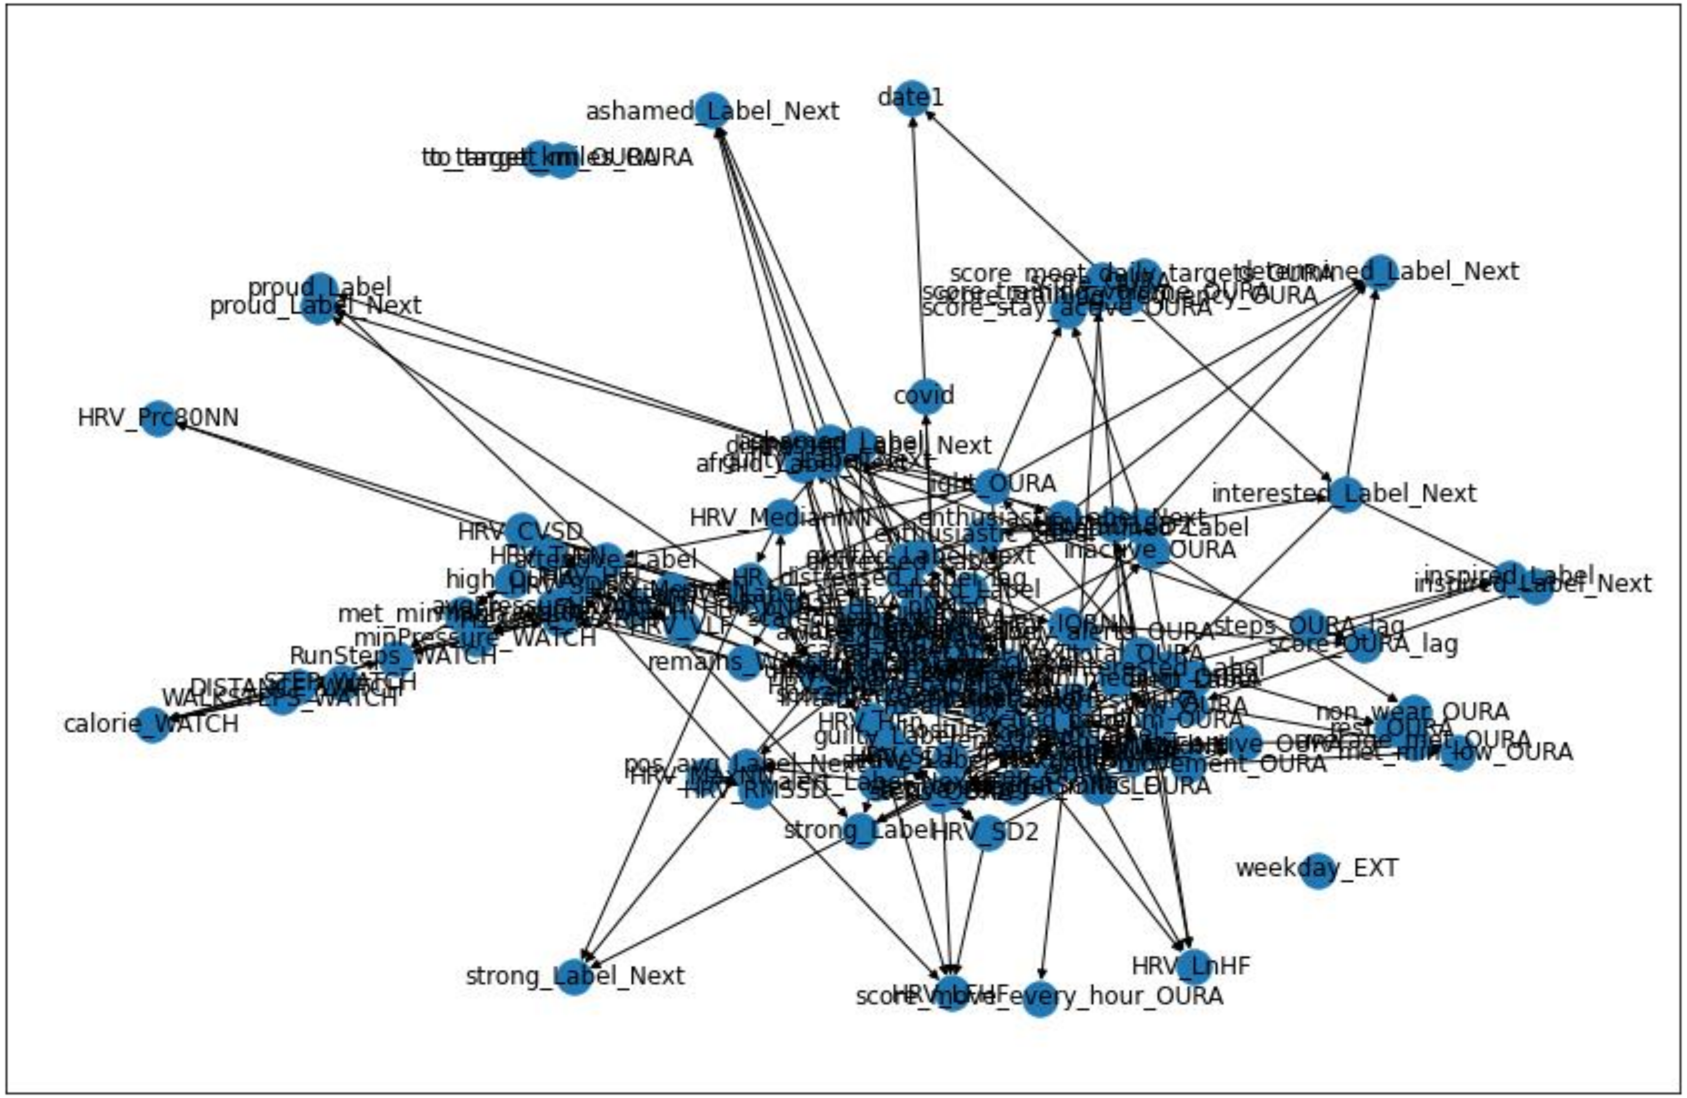

Causal Reasoning

Effect of COVID on Sleep quality


```
Increasing the treatment variable(s) [covid] from [0] to [1] causes an increase of -4.1337892753814245 in the expected value
of the outcome [score_OURA], over the data distribution/population represented by the dataset.
```

Causal Reasoning

Multi-treatment effect on Sleep quality

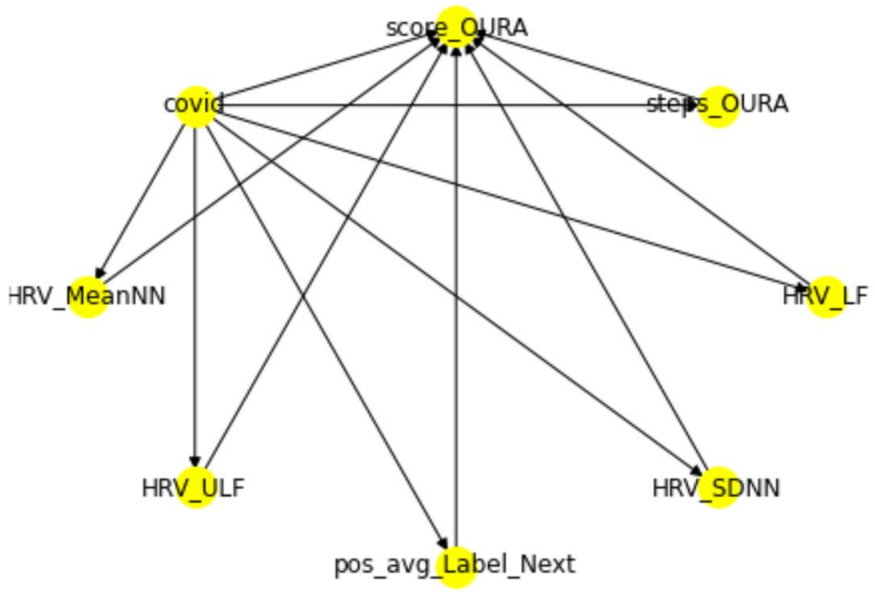

Causal Estimate is 4.125814211984448

Causal Reasoning

Multi-treatment effect on distressed

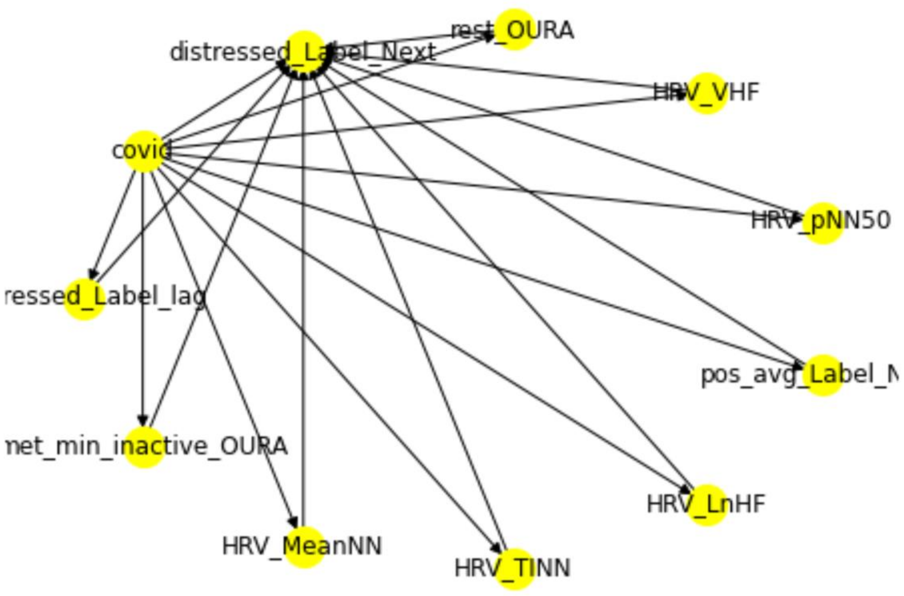

Causal Estimate is 14.946463892517530

Conclusion

- “Causal models have a unique combination of capabilities that could enable a deeper understanding of complex systems and allow us to better align decision systems with society’s values.”

DeepMind

- “Causal inference may be the new frontier as we migrate away from association-based analysis only.”

Swiss Re Institute

- “Causal inference and machine learning have a symbiotic relationship that is growing deeper… We have made a large investment in Causal AI.”

Netflix Research

Questions?

1. What do “Causal Invariant Features” mean?

2. What is Causal Reinforcement Learning area does? How Causality helps RL?

References

[1] Sanchez, Pedro, et al. "Causal machine learning for healthcare and precision medicine." *Royal Society Open Science* 9.8 (2022): 220638.

[2] Gao, Junyi, et al. "STAN: spatio-temporal attention network for pandemic prediction using real-world evidence." *Journal of the American Medical Informatics Association* 28.4 (2021): 733-743.

[3] Prosperi, Mattia, et al. "Causal inference and counterfactual prediction in machine learning for actionable healthcare." *Nature Machine Intelligence* 2.7 (2020): 369-375.

[4] Zhang, Wenhao, Ramin Ramezani, and Arash Naeim. "Causal Inference in Medicine and in Health Policy: A Summary." *HANDBOOK ON COMPUTER LEARNING AND INTELLIGENCE: Volume 2: Deep Learning, Intelligent Control and Evolutionary Computation*. 2022. 263-302.

[5] Wang, Lijing, et al. "CausalGNN: Causal-based Graph Neural Networks for Spatio-Temporal Epidemic Forecasting." (2022).

[6] Shen, Xinpeng, et al. "Challenges and opportunities with causal discovery algorithms: application to Alzheimer's pathophysiology." Scientific reports 10.1 (2020): 2975.

[7] Uleman, Jeroen F., et al. "Mapping the multicausality of Alzheimer's disease through group model building." GeroScience 43 (2021): 829-843.

[8] Xia, Tian, et al. "Learning to synthesise the ageing brain without longitudinal data." Medical Image Analysis 73 (2021): 102169.

[9] Sullivan, Edith V., et al. "Age-related decline in MRI volumes of temporal lobe gray matter but not hippocampus." Neurobiology of aging 16.4 (1995): 591-606.

[10] Schölkopf, Bernhard, et al. "On causal and anticausal learning." arXiv preprint arXiv:1206.6471 (2012).

[11] Kilbertus, Niki, Giambattista Parascandolo, and Bernhard Schölkopf. "Generalization in anti-causal learning." arXiv preprint arXiv:1812.00524 (2018).

[12] Heinze-Deml, Christina, and Nicolai Meinshausen. "Conditional variance penalties and domain shift robustness." Machine Learning 110.2 (2021): 303-348.


THANK YOU

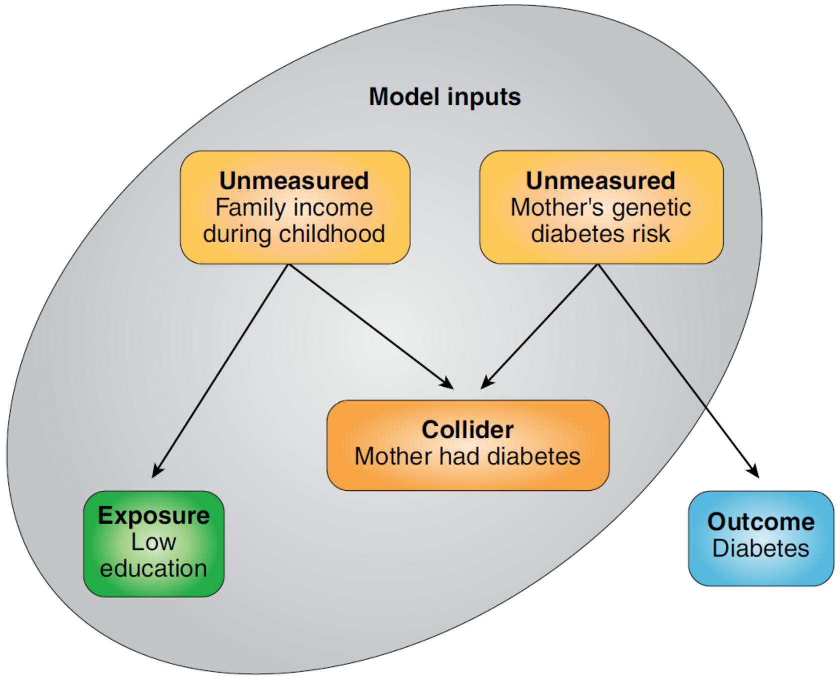

Fig. 3 | An example of M-bias. When estimating the effect of education level on diabetes risk, mother's history of diabetes could be mistaken as a confounder and included in a model, but it is a collider by the effect of history of family income and genetic risk.

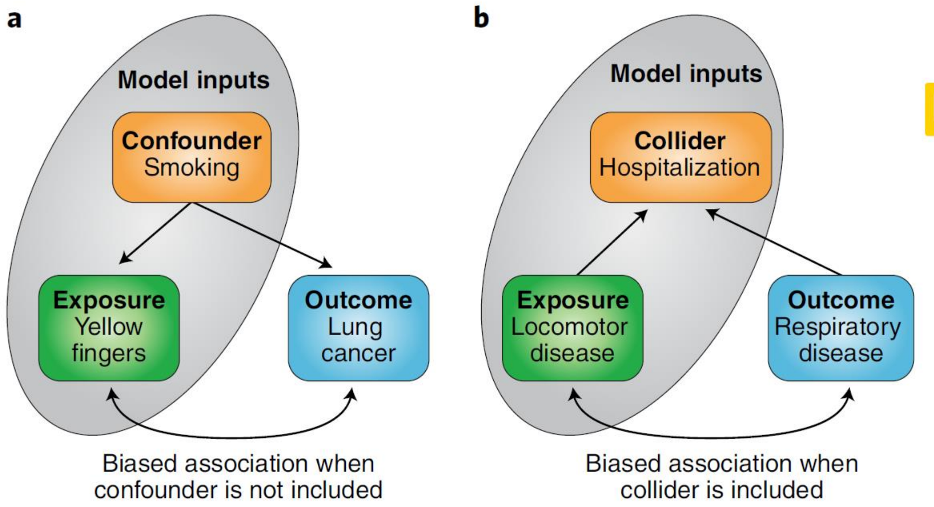

Fig. 2 | Examples of confounding bias and collider bias. a,**b**, Confounding (**a**) can occur when there exists a common cause for both exposure and outcome, while a collider (**b**) is a common effect of both exposure and outcome. Not including a confounder or including a collider in a model results in biased associations.