Learning Structural Causal Models through Deep Generative Models

Survey on Methods, Guarantees, and Challenges

Audrey Poinsot, Alessandro Leite, Nicolas Chesneau, Michèle Sébag, Marc Schoenauer
audrey.poinsot@ekimetrics.com, alessandro.leite@inria.fr

IJCAI 2024, 8th August 2024

Ekimetrics.

université PARIS-SACLAY

Inria

IJCAI JEJU 2024

SCAN ME

Correlation is not causation

Does chocolate make you smart?

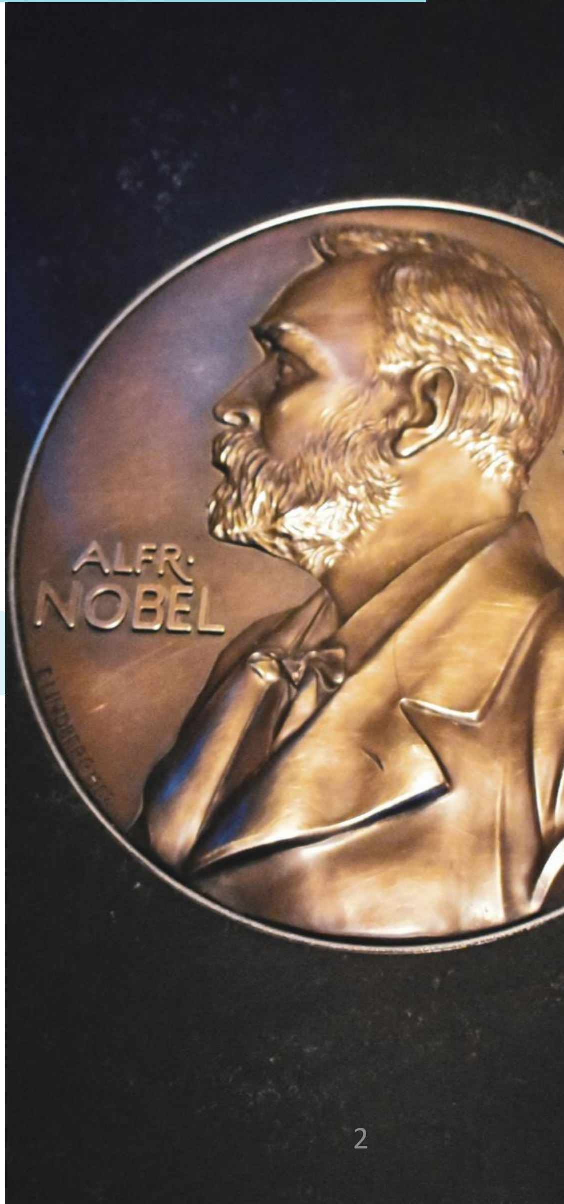

[Messerli, 2012]

Correlation is not causation

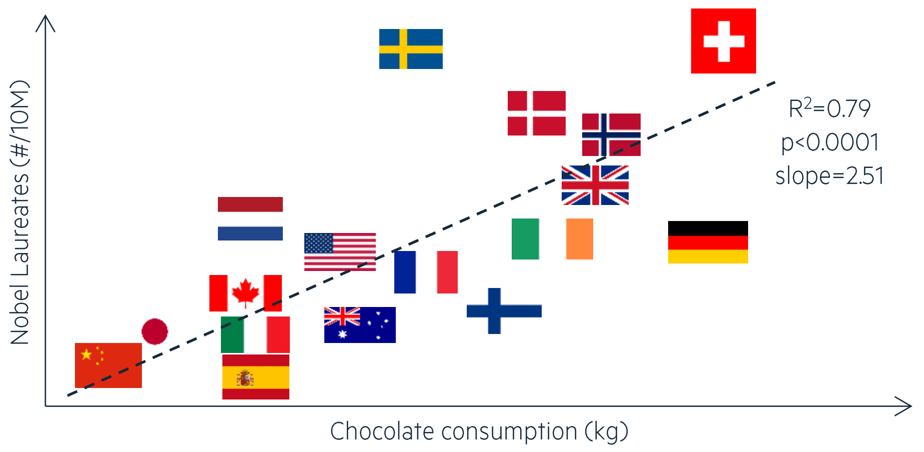

Does chocolate make you smart?

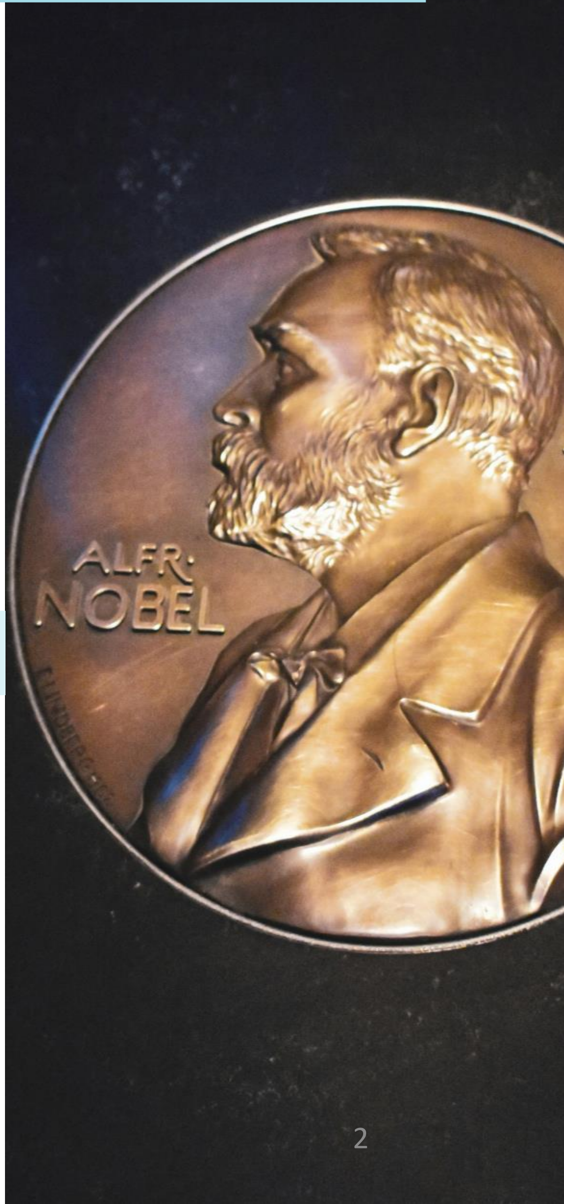

[Messerli, 2012]

Correlation is not causation

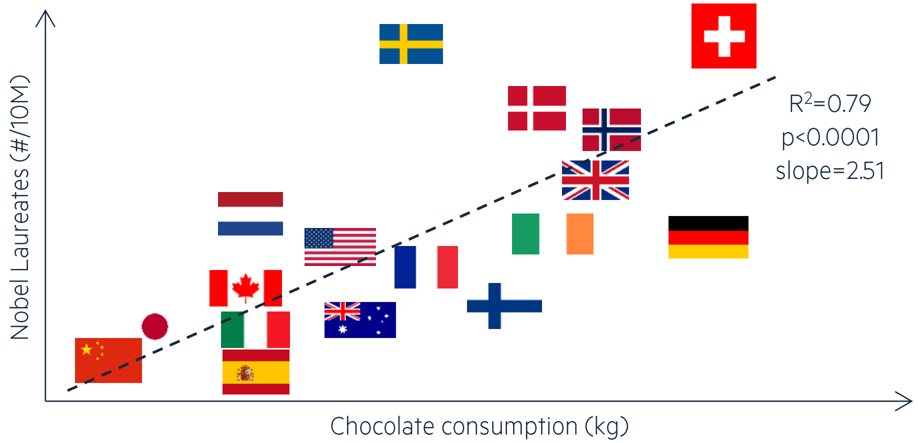

Does chocolate make you smart?

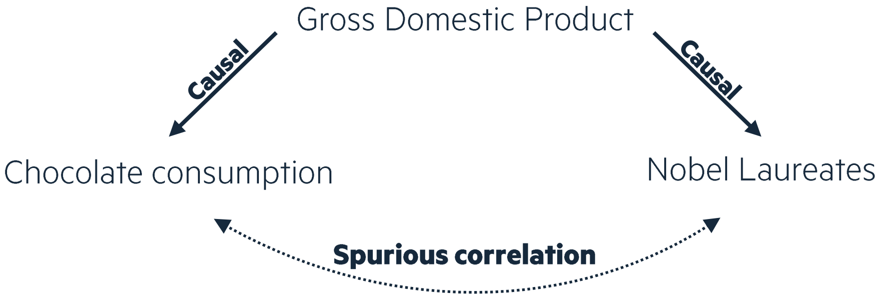

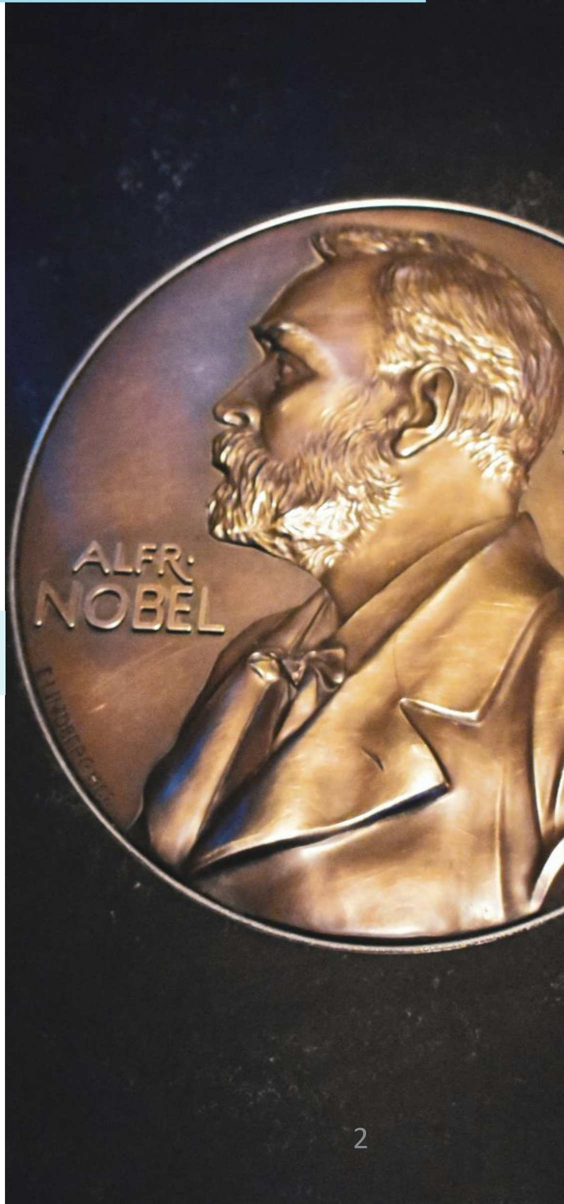

[Messerli, 2012]

Causal Inference to the rescue

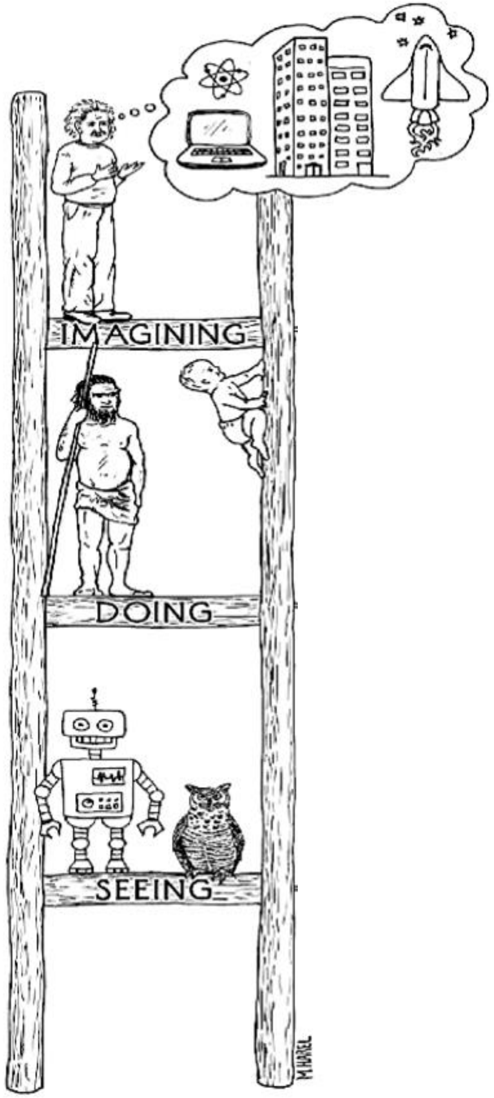

Figure inspired from [Pearl and Mackenzie, 2018]

Causal Inference to the rescue

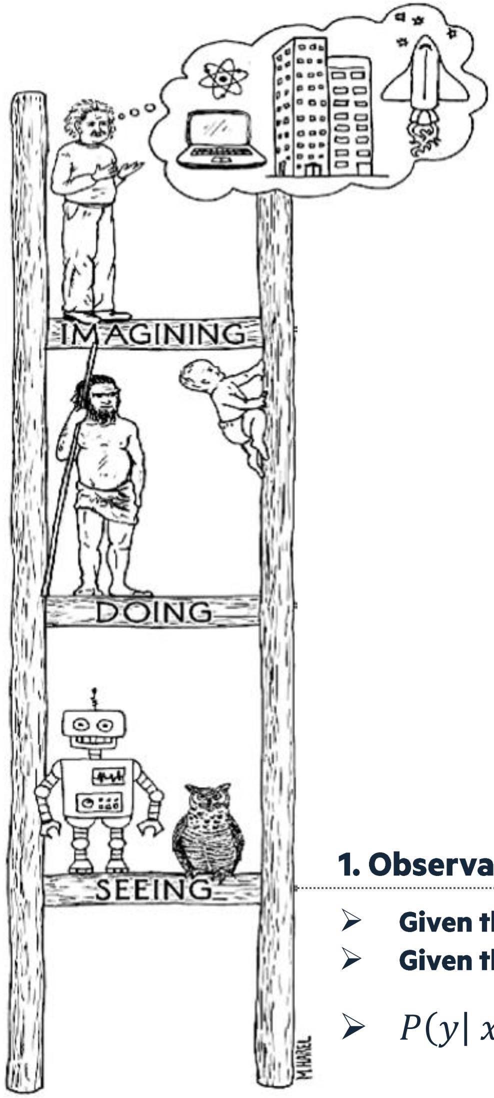

1. Observations – associations (L_1)

- **Given that** I eat chocolate, **how likely am I** to win a Nobel?
- **Given that** I am a woman, **how likely am I** to get a loan?
- $P(y|\, x)$

Figure inspired from [Pearl and Mackenzie, 2018]

Causal Inference to the rescue

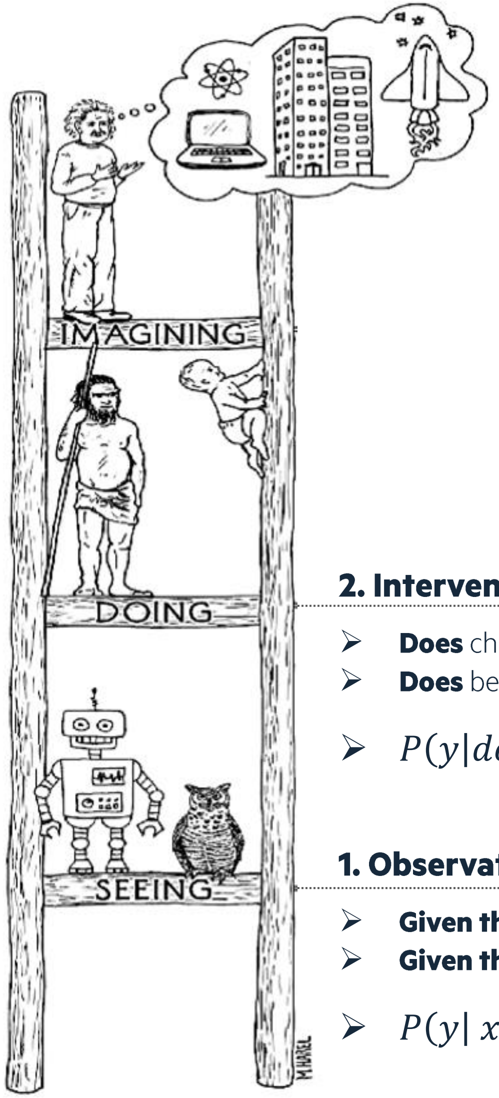

2. Interventions – action-guidance (L_2)

- **Does** chocolate **make you** smarter?
- **Does** being a woman **reduce your** chances of getting a loan?
- $P(y|do(x))$ → L_2 assumptions (e.g. causal graph, adjustment sets)

1. Observations – associations (L_1)

- **Given that** I eat chocolate, **how likely am I** to win a Nobel?
- **Given that** I am a woman, **how likely am I** to get a loan?
- $P(y| x)$

Figure inspired from [Pearl and Mackenzie, 2018]

Causal Inference to the rescue

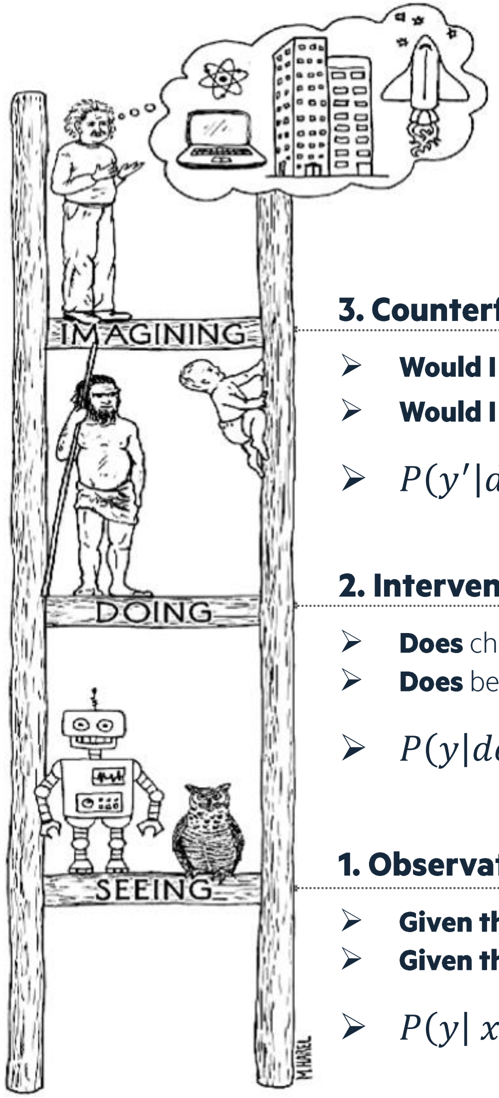

3. Counterfactuals – against existing observations (L_3)

- **Would I** have won the Nobel **if I had** eaten chocolate?
- **Would I** have got the loan **if I had** been a man?
- $P(y'|do(x'), y, x)$ → L_3 assumptions (e.g. exogeneous noise)

2. Interventions – action-guidance (L_2)

- **Does** chocolate **make you** smarter?
- **Does** being a woman **reduce your** chances of getting a loan?
- $P(y|do(x))$ → L_2 assumptions (e.g. causal graph, adjustment sets)

1. Observations – associations (L_1)

- **Given that** I eat chocolate, **how likely am I** to win a Nobel?
- **Given that** I am a woman, **how likely am I** to get a loan?
- $P(y|\ x)$

Figure inspired from [Pearl and Mackenzie, 2018]

Causal Inference to the rescue

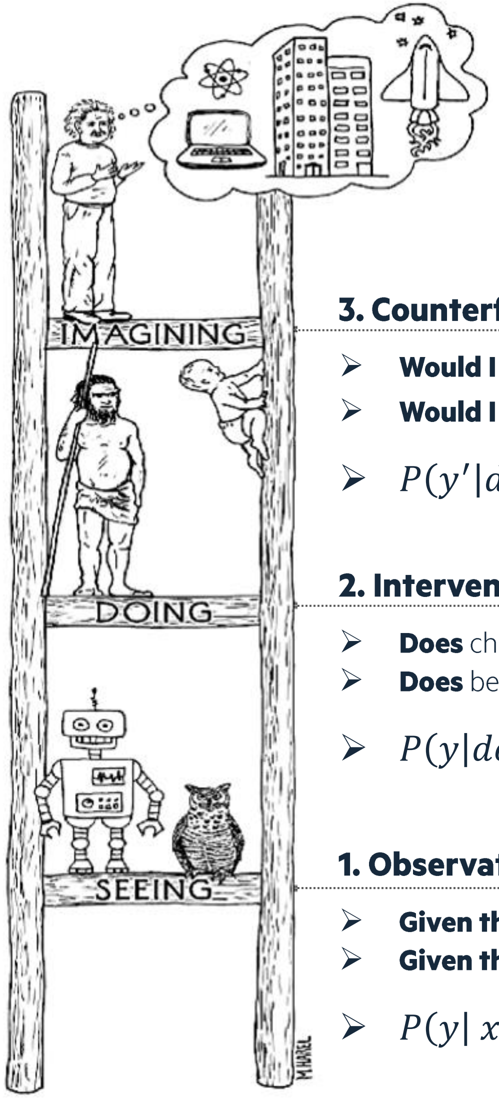

3. Counterfactuals – against existing observations (L_3)

- **Would I** have won the Nobel **if I had** eaten chocolate?
- **Would I** have got the loan **if I had** been a man?
- $P(y'|do(x'), y, x)$ → L_3 assumptions (e.g. exogeneous noise)

2. Interventions – action-guidance (L_2)

- **Does** chocolate **make you** smarter?
- **Does** being a woman **reduce your** chances of getting a loan?
- $P(y|do(x))$ → L_2 assumptions (e.g. causal graph, adjustment sets)

1. Observations – associations (L_1)

- **Given that** I eat chocolate, **how likely am I** to win a Nobel?
- **Given that** I am a woman, **how likely am I** to get a loan?
- $P(y|\, x)$

Definition

A **Structural Causal Model** (SCM) is a tuple $M \coloneqq (F, P(U))$ where F comprises a set of d structural equations, f_i, one for each endogenous random variable $X_i \in \boldsymbol{X}$:
$F = \{X_i \coloneqq f_i(PA(X_i), U_i)\}_{i\in[1,d]}$ with $PA(X_i)$ the parents of X_i and U_i the exogeneous noise

Definition from [Pearl, 2009]

Example

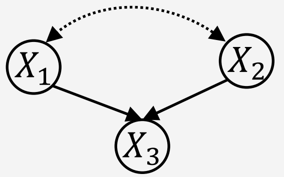

$$X_1 = f_1(U_1)$$
$$X_2 = f_2(U_2)$$
$$X_3 = f_3(X_1, X_2, U_3)$$

with $P(U)$ s.t. $U_3 \perp\!\!\!\perp U_1$ and $U_3 \perp\!\!\!\perp U_2$

Figure inspired from [Pearl and Mackenzie, 2018]

Causal Inference to the rescue

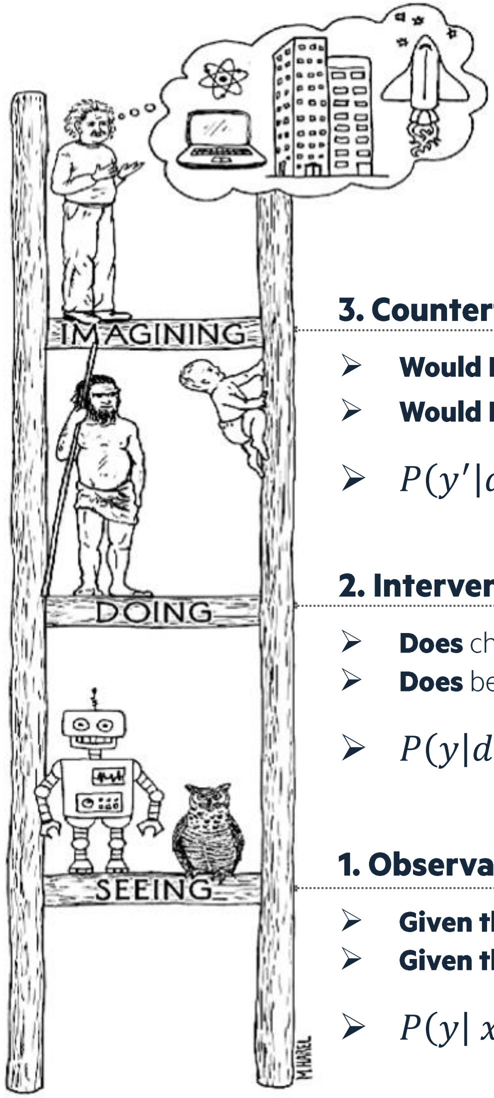

3. Counterfactuals – against existing observations (L_3)

- **Would I** have won the Nobel **if I had** eaten chocolate?
- **Would I** have got the loan **if I had** been a man?
- $P(y'|do(x'), y, x)$ → L_3 assumptions (e.g. exogeneous noise)

2. Interventions – action-guidance (L_2)

- **Does** chocolate **make you** smarter?
- **Does** being a woman **reduce your** chances of getting a loan?
- $P(y|do(x))$ → L_2 assumptions (e.g. causal graph, adjustment sets)

1. Observations – associations (L_1)

- **Given that** I eat chocolate, **how likely am I** to win a Nobel?
- **Given that** I am a woman, **how likely am I** to get a loan?
- $P(y|\, x)$

Figure inspired from [Pearl and Mackenzie, 2018]

Definition

A **Structural Causal Model** (SCM) is a tuple $M \coloneqq (F, P(U))$ where F comprises a set of d structural equations, f_i, one for each endogenous random variable $X_i \in \boldsymbol{X}$:
$F = \{X_i \coloneqq f_i(PA(X_i), U_i)\}_{i\in[1,d]}$ with $PA(X_i)$ the parents of X_i and U_i the exogeneous noise

Definition from [Pearl, 2009]

Example

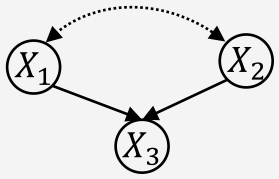

$X_1 = f_1(U_1)$
$X_2 = f_2(U_2)$
$X_3 = f_3(X_1, X_2, U_3)$

with $P(U)$ s.t. $U_3 \perp\!\!\!\perp U_1$ and $U_3 \perp\!\!\!\perp U_2$

Scope of the survey

Structural Causal Models
Numerous assumptions
Identification concerns

Deep Generative Models
Flexibility, few assumptions
Few guarantees

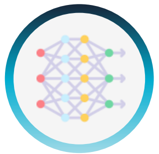

Learning Structural Causal Models through Deep Generative Models
Existing works, capabilities, and remaining open questions

Survey on methods, guarantees and challenges

Classification

Challenges & Opportunities

Survey on methods, guarantees and challenges

Classification

Learning an SCM with DGMs

- **SCM class**
  - Deep Structural Causal Models
    - Neural Causal Models
  - Bijective Generation Mechanism
- **DGM class**
  - Invertible Explicit
  - Amortised Explicit
  - Amortised Implicit

Challenges & Opportunities

Survey on methods, guarantees and challenges

Classification

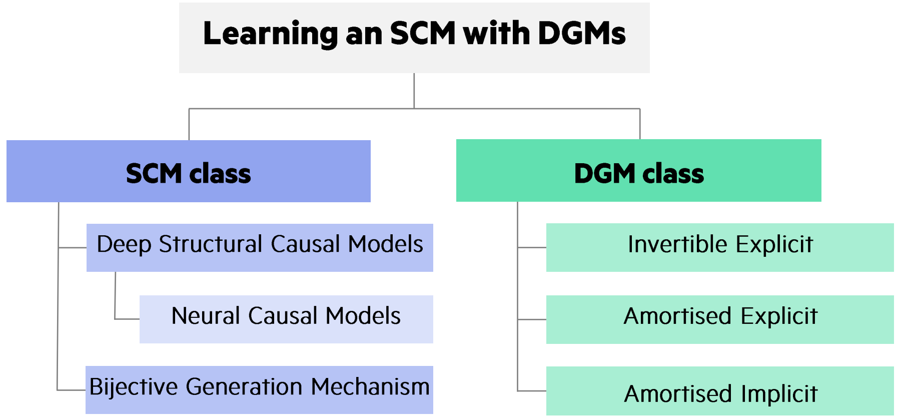

Challenges & Opportunities

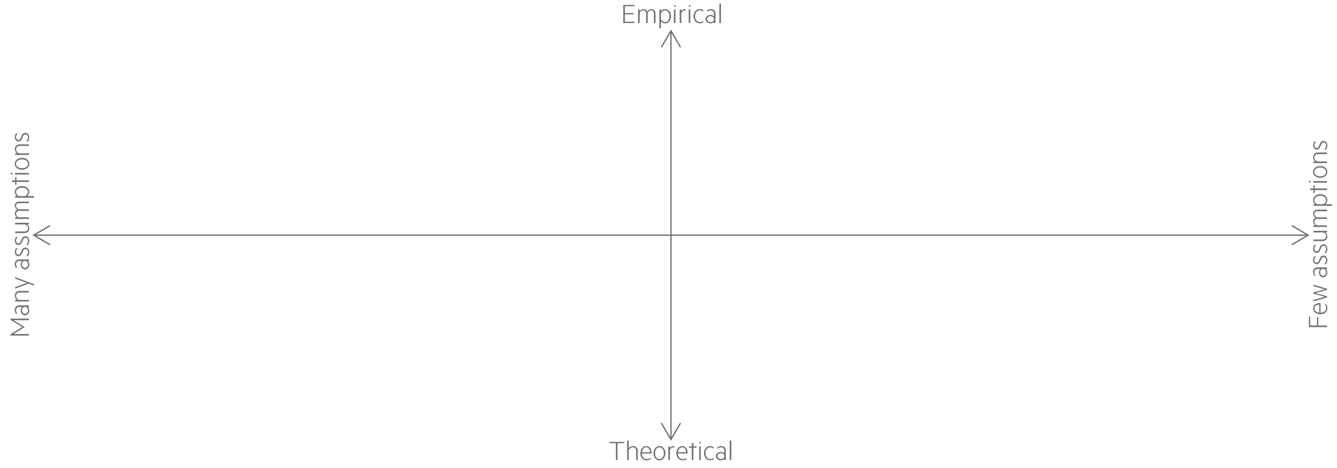

Survey on methods, guarantees and challenges

Classification

Learning an SCM with DGMs

- **SCM class**
  - Deep Structural Causal Models
    - Neural Causal Models
  - Bijective Generation Mechanism
- **DGM class**
  - Invertible Explicit
  - Amortised Explicit
  - Amortised Implicit

Challenges & Opportunities

Empirical

Many assumptions

Few assumptions

Theoretical

- Expressivity guarantees
- L_3-identifiability guarantees
- NeuralID to test for point identification

Survey on methods, guarantees and challenges

Classification

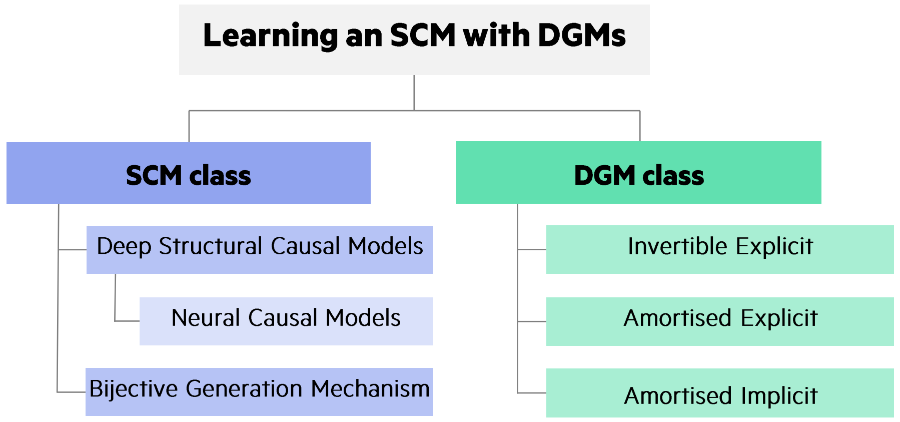

Challenges & Opportunities

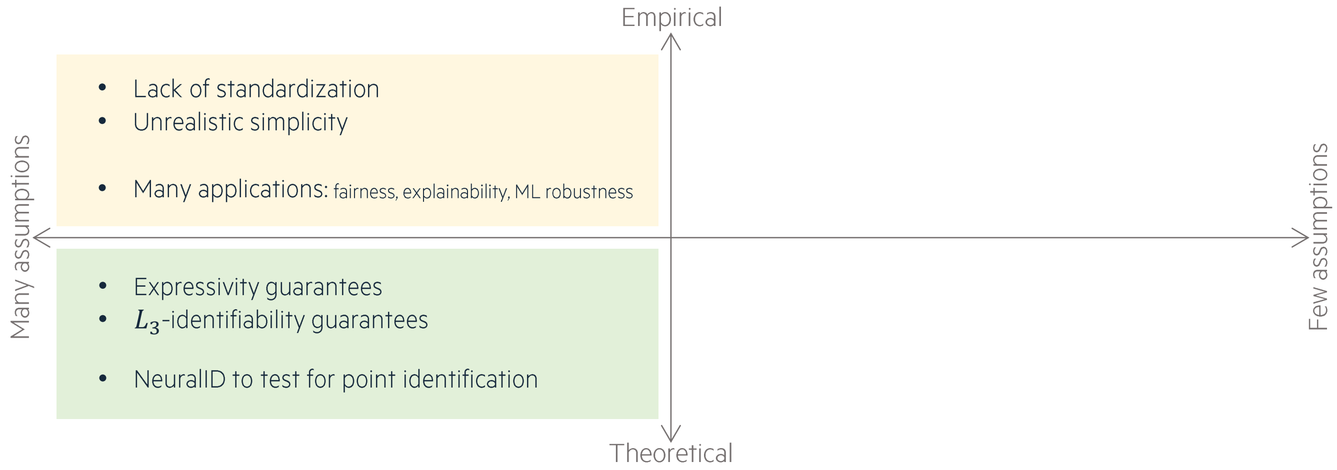

Survey on methods, guarantees and challenges

Classification

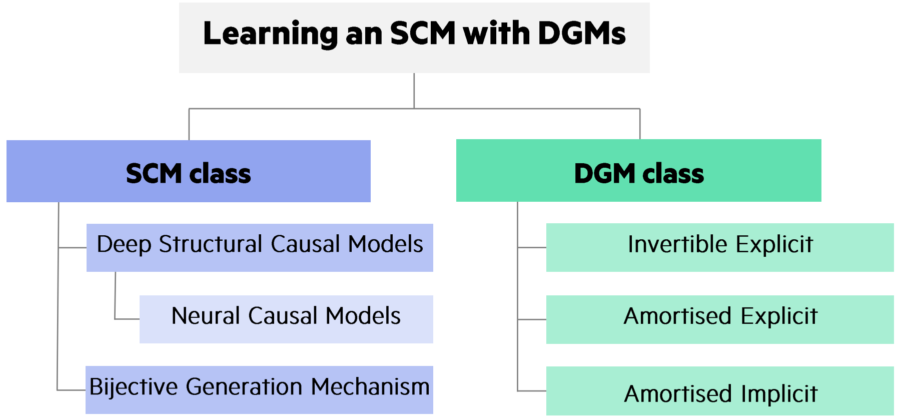

Challenges & Opportunities

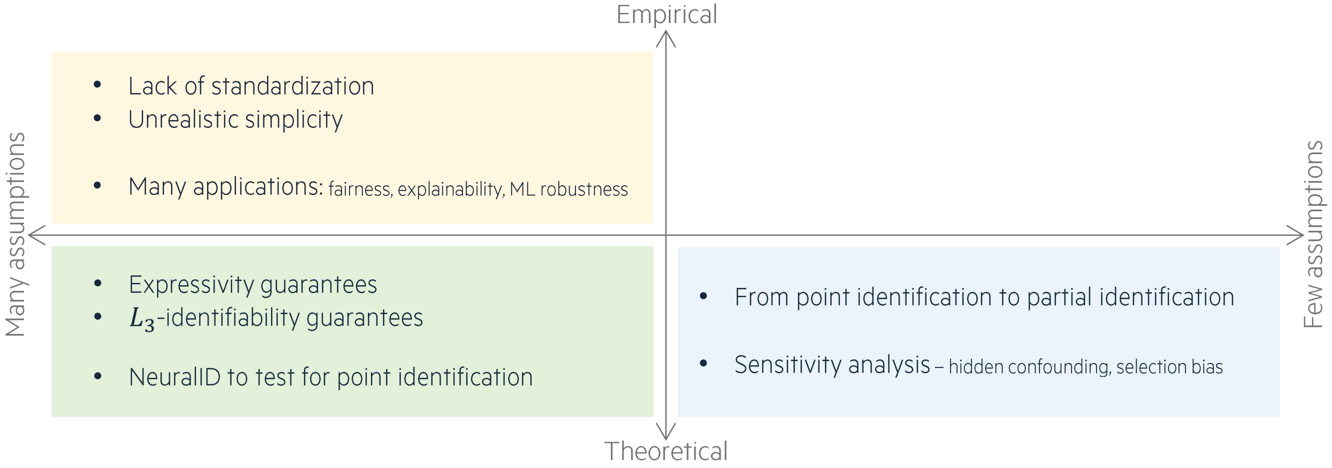

Survey on methods, guarantees and challenges

Classification

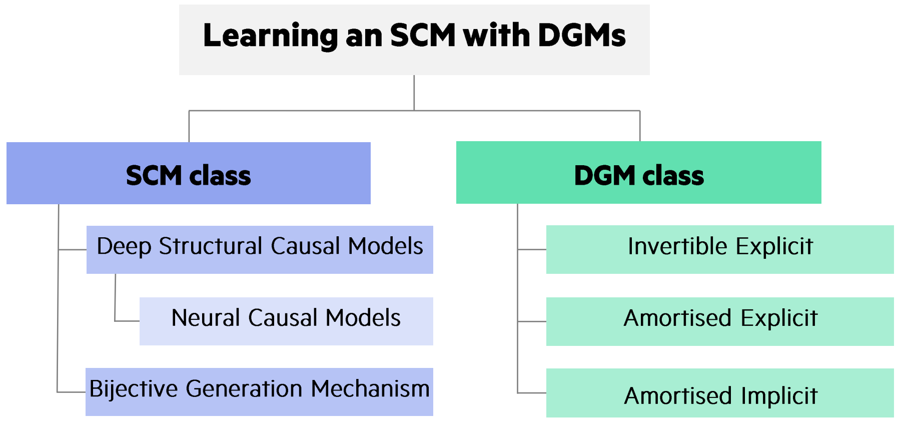

Challenges & Opportunities

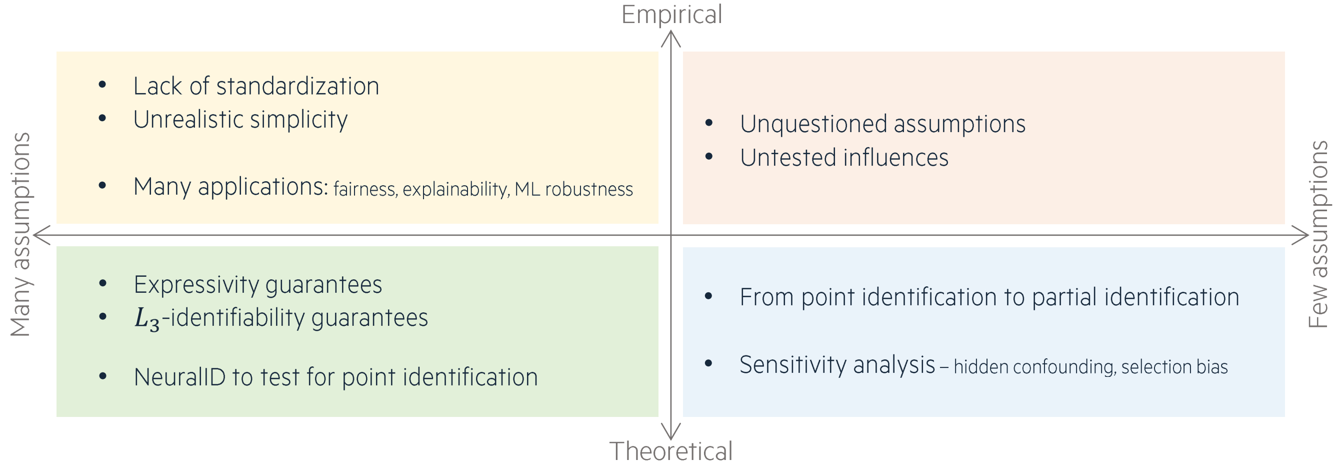

Survey on methods, guarantees and challenges

Classification

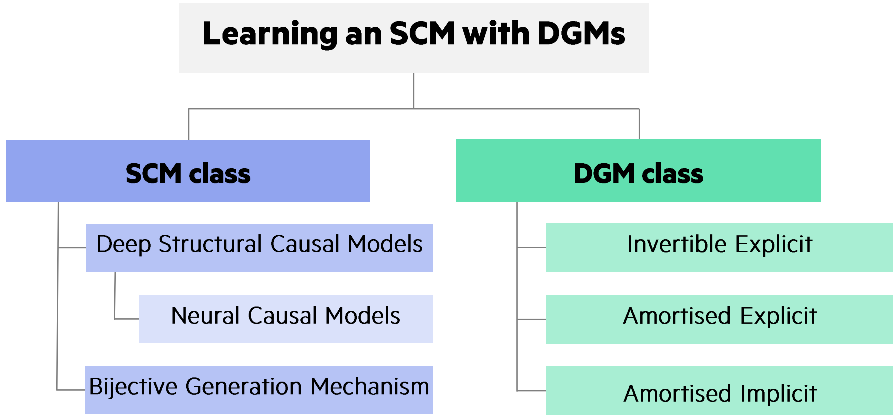

Challenges & Opportunities

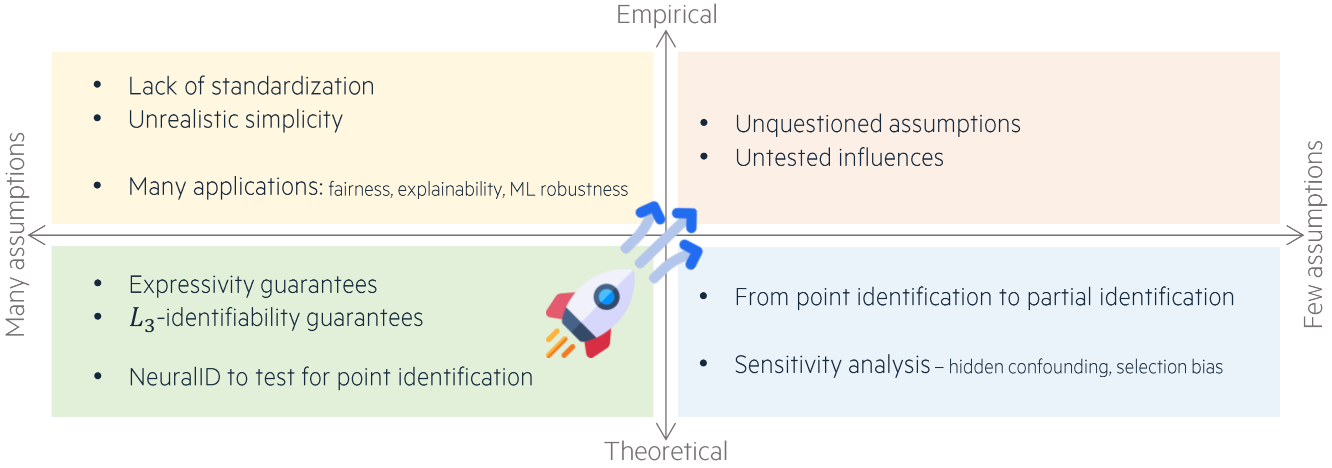

Survey on methods, guarantees and challenges

Classification

Learning an SCM with DGMs

- **SCM class**
  - Deep Structural Causal Models
    - Neural Causal Models
  - Bijective Generation Mechanism
- **DGM class**
  - Invertible Explicit
  - Amortised Explicit
  - Amortised Implicit

Challenges & Opportunities

Empirical

Many assumptions

- Lack of standardization
- Unrealistic simplicity
- Many applications: fairness, explainability, ML robustness

- Unquestioned assumptions
- Untested influences

- Expressivity guarantees
- L_3-identifiability guarantees
- NeuralID to test for point identification

- From point identification to partial identification
- Sensitivity analysis – hidden confounding, selection bias

Few assumptions

Theoretical

Come to chat at the poster session!

Poster location F29-40

Questions

IJCAI
JEJU 2024

université PARIS-SACLAY

Inria

Ekimetrics.

References

[Pearl, 2000] Judea Pearl. Models, Reasoning, and Inference. Cambridge University Press, 2000.

[Pearl, 2009] Judea Pearl. Models, Causal Inference in statistics: an overview. Statistics Surveys, 2009.

[Spirtes *et al.*, 2000] P. Spirtes, C. Glymour, and R. Scheines. Causation, Prediction, and Search. Springer-Verlag, 1993. (2nd ed. MIT Press 2000).

[Bollen, 1989] K. A. Bollen. Structural Equations with Latent Variables. John Wiley & Sons, 1989.

[Pearl and Mackenzie, 2018] Judea Pearl and Dana Mackenzie. The book of why: the new science of cause and effect. Basic books, 2018.

[Xia *et al.*, 2021] Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. In Advances in Neural Information Processing Systems, 2021.

[Zhou *et al.*, 2023] Guanglin Zhou, Shaoan Xie, Guangyuan Hao, Shiming Chen, Biwei Huang, Xiwei Xu, Chen Wang, Liming Zhu, Lina Yao, and Kun Zhang. Emerging synergies in causality and deep generative models: A survey. arXiv:2301.12351, 2023.

[Komanduri *et al.*, 2023] Aneesh Komanduri, Xintao Wu, Yongkai Wu, and Feng Chen. From identifiable causal representations to controllable counterfactual generation: A survey on causal generative modeling. arXiv:2310.11011, 2023.

[Kaddour *et al.*, 2022] Jean Kaddour, Aengus Lynch, Qi Liu, Matt J. Kusner, and Ricardo Silva. Causal machine learning: A survey and open problems. arXiv:2206.15475, 2022.

[Pawlowski *et al.*, 2020] Nick Pawlowski, Daniel Coelho de Castro, and Ben Glocker. Deep structural causal models for tractable counterfactual inference. In Advances in Neural Information Processing Systems, 2020.

[Xia *et al.*, 2023] Kevin Muyuan Xia, Yushu Pan, and Elias Bareinboim. Neural causal models for counterfactual identification and estimation. In 11th International Conference on Learning Representations, 2023.

[Nasr-Esfahany *et al.*, 2023] Arash Nasr-Esfahany, Mohammad Alizadeh, and Devavrat Shah. Counterfactual identifiability of bijective causal models. In 40th International Conference on Machine Learning, 2023.

References

[Poinsot *et al.*, 2024] Audrey Poinsot, Alessandro Leite, Nicolas Chesneau, Michèle Sébag, Marc Schoenauer. Learning Structural Causal Models through Deep Generative Models: Methods, Guarantees, and Challenges. In International Joint Conference on Artificial Intelligence, 2024.

[1] Arash Nasr-Esfahany, Mohammad Alizadeh, and Devavrat Shah. Counterfactual identifiability of bijective causal models. In 40th International Conference on Machine Learning, 2023.

[2] Nick Pawlowski, Daniel Coelho de Castro, and Ben Glocker. Deep structural causal models for tractable counterfactual inference. In Advances in Neural Information Processing Systems, 2020.

[3] Kevin Muyuan Xia, Yushu Pan, and Elias Bareinboim. Neural causal models for counterfactual identification and estimation. In 11th International Conference on Learning Representations, 2023.

[4] Adrian Javaloy, Pablo Sanchez-Martin, and Isabel Valera. Causal normalizing flows: from theory to practice. Advances in Neural Information Processing Systems, 2023.

[5] Alvaro Parafita and Jordi Vitria. Causal inference with deep causal graphs. arXiv:2006.08380, 2020.

[6] Ilyes Khemakhem, Ricardo Monti, Robert Leech, and Aapo Hyvarinen. Causal autoregressive flows. In 24th International Conference on Artificial Intelligence and Statistics, 2021.

[7] Matej Zecevic, Devendra Singh Dhami, Petar Velickovic, and Kristian Kersting. Relating graph neural networks to structural causal models. arXiv:2109.04173, 2021.

[8] Pablo Sanchez-Martin, Miriam Rateike, and Isabel Valera. VACA: designing variational graph autoencoders for causal queries. In AAAI Conference on Artificial Intelligence, 2022.

[9] Patrick Chao, Patrick Blöbaum, and Shiva Prasad Kasiviswanathan. Interventional and counterfactual inference with diffusion models. arxiv:2302.00860, 2023.

[10] Aneesh Komanduri, Yongkai Wu, Wen Huang, Feng Chen, and Xintao Wu. SCM-VAE: Learning identifiable causal representations via structural knowledge. In IEEE International Conference on Big Data, 2022.

[11] Bingyang Wen, Yupeng Cao, Fan Yang, Koduvayur Subbalakshmi, and Rajarathnam Chandramouli. Causal-TGAN: Modeling tabular data using causally-aware GAN. In ICLR Workshop on Deep Generative Models for Highly Structured Data, 2022.

[12] Murat Kocaoglu, Christopher Snyder, Alexandros G. Dimakis, and Sriram Vishwanath. CausalGAN: Learning causal implicit generative models with adversarial training. In International Conference on Learning Representations, 2018.

References

[13] Depeng Xu, Yongkai Wu, Shuhan Yuan, Lu Zhang, and Xintao Wu. Achieving causal fairness through generative adversarial networks. In 28th International Joint Conference on Artificial Intelligence, 2019.

[14] Boris van Breugel, Trent Kyono, Jeroen Berrevoets, and Mihaela van der Schaar. DECAF: Generating fair synthetic data using causally-aware generative networks. In Advances in Neural Information Processing Systems, 2021.

[15] Md Musfiqur Rahman and Murat Kocaoglu. Towards modular learning of deep causal generative models. In ICML Workshop on Structured Probabilistic Inference & Generative Modeling, 2023.

[16] Axel Sauer and Andreas Geiger. Counterfactual generative networks. In International Conference on Learning Representations, 2021.

[17] Xinwei Shen, Furui Liu, Hanze Dong, Qing Lian, Zhitang Chen, and Tong Zhang. Weakly supervised disentangled generative causal representation learning. Journal of Machine Learning Research, 2022.

[18] Patrik Hoyer, Dominik Janzing, Joris Mooij, Jonas Peters and Bernhard Schölkopf. Nonlinear causal discovery with additive noise modes. In Advances in Neural Information Processing Systems, 2008.

[19] Alexander Immer, Christoph Schultheiss, Julia E Vogt, Bernhard Schölkopf, Peter Bühlmann and Alexander Marx. On the Identifiability and Estimation of Causal Location-Scale Noise Models. In International Conference of Machine Learning, 2023.

[20] Kun Zhang and Aapo Hyvärinen. On the Identifiability of the Post-Nonlinear Causal Model. In Conference on Uncertainty in Artificial Intelligence, 2009.

[Zaffalon *et al.*, 2020] Marco Zaffalon, Alessandro Antonucci, and Rafael Cabanas. Structural causal models are (solvable by) credal networks. In 10th International Conference on Probabilistic Graphical Models, 2020.

[Zhang *et al.*, 2022] Junzhe Zhang, Jin Tian, and Elias Bareinboim. Partial counterfactual identification from observational and experimental data. In 39th International Conference on Machine Learning, 2022.

[Schrönder *et al.*, 2024] Maresa Schr¨oder, Dennis Frauen, and Stefan Feuerriegel. Causal fairness under unobserved confounding: a neural sensitivity framework. 12th International Conference on Learning Representations, 2023.

[Frauen *et al.*, 2024] Dennis Frauen, Fergus Imrie, Alicia Curth, Valentyn Melnychuk, Stefan Feuerriegel, and Mihaela van der Schaar. A neural framework for generalized causal sensitivity analysis. In 12th International Conference on Learning Representations, 2024.

[Messerli, 2012] Franz H. Messerli. Chocolate Consumption, Cognitive Function, and Nobel Laureates. The New England Journal of Medicine, 2012.