Causal, Casual, and Curious (2013-2020): A collage in the art of causal reasoning

Judea Pearl

Cognitive Systems Laboratory

Computer Science Department

University of California, Los Angeles, CA 90024 USA

judea@cs.ucla.edu

January 30, 2020

1   Introduction

This collection of 14 short articles represents adventurous ideas and semi-heretical thoughts that emerged when, in 2013, I was given the opportunity to edit a fun section of the *Journal of Causal Inference* called "Causal, Casual, and Curious."

   The direct contact with readers, unmediated by editors or reviewers, had a healthy liberating effect on me and has unleashed some of my best, perhaps most mischievous explorations. I thank the editors of the *Journal of Causal Inference* for giving me this opportunity to undertake this adventure and for trusting me to manage it as prudently as I could.

"Linear Models: A Useful "Microscope" for Causal Analysis," *Journal of Causal Inference*, 1(1): 155–170, May 2013.
https://ucla.in/2LcpmHz
Abstract: This note reviews basic techniques of linear path analysis and demonstrates, using simple examples, how causal phenomena of non-trivial character can be understood, exemplified and analyzed using diagrams and a few algebraic steps. The techniques allow for swift assessment of how various features of the model impact the phenomenon under investigation. This includes: Simpson's paradox, case-control bias, selection bias, missing data, collider bias, reverse regression, bias amplification, near instruments, and measurement errors.

"The Curse of Free-will and the Paradox of Inevitable Regret" *Journal of Causal Inference*, 1(2): 255-257, December 2013.
https://ucla.in/2N6x36Q
Abstract: The paradox described below aims to clarify the principles by which

population data can be harnessed to guide personal decision making. The logic that permits us to infer counterfactual quantities from a combination of experimental and observational studies gives rise to situations in which an agent knows he/she will regret whatever action is taken.

"Is Scientific Knowledge Useful for Policy Analysis? A Peculiar Theorem says: No," *Journal of Causal Inference*, 2(1): 109–112, March 2014.

https://ucla.in/2JfkxuY

Abstract: Conventional wisdom dictates that the more we know about a problem domain the easier it is to predict the effects of policies in that domain. Strangely, this wisdom is not sanctioned by formal analysis, when the notions of "knowledge" and "policy" are given concrete definitions in the context of nonparametric causal analysis. This note describes this peculiarity and speculates on its implications.

"Graphoids over counterfactuals" *Journal of Causal Inference*, 2(2): 243-248, September 2014.

https://ucla.in/2QpcGzS

Abstract: Augmenting the graphoid axioms with three additional rules enables us to handle independencies among observed as well as counterfactual variables. The augmented set of axioms facilitates the derivation of testable implications and ignorability conditions whenever modeling assumptions are articulated in the language of counterfactuals.

"Conditioning on Post-Treatment Variables," *Journal of Causal Inference*, 3(1): 131-137, March 2015. Includes Appendix (appended to published version).

https://ucla.in/2wtLB5h

Abstract: In this issue of the Causal, Casual, and Curious column, I compare several ways of extracting information from post-treatment variables and call attention to some peculiar relationships among them. In particular, I contrast *do*-calculus conditioning with counterfactual conditioning and discuss their interpretations and scopes of applications. These relationships have come up in conversations with readers, students and curious colleagues, so I will present them in a question–answers format.

"Generalizing experimental findings," *Journal of Causal Inference*, 3(2): 259-266, September 2015.

https://ucla.in/2L6yTzE

Abstract: This note examines one of the most crucial questions in causal inference: "How generalizable are randomized clinical trials?" The question has received a formal treatment recently, using a non-parametric setting, and has led to a simple and general solution. I will describe this solution and several of its ramifications, and compare it to the way researchers have attempted to tackle the problem using the language of ignorability. We will see that ignorability-type assumptions need

to be enriched with structural assumptions in order to capture the full spectrum of conditions that permit generalizations, and in order to judge their plausibility in specific applications.

"The Sure-Thing Principle," *Journal of Causal Inference*, 4(1): 81-86, March 2016.

https://ucla.in/2NTbnrS

Abstract: In 1954, Jim Savage introduced the Sure Thing Principle to demonstrate that preferences among actions could constitute an axiomatic basis for a Bayesian foundation of statistical inference. Here, we trace the history of the principle, discuss some of its nuances, and evaluate its significance in the light of modern understanding of causal reasoning.

"Lord's Paradox Revisited — (Oh Lord! Kumbaya!)," *Journal of Causal Inference*, Published Online 4(2): September 2016.

https://ucla.in/2JeJs1Q

Abstract: Among the many peculiarities that were dubbed "paradoxes" by well meaning statisticians, the one reported by Frederic M. Lord in 1967 has earned a special status. Although it can be viewed, formally, as a version of Simpson's paradox, its reputation has gone much worse. Unlike Simpson's reversal, Lord's is easier to state, harder to disentangle and, for some reason, it has been lingering for almost four decades, under several interpretations and re-interpretations, and it keeps coming up in new situations and under new lights. Most peculiar yet, while some of its variants have received a satisfactory resolution, the original version presented by Lord, to the best of my knowledge, has not been given a proper treatment, not to mention a resolution.

The purpose of this paper is to trace back Lord's paradox from its original formulation, resolve it using modern tools of causal analysis, explain why it resisted prior attempts at resolution and, finally, address the general methodological issue of whether adjustments for preexisting conditions is justified in group comparison applications.

"A Linear 'Microscope' for Interventions and Counterfactuals," *Journal of Causal Inference*, Published Online 5(1): 1-15, March 2017.

https://ucla.in/2L8jUFg

Abstract: This note illustrates, using simple examples, how causal questions of non-trivial character can be represented, analyzed and solved using linear analysis and path diagrams. By producing closed form solutions, linear analysis allows for swift assessment of how various features of the model impact the questions under investigation. We discuss conditions for identifying total and direct effects, representation and identification of counterfactual expressions, robustness to model misspecification, and generalization across populations.

"Physical and Metaphysical Counterfactuals" Revised version, *Journal of Causal Inference*, 5(2): September 2017.

https://ucla.in/2N9nSCV

Abstract: The structural interpretation of counterfactuals as formulated in Balke and Pearl (1994a,b) excludes disjunctive conditionals, such as "had X been x_1 or x_2," as well as disjunctive actions such as $do(X = x_1$ or $X = x_2)$. In contrast, the closest-world interpretation of counterfactuals (e.g. Lewis (1973)) assigns truth values to all counterfactual sentences, regardless of the logical form of the antecedent. This paper leverages "imaging"–a process of "mass-shifting" among possible worlds, to define disjunction in structural counterfactuals. We show that every imaging operation can be given an interpretation in terms of a stochastic policy in which agents choose actions with certain probabilities. This mapping, from the metaphysical to the physical, allows us to assess whether metaphysically-inspired extensions of interventional theories are warranted in a given decision making situation.

"What is Gained from Past Learning" *Journal of Causal Inference*, 6(1), Article 20180005, https://doi.org/10.1515/jci-2018-0005, March 2018.

https://ucla.in/2ocoWqq

Abstract: We consider ways of enabling systems to apply previously learned information to novel situations so as to minimize the need for retraining. We show that theoretical limitations exist on the amount of information that can be transported from previous learning, and that robustness to changing environments depends on a delicate balance between the relations to be learned and the causal structure of the underlying model. We demonstrate by examples how this robustness can be quantified.

"Does Obesity Shorten Life? Or is it the Soda? On Non-manipulable Causes," *Journal of Causal Inference*, 6(2), online, September 2018.

https://ucla.in/2EpxcNU

Abstract: Non-manipulable factors, such as gender or race have posed conceptual and practical challenges to causal analysts. On the one hand these factors do have consequences, and on the other hand, they do not fit into the experimentalist conception of causation. This paper addresses this challenge in the context of public debates over the health cost of obesity, and offers a new perspective, based on the theory of Structural Causal Models (SCM).

"On the interpretation of $do(x)$," *Journal of Causal Inference*, 7(1), online, March 2019.

https://ucla.in/2LY8M2X

Abstract: This paper provides empirical interpretation of the $do(x)$ operator when applied to non-manipulable variables such as race, obesity, or cholesterol level. We view $do(x)$ as an ideal intervention that provides valuable information on the effects of manipulable variables and is thus empirically testable. We draw parallels between

this interpretation and ways of enabling machines to learn effects of untried actions from those tried. We end with the conclusion that researchers need not distinguish manipulable from non-manipulable variables; both types are equally eligible to receive the $do(x)$ operator and to produce useful information for decision makers.

"Sufficient Causes: On Oxygen, Matches, and Fires," *Journal of Causal Inference*, AOP, https://doi.org/10.1515/jci-2019-0026, September 2019. https://ucla.in/2lEOJtO

Abstract: We demonstrate how counterfactuals can be used to compute the probability that one event was/is a sufficient cause of another, and how counterfactuals emerge organically from basic scientific knowledge, rather than manipulative experiments. We contrast this demonstration with the potential outcome framework and address the distinction between causes and enablers.

References

BALKE, A. and PEARL, J. (1994a). Counterfactual probabilities: Computational methods, bounds, and applications. In *Uncertainty in Artificial Intelligence 10* (R. L. de Mantaras and D. Poole, eds.). Morgan Kaufmann, San Mateo, CA, 46–54.

BALKE, A. and PEARL, J. (1994b). Probabilistic evaluation of counterfactual queries. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, vol. I. MIT Press, Menlo Park, CA, 230–237.

LEWIS, D. (1973). *Counterfactuals.* Harvard University Press, Cambridge, MA.

LORD, F. (1967). A paradox in the interpretation of group comparisons. *Psychological Bulletin* **68** 304–305.