The Importance of Causality in Decision Making: A Perspective on Recommender Systems

Emanuele Cavenaghi[1,*], Alessio Zanga[2,3], Fabio Stella[2] and Markus Zanker[1,5]

Abstract
Causality is receiving increasing attention from the artificial intelligence and machine learning communities. Similarly, growing attention to causality is currently going on in the Recommendation Systems (RSs) community, which has realised that RSs could greatly benefit from causality to transform accurate predictions into effective and explainable decisions. Indeed, the RS literature has repeatedly highlighted that, in real-world scenarios, recommendation algorithms suffer many types of biases since assumptions ensuring unbiasedness are likely not met. In this discussion paper, we formulate the RS problem in terms of causality, using potential outcomes and structural causal models, by giving formal definitions of the causal quantities to be estimated and a general causal graph to serve as a reference to foster future research and development.

Keywords
Causal Models, Decision Making, Recommender Systems

1. Introduction

Predicting and deciding are two fundamentally different tasks. As described by the *Ladder of Causation* [1], a decision manipulates the system which can react to our decision, while a prediction does not affect the system in any manner: the system is eventually affected only when we exploit the prediction to make a decision. Overlooking this difference usually leads to biased predictions that, in turn, result in wrong decisions. In this sense, the RSs community is facing several problems with biased estimates [2] to assess the effect of recommendations based on predictions. Indeed, according to [3, 4], the recommendation problem is usually framed as a prediction problem while, as pointed out in [5, 6], it is indeed a decision-making problem, since we have to decide which item(s) to recommend to which user(s).

Furthermore, human beings are not interested in mere correlations but in understanding the actual causes of the effects, manipulating the world to achieve the desired outcome. In fact, scientists are familiar with the phrase: "Correlation is not causation", that is, for example, "the rooster's crow is highly correlated with the sunrise; yet it does not cause the sunrise" [1]. Indeed, under general conditions, machine learning approaches do not allow to state that X is the cause of Y but only that they are "correlated" or "associated" to each other.

This is why causality becomes important: we need a way to translate cause-and-effect relations and interventions on a system using a mathematical formulation. To this end, in [6], we proposed a causal decision-making framework for RSs using Potential Outcomes (POs) [7] to define causal estimands of interest and Causal Graphs (CGs) [8] to propose a general probabilistic graphical model for RSs to encode the cause-and-effect relations among variables. Using this framework, we introduce the process, illustrated in Figure 1, which allows to make decisions by combining data and expert knowledge.

AIxIA 2024 Discussion Papers - 23rd International Conference of the Italian Association for Artificial Intelligence, Bolzano, Italy, November 25–28, 2024
*Corresponding author.
ecavenaghi@unibz.it (E. Cavenaghi); alessio.zanga@unimib.it (A. Zanga); fabio.stella@unimib.it (F. Stella); Markus.Zanker@unibz.it (M. Zanker)
0000-0002-0235-0421 (E. Cavenaghi); 0000-0003-4423-2121 (A. Zanga); 0000-0002-1394-0507 (F. Stella); 0000-0002-4805-5516 (M. Zanker)

CAUSAL DISCOVERY	CAUSAL ESTIMAND IDENTIFICATION	ESTIMATION	MAKING DECISIONS
Learn the CG that describes the data-generating process	Define the causal estimand(s) and identify an equivalent statistical estimand	Estimate using any model compatible with the outcome type	Make decisions using the extracted knowledge

Figure 1: From available data and expert knowledge to making decisions.

2. Making Decisions Through Causality

2.1. Causal Discovery

The first step is to have a CG that describes the data-generating process of the system under study. The CG can be learned by combining observational data with experts' knowledge through a process called *causal discovery*, which is enabled by several causal discovery algorithms [9, 10]. While the CG must be learned in each scenario, we proposed a reference CG for RSs [6] to guide the construction of a CG for specific RSs problems as done in [11, 12].

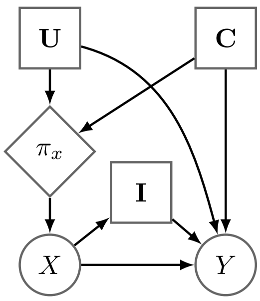

Figure 2: Causal Graph for item ID-based recommendation. The squared nodes represent clusters of nodes composed of multiple nodes.

The problem of recommending a *single* item with features $\mathbf{I}$ to a user $\mathbf{U}$ in context $\mathbf{C}$ is described by the CG of Figure 2 where the node X represents the action of recommending an item whose domain corresponds to the item set, i.e., $x \in dom(X)$. For example, in the context of film recommendation, $dom(X)$ is the catalogue of the films and our recommendation X is one of the films in the catalogue. To decide which item to recommend to the current user $\mathbf{U}$ in context $\mathbf{C}$, we use a policy π_x based on user and context features. Once we decide which item x to recommend, the corresponding item's features $\mathbf{I}$ are fixed and they mediate the effect of our recommendation X on the user feedback Y through the path $X \rightarrow \mathbf{I} \rightarrow Y$. For example, once we decide to recommend a film, its genre is fixed ($X \rightarrow genre$), and the genre (likely) affects a user's feedback ($genre \rightarrow Y$). It is worth noticing that not all the item's features have to influence the user's feedback, i.e., some item's features are not taken into account by the users. On the other hand, part of the effect of the recommendation X on the user's feedback Y may not be captured by the modelled features $\mathbf{I}$ and flows directly through the edge $X \rightarrow Y$. For example, if we are not able to model the film's popularity, since it is difficult to know or to model it, the effect of the film's popularity will flow through the edge $X \rightarrow Y$ as this feature is not included in our model.

2.2. Causal Estimand Identification

To exploit the potential of causality, we should frame the quantity to estimate as a *causal estimand* that encodes the notion of the causal effect of a variable (the cause) on another (the effect). Generally, we can define it, using the POs framework and the *do-operator* [13, 14], as $\mathbb{E}[Y|do(X=x), \mathbf{u}, \mathbf{i}, \mathbf{c}]$. This encodes the value of the expected feedback Y given by the user $\mathbf{u}$ in context $\mathbf{c}$ when we recommend item x with features $\mathbf{i}$. The difference that separates causal estimands from classical statistical estimands is the presence of the so-called *do-operator*, denoted with $do(X = x)$, that defines the intervention of fixing the value of X to x for the whole population of users. In contrast, conditioning on $X = x$ means that X takes a value x naturally, which simply translates to focusing only on the sub-population where X has been observed to be equal to x. In a decision-making problem, such as RSs, we are interested in estimating causal estimands since we actively decide which item(s) to recommend.

However, expressions with the *do-operator* can only be estimated in controlled experiments where the variables in the *do-terms* can be appropriately controlled. To estimate a causal estimand using only observational data, it is necessary to remove the *do-terms* and obtain an equivalent expression. To this end, the *adjustment formula* estimator [14] adopts a model-based approach to adjust for an *adjustment set* $\mathbf{Z}$ and obtain a statistical estimand:

$$P(Y = y|do(X = x)) = \sum_{\mathbf{z}} P(Y = y|X = x, \mathbf{Z} = \mathbf{z})P(\mathbf{Z} = \mathbf{z}) \tag{1}$$

To *identify* the variables that must be included in $\mathbf{Z}$, we can query the CG by evaluating an *identification criterion*, such as the *backdoor criterion* [14], *frontdoor criterion* [8] or *do-calculus* [13]. In particular, if no identification is possible with do-calculus, the causal effect is guaranteed to be *unidentifiable*. Thus, every estimate of the causal estimand will be biased.

2.3. Estimation

Once we have proved the identifiability of the causal estimand, i.e., once we have shown that the causal estimand is equal to a statistical estimand, this can be estimated using classical statistical estimators. To this end, any model that is compatible with the type of the outcome variable, e.g. linear regression for a continuous outcome or neural networks for non-linear relations, is suitable for this estimation. Clearly, the model should be chosen carefully for each problem by considering the available data characteristics to avoid estimation errors.

2.4. Making Decisions

Finally, with the estimated causal effects, e.g., the effect of our recommendation on the user's propensity to click on the recommended item(s), we can decide which items to recommend. This could be done in different ways: (i) greedy, (ii) ϵ-greedy and (iii) more sophisticated policies. To this end, in recent years, several works have exploited causality by linking it to *Multi-Armed Bandit* (MAB) [15, 16, 17] and *Reinforcement Learning* (RL) [18]. For example, in [19], the authors define the notion of Possibly-Optimal Minimal Intervention Set with the idea of determining the minimum set of variables on which a MAB agent should intervene to understand all the possible arms that are worth intervening on. Moreover, [20] extends the method by considering that some variables can not be manipulated. Using causality with RL, [21, 22] approached the Dynamic Treatment Regimes problem with confounded observational dataset.

3. Conclusions

In this paper, we proposed a causal view of the RS problem and highlighted the importance of framing the recommendation problem in terms of causality. The causality framework can, in our view, be considered as a single framework allowing researchers to wholistically define and address several problems widely acknowledged in the RSs community to bridge the gaps in future works.

However, we would like to stress that causality is not magic but ruthlessly honest and, differently from other approaches, it makes explicit assumptions, such as ignorability and unconfoundedness, leaving us with the burden of judging whether they are likely to be satisfied for the addressed context. Indeed, causality is not the sole ingredient to solve the RS problem while we are fully convinced that exploiting the body of knowledge generated over more than 30 years of research in RSs and users' behaviour remains fundamental.

References

[1] J. Pearl, D. Mackenzie, The book of why: the new science of cause and effect, Basic books, New York, United States, 2018.

[2] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, X. He, Bias and debias in recommender system: A survey and future directions, ACM Trans. Inf. Syst. 41 (2023).

[3] G. Adomavicius, A. Tuzhilin, Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions, IEEE Transactions on Knowledge and Data Engineering 17 (2005) 734–749.

[4] F. Ricci, L. Rokach, B. Shapira, Introduction to recommender systems handbook, in: Recommender systems handbook, Springer, Boston, MA, 2011, pp. 1–35.

[5] O. Jeunen, B. Goethals, Pessimistic decision-making for recommender systems, ACM Trans. Recomm. Syst. 1 (2022).

[6] E. Cavenaghi, A. Zanga, F. Stella, M. Zanker, Towards a causal decision-making framework for recommender systems, ACM Trans. Recomm. Syst. 2 (2024).

[7] D. B. Rubin, Estimating causal effects of treatments in randomized and nonrandomized studies., Journal of educational Psychology 66 (1974) 688.

[8] J. Pearl, Causality, Cambridge university press, Cambridge, United Kingdom, 2009.

[9] M. J. Vowels, N. C. Camgoz, R. Bowden, D'ya like dags? a survey on structure learning and causal discovery, ACM Computing Surveys (CSUR) (2021).

[10] A. Zanga, E. Ozkirimli, F. Stella, A survey on causal discovery: Theory and practice, International Journal of Approximate Reasoning 151 (2022) 101–129.

[11] E. Cavenaghi, A. Zanga, A. Rimoldi, P. Minasi, F. Stella, M. Zanker, Causal discovery in recommender systems: a case study in online hotel search 1 (2023).

[12] E. Cavenaghi, A. Zanga, A. Rimoldi, P. Minasi, M. Zanker, F. Stella, Analysis of relevant factors in online hotel recommendation through causal models 1 (2023).

[13] E. Bareinboim, J. D. Correa, D. Ibeling, T. Icard, On pearl's hierarchy and the foundations of causal inference, in: Probabilistic and Causal Inference: The Works of Judea Pearl, Association for Computing Machinery, New York, NY, USA, 2022, p. 507–556.

[14] M. Glymour, J. Pearl, N. P. Jewell, Causal inference in statistics: A primer, John Wiley & Sons, Hoboken, United States, 2016.

[15] F. Lattimore, T. Lattimore, M. D. Reid, Causal bandits: Learning good interventions via causal inference, in: D. Lee, M.Sugiyama, U.Luxburg, I.Guyon, R.Garnett (Eds.), Advances in Neural Information Processing Systems, volume 29, Curran Associates, Inc., Barcelona, Spain, 2016, p. 9.

[16] Y. Lu, A. Meisami, A. Tewari, W. Yan, Regret analysis of bandit problems with causal background knowledge, in: J. Peters, D. Sontag (Eds.), Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI), volume 124 of *Proceedings of Machine Learning Research*, PMLR, Virtual Event, 2020, pp. 141–150.

[17] V. Nair, V. Patil, G. Sinha, Budgeted and non-budgeted causal bandits, in: A. Banerjee, K. Fukumizu (Eds.), Proceedings of The 24th International Conference on Artificial Intelligence and Statistics, volume 130 of *Proceedings of Machine Learning Research*, PMLR, Virtual Event, 2021, pp. 2017–2025.

[18] Y. Lu, A. Meisami, A. Tewari, Causal markov decision processes: Learning good interventions efficiently, arXiv preprint arXiv:2102.07663 1 (2021).

[19] S. Lee, E. Bareinboim, Structural causal bandits: Where to intervene?, in: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 31, Curran Associates, Inc., Montréal, Canada, 2018, p. 11.

[20] S. Lee, E. Bareinboim, Structural causal bandits with non-manipulable variables, Proceedings of the AAAI Conference on Artificial Intelligence 33 (2019) 4164–4172.

[21] J. Zhang, Designing optimal dynamic treatment regimes: A causal reinforcement learning approach, in: H. D. III, A. Singh (Eds.), Proceedings of the 37th International Conference on Machine Learning, volume 119 of *Proceedings of Machine Learning Research*, PMLR, Vienna, Austria, 2020, pp. 11012–11022.

[22] J. Zhang, E. Bareinboim, Near-optimal reinforcement learning in dynamic treatment regimes, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 32, Curran Associates, Inc., Vancouver, Canada, 2019.