

A Common Misassumption in Online Experiments with Machine Learning Models

Olivier Jeunen

ShareChat

Edinburgh, UK

jeunen@sharechat.co

Abstract

Online experiments such as Randomised Controlled Trials (RCTs) or A/B-tests are the bread and butter of modern platforms on the web. They are conducted continuously to allow platforms to estimate the causal effect of replacing system variant “A” with variant “B”, on some metric of interest. These *variants* can differ in many aspects. In this paper, we focus on the common use-case where they correspond to machine learning models. The online experiment then serves as the final arbiter to decide which model is superior, and should thus be shipped. The statistical literature on causal effect estimation from RCTs has a substantial history, which contributes deservedly to the level of trust researchers and practitioners have in this “gold standard” of evaluation practices. Nevertheless, in the particular case of machine learning experiments, we remark that certain critical issues remain. Specifically, the assumptions that are required to ascertain that A/B-tests yield unbiased estimates of the causal effect, are seldom met in practical applications. We argue that, because variants typically learn using pooled data, a lack of *model interference* cannot be guaranteed. This undermines the conclusions we can draw from online experiments with machine learning models. We discuss the implications this has for practitioners, and for the research literature.

1 Randomised Controlled Trials and their Assumptions

Randomised experiments have existed in the scientific literature for close to 140 years, first introduced in psychology [Peirce and Jastrow, 1884]. Since then, they have been a popular topic of study in the statistical literature —a feat often ascribed to the seminal works of Fisher [1925, 1936]— and are generally well-understood [Imbens and Rubin, 2015]. Randomised Controlled Trials (RCTs) form the theoretical basis for the online experiments that modern web platforms run continuously [Gupta et al., 2019], colloquially known as A/B-tests [Kohavi et al., 2020].

Generally speaking, RCTs deal with *treatments* being applied to *units*, leading to certain *outcomes* [Rubin, 1974]. Typical examples from the early literature revolve around agricultural applications, where we have types of fertiliser we can apply to plots of land, which has an effect on crop yield. In an RCT, we randomly assign units to treatment/control, and as a result, the average measured outcomes for units under control C and treatment T give a finite-sample estimate of

the causal effect that the treatment has on the outcome, the Average Treatment Effect (ATE):

$$\text{ATE}(C \rightarrow T; Y) = \bar{Y}(T) - \bar{Y}(C). \quad (1)$$

Under seemingly reasonable and light assumptions, this estimate is *consistent* and *unbiased*. That is, given infinite samples, the expectation of the estimate converges to the true average causal effect of applying the treatment instead of the control [Imbens and Rubin, 2015].

The assumption that is needed to support this claim, is called the Stable Unit Treatment Value Assumption (SUTVA). The SUTVA implies that outcomes of units are independent of the outcomes of other units under different treatments. An agricultural example of when things can go wrong is given by Rubin [1974]: when “*the plots are in such close proximity that following rainfall plot i receives fertilizer [sic] from adjacent plots.*” In many classical cases, it is easily argued that this assumption holds due to experimental design choices and proper randomisation. In A/B-tests, it is often assumed that there are no “network effects” or “spillovers” among units [Gupta et al., 2019, §10]. When spillovers are known to exist, alternative experimental design frameworks have been proposed in the literature. These are typically proposed because of some shared resource among treatments (e.g. campaign budgets in online advertising [Liu et al., 2021]), or because multiple populations interplay (e.g. buyers and sellers interfering [Bajari et al., 2021]).

Kohavi et al. provide some related examples of problematic cases where interference violates the SUTVA, either through direct or indirect connections [Kohavi et al., 2020, Ch. 22].

In this article, we focus on the simplest possible setting, devoid of network effects, budgets, or multiple populations, and argue that the SUTVA is still violated in the large majority of online experiments that typical web platforms run on a daily basis to compare machine learning models.

2 Online Experiments with Machine Learning Models

Suppose we operate a platform on the web with a deployed recommender system. To make things concrete, we have a typical contextual bandit setup where *units* correspond to contexts $x \in \mathcal{X}$, in which we recommend items $a \in \mathcal{A}$ (which we refer to as *actions*). A recommendation policy maps contexts to actions, potentially stochastically:

$$\pi : \mathcal{X} \rightarrow \Delta^{|\mathcal{A}|}. \quad (2)$$

As is common practice, these policies are informed by user-item interactions we have collected on the platform, and they aim to optimise some metric of interest: the outcome Y . The recommendation policy itself is the *treatment*. We wish to deploy a recommendation policy, but have multiple contenders. To reliably estimate the performance that we would obtain from deploying a given recommendation policy to the entire user population, we run an A/B-test. That is, for every new context that the system is presented with, we randomly allocate it to either π_C or π_T . By assuming that users are allocated consistently, or simply that all contexts are independent, we ensure that no *traditional* network effects among units exist in this setup.

Nevertheless, we argue that interference *does* occur when information about logged context-action-outcome triplets under policy π_C is used to inform the recommendation policy π_T (or vice versa). Furthermore, this erratic experimental design is pervasive throughout the industry, either because practitioners and researchers are unaware that they are violating the SUTVA, or because

they choose to ignore it (implicitly assuming the effects will be negligible). To the best of our knowledge, this assumption has not been validated in the research literature, and should not be accepted at face value until these effects have been studied more rigorously. In any case, when A/B-tests are often seen as the *gold standard* of evaluation methods and several research papers rely on them to provide empirical validation of newly proposed methods, we believe that it is important for practitioners, authors and reviewers to be aware of this issue, and to deal with it explicitly when evaluating conclusions that are drawn from online experiments.

We highlight two common scenarios where this type of *model* interference occurs: (1) when policies learn from pooled data collected under different policies (i.e. shared training data), or (2) when features used to make decisions are informed by data collected under different policies (i.e. shared behavioural features). The first case is easy to detect, and already widespread. The second case can be more subtle, when e.g. embeddings are used to represent contexts and actions, and they are updated throughout the experiment (even if the recommendation policies themselves are not necessarily updated). In both these cases, a model’s output at a certain time t will not be independent of other policies in place, which violates the SUTVA explicitly: $\pi_T^t(x) \not\perp \pi_C^{t-1}(x')$.

Naturally, this insight does not automatically dismiss all experiments that are run with similar configurations, and it is by no means our intention to undermine the credibility of previously published work. It does, however, implore us to take a critical look at online experiments presented in the research literature, and we hope to incentivise authors to be explicit and clear when describing their experimental setup. Indeed, A/B-tests on large-scale platforms are inherently non-reproducible for the majority of the academic audience, and we believe that some additional scrutiny may be warranted in light of the arguments presented in this paper.

In what follows, we provide empirical evidence on synthetic data to highlight how a SUTVA violation can be problematic, and how it invalidates conclusions drawn from online experiments.

3 Empirical Evidence of Model Interference with Pooled Data

To empirically validate our insight that *model interference* violates the SUTVA when learning from pooled data, we provide an intuitive example and experimental proof on synthetic data.¹

In line with the setup we laid out above, we consider varying multi-armed bandit methods to learn a recommendation policy over time. We assume there are $|\mathcal{A}| = N$ items to recommend with binary rewards: $r \sim \text{Bernoulli}(\rho_a)$. Naturally, ρ_a is unknown to the learning methods. We simulate it to be uniformly spaced over a $[0.05, 0.15]$ interval, to represent “clicks” on items.

We consider Bayesian modelling methods with a Beta(2, 10) prior over ρ_a — emulating reasonable knowledge of the system (i.e. $\widehat{\rho}(a) = 0.10$ in absence of data). For every action $a \in \mathcal{A}$, we log the reward r . For convenience of notation, we will denote with $s(a)$ the number of clicks (successes) we observed for action a . Analogously, $f(a)$ denotes the number of non-clicks (failures). The data aggregation functions s and f can either be computed based on all available data (*pooled*), or only on data collected by the learning method itself (*siloed*). We consider the following classical multi-armed bandit methods, where ties are broken randomly:

¹A Python notebook reproducing the synthetic data and figures presented in this section has been made available at <https://github.com/olivierjeunen/pooled-data-experiments>.

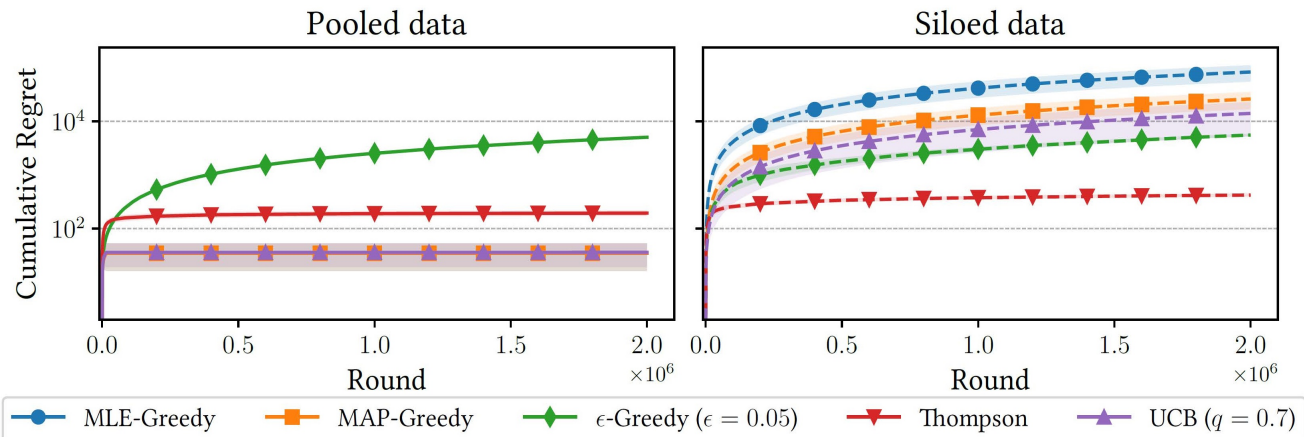


Figure 1. Cumulative regret for competing methods in an online experiment, where training data is either *pooled* or *siloed*. When siloed, the SUTVA holds, and online results are representative of what would happen if we would have shipped a method to the entire population. When pooled, models are faster to converge, but results are *not* representative because the SUTVA is violated.

MLE-Greedy acts according to the Maximum Likelihood Estimate: $\arg \max_{a \in \mathcal{A}} \frac{s(a)}{s(a)+f(a)}$,

MAP-Greedy acts according to the Maximum A Posteriori estimate: $\arg \max_{a \in \mathcal{A}} \frac{2+s(a)}{12+s(a)+f(a)}$,

epsilon-Greedy acts uniformly at random with probability ϵ , and according to MAP-Greedy otherwise,

Thompson follows a classical Thompson Sampling approach [Chapelle and Li, 2011, Alg. 2],

UCB acts according to the q^{th} quantile of the posterior distribution [Kaufmann et al., 2012].

We simulate 2×10^6 successive rounds where every method is presented with an opportunity to act and learn from the resulting observed reward, with $N = 11$ actions. To reduce the effects of statistical variation, this process is repeated from the onset 20 times, and we report 95% confidence intervals for the average cumulative regret. Note that this setup is overly simplified to represent any real-world system (2–10 million training samples for only 11 actions and no contextual dependence), and it is not our goal to simulate real-world behaviour as accurately as possible. Nevertheless, it provides a simple, intuitive and reproducible example of the consequences that occur due to a SUTVA violation in online experiments with machine learning models.

We repeat the above scenario with two configurations: *pooled* data, i.e. after every round, all models are updated with the observed action-reward pairs from all variants; and *siloed* data, i.e. after every round, models are only updated with the observed rewards for their own actions. Figure 1 visualises 95% confidence intervals over the cumulative regret for both configurations. As expected, we observe disparate relative performance of competing methods under the two regimes. Specifically, under the pooled data regime, the greedy methods (and UCB) perform exceedingly well in comparison with the other bandit algorithms (ϵ -Greedy and Thompson). This conclusion reverses when we consider the siloed data regime, where the Thompson sampling approach significantly outperforms others, and the ϵ -Greedy method has gone, relative to competing methods, from the worst to the 2nd best method (in terms of cumulative regret).

To researchers and practitioners familiar with the multi-armed bandit paradigm, it will be obvious that the *pooled* data scenario cannot be reconciled with the very motivation behind these

sequential decision-making methods [Robbins, 1952]. Indeed, we famously need to balance *exploitation* with *exploration*: considering both the immediate reward an action can bring, as well as the informational value we get from executing it. By pooling the data together, *all* competing methods can reap the benefits of exploration data generated by *any* of the variants. As a result, we observe a piggybacking effect, where greedy (exploitation-heavy) methods are disproportionately favoured over exploration-heavy methods. From the siloed data regime, however, we can clearly observe that pooled data experiments give a false representation of what would happen if the variants were shipped in isolation. The SUTVA is violated, and this type of experimental setup is unsuitable when the goal is to provide causal estimates of the effects that different machine learning methods have on the reward that we incur from running them on the platform.

As such, we are convinced that online experiments where competing machine learning models learn from pooled data, cannot be guaranteed to provide unbiased estimates of true causal effects. Nevertheless, we argue that this style of experiment is widespread in industry. Often, the differences between variants might not be as outspoken as in the intuitive example laid out above, where we have exploitation-heavy methods piggybacking off of exploration-heavy alternatives. In such cases, a SUTVA violation might not be as detrimental to the truthfulness of the online experiment as it is in our synthetic example. Nevertheless, the SUTVA is violated, and we have reasons to be critical about conclusions that are drawn from it. This holds not only for the platforms that are running the experiments themselves, but also for the scientific community. Indeed, a SUTVA violation is crucial information to have for practitioners and researchers (readers and reviewers) that wish to assess the scientific contributions of such experiments.

4 Why is the pooled data setup dominant in practice?

To consider the reasons why the pooled data setup occurs so often in practice, we must consider what makes it attractive. A first argument in favour, comes from its engineering simplicity. Indeed, interactions between users and items that are generated over time on the platform can be logged in a shared database, and they need no further special treatment down the line. That is, the training data for machine learning methods can simply be obtained through a scan of the database.

A second and related argument in favour, is that of training data size. In the age of “*data hungry*” deep learning methods [Marcus, 2018], it is hard to make a case for —optimistically— halving the training data size of all models, or —more realistically— learning from only a fraction of the available training interactions, when we have multiple treatments being tested. This is especially true when we consider a recently popular class of off-policy learning approaches that typically suffer from high variance in their objective functions [Vasile et al., 2020; Saito and Joachims, 2022]. In these use-cases, it comes natural to try and combat variance with an increased training sample size. Whilst natural, we must be aware that it violates the SUTVA.

Finally, there are differing incentives when running a platform successfully as a business, compared to measuring experimental outcomes in a scientifically sound manner. Often, the long-term goals of researchers and practitioners running online experiments in industry will be focused on *improving* a given metric of interest (the outcome), and not necessarily on *measuring* the effect of the treatment on the outcome as truthfully as possible. Online experiments then become a means to an end, and minimising the number of false positives that they generate might become more important than having unbiased causal effect estimates. Understandably so, there is a rich

literature on avoiding pitfalls like false positives, related to statistical mis-reasoning in A/B-testing, and focusing on p -values and hypothesis testing [Kohavi et al., 2022].

Nevertheless, when results from online experiments are adopted as empirical proof in the scientific literature, it is critical that we are explicit about the assumptions that underlie the methods that are used, and to which extent they are likely to be violated.

5 Conclusions and Outlook

Offline experiments to validate the effectiveness of machine learning models in applications that rely heavily on user interactions, are notoriously hard to get right [Gilotte et al., 2018; Jeunen, 2019; Diaz, 2021; Deffayet et al., 2023]. This is a well-known issue in the recommender systems research community, and contrasting results between on- and off-line experiments have been reported in the literature time and again [Beel et al., 2013; Garcin et al., 2014; Rossetti et al., 2016; Jeunen et al., 2018]. The reasons underlying these diverging results, are often (rightfully) ascribed to problematic offline evaluation practices, in that they do not produce unbiased offline estimators of online metrics [Jeunen, 2021]. Online experiments, on the other hand, are typically seen as the gold standard, owing to the elegant statistical theory that motivates their use.

In this short opinion paper, we have argued that this confidence might be misplaced in many practical A/B-testing scenarios. We have briefly reviewed the main assumption that makes RCTs tick (i.e. the Stable Unit Treatment Value Assumption (SUTVA)), and have motivated how it is easily violated in the prevalent online experiment setup where we have competing machine learning methods that learn from a shared *information* resource — be it training samples or features. When this information is somehow influenced by other machine learning methods evaluated in the experiment, this is a clear SUTVA violation. It occurs when either the models themselves or the features they consume are updated with new (pooled) data throughout the experiment. Through a classical multi-armed bandit setup, we have provided an empirical example that highlights how this practice can invalidate conclusions drawn from online experiments with similar violations.

As we have mentioned earlier, it is not the intention of this paper to undermine the credibility of existing work, or to dissuade the use of A/B-tests as an evaluation tool. We do wish to cultivate an awareness of the assumptions that we (often implicitly) make when running online experiments, and point out that they are easily violated in typical modern use-cases. As we have argued, there can be good reasons for adopting an experimental design where SUTVA is violated (e.g. to increase the training sample size for competing methods). Nevertheless, in such cases, results obtained through and conclusions drawn from online experiments, should be presented with an explicit caveat. This ensures that readers have access to the information that is necessary to properly assess their scientific value, and possible contributions to the state-of-the-art.

This paper raises more questions than it answers. It is our hope that it can contribute to conversations about future work, to (1) qualitatively and quantitatively assess to what extent SUTVA violations affect experimental outcomes, (2) to develop tools and metrics that can identify problematic cases, and (3) to provide guidance and improved methods for problematic scenarios.

In particular, future work that studies the relation between the SUTVA violations we describe and the bias they impose on experimental outcomes, will be hugely impactful for both the academic and industrial research communities. This can include both theoretical works that aim to prove

bounds on the statistical bias of the estimator, as well as empirical works that quantify the extent to which different types of common online experiments with SUTVA violations can be trusted.

Acknowledgements

The opinions voiced in this article have benefited from discussions with many colleagues and collaborators. The author would like to explicitly thank Andrés Ferraro, Thorsten Joachims and Lien Michiels for their feedback on earlier drafts.

References

- Patrick Bajari, Brian Burdick, Guido W. Imbens, Lorenzo Masoero, James McQueen, Thomas Richardson, and Ido M. Rosen. Multiple randomization designs, 2021. URL <https://arxiv.org/abs/2112.13495>.
- Joeran Beel, Marcel Genzmehr, Stefan Langer, Andreas Nürnberger, and Bela Gipp. A comparative analysis of offline and online evaluations and discussion of research paper recommender system evaluation. In *Proceedings of the International Workshop on Reproducibility and Replication in Recommender Systems Evaluation*, RepSys '13, page 7–14. ACM, 2013. ISBN 9781450324656. URL <https://doi.org/10.1145/2532508.2532511>.
- Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/e53a0a2978c28872a4505bdb51db06dc-Paper.pdf>.
- Romain Deffayet, Thibaut Thonet, Jean-Michel Renders, and Maarten de Rijke. Offline evaluation for reinforcement learning-based recommendation: A critical issue and some alternatives. *SIGIR Forum*, 56(2), jan 2023. ISSN 0163-5840. URL <https://doi.org/10.1145/3582900.3582905>.
- Fernando Diaz. On evaluating session-based recommendation with implicit feedback. In *Workshop on Perspectives on Offline Evaluation for Recommender Systems at RecSys '21*, PERSPECTIVES '21, 2021.
- Ronald Aylmer Fisher. *Statistical methods for research workers*. Springer, 1925.
- Ronald Aylmer Fisher. Design of experiments. *British Medical Journal*, 1(3923):554, 1936.
- Florent Garcin, Boi Faltings, Olivier Donatsch, Ayar Alazzawi, Christophe Bruttin, and Amr Huber. Offline and online evaluation of news recommender systems at swissinfo.ch. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, page 169–176. ACM, 2014. ISBN 9781450326681. URL <https://doi.org/10.1145/2645710.2645745>.
- Alexandre Gilotte, Clément Calauzènes, Thomas Nedelec, Alexandre Abraham, and Simon Dollé. Offline A/B testing for recommender systems. In *Proceedings of the Eleventh ACM International*

-
- Conference on Web Search and Data Mining, WSDM '18*, page 198–206. ACM, 2018. ISBN 9781450355810. URL <https://doi.org/10.1145/3159652.3159687>.
- Somit Gupta, Ronny Kohavi, Diane Tang, Ya Xu, Reid Andersen, Eytan Bakshy, Niall Cardin, Sumita Chandran, Nanyu Chen, Dominic Coey, Mike Curtis, Alex Deng, Weitao Duan, Peter Forbes, Brian Frasca, Tommy Guy, Guido W. Imbens, Guillaume Saint Jacques, Pranav Kantawala, Ilya Katsev, Moshe Katzwer, Mikael Konutgan, Elena Kunakova, Minyong Lee, MJ Lee, Joseph Liu, James McQueen, Amir Najmi, Brent Smith, Vivek Trehan, Lukas Vermeer, Toby Walker, Jeffrey Wong, and Igor Yashkov. Top challenges from the first practical online controlled experiments summit. *SIGKDD Explor. Newsl.*, 21(1):20–35, may 2019. ISSN 1931-0145. URL <https://doi.org/10.1145/3331651.3331655>.
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Olivier Jeunen. Revisiting offline evaluation for implicit-feedback recommender systems. In *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys '19*, page 596–600. ACM, 2019. ISBN 9781450362436. URL <https://doi.org/10.1145/3298689.3347069>.
- Olivier Jeunen. *Offline Approaches to Recommendation with Online Success*. PhD thesis, University of Antwerp, 2021.
- Olivier Jeunen, Koen Verstrepen, and Bart Goethals. Fair offline evaluation methodologies for implicit-feedback recommender systems with mmar data. In *Workshop on Offline Evaluation for Recommender Systems at RecSys '18, REVEAL '18*, 2018.
- Emilie Kaufmann, Olivier Cappe, and Aurelien Garivier. On bayesian upper confidence bounds for bandit problems. In Neil D. Lawrence and Mark Girolami, editors, *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *ICML '12*, pages 592–600. PMLR, 21–23 Apr 2012. URL <https://proceedings.mlr.press/v22/kaufmann12.html>.
- Ron Kohavi, Diane Tang, and Ya Xu. *Trustworthy online controlled experiments: A practical guide to A/B testing*. Cambridge University Press, 2020.
- Ron Kohavi, Alex Deng, and Lukas Vermeer. A/b testing intuition busters: Common misunderstandings in online controlled experiments. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, page 3168–3177. ACM, 2022. ISBN 9781450393850. URL <https://doi.org/10.1145/3534678.3539160>.
- Min Liu, Jialiang Mao, and Kang Kang. Trustworthy and powerful online marketplace experimentation with budget-split design. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, page 3319–3329. ACM, 2021. ISBN 9781450383325. URL <https://doi.org/10.1145/3447548.3467193>.
- Gary Marcus. Deep learning: A critical appraisal, 2018. URL <https://arxiv.org/abs/1801.00631>.

-
- Charles Sanders Peirce and Joseph Jastrow. On small differences in sensation. *Memoirs of the National Academy of Sciences*, 3:73–83, 1884.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.*, 58:527–535, 1952.
- Marco Rossetti, Fabio Stella, and Markus Zanker. Contrasting offline and online results when evaluating recommendation algorithms. In *Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16*, page 31–34. ACM, 2016. ISBN 9781450340359. URL <https://doi.org/10.1145/2959100.2959176>.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Yuta Saito and Thorsten Joachims. Counterfactual evaluation and learning for interactive systems: Foundations, implementations, and recent advances. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, page 4824–4825. ACM, 2022. ISBN 9781450393850. URL <https://doi.org/10.1145/3534678.3542601>.
- Flavian Vasile, David Rohde, Olivier Jeunen, and Amine Benhalloum. A gentle introduction to recommendation as counterfactual policy learning. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '20*, page 392–393. ACM, 2020. ISBN 9781450368612. URL <https://doi.org/10.1145/3340631.3398666>.