

Internship Proposal (Research): “Online Algorithms for Fair Ad Auctions”

Keywords: Online Learning, Ad auctions, Fairness, Machine learning, Matching

Host lab & team: Laboratoire d’Informatique de Grenoble (LIG), Grenoble, France; POLARIS team

Supervisors:

Nicolas Gast (Inria) – <http://polaris.imag.fr/nicolas.gast/> – nicolas.gast@inria.fr

Patrick Loiseau (Inria & Ecole Polytechnique) – <http://lig-membres.imag.fr/loiseapa/> – patrick.loiseau@inria.fr

Background

Ad exchange platforms are used to match advertisers and users. Typically, each time a user comes, bids are collected from various relevant advertisers and an auction is conducted in order to decide which ad should be shown to the user, and how much the chosen advertisers should pay (e.g., first price, second price auctions). In practice, many platforms use automatic bidding platforms. Each advertiser is asked to provide general information about its campaign (such as target audience and budget) and the platform places bids on their behalf [10]. This allows the platform to provide better targeted advertisement (for instance, higher *click-through-rate*) than if bids were chosen only by the advertisers. To summarize, the ads that are shown to users are chosen according to a complex mechanism that involves the audience selection, auctions and some inside action by the platform (such as targeting or other mechanisms like pacing).

While this process is quite efficient, it has been shown in the literature that targeted advertisement can lead to discriminations [2], such as gender discrimination, that are often illegal [1]. For instance, it has been shown in [7, 9] that ads for highly qualified jobs are more displayed to men than to women. In fact, even if the audience defined by the advertisers is not discriminating, it has been shown that the actual population that sees an ad can be highly biased against some sub-population. This can be due to the targeting mechanism done by the platform [2] or to the auctions mechanisms and the competition between different advertisers [2, 6, 3]. The crux of the problem is that while most fairness notions (called group fairness) are defined per ad based on the entire population seeing the ad, the mechanism to decide ad impression is on the contrary defined per user based on the set of available ads; it is therefore not straightforward how one can modify current ad delivery mechanisms to guarantee fairness notions.

Goal of the internship

The main objective of this internship is to formulate the notion of an online fair advertisement system and to develop online learning algorithms that choose which ads to display in a fair manner. We will focus on the notion of group fairness.¹ The main idea is that, considering a given auction mechanism, the platform can adapt dynamically the bids of each advertisers to achieve fairness guarantees. Our proposed approach will be similar to the pacing multipliers, described for instance in [4, 5], but using different multipliers for the different subgroup of the population (i.e., the groups with different values of the sensitive attributes).

The starting point will be to study fairness in a static auction problem where the targeted population of all ads is known in advance along with all their characteristics and in which an auction mechanism is fixed (e.g., to first price). The goal there will be to introduce the notion of auction equilibrium under fairness constraints. Seeing the fairness as constraints of an optimization problem will allow us to consider the dual version of the problem which we intend to use to derive optimal offline multipliers that can be used to adapt the bids of each advertiser. This static problem will serve as a benchmark for the online case. It will also be useful to study how fairness affects ad performance (measured for instance in terms of revenue, advertisers welfare or quality of targeting).

Following the understanding of the static problem, the intern will use this formulation in an online setting in which users (but also advertisers) are dynamic and in which decisions must be taken online. Our view is that the fairness constraints could be learned by using online Lagrange multipliers that would modify the prices to show an ad to a category of users. This could be done similarly to pacing mechanisms [4, 5], by considering one multiplier per group and per fairness constraint that varies with time. The goal is to make these virtual multipliers evolve fast enough in order to obtain envy-free (or no-regret) algorithms.

Expected ability of the student and additional information

A strong background in probability is necessary, exposure to (online) learning and programming are a plus. The internship is part of the Explainable and Responsible AI chair of the MIAI@Grenoble Alpes institute and will be hosted in the POLARIS team, a joint team between Inria and LIG (Grenoble CS lab). It may be continued as a PhD. For more information, please contact patrick.loiseau@inria.fr and nicolas.gast@inria.fr.

¹Intuitively, this notion prescribes that different demographic groups receive similar outcomes [8].

References

- [1] “24 CFR § 100.75 - Discriminatory advertisements, statements and notices.” URL: <https://www.law.cornell.edu/cfr/text/24/100.75>.
- [2] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke. “Discrimination through optimization: How Facebook’s ad delivery can lead to skewed outcomes”. In: *ACM Conference on Computer Supported Cooperative Work (CSCW)*. 2019.
- [3] L. E. Celis, A. Mehrotra, and N. K. Vishnoi. “Toward Controlling Discrimination in Online Ad Auctions”. In: *International Conference on Machine Learning (ICML)*. 2019.
- [4] V. Conitzer, C. Kroer, E. Sodomka, and N. E. Stier-Moses. “Multiplicative pacing equilibria in auction markets”. In: *Conference on Web and Internet Economics (WINE)*. 2018.
- [5] V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. E. Stier-Moses, and C. Wilkens. “Pacing equilibrium in first-price auction markets”. In: *Conference on Economics and Computation (EC)*. 2019.
- [6] C. Dwork and C. Ilvento. “Fairness under composition”. In: *Innovations in Theoretical Computer Science (ITCS)*. 2019.
- [7] A. Hannák, C. Wagner, D. Garcia, A. Mislove, M. Strohmaier, and C. Wilson. “Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr”. In: *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM. 2017, pp. 1914–1933.
- [8] M. Hardt, E. Price, N. Srebro, et al. “Equality of opportunity in supervised learning”. In: *Advances in neural information processing systems*. 2016, pp. 3315–3323.
- [9] A. Lambrecht and C. Tucker. “Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads”. In: *Management Science* (2019).
- [10] “Mobile Advertising assessment study guide. About automated bidding”. URL: <https://support.google.com/google-ads/answer/2979071?hl=en>.