

Fairness in Online Social Network Timelines: Measurements, Models and Mechanism Design

Eduardo Hargreaves, Claudio Agosti, Daniel Menasché, Giovanni Neglia,
Alexandre Reiffers-Masson, Eitan Altman

► **To cite this version:**

Eduardo Hargreaves, Claudio Agosti, Daniel Menasché, Giovanni Neglia, Alexandre Reiffers-Masson, et al.. Fairness in Online Social Network Timelines: Measurements, Models and Mechanism Design. IFIP Performance, Dec 2018, Toulouse, France. hal-01910462

HAL Id: hal-01910462

<https://hal.inria.fr/hal-01910462>

Submitted on 1 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fairness in Online Social Network Timelines: Measurements, Models and Mechanism Design*

Eduardo Hargreaves
Federal University of Rio de Janeiro
Rio de Janeiro, Brazil

Claudio Agosti
University of Amsterdam
Amsterdam, The Netherlands

Daniel Menasché
Federal University of Rio de Janeiro
Rio de Janeiro, Brazil

Giovanni Neglia
Université Côte d'Azur, Inria
Sophia Antipolis, France

Alexandre Reiffers-Masson
Indian Institute of Science
Bangalore, India

Eitan Altman
Université Côte d'Azur, Inria
Sophia Antipolis, France

ABSTRACT

Facebook News Feed personalization algorithm has a significant impact, on a daily basis, on the lifestyle, mood and opinion of millions of Internet users. Nonetheless, the behavior of such algorithm lacks transparency, motivating measurements, modeling and analysis in order to understand and improve its properties. In this paper, we propose a reproducible methodology encompassing measurements, an analytical model and a fairness-based News Feed design. The model leverages the versatility and analytical tractability of time-to-live (TTL) counters to capture the visibility and occupancy of publishers over a News Feed. Measurements are used to parameterize and to validate the expressive power of the proposed model. Then, we conduct a what-if analysis to assess the visibility and occupancy bias incurred by users against a baseline derived from the model. Our results indicate that a significant bias exists and it is more prominent at the top position of the News Feed. In addition, we find that the bias is non-negligible even for users that are deliberately set as neutral with respect to their political views, motivating the proposal of a novel and more transparent fairness-based News Feed design.

1. INTRODUCTION

Background. Online social networks (OSNs) have an increasingly important influence in the life of millions of Internet users, shaping their mood, tastes and political views [3]. In essence, the goal of OSNs is to allow users to connect and to efficiently share meaningful information. To this aim, one of the key building blocks of OSNs is their filtering algorithm, which personalizes content made available to each user of the system. Facebook, for instance, developed the *News Feed* algorithm for that purpose.

However, Facebook users may be unaware of the influence of the filtering they are subject to [5]. Such lack of awareness favors the creation of a *filter bubble* that reinforces users' opinions by selecting the users information diets [11].

*The full version of this paper is available online [6].

Challenges. The News Feed algorithm is a recommendation system that shows posts to users based on inferred users' preferences, trading among possibly conflicting factors while prioritizing posts [10]. Hence, the News Feed algorithm shares common features with traditional recommender systems, such as those used by Netflix and Spotify, to recommend movies and music. For instance, in all such systems users typically do not provide explicit feedback about recommendations. Nonetheless, due to the nature of OSNs, the News Feed algorithm also poses its own set of challenges, related to the measurement, modeling and control of publishers' visibilities.

Models are instrumental to perform what-if analysis, e.g., to understand how the News Feed would behave in the presence of different filtering algorithms. In addition, analytical models can also serve as building blocks towards novel mechanisms to populate the News Feed. Such foundational development of principled mechanisms to populate the News Feed is key to build transparency into the system.

Prior art. Most of the prior work on the News Feed algorithm relies on measurements obtained through restrictive non-disclosure agreements that are not made publicly available to other researchers and practitioners [1, 2, 8].

Goals. Our goal is to provide insights on the filtering that occurs in OSNs through a reproducible methodology and a dataset made publicly available. Given such measurements, we pose the following questions:

1. what would be the occupancies of the various sources under alternative scenarios wherein different filtering algorithms are in place?
2. how to design mechanisms to populate timelines in a principled fashion, accounting for users preferences and providing content diversity, e.g., under a fairness-based framework?

To address the first question above, we propose an analytical model for the News Feed. The model allows us to derive the occupancy and visibility of each publisher at users' timelines, as a function of the considered filtering process. Using the model, we conduct what-if analysis, e.g., to assess publishers' visibilities in a scenario without filters.

Then, we use the proposed model to build fairness-driven mechanisms to populate timelines. Utilities are used to cap-

ture the preferences of users with respect to the exposure to posts from different publishers. The mechanism leverages results on utility-based cache design [4], and accounts for fairness among publishers through utility functions.

Contributions. In this paper we take important steps towards measuring, modeling, auditing and designing timelines. Our key contributions are summarized as follows.

A measurement methodology is introduced to publicly and transparently audit the OSN ecosystem, focusing on the Facebook News Feed algorithm. The methodology encompasses an Internet browser extension to autonomously and independently collect information on the posts presented to users by the News Feed algorithm. Such information is not available through the Facebook API.

Empirical findings are reported using data collected from measurements conducted during the 2018 Italian presidential elections. We observed that *a*) the filtering algorithm is impacted by the profile of pages that a user “likes”, *b*) this effect is stronger at the topmost News Feed position and *c*) neutral users are also exposed to non-uniform filtering.

An analytical model is proposed to quantify the visibility and occupancy of publishers in the users’ News Feeds. The model allows us to conduct a what-if analysis, to assess the metrics of interest under different filtering mechanisms and is validated using data from the experiment.

A fairness-driven mechanism design is proposed, leveraging the proposed model and measurements. Given a user profile, that “likes” a certain subset of publishers, the measurements are used to parameterize a simple instance of the model. Then, a family of α -fair utility functions are used to allocate resources to publishers subject to users preferences under a utility maximization framework.

A model-based bias assessment is conducted where the News Feed occupancies are contrasted against an unfiltered baseline to quantify the *bias*, i.e., how publishers’ occupancies are affected by the News Feed algorithm.

2. A FAIRNESS-BASED MECHANISM

In order to fulfill the proposed goals, we formulate an utility maximization problem. Then, multiple notions of fairness are introduced using a family of α -fair utility functions.

2.1 Utility Maximization Problem formulation

Let N_{ij} be the average number of posts of publisher j in top K positions of the News Feed of user i . We associate to each user-publisher pair a utility function $U_{ij}(h_{ij})$, which is an increasing, strictly concave and continuously differentiable function of h_{ij} over the range $0 \leq h_{ij} \leq 1$. In the simplest case, h_{ij} equals the normalized occupancy of publisher j at the News Feed of user i , i.e. $h_{ij} = N_{ij}/K$.

We further assume that utilities are additive. The goal is to maximize the sum of the utilities for each of the individual publishers. The optimization problem is posed as follows,

$$\max \sum_{j=1}^J U_{ij}(h_{ij}) \quad (1)$$

$$\text{s.t.} \quad \sum_{j=1}^J N_{ij} = K, \quad N_{ij} \geq 0 \quad (2)$$

where h_{ij} are concave and non-decreasing functions of N_{ij} .

2.2 News Feed fairness: α -fair utilities

The optimal allocation of News Feed space to publishers through equations (1) and (2) depends on the shape of the utility functions. Let w_{ij} be a weight that gives the preference that user i has for publisher j . The use of a given family of utility functions corresponds to the selection of a fairness criterion. As special case, we study the optimal allocation under the concept of α -fairness as it is considered in communication networks [7].

$$U_{ij}(h_{ij}) = \begin{cases} w_{ij} \frac{h_{ij}^{1-\alpha}}{1-\alpha}, & \alpha \leq 0, \alpha \neq 1. \\ w_{ij} \log(h_{ij}), & \alpha = 1. \end{cases} \quad (3)$$

The solution of (1)-(3) is presented and discussed in [6]. In [6] we also compare the allocation derived from the optimization problem above with that observed in our measurements of Facebook News Feeds collected during the 2018 Italian elections.

3. CONCLUSION

We believe that the measurements, model and design mechanisms presented in this work are one step further towards algorithmic transparency, as they promote user awareness about the filtering process they are submitted to. Ultimately, such awareness is key to protect and empower online social networks users, communities, society and democracy as a whole [9].

4. REFERENCES

- [1] E. Bakshy, S. Messing, and L. A. Adamic. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239):1130–1132, 2015.
- [2] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic. The Role of Social Networks in Information Diffusion. In *WWW '12*, 2012.
- [3] R. M. Bond, C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295, 2012.
- [4] M. Dehghan, L. Massoulie, D. Towsley, D. Menasche, and Y. C. Tay. A utility optimization approach to network cache design. In *IEEE INFOCOM*, 2016.
- [5] M. Eslami, A. Rickman, K. Vaccaro, A. Aleyasen, A. Vuong, K. Karahalios, K. Hamilton, and C. Sandvig. “I always assumed that I wasn’t really that close to [her]”. *CHI '15*, 2015.
- [6] E. Hargreaves, C. Agosti, D. Menasché, G. Neglia, A. Reiffers-Masson, and E. Altman. Fairness in Online Social Network Timelines: Measurements, Models and Mechanism Design. <http://bit.ly/2NI925P>, 2018.
- [7] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 1997.
- [8] A. Kramer, J. E. Guillory, and J. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, (29), 2014.
- [9] G. Lichfield. Technology is threatening our democracy. how do we save it? *MIT Technology Review*, September/October 2018.
- [10] A. Mosseri. News feed ranking in three minutes flat, 2018. <http://bit.ly/2QneMQG>.
- [11] E. Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. The Penguin Group, 2011.