

# Balanced News Using Constrained Bandit-based Personalization

Sayash Kapoor<sup>\*†</sup>, Vijay Keswani<sup>\*</sup>, Nisheeth K. Vishnoi<sup>\*</sup>, L. Elisa Celis<sup>\*</sup>

<sup>\*</sup> École Polytechnique Fédérale de Lausanne

<sup>†</sup> Indian Institute of Technology, Kanpur

{sayash.kapoor, vijay.keswani, nisheeth.vishnoi, elisa.celis}@epfl.ch

## Abstract

We present a prototype for a news search engine that presents *balanced* viewpoints across liberal and conservative articles with the goal of depolarizing content and allowing users to escape their filter bubble. The balancing is done according to flexible user-defined constraints, and leverages recent advances in constrained bandit optimization. We showcase our balanced news feed by displaying it side-by-side with the news feed produced by a traditional (polarized) feed.

## 1 Introduction

The personalization of content online can lead to opinion bubbles and polarization [Flaxman *et al.*, 2016]. As a user clicks on content of a particular type (e.g., links from liberal media outlets), the content delivery system learns these preferences, and in the future presents the user with more items of that type. In doing so, many content delivery engines polarize completely, displaying a very narrow range of content to each user. This can lead to polarization and the formation of opinion bubbles in which people are only exposed to content that matches their pre-existing beliefs [Matakos *et al.*, 2017], and can lead people to mistrust information that does not match their opinion [H. Ribeiro *et al.*, 2017]. Personalization leading to polarization has been observed in social media feeds [Flaxman *et al.*, 2016], product recommendation [De Myttenaere *et al.*, 2015], online advertising [Datta *et al.*, 2015] and other media that pervades the internet. Moreover, diverse search results might be desirable from the point-of-view of user satisfaction [Celis and Tekriwal, 2017]. Our goal in this work is to show that an alternative exists: Balanced content delivery systems are feasible, and one can even give the user the information and freedom to control their balance of content. In this demo, as an example, we take news search in the US diversified across liberal and conservative content.

**Balanced News** is our prototype <sup>1 2</sup> aimed at demonstrating novel bandit-based *constrained* personalization algorithms that can be used for delivering diverse content while still allowing for personalization. <sup>2</sup> A feature of this recent

<sup>1</sup>Demo video: <http://y2u.be/Zvum0t1EYtQ>

<sup>2</sup>Demo website: <https://bit.ly/2H1vroP>

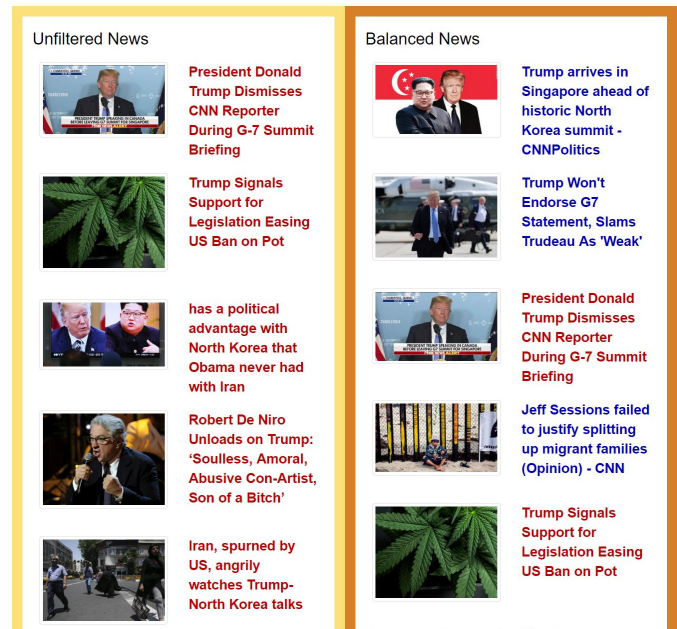


Figure 1: Balanced news search demo for a liberal-favoring user. A traditional system (unfiltered news) on the left would learn about the user and only display liberal-favoring articles. Our constrained system (balanced news) on the right similarly learns the user preferences, but still displays some conservative articles in order to provide more diversified content. (*Liberal-favoring articles in blue, conservative-favoring articles in red.*)

work is that the algorithms are scalable despite constraints, allowing for real-time systems to be balanced in this manner. The algorithms optimize content in the presence of constraints, that is, limits on how much or how little content of a particular *type* is displayed. We consider news search in the US and classify the content into two *types*: liberal-favoring and conservative-favoring. Figure 1 shows a screenshot of the demo website. For showcasing the types that each article belongs to, articles belonging to the liberal type have blue hyperlinks, whereas articles belonging to the conservative type have red hyperlinks. There are two news feeds shown in the demo – the *unfiltered* feed that runs a standard Bandit algorithm used for personalization, and the *balanced* feed that runs the constrained Bandit algorithm from [Celis

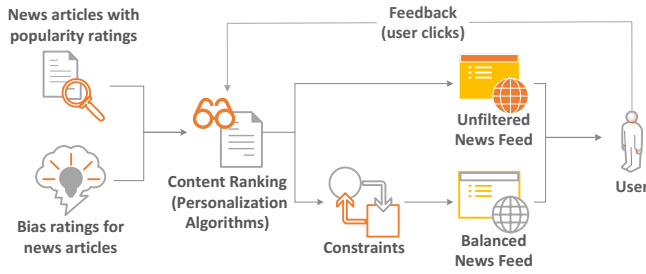


Figure 2: Flowchart of the balanced search prototype.

*et al.*, 2018b]. In both feeds, clicking on conservative articles decreases the proportion of liberal articles in the next iteration. Similarly, clicking on liberal articles increases the proportion of liberal articles. The main difference between the unfiltered and balanced news feeds is that while the unfiltered feed can continue getting polarized in a direction until *all* articles belong to only one type, the balanced news feed can only specialize up to user-defined constraints – once the constraints have been reached, the proportion of articles of the type favored by the user can not increase.

On a high-level, similar “group-fairness” constraints have also been considered for other fundamental problems, including classification [Dwork *et al.*, 2012; Zafar *et al.*, 2017], data summarization [Celis *et al.*, 2016; 2018c], ranking [Celis *et al.*, 2018d; Yang and Stoyanovich, 2017], and voting [Celis *et al.*, 2018a].

## 2 Implementation

In our demo we display two different news feeds side-by-side in order to contrast the difference between a traditional (unfiltered) news feed, and our (balanced) news feed. The former is implemented using standard bandit algorithms, and the latter is implemented using constrained bandit algorithms as inspired by [Celis *et al.*, 2018b]. For implementing the algorithms, for each article we require article a *rating*, a *type* (either liberal or conservative) and a *reward* function. For the balanced algorithm, we also require *constraints* on each type of content.

**Ratings.** We used the *WebHose News Search API*,<sup>3</sup> which ranks articles by popularity. This gives a warm-start to the Bandit algorithm; in the beginning we display articles based on popularity, and gradually personalize to the user’s preferences as they are learned.

**Type.** For each article, the *type* – either liberal or conservative – is determined based on bias ratings from *AllSides*.<sup>4</sup>

**Reward.** We take every user click as an indicator of a user’s preference for viewing more articles of that type, and use this as the *reward* for the bandit algorithm. An article  $a$  that the user clicks at the  $t$ -th iteration receives a reward  $r_t^a = 1$ , all other articles  $a'$  receive a reward  $r_t^{a'} = 0$ .

**Constraints.** We take upper and lower constraints on the percentage of liberal articles as input through the slider to the right of the graph (see Figure 3). By default, the constraints




Figure 3: Dashboard displaying the percentage of liberal articles in the unfiltered and balanced news feeds. Unfiltered news in yellow, balanced news in orange, constraints in black. While the percentage of liberal articles can increase (or decrease) to 100% (0%) in the case of unfiltered news, they are constrained to lie within the given region for balanced news (*orange*). The sliders on the right allow the user to change the constraints.

are set to 80% and 20% respectively; i.e., the balanced news feed would always contain *at least* 20% liberal articles, and 20% conservative articles – the remaining 60% would depend on the given user’s preference history.

The difference between the unfiltered and balanced news feeds becomes apparent once a user has clicked on enough articles of just one type: While unfiltered news completely polarizes to display articles largely of a single type (either conservative or liberal, depending on user preferences), balanced news always has a certain minimum percentage of articles of both types (as defined by the constraints) (see Figure 1).

Figure 3 demonstrates an instance of the constraints in action by showing the outcome for a new user who, upon entering the system, clicks on 5 consecutive liberal articles, and no conservative articles. The yellow line depicts the percentage of liberal articles in the unfiltered news feed over time, the orange line depicts the percentage of liberal articles in the balanced news feed over time, and the black lines depict the constraints over time. For a user who has clicked a large number of liberal-favoring articles, the unfiltered news feed has high chances of displaying liberal-favoring articles, whereas this probability is constrained by a certain upper (and lower) bound for the balanced news.

## 3 Conclusion

Using constrained personalization algorithms allows us to increase diversity and reduce polarization in content delivery systems, which can help combat the formation of opinion bubbles. Our balanced search prototype<sup>2</sup> is aimed at showcasing what such content delivery systems might look like. An important open problem is to determine how to set constraints on articles being delivered to users as soliciting constraints from the user directly may not always be feasible or desirable. In particular, it would be important to determine which constraints provide enough diversity to prevent opinion bubbles from forming or content from having disparate impact on specific populations, while at the same time not compromising user satisfaction.

<sup>3</sup><https://webhose.io>

<sup>4</sup><https://bit.ly/2GX11Rx>

## References

- [Celis and Tekriwal, 2017] L. Elisa Celis and Siddhartha Tekriwal. What Do Users Want in Q&A Sites: Quality or Diversity? In *International Conference on Computational Social Science (IC2S2)*, 2017.
- [Celis *et al.*, 2016] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. How to be fair and diverse? *Fairness, Accountability and Transparency in Machine Learning 2016*, 2016.
- [Celis *et al.*, 2018a] L. Elisa Celis, Lingxiao Huang, and Nisheeth K. Vishnoi. Multiwinner voting with fairness constraints. In *International Joint Conference on Artificial Intelligence*, 2018.
- [Celis *et al.*, 2018b] L. Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth K. Vishnoi. An Algorithmic Framework to Control Bias in Bandit-based Personalization. *ArXiv e-prints*, February 2018.
- [Celis *et al.*, 2018c] L. Elisa Celis, Vijay Keswani, Damian Straszak, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. Fair and diverse data summarization. In *International Conference on Machine Learning*, 2018.
- [Celis *et al.*, 2018d] L. Elisa Celis, Damian Straszak, and Nisheeth K. Vishnoi. Ranking with fairness constraints. In *International Colloquium on Automata, Languages and Programming*, 2018.
- [Datta *et al.*, 2015] Amit Datta, Michael C. Tschantz, and Anupam Datta. Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies*, 2015(1):92–112, 2015.
- [De Myttenaere *et al.*, 2015] Arnaud De Myttenaere, Boris Golden, Bénédicte Le Grand, and Fabrice Rossi. Study of a bias in the offline evaluation of a recommendation algorithm. *ArXiv e-prints*, November 2015.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012*, pages 214–226, 2012.
- [Flaxman *et al.*, 2016] Seth Flaxman, Sharad Goel, and Justin M. Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1):298–320, 2016.
- [H. Ribeiro *et al.*, 2017] Manoel H. Ribeiro, Pedro H. Calais, Virgílio A. F. Almeida, and Wagner Meira, Jr. “Everything I Disagree With is #FakeNews”: Correlating Political Polarization and Spread of Misinformation. *ArXiv e-prints*, June 2017.
- [Matakos *et al.*, 2017] Antonis Matakos, Evimaria Terzi, and Panayiotis Tsaparas. Measuring and moderating opinion polarization in social networks. *Data Mining and Knowledge Discovery*, 31(5):1480–1505, Sep 2017.
- [Yang and Stoyanovich, 2017] Ke Yang and Julia Stoyanovich. Measuring fairness in ranked outputs. In *Scientific and Statistical Database Management*, pages 22:1–22:6, 2017.
- [Zafar *et al.*, 2017] Muhammad B. Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In *International Conference on Artificial Intelligence and Statistics 2017*, pages 962–970, 2017.