Dealing with Item Cold-Start in News Recommender at Globo

Joel Pinho Lucas* Globo Rio de Janeiro, RJ, Brazil joel.pinho@g.globo Leticia Freire de Figueiredo[†]
Globo
Rio de Janeiro, RJ, Brazil
leticia.figueiredo@g.globo

Felipe Alves Ferreira Globo Rio de Janeiro, RJ, Brazil faferreira@inf.puc-rio.br

Abstract

Globo is the largest Latin American mass media group, where its vertical information portals play an important role in content distribution. Among such portals, G1 is Globo's journalism portal, being the most popular news portal in Brazil and responsible for delivering informative content to more than 100 million unique users per day. In this context, recommender systems play an important role in achieving a good user experience, offering personalized content. In this paper, we discuss how G1's recommender system identifies and deals with the item cold-start problem, describing the recommendation scenarios and how the applied improvements in the currently deployed algorithms led to a decreased processing time and an increased CTR in the context of news recommendations.

Keywords: recommender systems, collaborative-filtering, cold-start, news recommender

1 Introduction

In this paper, we describe a challenging scenario related to cold-start in Globo's Recommender System. Globo is the largest Latin American mass media group, where its vertical information portals play an important role in content distribution. Among such portals, G1 is Globo's journalism portal, being the most popular news portal in Brazil and responsible for delivering informative content to more than 100 million unique users per day. In this context, recommender systems play an important role in achieving a good user experience and offering personalized content.

Besides addressing recommendations to millions of users encompassing diverse profiles in terms of engagement and the type of content they consume, providing fresh and context-aware recommendations is also challenging, since thousands of news articles and videos are published every day.

In this paper, we discuss how G1's recommender system identifies and deals with the item cold-start problem. In this

In: WebMedia in Practice, Curitiba, Brasil. Anais Estendidos do Simpósio Brasileiro de Sistemas Multimídia e Web (WebMedia). Porto Alegre: Sociedade Brasileira de Computação, 2022.

© 2022 SBC – Sociedade Brasileira de Computação. ISSN 2596-1683 way, in section 2, we describe the recommendations scenario in G1 and the foundations and magnitude of such a problem. Secondly, in section 3 we propose an approach for measuring item cold-start impact by tracking recommendations coverage and architectural capability of responsiveness to new items published. Finally, in section 4, we describe a case study in order to analyze the impacts of efforts conducted last year to address this scenario.

2 G1's Recommendation Scenario

Observing specifically the G1 scenario, in addition, to reach millions of users, recommendations must deal with the dynamic environment of editorial news publication. In 2021, G1 published more than 350k news content, which resulted in a daily average of almost 1k new items ingested in Globo's recommender system. This scenario poses scalability and personalization challenges due to the diversity of content segments and the intensity of user interests.

In most recommendation scenarios, including e-commerce and video streaming platforms, new items usually represent a small piece of the relevant recommendation candidates. On the other hand, in news recommender systems the magnitude of recent items' impact is way more challenging. [1] states the news domain poses some challenges for recommender systems due to the fast-growing number of items, accelerated decay of item's value, and user preferences shift. Figure 1 depicts the total volume of G1's news page visits in April 2022, where the x-axis corresponds to the number of hours between the news publication and the page visit event.

The histogram clearly shows that items (news pages) are consumed predominantly on the same day they become available. This scenario, inevitably, results in a significant number of cold-start items. In order to deal with cold-start items, Globo's recommender systems architecture provides a fallback strategy, where a secondary, and typically more straightforward, model is requested for providing items to the recommendation list.

The scenario we describe in this paper refers to an itembased recommendation scenario in the G1 portal, where the article news displayed on the page is the anchor, or pivot item, for the recommendation. In G1, every news page encompasses a newsfeed component displaying other news related to the pivot, which is selected by recommendation models on the top of the architecture mentioned before. In

^{*}Both authors contributed equally to this research.

[†]Both authors contributed equally to this research.

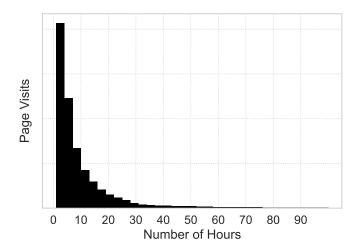


Figure 1. Number of hours between content publication and page visit event.

this item-based component, the primary recommendation algorithm is an item-based collaborative filtering (CFI) approach and the fallback is merely a list of the topmost visited news. Despite the highest conversion metrics, the CF algorithm depends on co-visitation events relating to the pivot. In this way, the high intensity of G1 content publication implies challenges related to item cold-start, which, along with the massive volume of user events, are intrinsically related to software engineering. This intensifies the cold-start problem, as for fresh items you cannot count on lots of interactions before starting to recommend them [2]. Intuitively, the faster the CF algorithm processes user events, the less the fallback algorithm is employed and the more page visits will occur. Thus, finding a strategy for measuring cold-start and also the primary algorithm coverage is crucial for facing this challenging scenario.

3 Measuring Item cold-start in Collaborative Filtering

In Globo's vertical portals, instead of considering the user as the starting point, our collaborative-filtering approaches take the item as a target. In G1, such methods provide recommendations that are related to the pivot item of the newsfeed. In this context, we employ two methods: association rules - to get items co-occurring together - and the cosine-based similarity. Suppose a set of P products, with m items. The association rule will define that an item $x_i \in X$ will co-occur with an item $y_j \in Y$ if it was seen by the same user. The set of items $X, Y \in P$ and $X \cap Y = \emptyset$ [4]. The cosine-based similarity calculate the angle between two items p_i and p_j [5].

The cold-start problem arises when the user or item just arrived and the model does not have information, such as how many users viewed the item or what items the user liked [3, 6]. Collaborative-filtering models cannot handle

this drawback, since they need prior information about both, the user and the item, to indicate what to recommend. Since G1 faces accelerated decay of item's value, recommending items has the challenge of dealing with items that will lose their interest in a few hours. Thus, the algorithm processing time is critical for computing items' co-occurrence information. News recommenders that are heavily exposed to cold-starts need to measure the velocity they respond to new items. Therefore, tracking the algorithm's response time to a new pivot item might also determine the amount of effort spent on software engineering improvements. In this sense, we propose to track the difference of time (number of hours) between the news publication event and the first time the algorithm outputs similar items to the pivot, meaning that significant co-occurred items were able to be tracked. We defined an intuition for the algorithm's Item-cold Start Responsiveness (ICSR) as metric accounting for the mean number of hours between those two events in a given period of observation.

As described in the previous section, collaborative filtering is the primary, and until the moment the most effective, approach employed in G1's newsfeed component. In this way, measuring the proportion of items, among all recommendations, were originated from the primary algorithm (the collaborative filtering approach in this case) might also determine the amount of effort spent on software engineering improvements. Therefore, we compare the primary algorithm's *coverage* against the fallback algorithms in a given period of observation.

4 Case Study

This case study concerns G1's recommendation scenario described in section 2, where Globo's recommendation architecture is responsible for populating the newsfeed component available on every G1's news page. The CFI algorithm is the foundation of recommendations provided in this component, commonly resulting in higher CTR (click-through rate). Since G1's revenue comes mostly from advertising displayed on page view events, the CTR is one of the main online metrics pursued by the product.

As stated previously, the high volume of publications along with the user interest in recent content turns item cold start into an arduous challenge for G1. In this context, algorithms runtime plays a crucial role in processing as many items as possible in order to keep user preferences up to date. Therefore, in this scenario, software engineering efforts have been taken to reduce the CFI runtime. During this process, collecting appropriate metrics was essential to identify pain points and to continuously increase performance. Table 1 describes a comparison between two versions of CFI populating G1's newsfeed recommendations: v.1 from March 2021 and v.2 from November 2021. The latter contains engineering improvements responsible for reducing CFI's processing

time and, consequently, considering item interaction data earlier. Such engineering improvements consisted mainly of optimizations in our Hadoop cluster, which is beyond the scope of this study. Table 1's data consists of CTR, algorithms' coverage, and the item cold-start responsiveness (ICSR) from both CFI versions.

Table 1. Measuring the engineering improvements

	CTR	ICSR	Algorithm Coverage
CFI v.1	4.14	4.47h	22.2%
CFI v.2	4.57	3.01h	32.96%

Table 1 shows that CTR is directly proportional to CFI coverage and inversely proportional to its responsiveness to item cold-start. This scenario demonstrates improvements in our recommender architecture resulted in higher CTR as a consequence of more responsiveness to cold-starts.

5 Conclusion

At G1, Globo's journalism vehicle, we experience millions of unique users per day and thousands of news published yearly. Due to the lack of prior data on item correlations, our algorithms, inevitably, have to face the cold-start problem. In this scenario, measuring algorithms' responsiveness was an effective strategy for planning and conducting software engineering improvements.

As described in the case study in section 4, after performing changes in the recommendation architecture, we could observe higher CTR, 10.38% of lift, as a consequence of the collaborative filtering coverage increase. Such CTR lift provided additional page visits to G1's newsfeed, generating more advertising revenue for the product. Therefore, the business impact of the tracking approaches presented in this paper contributed to G1 key results and, consequently, generated more revenue from advertising. It is important to highlight that more recent items are not directly related to CTR, and other conversion metrics, in most recommender systems, however, in the news domain, in most scenarios, it does provide a high benefit.

References

- [1] Gabriel de Souza Pereira Moreira, Felipe Ferreira, and Adilson Marques da Cunha. 2018. News Session-Based Recommendations using Deep Neural Networks. In *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems*. ACM. https://doi.org/10.1145/3270323.3270328
- [2] Jorge Díez Peláez, David Martínez Rego, Amparo Alonso-Betanzos, Óscar Luaces Rodríguez, Antonio Bahamonde Rionda, et al. 2016. Metrical representation of readers and articles in a digital newspaper. In 10th ACM Conference on Recommender Systems (RecSys 2016). ACM.
- [3] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2011. Content-based Recommender Systems: State of the Art and Trends. Springer US, Boston, MA, 73–105. https://doi.org/10.1007/978-0-387-85820-3_3

- [4] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2000. Analysis of recommendation algorithms for e-commerce. In Proceedings of the 2nd ACM Conference on Electronic Commerce. 158–167.
- [5] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In Proceedings of the 10th international conference on World Wide Web. 285–295.
- [6] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval. 253–260.