



Measuring the User Satisfaction in a Recommendation Interface with Multiple Carousels

Nicolò Felicioni
Politecnico di Milano
Italy
nicolo.felicioni@polimi.it

Maurizio Ferrari Dacrema
Politecnico di Milano
Italy
maurizio.ferrari@polimi.it

Paolo Cremonesi
Politecnico di Milano
Italy
paolo.cremonesi@polimi.it

ABSTRACT

It is common for video-on-demand and music streaming services to adopt a user interface composed of several recommendation lists, i.e., *widgets* or *swipeable carousels*, each generated according to a specific criterion or algorithm (e.g., most recent, top popular, recommended for you, editors' choice, etc.). Selecting the appropriate combination of carousel has significant impact on user satisfaction. A crucial aspect of this user interface is that to measure the relevance a new carousel for the user it is not sufficient to account solely for its individual quality. Instead, it should be considered that other carousels will already be present in the interface. This is not considered by traditional evaluation protocols for recommenders systems, in which each carousel is evaluated in isolation, regardless of (i) which other carousels are displayed to the user and (ii) the relative position of the carousel with respect to other carousels. Hence, we propose a two-dimensional *evaluation protocol for a carousel setting* that will measure the quality of a recommendation carousel based on how much it improves upon the quality of an already available set of carousels. Our evaluation protocol takes into account also the position bias, i.e., users do not explore the carousels sequentially, but rather concentrate on the top-left corner of the screen.

We report experiments on the movie domain and notice that under a carousel setting the definition of which criteria has to be preferred to generate a list of recommended items changes with respect to what is commonly understood.

CCS CONCEPTS

• **Information systems** → **Collaborative filtering; Recommender systems**; • **General and reference** → **Evaluation**.

KEYWORDS

Recommender Systems; User Interface; Evaluation

ACM Reference Format:

Nicolò Felicioni, Maurizio Ferrari Dacrema, and Paolo Cremonesi. 2021. Measuring the User Satisfaction in a Recommendation Interface with Multiple Carousels. In *ACM International Conference on Interactive Media Experiences (IMX '21)*, June 21–23, 2021, Virtual Event, NY, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3452918.3465493>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

IMX '21, June 21–23, 2021, Virtual Event, NY, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8389-9/21/06.

<https://doi.org/10.1145/3452918.3465493>

1 INTRODUCTION

Recommender systems have become a widespread technology, which is present in a multitude of the services we use today. Their main goal is to help the user explore the vast catalogs at their disposal and are very important tools to ensure and improve the user satisfaction. The ability to provide high quality personalized recommendation is central in many business models, among which video on demand and music streaming services. Most of the research that proposes new recommendation algorithms is evaluated with offline experiments which are known to be often insufficient to accurately predict user satisfaction due to a number of factors [1].

In a traditional *offline* evaluation scenario, the systems is asked to generate a list of recommended items to a set of users, one list for each user. Each list is then compared with the historical opinions of the user and quality is measured with standard accuracy metrics (such as NDCG) developed in the fields of information retrieval and machine learning.

This approach, however, disregards an essential component of the user experience, that is the user interface. In fact, in video-on-demand services (e.g., Netflix, Amazon Prime Video) and music streaming platforms (e.g., Spotify) users are frequently provided with several recommendation lists or *carousels*, each with a certain theme e.g., recently added, originals, trending, editorially curated (see Figure 1). In this scenario, the user satisfaction depends on all the carousels that are shown, not just one, and there is significant industrial interest in finding effective strategies to select which carousels to display [2, 5, 15]. In order to represent such a scenario more closely, it is necessary to take into account how the recommendations in the various carousels complement each other and how the user explore a two-dimensional interface. There seems to be no agreed evaluation protocol to address both cases. Sometimes the recommendation lists are simply concatenated one after the other, or the average of the recommendation quality is computed (see Figure 2). Alternatively, the evaluation is done not in terms of correct recommendations, but rather in terms of which carousel contains at least a correct recommendation [15]. These approaches we argue are insufficient because they do not take into account how the two-dimensional structure of the page affects the user behavior and how the recommendation lists complement each other. Consider, for example, a scenario where there is a page containing a carousel and we need to select the second one. Clearly, we should not select as second carousel one which contains recommendations very similar to the first one (even if it has, taken individually, a high recommendation quality), rather we should balance the recommendation quality and diversity of the recommendations [16]. In particular, we should compute the recommendation quality of the

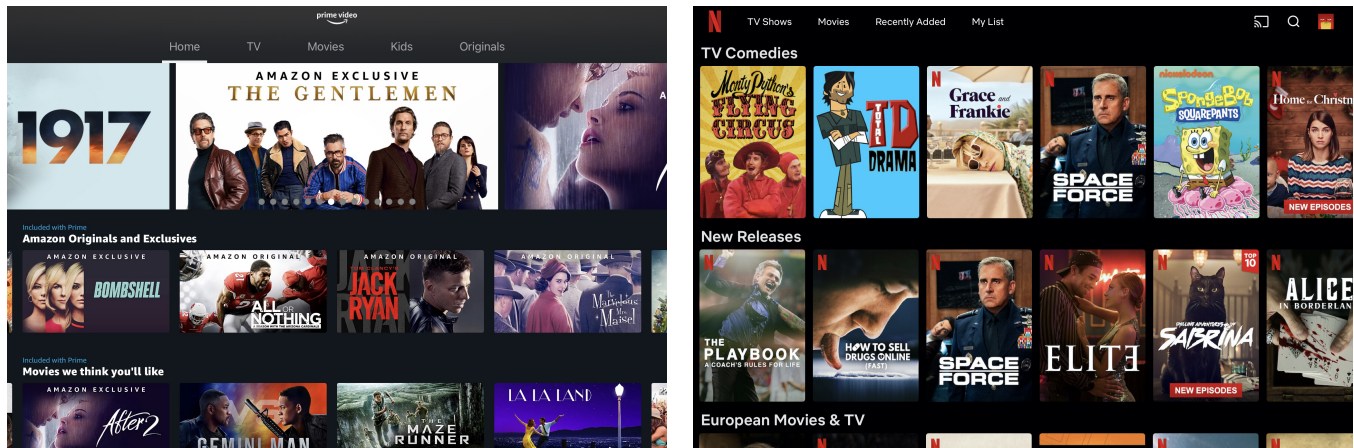


Figure 1: Examples of carousel user interfaces in the multimedia streaming domain. On the left, there is the Amazon Prime Video homepage, while on the right, there is the Netflix homepage.

second carousel by accounting only for the new correct recommendations. Furthermore, to the best of our knowledge no metric exists that measures the ranking quality of the recommendations based on the way users tend to explore a two-dimensional interface, by focusing on the top-left corner of the screen and then exploring the recommendations both to the right and to the bottom. We also note that most of the few articles discussing recommendation strategies for carousel-based user interfaces are evaluated in an online setting (e.g. with surveys or A/B testing) and the absence of a standardized offline evaluation methodology makes it difficult for researchers to explore a carousel-based user interface.

In this paper we propose a definition for this carousel setting highlighting the specific characteristics that distinguish both the user interface and the user behavior from the traditional single list evaluation. We also propose a 2-D strategy for the offline evaluation tailored for a carousel setting as well as a ranking metric which takes into account how users explore such interfaces (see Figure 3). We finally report experimental results indicating that accounting for the carousel interface and user behavior changes the relative performance of recommendation models such that they lead to a different conclusion on which should be the carousel to display.

The rest of the paper is organized as follows. In Section 2 we presents related works on carousel interfaces, in Section 3 we define and describe the carousel scenario as well as the two-dimensional ranking metric. Finally, in Section 4 we report the results of our experimental analysis and in Section 5 we draw the conclusions and possible future works.

2 RELATED WORK

Many research works tried to understand user behavior in a two-dimensional setting, usually by employing eye-tracking. For example, Kammerer and Gerjets [10] examine the differences in user behavior when presented the same 9 results of a search engine result page in two different arrangements. The page is displayed as a plain list to some of the users involved, while it is displayed as a 3x3 two-dimensional grid to others. The study indicated that users navigate the list sequentially from top to bottom, while the

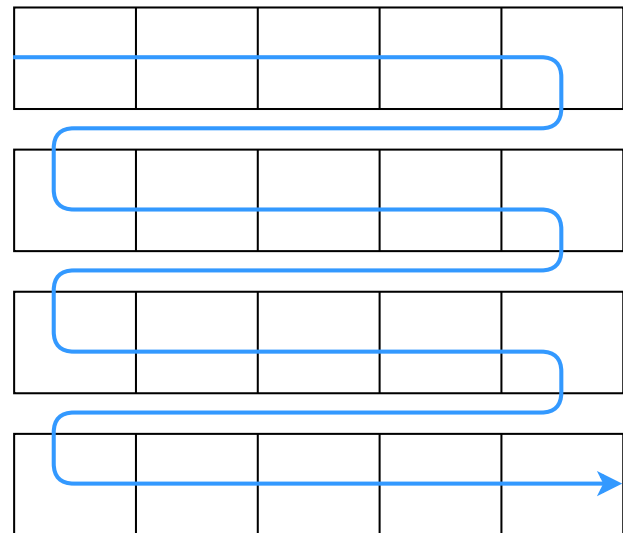


Figure 2: Traditional offline evaluation strategy in which all carousels are concatenated, assuming the user will explore them sequentially one at a time.

grid navigation traverses both rows and columns according to a two-dimensional pattern. Zhao et al. [17] analyze the behavior of users of a two-dimensional recommendation engine. They make the examples of YouTube, MovieLens, Hulu (three video streaming services), and Google Apps (a popular mobile application store). From their analyses, the authors conclude that there is an "F-pattern" of the users' gaze behavior. In other words, users are more focused on the top-left corner of a two-dimensional interface and that the attention is lower both to the bottom and to the right. This phenomenon is also called the "golden triangle" and it can be visualized by looking at Figure 5.

While there are industrial research papers focusing on the carousel scenario, indicating the importance of this topic for the

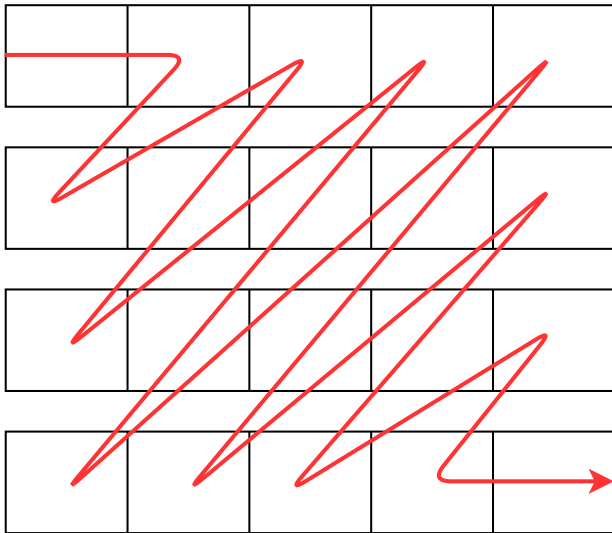


Figure 3: The proposed 2-D offline evaluation strategy which accounts for the way users explore a two-dimensional recommendation carousel by concentrating on the top-left corner and then proceeding both to the right and to the bottom across different carousels.

industry, there seems to be much less papers from the academia. This may be due to the requirement that such studies be evaluated in an online setting and to the lack of an offline evaluation protocol for the carousel setting. Such an evaluation protocol could allow researchers to more easily work on the topic even when they do not have easy access to a real user base. Wu et al. [15] try to optimize the ranking of the carousels on the user interface of Netflix, a popular online video streaming service. In this work, they focus on a carousel interface that Netflix uses, with a fixed recommendation on top of the screen and a set of thematically coherent carousels. Their goal is not to design a novel UI layout, but to include navigation feedback to understand better the user satisfaction on their services. Gruson et al. [8] optimize the homepage of Spotify, an online music streaming service. In this work, they focus on a carousel interface visualized on a mobile device, with users that can scroll left and right within carousels to reveal more items, and up and down to reveal more carousels. In their offline evaluation, the carousels are sequentially concatenated as a single long recommendation list. Evaluating in this way, though, does not reflect the actual user behavior in a two-dimensional layout. Also, the authors perform an online evaluation and find discrepancies between online and offline evaluations. Bendada et al. [2] try to optimize a carousel layout of the homepage of Deezer, a music streaming service. In particular, they design a real-time model to understand the user preferences and adjust the carousels displayed. Ding et al. [5] target the problem of optimizing the carousel homepage of Amazon Prime Video, a video streaming service. They point out that business constraints may fix carousels in certain positions of the grid, so the evaluation should take into account how the other carousel intersect with the fixed ones and how they improve the user experience.

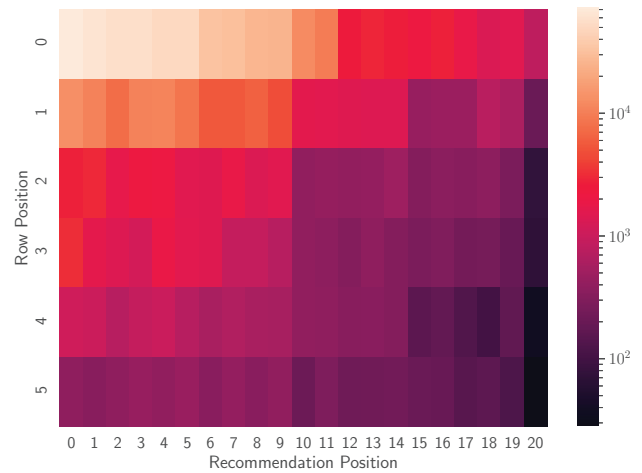


Figure 4: Heatmap of the number of interactions per position on the screen taken from the dataset presented in [11]. Values are log-scaled.

3 EVALUATION IN A CAROUSEL LAYOUT

In this section we describe common characteristics of a real industrial service that employs a carousel layout.

First of all, the *interface* is two-dimensional. The user is presented with a grid of recommendations, divided in carousels, which are the rows of this grid. Each carousel may be generated by different providers. In general the provider of a certain carousel is not aware of the recommendations contained in the others. This means that among different carousels there may be duplicate recommendations¹, while inside a single carousel this is avoided.

In terms of *user behavior*, it is known that users do not explore a grid of recommendations sequentially, one carousel at a time, but instead they concentrate on the top-left corner. This should be taken into account when calculating the quality of the provided ranking.

Furthermore, in this scenario a content provider will need to account for as much information as possible on the user behavior and interface characteristics in order to be able to select which are the best recommendation carousels to display.

3.1 NDCG in two dimensions

One of the most used metrics to measure the quality of a ranked one-dimensional list of results is the *Normalized Discounted Cumulative Gain* (NDCG) [9]. This metric is defined based on the DCG, computed as follows:

$$DCG = \sum_{i=1}^c \frac{rel(i)}{\log_2(i+1)}$$

Where c is the length of the recommendation list and $rel(i)$ a function that computes the relevance of the item in position i , e.g., the rating or just a binary 0-1 value. In this formulation, the

¹For example, in the Netflix homepage shown in Figure 1 the TV series *Space Force* appears both in the *TV Comedies* and *New Releases* carousels.

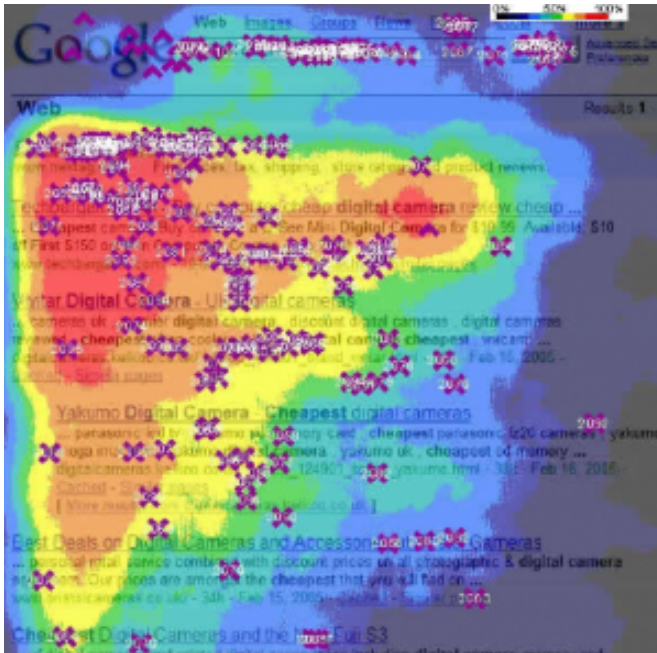


Figure 5: A known visualization of how users navigate a simple search result page, the attention is focused on a top-left golden triangle and progressively decreases moving to the right and the bottom [3].

numerator is responsible for rewarding items that are relevant for the user (the function $rel(i)$), while the denominator should penalize the items proportionally to how far they are from the beginning of the list. Clearly, it is better if a correct recommendation is positioned at the beginning of the list rather than at the end. This reasoning is sound in a layout where a list of results is presented because users will explore the list sequentially [10]. In the case of a two-dimensional layout, this metric could be simply adapted, by assuming that the grid of results will be explored as a long list, with the various rows concatenated one after the other. A visualization of the penalty imposed in this case is given in Figure 6a.

This assumption, though, is not realistic, as many research studies on user behavior suggest [10, 17]. Hence, we propose to adapt the penalization term to follow the "golden triangle" assumption, i.e., the fact that users tend to explore a two-dimensional interface focusing on the top-left corner, with decreasing attention both to the right and to the bottom. The penalty that we propose can be computed with the logarithm of the sum of the row and column of each recommendation and is visualized in Figure 6. This penalty can also be modified depending on the particular interface. For example, in a mobile application on a smartphone, we can penalize less the vertical dimension. This can be easily achieved by weighting the row and column coordinates with different weight and therefore can be easily adapted to the desired interface layout. An important aspect is that duplicate recommendations should be taken into account when computing the numerator. We propose to reward the correct recommendations only the first time the user sees them,

but not the following ones since repeating the same recommendation multiple times not enhance user satisfaction. With this new definition of the reward and the penalty, we call this metric $NDCG2D$.

4 RESULTS

We now compare the proposed carousel evaluation methodology and the two-dimensional NDCG metric with the traditional evaluation done on a single list at a time.

We report the results for *MovieLens 10M*, a dataset from the video-on-demand domain, which is among the domains that use the carousel user interface. *MovieLens 10M* contains almost 70 thousand users and 10 thousand items, for a total of 10 million ratings in the range 1-5. The dataset also contains item features such as tags, provided by users, the genre of the movie and the year of release. We randomly selected 80% of interactions for the training set and 10% for both validation and test set.

4.1 Algorithms

We select different recommendation algorithms from different families of models.

TopPopular: is a frequently used recommendation technique that recommends to all users the same set of the most popular items.

ItemKNN Hybrid: is a nearest-neighbor technique that computes the similarity between items using both the user interactions and the item metadata, using cosine with shrinkage [12].

RP^{3β}: is a graph-based technique that represents the data as a graph. Users and items are represented as nodes and a connection exists between them if the user has interacted with that item. In RP^{3β} the similarity between items is computed by simulating a random walk on the graph [13].

EASE^R: is a machine learning model which was recently proposed [14]. EASE^R is equivalent to a shallow autoencoder, since it has a closed-form solution the model can be computed very rapidly.

FunkSVD: is a known matrix factorization algorithm² created during the Netflix Prize. FunkSVD decomposes the user-item interactions as the product of two lower dimensional matrices.

NMF: Non-negative Matrix Factorization [4] is another matrix factorization algorithm that uses a known matrix-decomposition technique applied on the user-item interactions.

In order to ensure a fair comparison between the recommendation quality of all the algorithms, we thoroughly optimized all of them by following the guidelines of Ferrari Dacrema et al. [7]. In particular, we optimized the hyperparameters of all models using a Bayesian search [6] exploring 50 cases, the first 15 of them are selected at random while the remaining 35 are selected according to the Gaussian process.³ We optimized the ranking quality with a recommendation list of 10 elements.

²<https://sifter.org/~simon/journal/20061211.html>

³We relied on the implementation provided in the scikit-optimize python library https://scikit-optimize.github.io/stable/modules/generated/skopt.gp_minimize.html

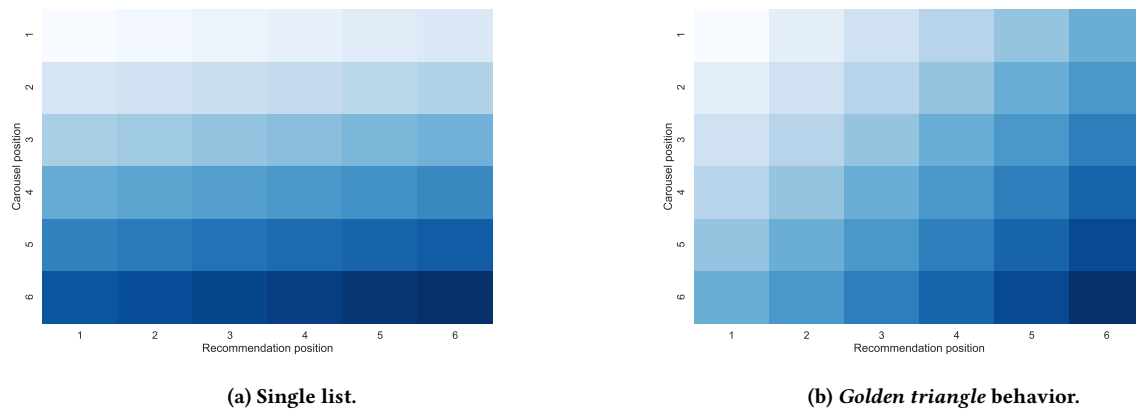


Figure 6: A visual comparison of the penalty for each recommended item under different assumptions. Figure 6a refers to carousels concatenated in a single list. Figure 6b refers to the two-dimensional penalty which accounts for the golden triangle behavior.

| | Individual NDCG | Carousel NDCG 2D | Individual NDCG rank | Carousel NDCG 2D rank | Δ rank |
|----------------|-----------------|------------------|----------------------|-----------------------|---------------|
| TopPopular | 0.0983 | – | – | – | – |
| ItemKNN hybrid | 0.2174 | 0.2148 | 3 | 4 | -1 |
| $RP^3\beta$ | 0.2160 | 0.2035 | 4 | 5 | -1 |
| $EASE^R$ | 0.2566 | 0.2293 | 1 | 2 | -1 |
| MF FunkSVD | 0.2307 | 0.2373 | 2 | 1 | +1 |
| NMF | 0.1974 | 0.2281 | 5 | 3 | +2 |

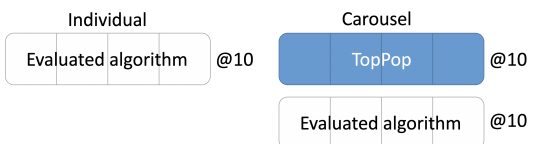


Table 1: Comparison of the MovieLens10M accuracy metrics with individual and carousel evaluation (with TopPop fixed as the first carousel) at recommendation list length of 10. Note that in the carousel evaluation there will be two recommendation lists.

4.2 Discussion

Each algorithm is evaluated in two settings. In the first setting each model is evaluated independently with the traditional approach. In the second, it is evaluated in a carousel setting where the first carousel is a TopPopular and the evaluated model is the second carousel in the layout (see Table 1). In both cases each model generates a recommendation list of 10 elements, but, in the carousel interface two lists are generated, one for each model.

For the two-dimensional NDCG we considered a case where the user explores both vertically and horizontally evenly. Depending on the device and user interface design, in other cases it may be more realistic to attribute a higher importance to the vertical or horizontal exploration.

We report the results in Table 1, the ranking quality of the models is evaluated with NDCG and NDCG2D. We can see how there are several differences between the relative performance of the recommender models. For example, following the individual NDCG, the best model is $EASE^R$, while according to the carousel NDCG2D, FunkSVD is the best. Both matrix factorization techniques tend to gain positions in the NDCG2D evaluation, while all other models lose positions. This result indicates that the algorithm that exhibited

the best recommendation quality taken independently, $EASE^R$, contains recommendations which are already present in a TopPopular and therefore would generate carousels with duplicated recommendations. Matrix factorization models, on the other hand, are less effective on their own but become more effective in a carousel layout because they better complement the recommendations provided by the TopPopular recommender.

These results indicate the importance of taking into account the user behavior when deciding which is the most appropriate model to be added in a carousel layout in order to ensure the user is provided with the most effective recommendations.

5 CONCLUSIONS AND FUTURE WORK

In this paper we proposed a novel approach for the evaluation of the recommendation quality when a carousel user interface is used in which the recommendation quality of a carousel depends on the entire page.

We also extended the NDCG metric to account for the user navigation behavior in a two-dimensional page.

The experimental results indicate that taking into account both the page layout and the user behavior changes the outcome of

an offline evaluation leading to prefer a different carousel. This indicates the importance of taking both the page layout and the user behavior into account during the offline evaluation phase. Important future work is an online study to assess the extent to which this evaluation protocol is able to more closely mirror the user satisfaction, using surveys or A/B testing. In the end, this evaluation protocol could constitute a viable option for the offline evaluation of algorithmic strategies to select the carousel layout of a page in order to provide the user with more effective carousels and improve user satisfaction.

REFERENCES

- [1] Joeran Beel, Marcel Genzmehr, Stefan Langer, Andreas Nürnberger, and Bela Gipp. 2013. 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* (Hong Kong, China) (RepSys '13). Association for Computing Machinery, New York, NY, USA, 7–14. <https://doi.org/10.1145/2532508.2532511>
- [2] Walid Bendada, Guillaume Salha, and Théo Bontempelli. 2020. Carousel Personalization in Music Streaming Apps with Contextual Bandits. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22–26, 2020*, Rodrygo L. T. Santos, Leandro Balby Marinho, Elizabeth M. Daly, Li Chen, Kim Falk, Noam Koenigstein, and Edleno Silva de Moura (Eds.). ACM, 420–425. <https://doi.org/10.1145/3383313.3412217>
- [3] Flavio Chierichetti, Ravi Kumar, and Prabhakar Raghavan. 2011. Optimizing two-dimensional search results presentation. In *Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM 2011, Hong Kong, China, February 9–12, 2011*, Irwin King, Wolfgang Nejdl, and Hang Li (Eds.). ACM, 257–266. <https://doi.org/10.1145/1935826.1935873>
- [4] Andrzej Cichocki and Anh Huy Phan. 2009. Fast Local Algorithms for Large Scale Nonnegative Matrix and Tensor Factorizations. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* 92-A, 3 (2009), 708–721. <https://doi.org/10.1587/transfun.E92.A.708>
- [5] Weicong Ding, Dinesh Govindaraj, and S. V. N. Vishwanathan. 2019. Whole Page Optimization with Global Constraints. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4–8, 2019*, Ankur Teredesai, Vipin Kumar, Ying Li, Römer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 3153–3161. <https://doi.org/10.1145/3292500.3330675>
- [6] Nicolò Felicioni, Andrea Donati, Luca Conterio, Luca Bartoccioni, Davide Yi Xian Hu, Cesare Bernardis, and Maurizio Ferrari Dacrema. 2020. Multi-Objective Blended Ensemble For Highly Imbalanced Sequence Aware Tweet Engagement Prediction. In *RecSys Challenge '20: Proceedings of the Recommender Systems Challenge 2020, Virtual Event Brazil, September, 2020*. ACM, 29–33. <https://dl.acm.org/doi/10.1145/3415959.3415998>
- [7] Maurizio Ferrari Dacrema, Simone Boglio, Paolo Cremonesi, and Dietmar Jannach. 2021. A Troubling Analysis of Reproducibility and Progress in Recommender Systems Research. *ACM Trans. Inf. Syst.* 39, 2, Article 20 (Jan. 2021), 49 pages. <https://doi.org/10.1145/3434185>
- [8] Alois Gruson, Praveen Chandar, Christophe Charbuillet, James McInerney, Samantha Hansen, Damien Tardieu, and Ben Carterette. 2019. Offline Evaluation to Make Decisions About Playlist Recommendation Algorithms. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11–15, 2019*, J. Shane Culpepper, Alistair Moffat, Paul N. Bennett, and Kristina Lerman (Eds.). ACM, 420–428. <https://doi.org/10.1145/3289600.3291027>
- [9] Kalervo Järvelin and Jaana Kekäläinen. 2000. IR evaluation methods for retrieving highly relevant documents. In *SIGIR 2000: Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, July 24–28, 2000, Athens, Greece*, Emmanuel J. Yannakoudakis, Nicholas J. Belkin, Peter Ingwersen, and Mun-Kew Leong (Eds.). ACM, 41–48. <https://doi.org/10.1145/345508.345545>
- [10] Yvonne Kammerer and Peter Gerjets. 2010. How the interface design influences users' spontaneous trustworthiness evaluations of web search results: comparing a list and a grid interface. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications, ETRA 2010, Austin, Texas, USA, March 22–24, 2010*, Carlos Hitoshi Morimoto, Howell O. Istance, Aulikki Hyrskykari, and Qiang Ji (Eds.). ACM, 299–306. <https://doi.org/10.1145/1743666.1743736>
- [11] Fernando Benjamin Pérez Maurera, Maurizio Ferrari Dacrema, Lorenzo Saule, Mario Scriminaci, and Paolo Cremonesi. 2020. ContentWise Impressions: An Industrial Dataset with Impressions Included. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19–23, 2020*, Mathieu d'Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux (Eds.). ACM, 3093–3100. <https://doi.org/10.1145/3340531.3412774>
- [12] Bamshad Mobasher, Xin Jin, and Yanzan Zhou. 2003. Semantically Enhanced Collaborative Filtering on the Web. In *Web Mining: From Web to Semantic Web, First European Web Mining Forum, EMWF 2003, Cavtat-Dubrovnik, Croatia, September 22, 2003, Revised Selected and Invited Papers (Lecture Notes in Computer Science)*, Bettina Berendt, Andreas Hotho, Dunja Mladenic, Maarten van Someren, Myra Spiliopoulou, and Gerd Stumme (Eds.), Vol. 3209. Springer, 57–76. https://doi.org/10.1007/978-3-540-30123-3_4
- [13] Bibek Paudel, Fabian Christoffel, Chris Newell, and Abraham Bernstein. 2017. Updatable, Accurate, Diverse, and Scalable Recommendations for Interactive Applications. *ACM Trans. Interact. Intell. Syst.* 7, 1 (2017), 1:1–1:34. <https://doi.org/10.1145/2955101>
- [14] Harald Steck. 2019. Embarrassingly Shallow Autoencoders for Sparse Data. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13–17, 2019*, Ling Liu, Ryan W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 3251–3257. <https://doi.org/10.1145/3308558.3313710>
- [15] Chao-Yuan Wu, Christopher V. Alvino, Alexander J. Smola, and Justin Basilico. 2016. Using Navigation to Improve Recommendations in Real-Time. In *Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15–19, 2016*, Shilad Sen, Werner Geyer, Jill Freyne, and Pablo Castells (Eds.). ACM, 341–348. <https://doi.org/10.1145/2959100.2959174>
- [16] Mi Zhang and Neil Hurley. 2008. Avoiding monotony: improving the diversity of recommendation lists. In *Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys 2008, Lausanne, Switzerland, October 23–25, 2008*, Pearl Pu, Derek G. Bridge, Bamshad Mobasher, and Francesco Ricci (Eds.). ACM, 123–130. <https://doi.org/10.1145/1454008.1454030>
- [17] Qian Zhao, Shuo Chang, F. Maxwell Harper, and Joseph A. Konstan. 2016. Gaze Prediction for Recommender Systems. In *Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15–19, 2016*, Shilad Sen, Werner Geyer, Jill Freyne, and Pablo Castells (Eds.). ACM, 131–138. <https://doi.org/10.1145/2959100.2959150>