# Let's Get It Started (Again!): Fostering the Discoverability of New Releases on Deezer

Léa Briand<sup>1</sup>, Théo Bontempelli<sup>1</sup>, Walid Bendada<sup>1,2</sup>, Mathieu Morlon<sup>1</sup>, François Rigaud<sup>1</sup>, Benjamin Chapus<sup>1</sup>, Thomas Bouabça<sup>1</sup> and Guillaume Salha-Galvan<sup>1</sup>

<sup>1</sup>Deezer Research, Paris, France <sup>2</sup>LAMSADE, Université Paris Dauphine, PSL, Paris, France

#### Abstract

This workshop paper presents an extended version of a work previously been accepted for oral presentation as an Industry Talk at the 46th European Conference on Information Retrieval (ECIR 2024), under the title *"Let's Get It Started: Fostering the Discoverability of New Releases on Deezer"* [1]. This conference version received the ECIR 2024 Best Industry Track Paper award. We present our recent initiatives to foster the discoverability of new releases on the music streaming service Deezer. After introducing our search and recommendation features dedicated to new releases, we outline our shift from editorial to personalized release suggestions using cold start embeddings and contextual bandits. Backed by online experiments, we discuss the advantages of this shift in terms of recommendation quality and exposure of new releases on the service.

#### Keywords

New Releases, Music Recommendation, Music Discovery, Music Streaming Service, A/B Testing.

## 1. Introduction

Music artists release hundreds of thousands of new albums weekly on the music streaming service Deezer [2]. The swift integration of this content, along with its discoverability through recommender systems and search engines, holds significant importance. For Deezer, it ensures that users have immediate access to the latest music of their favorite artists while also easily coming upon new ones they might like, which improves the user experience [3, 4]. The proper exposure of new releases also benefits artists by amplifying their visibility, which can contribute to their success, boost their revenues, and foster the emergence of new talents [5, 6]. Nonetheless, displaying the right releases to the right users remains challenging due to the limited prior information on this fresh content, especially for new artists unknown from the service [4, 7].

In this paper, we present our recent efforts to better showcase new releases on Deezer, both in terms of recommendation performance and number of new releases exposed to users. Specifically, we first focus on the product context. We describe in Section 2 our search and recommendation features dedicated to new releases, along with their objectives and differences. We also dive into our historical semi-personalized solution for new release recommendation, based on editorial pre-selections, and its limitations in terms of catalog coverage and adaptability. Then, in Section 3, we detail the fully personalized systems we deployed in 2023 to overcome

MuRS'24: Music Recommender Systems Workshop 2024, October 14−18, 2024, Bari, Italy research@deezer.com (L. Briand)

<sup>© 2024</sup> Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



Figure 1: Interface of the "New releases for you" carousel on the website version of Deezer, recommending a personalized shortlist of recently released albums to each user.

these limitations, involving a cold start neural embedding model along with contextual bandits [8, 9]. We provide an online evaluation of our approach on Deezer's "New releases for you" carousel of recommended new albums. Finally, we conclude in Section 4 by discussing our ongoing efforts to improve our models and how we present new releases on the Deezer homepage.

# 2. Listening to New Releases on Deezer

Hundreds of thousands of albums and tracks are released on Deezer weekly. Users can instantly access this content using our *search engine*, based on an Elasticsearch index plugged into a Kafka topic [10, 11] refreshed in real-time. They also receive *notifications* informing them about updates from their favorite artists, which we refer to as *unmissable* releases. In addition, on the Deezer homepage, we propose "*Friday cards*", i.e., weekly playlists mixing unmissable tracks.

All these features help users retrieve releases from artists they already know. However, they hardly permit *discovering* new music from emergent or unfamiliar artists [7]. For this distinct purpose, the Deezer homepage relies on two album carousels<sup>1</sup> [12]. The first one, known as "*Fresh picks of the week*", presents the biggest releases according to professional editors from Deezer. The second one, called "*New releases for you*" and illustrated in Figure 1, aims to provide personalized recommendations. Up to early 2023, it first showcased unmissable albums for each user, followed by new albums from their favorite music genres. As shown in Figure ??, our data scientists were identifying each user's favorite genres, but the lists of recommended albums for each genre were manually pre-selected by Deezer editors.

This historical method for discovering new releases had limitations that motivated the work presented in this paper. To begin with, "New releases for you" carousels were only partially personalized. All fans from a genre used to receive the same (non-unmissable) recommendations, regardless of their individual subgenre preferences. Moreover, editorial lists were only updated once a week, on Fridays. Carousels remained static for seven days, making the feature less

<sup>&</sup>lt;sup>1</sup>Our carousels display swipeable ranked lists of 12 albums, followed by a "View all" button leading to a page displaying up to 100 recommended albums.



**Figure 2:** How "New releases for you" used to work until 2023. Each user's favorite music genres were computed by our data scientists, but lists of new albums per genre were pre-selected by our editors.

engaging. New releases added during the week went unrecommended for days, and user interactions were not factored in for improved recommendations. Last but not least, only the small minority of albums handpicked by editors were recommended. Most new releases never got the chance to even be exposed to users.

# 3. Improving the Discoverability of New Releases

This section outlines our recent initiatives to overcome these limitations. For the sake of clarity, we focus our presentation on the discovery of *new albums*.

#### 3.1. Cold Start Embedding Representations of New Releases

Most of our personalized recommender systems [8, 12, 13, 14, 15] leverage latent models for collaborative filtering (CF) [16, 17]. By analyzing usage data on Deezer, they learn lowdimensional *embedding* vector representations of users and musical items including albums, in a vector space where proximity reflects preferences. Then, they offer recommendations based on embedding similarity metrics [8, 14]. Unfortunately, such CF models often struggle to incorporate new items lacking usage data [4, 18, 19]. This issue, known as the *cold start* problem, explains why Deezer historically favored an editorial approach for suggesting new releases.

In 2023, we developed CF-Cold-Start. For any existing CF embedding space, CF-Cold-Start predicts embedding vectors of items such as new albums *as early as their release date*, without requiring extensive usage data. CF-Cold-Start draws inspiration from advances in cold start research [19, 20, 21], including a Deezer system for *user* cold start [8]. It consists in a 3-layer neural network [22], predicting "ground truth" embedding vectors that would be computed by the CF model using a week of usage data. In the case of a new album, the input layer processes various album metadata, including information on labels and artists, as well as usage data (streams, likes...) at the release date when available. The output layer returns a *predicted embedding vector* for this album. We train CF-Cold-Start in a supervised manner using musical



**Figure 3:** Overview of the CF-Cold-Start new album recommendation framework from Section 3.1, available for online requests in our production environment on Deezer.

					- 11	
		Top-K new albums recommended by	ums	aı	album	
			d by	a2	<u>album</u>	
	CF-Cold-Start	rt →	a3	<u>album</u>		
				a4	<u>album</u>	
"New releas	ses for you"	carousel				
Unmissable 1	Unmissable 2	Unmissable 3	al		a2	a3

Figure 4: CF-Cold-Start replaces editorial selections from Figure 2 with personalized recommendations.

a4

items for which both input data and ground truth embedding vectors are available, by minimizing the mean squared error between predicted and ground truth vectors using gradient descent [22].

We deployed CF-Cold-Start on Deezer in our "New releases for you" feature. While carousels still begin with unmissable new albums, CF-Cold-Start replaces editorial pre-selections with personalized recommendations, consisting in ranked lists of albums whose predicted embedding vectors are the most similar to each user<sup>2</sup> in the CF embedding space. We update embedding vector predictions every four hours through a forward pass in CF-Cold-Start, to refine recommendations as we collect more usage data on new albums. We use approximate nearest neighbors methods via a Golang application incorporating the Faiss library [23]. Figure 3 and Figure 4 provide an illustrated summary of our CF-Cold-Start new album recommendation pipeline in production.

<sup>&</sup>lt;sup>2</sup>Precisely, this feature uses dot product similarities and the SVD embeddings from Bendada et al. [13].



**Figure 5:** Overview of the TS-CF-Cold-Start new album recommendation framework from Section 3.2, available for online requests in our production environment on Deezer.

### 3.2. Carousel Personalization with Contextual Bandits

CF-Cold-Start offers fully personalized recommendations from the entire set of new albums, with refreshes every four hours instead of once a week. Nonetheless, it solely *exploits* the most tailored albums based on usage data. In this work, we also wanted to occasionally deviate from them, to *explore* less popular but promising albums which could be overlooked by relying on the existing usage.

To this end, we developed TS-CF-Cold-Start, a CF-Cold-Start variant with a *multi-arm bandit* component for adaptive "New releases for you" album ranking. Bandit algorithms effectively handle exploration-exploitation face-offs [9, 24, 25]. Here, we use the *Thompson Sampling* (TS) [26] extension for contextual bandits [9]. Our *arms* are new albums, recommended by ranked batches at each bandit *round*, i.e., each carousel display on Deezer. The set of arms evolves over time, as new albums enter daily and exit the new release set after seven days.

Album embedding vectors predicted by CF-Cold-Start act as the *prior expected arm representations* [9] used for arm selection by the bandit. They are updated every four hours using TS with Gaussian distributions. Each album's mean and variance are exported in an ONNX model [27] in our service for inference and sampling.

Updates are based on user click *rewards* and a *cascade-based* scheme [12] accounting for the album position in the carousel. We note that, while carousels still start with fixed unmissable releases, the stochastic nature of TS-based exploration-exploitation permits recommending different non-unmissable albums with each homepage refresh, even between two arm updates. Figure 5 summarizes our TS-CF-Cold-Start pipeline in production.



#### Editorial playlists carousel

TS Reco 1	TS Reco 2	TS Reco 3	TS Reco 4	TS Reco 5



"New releases for you" carousel							
	Unmissable 1	Unmissable 2	Unmissable 3	TS Reco 1	TS Reco 2		

**Figure 6:** Unlike the editorial playlist carousels discussed in Bendada et al. [12] (*top*), "New releases for you" carousels (*bottom*) includes fixed unmissable albums, limiting Thompson Sampling (TS) bandit actions to less visible carousel slots.

### 3.3. Online Experimental Evaluation on Deezer

We conducted large-scale A/B tests on Deezer between March and May 2023. On average, replacing editorial selections with CF-Cold-Start recommendations in "New releases for you" not only improved display-to-click rates by 6%, but also multiplied by ×3 and ×1.5 the weekly numbers of new albums displayed and clicked in carousels, respectively. Hence, CF-Cold-Start boosted their exposure.

Using TS-CF-Cold-Start added dynamism to carousels by sampling albums with each app refresh. However, its performance was on par with CF-Cold-Start in both clicks and exposure. This mixed result contrasts with a previous study [12] where bandits improved our *editorial playlist* carousels. We attribute our lower performance to two factors. Firstly, as illustrated in Figure 6, unmissable albums restrict bandit actions to less visible carousel slots, unlike in playlist carousels [12]. Secondly, our set of new albums constantly evolves, potentially hindering model convergence.

Note: our experiments will be further illustrated and discussed in the presentation associated with this article. Resources related to this talk are available on: https://github.com/deezer/new-releases-ecir2024.

# 4. Conclusion and Future Work

In conclusion, both CF-Cold-Start and TS-CF-Cold-Start enhanced our "New releases for you" recommendations and the exposure of new albums, all of this with more dynamic updates. These initiatives foster the overall discoverability of new releases on Deezer. They also open up interesting avenues for future work. We believe our problem is conducive to the study of bandits explicitly accounting for *equity of exposure*, for improved fairness in release suggestions [28]. From a product perspective, our mixed results with TS-CF-Cold-Start also prompt us to separate unmissable albums from discoveries into two carousels, with the latter fully controlled by bandits. Lastly, we will use CS-Cold-Start to better incorporate releases in other Deezer features, including our "Flow" personalized radio [14].

# References

- L. Briand, T. Bontempelli, W. Bendada, M. Morlon, F. Rigaud, B. Chapus, T. Bouabça, G. Salha-Galvan, Let's get it started: Fostering the discoverability of new releases on deezer, in: Proceedings of the 46th European Conference on Information Retrieval, Springer, 2024, pp. 286–291.
- [2] Deezer, https://www.deezer.com, 2023.
- [3] Y. Li, K. Liu, R. Satapathy, S. Wang, E. Cambria, Recent Developments in Recommender Systems: A Survey, arXiv Preprint arXiv:2306.12680 (2023).
- [4] M. Schedl, H. Zamani, C.-W. Chen, Y. Deldjoo, M. Elahi, Current Challenges and Visions in Music Recommender Systems Research, International Journal of Multimedia Information Retrieval 7 (2018) 95–116.
- [5] L. Aguiar, J. Waldfogel, Platforms, Promotion, and Product Discovery: Evidence from Spotify Playlists, Technical Report, National Bureau of Economic Research (2018).
- [6] A. Ferraro, X. Serra, C. Bauer, What is Fair? Exploring the Artists' Perspective on the Fairness of Music Streaming Platforms, in: IFIP Conference on Human-Computer Interaction, Springer, 2021, pp. 562–584.
- [7] O. Celma, O. Celma, The Long Tail in Recommender Systems, Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space (2010) 87–107.
- [8] L. Briand, G. Salha-Galvan, W. Bendada, M. Morlon, V.-A. Tran, A Semi-Personalized System for User Cold Start Recommendation on Music Streaming Apps, in: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2021, pp. 2601–2609.
- [9] O. Chapelle, L. Li, An Empirical Evaluation of Thompson Sampling, Advances in Neural Information Processing Systems 24 (2011) 2249–2257.
- [10] Elasticsearch, https://www.elastic.co/elasticsearch, 2023.
- [11] Kafka, https://kafka.apache.org/, 2023.
- [12] W. Bendada, G. Salha, T. Bontempelli, Carousel Personalization in Music Streaming Apps with Contextual Bandits, in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020, pp. 420–425.
- [13] W. Bendada, T. Bontempelli, M. Morlon, B. Chapus, T. Cador, T. Bouabça, G. Salha-Galvan, Track Mix Generation on Music Streaming Services using Transformers, in: Proceedings of the 17th ACM Conference on Recommender Systems, 2023, pp. 112–115.
- [14] T. Bontempelli, B. Chapus, F. Rigaud, M. Morlon, M. Lorant, G. Salha-Galvan, Flow Moods: Recommending Music by Moods on Deezer, in: Proceedings of the 16th ACM Conference on Recommender Systems, 2022, pp. 452–455.
- [15] G. Salha-Galvan, R. Hennequin, B. Chapus, V.-A. Tran, M. Vazirgiannis, Cold Start Similar Artists Ranking with Gravity-Inspired Graph Autoencoders, in: Proceedings of the 15th ACM Conference on Recommender Systems, 2021, pp. 443–452.
- [16] D. Bokde, S. Girase, D. Mukhopadhyay, Matrix Factorization Model in Collaborative Filtering Algorithms: A Survey, Procedia Computer Science 49 (2015) 136–146.
- [17] Y. Koren, R. Bell, Advances in Collaborative Filtering, Recommender Systems Handbook (2015) 77–118.

- [18] R. Mu, A Survey of Recommender Systems based on Deep Learning, IEEE Access 6 (2018) 69009–69022.
- [19] J. Wang, P. Huang, H. Zhao, Z. Zhang, B. Zhao, D. L. Lee, Billion-Scale Commodity Embedding for E-Commerce Recommendation in Alibaba, in: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2018, pp. 839–848.
- [20] P. Covington, J. Adams, E. Sargin, Deep Neural Networks for Youtube Recommendations, in: Proceedings of the 10th ACM Conference on Recommender Systems, 2016, pp. 191–198.
- [21] H. Lee, J. Im, S. Jang, H. Cho, S. Chung, MeLU: Meta-Learned User Preference Estimator for Cold-Start Recommendation, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2019, pp. 1073–1082.
- [22] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT Press, 2016.
- [23] J. Johnson, M. Douze, H. Jégou, Billion-Scale Similarity Search with GPUs, IEEE Transactions on Big Data 7 (2019) 535–547.
- [24] L. Li, W. Chu, J. Langford, R. E. Schapire, A Contextual-Bandit Approach to Personalized News Article Recommendation, in: Proceedings of the 19th International Conference on World Wide Web, 2010, pp. 661–670.
- [25] J. McInerney, B. Lacker, S. Hansen, K. Higley, H. Bouchard, A. Gruson, R. Mehrotra, Explore, Exploit, and Explain: Personalizing Explainable Recommendations with Bandits, in: Proceedings of the 12th ACM Conference on Recommender Systems, 2018, pp. 31–39.
- [26] W. R. Thompson, On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples, Biometrika 25 (1933) 285–294.
- [27] Open Neural Network Exchange (ONNX), https://github.com/onnx/onnx, 2023.
- [28] O. Jeunen, B. Goethals, Top-K Contextual Bandits with Equity of Exposure, in: Proceedings of the 15th ACM Conference on Recommender Systems, 2021, pp. 310–320.