



Neural content-aware collaborative filtering for cold-start music recommendation

Paul Magron, Cédric Févotte

► To cite this version:

Paul Magron, Cédric Févotte. Neural content-aware collaborative filtering for cold-start music recommendation. Data Mining and Knowledge Discovery, 2022, 10.1007/s10618-022-00859-8 . hal-03152158v2

HAL Id: hal-03152158

<https://inria.hal.science/hal-03152158v2>

Submitted on 20 Jul 2022

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.

Neural content-aware collaborative filtering for cold-start music recommendation*

Paul Magron[†], Cédric Févotte[‡]

Abstract

State-of-the-art music recommender systems are based on collaborative filtering, which builds upon learning similarities between users and songs from the available listening data. These approaches inherently face the cold-start problem, as they cannot recommend novel songs with no listening history. Content-aware recommendation addresses this issue by incorporating content information about the songs on top of collaborative filtering. However, methods falling in this category rely on a shallow user/item interaction that originates from a matrix factorization framework. In this work, we introduce neural content-aware collaborative filtering, a unified framework which alleviates these limits, and extends the recently introduced neural collaborative filtering to its content-aware counterpart. We propose a generative model which leverages deep learning for both extracting content information from low-level acoustic features and for modeling the interaction between users and songs embeddings. The deep content feature extractor can either directly predict the item embedding, or serve as a regularization prior, yielding two variants (*strict* and *relaxed*) of our model. Experimental results show that the proposed method reaches state-of-the-art results for a cold-start music recommendation task. We notably observe that exploiting deep neural networks for learning refined user/item interactions outperforms approaches using a more simple interaction model in a content-aware framework.

Keywords — Content-aware recommendation, neural collaborative filtering, matrix factorization, cold-start problem.

1 Introduction

Music recommendation consists in predicting users’ listening habits in order to suggest them novel tracks that they might enjoy (Schedl et al., 2015). This task, which is at the core of many commercial platforms, has been extensively investigated, but remains challenging due to the complexity of music and to the lack of explicit and reliable users feedback (Hu et al., 2008). In particular, leveraging the musical *function* at hand is necessary to perform recommendation that are tailored for a specific usage (Schedl et al., 2017). More generally, the usage of contextual information (Pichl and Zangerle, 2021) is of paramount importance in order to adapt the recommendation to a particular location or event (Gillhofer and Schedl, 2015; Cheng and Shen, 2016). Besides, it has been shown that psychological cues such as personality and emotional response play a major role in musical tastes and users behavior (Ferwerda et al., 2015). Recommender systems could then benefit from such findings in music psychology research (Laplante, 2014; Soleymani et al., 2015; Magron and Févotte, 2021). Finally, music recommender systems face the *cold-start* problem (Schein et al., 2002), which is the topic of investigation of this paper: when a new song is added to a music streaming platform, it has no interaction data with the set of users. Consequently, the system cannot properly recommend this novel item to existing users. Even though this problem has fueled an important amount of studies, it is still considered as a major challenge in music recommendation research (Schedl et al., 2018).

State-of-the-art approaches for music recommendation are based on collaborative filtering (Hu et al., 2008), a family of techniques which rely solely on users’ listening history: the interest of a given user for a given song is predicted using similarities between various user profiles. The users’ feedback are most often implicit and in the form of *playcounts*, that is, how many times a given user has listened to a particular song. Collaborative filtering techniques consist in learning user and item embeddings from the data, where these embeddings respectively characterize the users’ preferences and the items’ attributes. This listening history is however noisy and lacks negative

*The work of P. Magron was conducted while he was with IRT, Université de Toulouse, CNRS, Toulouse, France. This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 681839 (project FACTORY). Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

[†]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France (e-mail: firstname.lastname@inria.fr).

[‡]IRIT, Université de Toulouse, CNRS, Toulouse, France (e-mail: firstname.lastname@irit.fr).

feedback data (Hu et al., 2008; Rendle et al., 2009). To alleviate this problem, weighted matrix factorization (WMF) techniques (Salakhutdinov and Mnih, 2007; Koren et al., 2009) process binarized data computed from the raw playcounts, and associate a measure of *confidence* to these binarized feedback. This family of techniques has shown good performance in music recommender systems (Liang et al., 2015). More recently, deep learning has been leveraged in collaborative filtering techniques with promising results (Cheng et al., 2016; Zhang et al., 2019; Liang et al., 2018). While some earlier models rely on shallow architectures (Salakhutdinov et al., 2007; Wu et al., 2016), deep neural networks (DNNs) are now exploited for learning deep user/item embeddings and interaction models to replace the matrix product in WMF (Xue et al., 2017; He et al., 2017, 2018; Chen et al., 2019).

However, these approaches are agnostic to any form of item-related content, which becomes a major issue for new items. Indeed, since collaborative filtering methods only exploit the available user/item interaction data, they are not able to recommend a song without listening history, and therefore face the cold-start problem. This problem has been tackled with content-based methods, which aim to exploit additional information about the items for recommendation (Yoshii et al., 2006; Wang and Blei, 2011; Fang and Si, 2011). Deep learning has been used as a tool to learn features from the content that can help the collaborative filtering. In (Liang et al., 2015), the last hidden layer of an auto-tagging DNN is used as content feature. Conversely, in (Oord et al., 2013; Wang and Wang, 2014), acoustic features are mapped to the learned item attribute matrix in order to be further used for cold-start recommendation. These are however limited in performance since the user/item embeddings and the deep content feature extractor are learned in two distinct stages. To alleviate this issue, several recent works have proposed to combine these steps into a single-stage approach, where the user/item embeddings and the content feature extractor are jointly learned. Examples of architectures for the deep content feature extractor include stacked denoising auto-encoders (Wang et al., 2015; Li et al., 2015), variational auto-encoders (Li and She, 2017), recurrent neural networks (Wang et al., 2016) and convolutional networks (Huan et al., 2017; Lee et al., 2018). However, these approaches still rely on a simple matrix product for modeling the interaction between users and items, and therefore do not leverage DNNs for learning more complex interaction models.

In this work, we introduce neural content-aware collaborative filtering (NCACF), a unified framework which overcomes these limits, and extends the recently introduced neural collaborative filtering (NCF) (He et al., 2017) to its content-aware counterpart. We propose a generative model which leverages DNNs for both extracting content information from low-level acoustic features and for learning refined interaction models between user and item embeddings. Two variants of the model are considered. In the *strict* variant, the deep content feature extractor directly predicts the item embedding. Conversely, in the *relaxed* variant, it serves as a regularization prior for the item embedding. These two approaches allow for more flexibility and for generalizing previous models from the literature. In particular, when the interaction model reduces to a simple product (which is a building block of matrix factorization models), we derive an estimation algorithm which hybridizes closed-form updates using alternating least squares (ALS) for the embeddings and gradient descent (GD) for the deep content feature extractor. We further incorporate the embeddings in the network in order to estimate the whole model jointly using a single GD algorithm. Finally, we use this technique to estimate the more general NCACF model which uses a deep user/item interaction model. To assess the potential of these methods, we conduct experiments on the Million Song Dataset, a publicly available database for music information retrieval tasks. We observe that the proposed NCACF outperforms the baselines using a shallow user/item interaction model, but also recent content-free approaches using a deep interaction model. This reveals its interest for cold-start music recommendation applications, but also for traditional collaborative filtering tasks, where leveraging side-content information is shown beneficial. To summarize, the main contributions of our work are as follows:

1. We introduce NCACF, a deep learning-based recommendation system. This model allows for learning refined user/item interactions as well as leveraging content information, which makes it suitable for cold-start recommendation.
2. By considering several variants of this model’s architecture, underlying generative processes, and training strategies, we show that many methods from the recent literature (both content-free and content-aware) can be cast as particular instances of the NCACF unified framework.
3. We conduct extensive experiments on the largest publicly available music recommendation dataset, where NCACF reaches state-of-the-art performance. Carefully evaluating its many particular cases notably allows to evaluate the potential of several recent methods for cold-start recommendation, which is complementary to the warm-start setting they were initially designed for.

The rest of this paper is structured as follows. Section 2 presents the work related to our approach. The proposed NCACF model is then introduced in Section 3. The experimental protocol is detailed in Section 4 and the results are presented and discussed in Section 5. Finally, Section 6 draws some concluding remarks.

2 Related work

In this section, we present the related work upon which our method builds. We first describe the baseline WMF model, and then present approaches that incorporate deep content feature extraction in WMF. Finally, we present

Table 1: Mathematical notations.

U	Number of users.
I	Number of items (= songs).
K	Dimension of the user/item embeddings.
L	Dimension of the content vectors (= acoustic features).
P	Number of layers in the deep content feature extractor.
Q	Number of hidden layers in the deep interaction model.
$\mathbf{Y}$	User/item interaction matrix (raw playcounts).
$\mathbf{R}$	Binarized playcount matrix.
$\mathbf{r}_u, \mathbf{r}_i$	Binarized playcounts for user u and item i .
$\mathbf{C}$	Confidence matrix.
$\mathbf{c}_u, \mathbf{c}_i$	Confidence for user u and item i .
$\mathbf{x}_i$	Acoustic feature vector for item i .
$\mathbf{W}, \mathbf{H}$	User and item embedding matrices.
$\mathbf{w}_u, \mathbf{h}_i$	Embedding vectors for user u and item i .
ϕ_θ	Deep content feature extractor, with parameters θ .
ψ_γ	User/item interaction model, with parameters γ .
-1	Matrix inverse.
$\odot$	Element-wise vector or matrix multiplication.
$\cdot^\top$	Vector or matrix transpose.
$\mathbf{I}_K$	Identity matrix of size K .
$\mathbf{1}$	Vector or matrix whose entries are all equal to 1.
$\text{diag}(\mathbf{a})$	Diagonal matrix whose entries are given by the vector $\mathbf{a}$.
$\ \cdot\ $	Euclidean norm.

the collaborative filtering methods using deep learning for user/item interaction modeling. The notations used throughout this paper are summarized in Table 1.

2.1 Weighted matrix factorization

Let us consider a large data matrix $\mathbf{Y} \in \mathbb{R}^{U \times I}$ representing interactions between a set of U users and I items, where $y_{u,i}$ denotes the interaction between user u and item i . Matrix factorization (Marlin and Zemel, 2004; Hu et al., 2008) consists in decomposing the data $\mathbf{Y}$ as the product of two low-dimensional factors: a (transposed) user preferences matrix $\mathbf{W} \in \mathbb{R}^{K \times U}$ and an item attributes matrix $\mathbf{H} \in \mathbb{R}^{K \times I}$, such that $\mathbf{Y} \approx \mathbf{W}^\top \mathbf{H}$. The rank K of the decomposition is chosen such that $K(U+I) \ll UI$ to ensure dimensionality reduction. More specifically, weighted matrix factorization (WMF) (Salakhutdinov and Mnih, 2008; Koren et al., 2009) is a variant of matrix factorization that is appropriate for handling implicit feedback data. In this variant, the factorization is performed on a binarized interaction data matrix $\mathbf{R}$ (that is, $r_{u,i} = 1$ if user u has interacted with item i and 0 otherwise).

The WMF model can be estimated using alternating least square (ALS) (Hu et al., 2008; He et al., 2016b), which yields closed-form updates for the user and item factors. Once these have been estimated, recommendation can be performed using the predicted interaction defined as $\hat{r}_{u,i} = \mathbf{w}_u^\top \mathbf{h}_i$, and constructing a recommendation list for each user by picking the items with the highest predicted interaction for this user. Even though WMF has shown good performance for recommender systems (Hu et al., 2008; Oord et al., 2013), it only applies to songs for which some listening history is available, hence facing the cold-start problem.

2.2 Combining WMF and deep content features

To alleviate the aforementioned cold-start problem, several works have proposed to exploit side content information about the songs in a factorization-based framework. While some of these works rely on closed-form models (Jeunen et al., 2020), e.g., based on matrix co-factorization (Gouvert et al., 2018), we focus here on approaches that leverage deep learning in conjunction with WMF (Liang et al., 2015; Oord et al., 2013). These are illustrated in Fig. 1 and described hereafter.

In (Liang et al., 2015), the authors propose to incorporate some content information in WMF. More precisely, they consider a prior on the item attribute matrix of the form of $\mathbf{h}_i \approx \Phi \mathbf{z}_i$, where $\mathbf{z}_i$ is the content feature vector, and Φ linearly maps this vector to the item embedding space. As for the content feature vector $\mathbf{z}_i$, the authors extract the last hidden representation of an auto-tagging DNN that is trained beforehand. This approach allows to alleviate the cold-start problem: for novel items without listening history, binarized playcounts can




Figure 1: Content-aware WMF methods using deep content features. Liang et al. (2015) (left) extract content features using an auto-tagging DNN and then incorporated it in WMF. Conversely, Oord et al. (2013) (right) first obtain the WMF decomposition, and then use the learned attribute matrix as target for training the deep content extractor.

still be predicted through $\hat{r}_{u,i} = \mathbf{w}_u^T \Phi \mathbf{z}_i$, from which recommendation can be made. However, the quality of the content features is highly dependent on the performance of the auto-tagging network. While these methods have significantly improved in the recent years (Kim et al., 2018; Pons et al., 2018), they are still limited by the noisy nature of the tags, which are mostly user-annotated. Besides, the learned features are not necessarily relevant for the task at hand (i.e., recommendation instead of tagging), since they are neither estimated jointly with the user/item embeddings nor fine-tuned for this task.

Oord et al. (2013) adopt a different strategy in order to leverage WMF for content-based recommendation. First, the factors $\mathbf{W}$ and $\mathbf{H}$ are estimated from the listening data using ALS updates, as mentioned in Section 2.1. Then, in a second stage, they train a DNN ϕ_θ with parameters θ whose purpose is to map input acoustic features $\mathbf{x}_i$ to the estimated attributes $\mathbf{h}_i$. This DNN can then be used for predicting item attributes for novel songs and therefore perform recommendation in a cold-start setting through $\hat{r}_{u,i} = \mathbf{w}_u^T \phi_\theta(\mathbf{x}_i)$. However, this approach strongly relies on the proper estimation of the WMF factors $\mathbf{W}$ and $\mathbf{H}$, which are needed for training ϕ_θ . Since the authors do not alternate the updates on the factors and the network parameters, the quality of the prior factorization sets a performance limit to this approach.

In a nutshell, these methods are mainly limited by that they operate in two stages, which means that the content feature extractor and the factorization model are not trained jointly. As a result, in (Liang et al., 2015) the extracted features are not necessarily relevant for a recommendation task, and in (Oord et al., 2013) the capacity of the deep content extractor is inherently limited by the prior factorization.

To alleviate these drawbacks, several recent works have proposed to combine these steps into a single-stage approach termed collaborative deep learning (CDL) (Wang et al., 2015). This allows for jointly learning the user/item embeddings and the deep content feature extractor. Various architectures have been proposed for modeling the DNN ϕ_θ , such as denoising (Wang et al., 2015) or variational (Li and She, 2017) auto-encoders. These approaches propose to estimate the whole model jointly, that is, alternating updates on the factors and the network parameters (this will be described in details in Section 3.2). Predictions are finally given by:

$$\hat{r}_{u,i} = \mathbf{w}_u^T (\mathbf{h}_i + \phi_\theta(\mathbf{x}_i)), \quad (1)$$

where $\mathbf{h}_i$ corresponds to the item embedding obtained in a collaborative filtering framework. However, this approach relies on a shallow user/item interaction model, since the collaborative filtering part remains matrix factorization-based.

More recently, knowledge graph-enhanced recommender systems have been proposed (Wang et al., 2019, 2020; Li et al., 2020). These consist in exploiting some latent information that might be shared between items (e.g., the composer of a given song might also have produced other records) in order to yield more explainable content feature extraction, and consequently improving recommendation performance. However, these approaches also consider a shallow user/item interaction model since predictions are made through the dot product of the embeddings.

2.3 Deep learning-based collaborative filtering

In the recent years, leveraging deep learning in collaborative filtering-based recommender systems has attracted a lot of attention (Zhang et al., 2019). The core idea of such approaches is to replace a shallow matrix factorization model such as presented in Section 2.1 with a DNN. Such a framework allows for learning refined user/item embeddings (Xue et al., 2017) and more complex interaction between these factors (He et al., 2017; Chen et al., 2019). More specifically, (He et al., 2017) proposes to integrate the user and item factors as embeddings in a neural network, which are then jointly learned using a gradient descent (GD) algorithm. First, they propose a model termed generalized matrix factorization (GMF), in which the predicted binarized playcounts are modeled as:

$$\hat{r}_{u,i} = \sigma(\mathbf{a}^\top (\mathbf{w}_u \odot \mathbf{h}_i)), \quad (2)$$

where σ is a non-linear activation function which replaces the identity function used in matrix factorization (5), and $\mathbf{a}$ is a weight vector. Alternatively, they consider a concatenation of the user and item factors (instead of their dot product), on top of which a multilayer perceptron (MLP) is applied to learn complex interactions between these factors. They also combine both approaches, yielding a model termed neural matrix factorization. This approach based on learning deep user/item interaction models is shown to outperform the shallow matrix factorization model, which has further been confirmed in (Chen et al., 2019). However, this so-called neural collaborative filtering (NCF) framework does not use any extra content information, and as such, face the cold-start problem. Several attempts were made to incorporate content information in such a framework (Lian et al., 2017), but do not leverage the full potential of deep learning. Besides, their potential has been shown for a traditional collaborative filtering task, but not for cold-start recommendation.

We propose to overcome this issue by introducing a unified model in which deep learning is leveraged for both learning deep user/item interactions (as presented in this section) and extracting content features (as described in Section 2.2 for cold-start recommendation).

3 Neural content-aware collaborative filtering

3.1 Model overview

3.1.1 Data

We consider implicit feedback data in the form of a playcount matrix $\mathbf{Y} \in \mathbb{R}^{U \times I}$. As recalled in Section 2.1, in order to better account for the over-dispersed nature of such data, it is common to consider the binarized playcount matrix $\mathbf{R}$ rather than the raw listening history. $\mathbf{R}$ indicates whether a user has listened to a song more than a certain amount of times τ or not (Liang et al., 2015, 2016; Tran et al., 2019):

$$\forall(u, i), r_{u,i} = \begin{cases} 1 & \text{if } y_{u,i} \geq \tau, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

In order to keep some information about the raw playcounts (which is lost when applying binarization), we also compute a confidence weight defined as:

$$\forall(u, i), c_{u,i} = 1 + \alpha \log \left(1 + \frac{y_{u,i}}{\epsilon} \right), \quad (4)$$

whose parameters are commonly set at $\alpha = 2$ and $\epsilon = 10^{-6}$ in the literature (Hu et al., 2008; Liang et al., 2015). Note that alternative schemes exist for defining the confidence, such as using constant values (Wang et al., 2015; Li and She, 2017). We resorted to using (4) as it yielded slightly better results in our preliminary experiments.

3.1.2 Generative process

We model the binarized playcounts as the result of the interaction between the user and item factors. Similarly as in WMF, we consider a Gaussian generative model:

$$\forall(u, i), r_{u,i} \sim \mathcal{N}(\psi_\gamma(\mathbf{w}_u, \mathbf{h}_i), c_{u,i}^{-1}), \quad (5)$$

where $c_{u,i}$ is the confidence defined in (4) and ψ_γ is the interaction model, which might depend on some parameters γ . The structure of this interaction model will be specifically described in Sections 3.2 and 3.3.

As in (Liang et al., 2015), we consider the following prior on the user factor:

$$\forall u \in \{1, \dots, U\}, \mathbf{w}_u \sim \mathcal{N}(0, \lambda_W^{-1} \mathbf{I}_K), \quad (6)$$

where λ_W is a regularization hyperparameter. In order to exploit some content information, one needs to consider an additional assumption about the item factor $\mathbf{h}_i$. Let us first consider a *relaxed* formulation, where this content information is incorporated in the model in the form of a prior as:

$$\forall i \in \{1, \dots, I\}, \mathbf{h}_i \sim \mathcal{N}(\phi_\theta(\mathbf{x}_i), \lambda_H^{-1} \mathbf{I}_K), \quad (7)$$

where ϕ_θ is a deep content feature extractor with parameters θ and $\mathbf{x}_i$ is a vector of low-level acoustic features for item i . We also consider a *strict* formulation, where the item attribute is directly predicted by the deep content feature extractor:

$$\forall i \in \{1, \dots, I\}, \mathbf{h}_i = \phi_\theta(\mathbf{x}_i), \quad (8)$$

which corresponds to (7) with $\lambda_H \rightarrow \infty$.

3.1.3 Deep content feature extractor

As for the deep content feature extractor ϕ_θ , we consider an MLP architecture with P layers, that is:

$$\phi_\theta(\mathbf{x}_i) = \phi^P(\phi^{P-1}(\dots\phi^1(\mathbf{x}_i))), \quad (9)$$

where $\mathbf{x}_i$ is a set of low-level acoustic features (described in Section 4.3), and ϕ^p denotes the p -th layer of ϕ_θ (for a clarity purpose, we do not explicitly write that ϕ^p depends on some parameters θ^p), such that:

$$\forall p \in \{1, \dots, P\}, \phi^p(\mathbf{z}) = \sigma^p(\mathbf{A}^p \mathbf{z} + \mathbf{b}^p), \quad (10)$$

where $\mathbf{A}^p$, $\mathbf{b}^p$ and σ^p respectively denote the weights, biases and activation function of the p -th layer. The activation function σ^p is chosen as the rectified linear unit (ReLU) function for all layers except for the last one, which uses the identity function. The choice of ReLU over alternative non-linear activation functions (such as the hyperbolic tangent or the sigmoid) is motivated by their consistent performance in the literature, notably their capability to reduce the gradient vanishing and overfitting problems (Glorot et al., 2011). Based on preliminary validation experiments and on prior work (Liang et al., 2015), each layer uses 1024 neurons, except for the last one which outputs the item factor of dimension K . We consider a total of $P = 3$ layers in our experiments.

The goal of this paper is to assess the potential of combining deep content feature extraction with deep interaction modeling in a unified content-aware collaborative filtering framework, rather than specifically optimizing the networks architectures. As such, we leave the usage of more advanced networks (e.g., convolutional layers potentially acting on raw audio waveforms (Lee et al., 2018), knowledge graph-based architectures (Wang et al., 2020)) to future work.

3.1.4 Estimation

Estimating the whole model in the maximum a posteriori sense results in the following optimization problem for the relaxed variant:

$$\min_{\xi_R} \mathcal{L}_{\text{NCACF-R}}(\xi_R) := \sum_{u,i} c_{u,i} (r_{u,i} - \psi_\gamma(\mathbf{w}_u, \mathbf{h}_i))^2 + \lambda_W \sum_u \|\mathbf{w}_u\|^2 + \lambda_H \sum_i \|\mathbf{h}_i - \phi_\theta(\mathbf{x}_i)\|^2, \quad (11)$$

where $\xi_R = \{\theta, \gamma, \mathbf{W}, \mathbf{H}\}$ is the whole set of parameters. The problem for the strict variant writes similarly:

$$\min_{\xi_S} \mathcal{L}_{\text{NCACF-S}}(\xi_S) := \sum_{u,i} c_{u,i} (r_{u,i} - \psi_\gamma(\mathbf{w}_u, \phi_\theta(\mathbf{x}_i)))^2 + \lambda_W \sum_u \|\mathbf{w}_u\|^2. \quad (12)$$

with $\xi_S = \{\theta, \gamma, \mathbf{W}\}$. Depending on the design choice for the interaction model ψ_γ , we can propose several optimization schemes for problems (11) and (12). In general, we will rely on a gradient descent strategy, but when the interaction model reduces to a shallow dot product, it becomes possible to leverage closed-form updates for estimating $\mathbf{W}$ and $\mathbf{H}$ (see Section 3.2.1).

3.1.5 Recommendation

Once the model is trained, cold-start recommendation is performed by computing the predicted binarized play-counts for all user/item pairs through:

$$\hat{r}_{u,i} = \psi_\gamma(\mathbf{w}_u, \phi_\theta(\mathbf{x}_i)). \quad (13)$$

The proposed model can also be used for a traditional collaborative filtering task which does not suffer from the cold-start problem. Warm-start recommendations is still performed using (13) for the strict variant. However, for the relaxed variant, it is usually preferred (Liang et al., 2015) to directly use the item embedding extracted from collaborative filtering $\mathbf{h}_i$, and to perform recommendation through:

$$\hat{r}_{u,i} = \psi_\gamma(\mathbf{w}_u, \mathbf{h}_i), \quad (14)$$

which we will use in our experiments in the warm-start setting.




Figure 2: Proposed NCACF models in the specific case where the interaction model reduces to a dot product between the user and item embeddings, in its relaxed (left) and strict (right) variants.

3.2 Linear interaction model

Let us first consider a particular case of the model described above, where the user/item interaction ψ_γ reduces to a matrix factorization model:

$$\psi_\gamma(\mathbf{w}_u, \mathbf{h}_i) = \mathbf{w}_u^\top \mathbf{h}_i. \quad (15)$$

The corresponding model is illustrated in Fig. 2 in its relaxed and strict variants. In such a scenario, the sets of parameters reduce to $\xi_R = \{\theta, \mathbf{W}, \mathbf{H}\}$ and $\xi_S = \{\theta, \mathbf{W}\}$, and problems (11) and (12) respectively rewrite:

$$\min_{\xi_R} \mathcal{L}_{\text{MF-R}}(\xi_R) := \sum_{u,i} c_{u,i} (r_{u,i} - \mathbf{w}_u^\top \mathbf{h}_i)^2 + \lambda_W \sum_u \|\mathbf{w}_u\|^2 + \lambda_H \sum_i \|\mathbf{h}_i - \phi_\theta(\mathbf{x}_i)\|^2, \quad (16)$$

and

$$\min_{\xi_S} \mathcal{L}_{\text{MF-S}}(\xi_S) := \sum_{u,i} c_{u,i} (r_{u,i} - \mathbf{w}_u^\top \phi_\theta(\mathbf{x}_i))^2 + \lambda_W \sum_u \|\mathbf{w}_u\|^2. \quad (17)$$

To address (16) and (17), we propose two optimization strategies, which we describe hereafter.

3.2.1 Hybrid algorithm

First, we propose a hybrid approach termed MF-Hybrid which boils down to combining two different algorithms in an alternating minimization framework. Similarly to the WMF model and its content-aware counterpart presented in Section 2.1 and 2.2, this strategy consists in estimating the user and item factors using an ALS scheme, which yields closed-form updates for $\mathbf{W}$ and $\mathbf{H}$ at each iteration. On the other hand, the parameters θ of the deep content network ϕ are estimated using a GD algorithm. This yields the following update scheme for the relaxed version:

$$\forall u \in \{1, \dots, U\}, \mathbf{w}_u \leftarrow (\mathbf{H} \text{diag}(\mathbf{c}_u) \mathbf{H}^\top + \lambda_W \mathbf{I}_K)^{-1} \mathbf{H} \text{diag}(\mathbf{c}_u) \mathbf{r}_u, \quad (18)$$

$$\forall i \in \{1, \dots, I\}, \mathbf{h}_i \leftarrow (\mathbf{W} \text{diag}(\mathbf{c}_i) \mathbf{W}^\top + \lambda_H \mathbf{I}_K)^{-1} (\mathbf{W} \text{diag}(\mathbf{c}_i) \mathbf{r}_i + \lambda_H \phi_\theta(\mathbf{x}_i)), \quad (19)$$

$$\theta \leftarrow \theta - \eta \nabla_\theta \left(\sum_i \|\mathbf{h}_i - \phi_\theta(\mathbf{x}_i)\|^2 \right), \quad (20)$$

where ∇ denotes the gradient operator, η is the learning rate, $\mathbf{r}_u = [r_{u,1}, \dots, r_{u,I}]^\top$, $\mathbf{r}_i = [r_{1,i}, \dots, r_{U,i}]^\top$, and similarly for $\mathbf{c}_u$ and $\mathbf{c}_i$.¹ Using a similar approach for the strict variant yields:

$$\forall u \in \{1, \dots, U\}, \mathbf{w}_u \leftarrow (\phi_\theta(\mathbf{X}) \text{diag}(\mathbf{c}_i) \phi_\theta(\mathbf{X})^\top + \lambda_W \mathbf{I}_K)^{-1} \phi_\theta(\mathbf{X}) \text{diag}(\mathbf{c}_i) \mathbf{r}_u, \quad (21)$$

$$\theta \leftarrow \theta - \eta \nabla_\theta \left(\sum_{u,i} c_{u,i} (r_{u,i} - \mathbf{w}_u^\top \phi_\theta(\mathbf{x}_i))^2 \right). \quad (22)$$

¹While $\mathbf{r}_i$ denotes the i -th column of $\mathbf{R}$, $\mathbf{r}_u$ denotes its u -th row, which might appear as a slight notation abuse. Indeed, using the same notation convention, the u -th row should be denoted by $[\mathbf{R}^\top]_u^\top$. Nonetheless, we decided to keep the notation $\mathbf{r}_u$ for brevity.

Algorithm 1: MF-Hybrid-Relaxed

```
1 Inputs: Binarized playcounts  $\mathbf{R} \in [0, 1]^{U \times I}$  and confidence  $\mathbf{C} \in \mathbb{R}_+^{U \times I}$ 
2 Initial user preferences  $\mathbf{W} \in \mathbb{R}^{K \times U}$  and item attributes  $\mathbf{H} \in \mathbb{R}^{K \times I}$  matrices
3 Initial deep content feature network parameters  $\theta$ 
4 Number of iterations  $N$ , and number of epochs per iteration  $N_{\text{gd}}$ 
5 for  $j = 1$  to  $N$  do
6   | Update  $\mathbf{W}$  using (18)
7   | Update  $\mathbf{H}$  using (19)
8   | for  $j' = 1$  to  $N_{\text{gd}}$  do
9   |   | Update  $\theta$  using (20)
10  | end
11 end
12 Output:  $\mathbf{W}$ ,  $\mathbf{H}$  and  $\theta$ .
```

Algorithm 2: MF-Hybrid-Strict

```
1 Inputs: Binarized playcounts  $\mathbf{R} \in [0, 1]^{U \times I}$  and confidence  $\mathbf{C} \in \mathbb{R}_+^{U \times I}$ 
2 Initial user preferences matrix  $\mathbf{W} \in \mathbb{R}^{K \times U}$ 
3 Initial deep content feature network parameters  $\theta$ 
4 Number of iterations  $N$ , and number of epochs per iteration  $N_{\text{gd}}$ 
5 for  $j = 1$  to  $N$  do
6   | Update  $\mathbf{W}$  using (21)
7   | for  $j' = 1$  to  $N_{\text{gd}}$  do
8   |   | Update  $\theta$  using (22)
9   | end
10 end
11 Output:  $\mathbf{W}$  and  $\theta$ .
```

Note that in practice, both (20) and (22) might be implemented differently as one might use a stochastic variant of a GD algorithm with momentum. The proposed procedures are termed MF-Hybrid-Relaxed and MF-Hybrid-Strict and summarized in Algorithms 1 and 2, respectively. These are iterative schemes in which the first stage of each iteration consists in updating the embeddings using ALS updates, and the second stage consists in updating the network parameters. Note that the DNN parameters θ update uses an arbitrary number N_{gd} of GD steps at each iteration. Indeed, in practice, one does not need to fully minimize the objective with respect to θ (at $\mathbf{W}$ and $\mathbf{H}$ fixed), but rather seeks for a decrease of the objective function. We will investigate the impact of this parameter in our experiments.

Therefore, these algorithms generalize the related work (Oord et al., 2013), which consists in first estimating the WMF model in a content-free scenario (that is, using (18) and (19) with $\phi_\theta = 0$) and then learning the mapping ϕ_θ through either (20) or (22) (the corresponding losses were termed “mean square error” and “weighted prediction error” in (Oord et al., 2013), respectively).

3.2.2 Unified learning

We now address the optimization problems (16) and (17) using a single GD algorithm. We integrate the user and item factors as embedding layers within a DNN, as illustrated in Fig. 2. This network consists of a collaborative filtering part and a content feature extractor part. The collaborative filtering part is fed with user and item indices to yield user and item factors through an embedding layer. These are then combined using an interaction model (which here reduces to a dot product) to yield predicted binarized playcounts. The item factor can be regularized using the deep content feature extractor branch (in the relaxed variant) or directly predicted by this branch (in the strict variant).

Even though the resulting models are equivalent to those presented in Section 3.2.1, this now allows for a unified learning approach where only the GD algorithm is used to train all the parameters, by considering as objective functions the losses given in (16) and (17). As a result, the gradient updates are:

$$\xi_{\text{R}} \leftarrow \xi_{\text{R}} - \eta \nabla_{\xi_{\text{R}}} \mathcal{L}_{\text{MF-R}}(\xi_{\text{R}}), \quad (23)$$

and

$$\xi_{\text{S}} \leftarrow \xi_{\text{S}} - \eta \nabla_{\xi_{\text{S}}} \mathcal{L}_{\text{MF-S}}(\xi_{\text{S}}). \quad (24)$$

The corresponding procedures will be referred to as MF-Uni-Relaxed and MF-Uni-Strict, respectively.




Figure 3: Proposed NCACF model in the general case, where a deep interaction model is used in addition to the deep content feature extractor, in its relaxed (left) and strict (right) variants.

3.3 Deep interaction model

We now propose a more general version of the network where we leverage a deep interaction model in order to learn more refined interactions between embedding vectors (He et al., 2017; Xue et al., 2019). This approach is illustrated in Fig. 3. First, the two embeddings are combined into a single vector. Drawing on prior work (Liu et al., 2015; He et al., 2017; Xue et al., 2019), we propose two embedding combinations, based on their multiplication or concatenation:

$$\forall(u, i), \mathbf{v}_{u,i} = \begin{cases} \mathbf{w}_u \odot \mathbf{h}_i & (\text{multiplication}), \\ \begin{bmatrix} \mathbf{w}_u \\ \mathbf{h}_i \end{bmatrix} & (\text{concatenation}). \end{cases} \quad (25)$$

The resulting combined vector $\mathbf{v}_{u,i}$ is of length K' , which is equal to K or $2K$ in the case of a multiplication or concatenation, respectively. Previous studies such as (Chen et al., 2019) have shown that using a concatenation of the user and item factors overall outperforms using a multiplication, regardless of the size and number of deep layers of the subsequent network. However, since our framework is different (we consider the cold-start scenario and we use an additional deep content extractor), we will evaluate both combination methods experimentally. This combined vector is then fed as input to an MLP with Q hidden layers to output the predicted binarized playcounts, which is defined similarly as the deep content feature extractor (10):

$$\hat{r}_{u,i} = \psi^{Q+1}(\psi^Q(\dots\psi^1(\mathbf{v}_{u,i}))), \quad (26)$$

where ψ^q denotes the q -th layer of the network ψ_γ , such that:

$$\forall q \in \{1, \dots, Q+1\}, \psi^q(\mathbf{z}) = \sigma^q(\mathbf{A}^q \mathbf{z} + \mathbf{b}^q), \quad (27)$$

where $\mathbf{A}^q$, $\mathbf{b}^q$ and σ^q are the weights, biases and activation function of the q -th layer, respectively (we use the same notations as for the deep content feature extractor ϕ_θ for clarity, but the corresponding parameters are different). The last layer (i.e., the output layer) has a particular structure, since it does not use biases (that is, $\mathbf{b}^{Q+1} = 0$) and uses a single neuron in order to output the predicted binarized playcount of dimension 1. As activation functions, we choose ReLU for the Q hidden layers and the sigmoid function for the output layer, which is chosen to enforce the predicted binarized playcounts to range between 0 and 1 (He et al., 2017; Chen et al., 2019).

In order to choose an appropriate number of neurons for each hidden layer, we resort to using the so-called “tower” structure, which has shown appropriate in deep collaborative filtering models (He et al., 2016a, 2017; Chen et al., 2019). Specifically, this structure consists in halving the size of each successive higher layers, as this

allows for learning more abstract representations in higher layers of the network (He et al., 2017). In practice, the q -th hidden layer therefore uses $K'/2^{q-1}$ neurons.

These models will be referred to as NCACF-Relaxed and NCACF-Strict, and are estimated by addressing the optimization problems (11) and (12) using a GD algorithm, which yields the following updates:

$$\xi_R \leftarrow \xi_R - \eta \nabla_{\xi_R} \mathcal{L}_{\text{NCACF-R}}(\xi_R), \quad (28)$$

and

$$\xi_S \leftarrow \xi_S - \eta \nabla_{\xi_S} \mathcal{L}_{\text{NCACF-S}}(\xi_S), \quad (29)$$

where $\mathcal{L}_{\text{NCACF-R}}$ and $\mathcal{L}_{\text{NCACF-S}}$ are the losses defined in (11) and (12), respectively. In particular, when the interaction model is based on a multiplication of the embeddings and uses a single output layer with a fixed weight vector $\mathbf{a} = \mathbf{1}$ and no activation function (i.e., the sigmoid is replaced with the identity), this model becomes equivalent to the MF-Uni model presented in Section 3.2.2.

3.4 Training considerations

3.4.1 Loss functions

The proposed NCACF model and its variants are based on a Gaussian generative process for the playcounts, from which a weighted prediction error naturally arises as loss function (*cf.* (11) and (12)). However, this loss is known not to be the most appropriate choice when handling implicit feedback data such as playcounts (Salakhutdinov et al., 2007). Consequently, other losses such as the log loss have been preferred in such tasks (He et al., 2017; Xue et al., 2017; Chen et al., 2019). Nonetheless, we decided to use the weighted prediction error in this work since the underlying Gaussian generative model allows to easily incorporate prior information, which has motivated previous work on content-aware collaborative filtering such as (Wang et al., 2015, 2016; Huan et al., 2017) to adopt this framework. Besides, alternative losses might improve the overall performance of all proposed methods, but it would make the comparison unfair with baselines such as (Oord et al., 2013).

Therefore, we leave the usage of alternative losses to future work. Note that this would be consistent with the design of alternative generative models, for instance by directly modeling the raw interaction data instead of the binarized playcounts. In particular, refined statistical models based on the Poisson or compound Poisson distributions (Gopalan et al., 2014a,b; Basbug and Engelhardt, 2016; Gouvert et al., 2019) allow for robust modeling of such over-dispersed implicit feedbacks, and might yield more appropriate losses for this task.

3.4.2 Sampling

Except for MF-Hybrid-Relaxed, learning the proposed models involves to perform GD updates for minimizing a weighted prediction error on the form of:

$$\sum_{u=1}^U \sum_{i=1}^I c_{u,i} (r_{u,i} - \hat{r}_{u,i})^2. \quad (30)$$

In practice, this problem is split over batches of data $\mathcal{B} \subset \{1, \dots, U\} \times \{1, \dots, I\}$, each of which being processed at each iteration of the GD algorithm. The loss then rewrites for each batch:

$$\sum_{(u,i) \in \mathcal{B}} c_{u,i} (r_{u,i} - \hat{r}_{u,i})^2. \quad (31)$$

Processing all the available data (both non-zero and zero playcounts) results in a relatively high computational burden, which might become cumbersome for very large datasets. To alleviate this issue, and to account for the sparsity of the dataset, it is common to resort to a negative sampling strategy (Hsieh et al., 2017; Lee et al., 2018; Chen et al., 2019), where only part of the null playcounts is considered, that is:

$$\sum_{(u,i) \in \mathcal{B}^+} \left[c_{u,i} (r_{u,i} - \hat{r}_{u,i})^2 + \sum_{j \in \mathcal{B}_u^-} c_{u,j} (r_{u,j} - \hat{r}_{u,j})^2 \right] \quad (32)$$

where $\mathcal{B}^+$ is a batch of positive samples (that is, a set of user/item pairs with a non-zero playcount) and $\mathcal{B}_u^- = \{j = 1, \dots, N^- | r_{u,j} = 0\}$ is a set of N^- negative samples for user u . We tested this strategy in preliminary experiments, but we obtained a relatively poor performance for all methods. Consequently, we rather consider a more traditional sampling strategy, where we considered batches of items:²

$$\sum_{i \in \mathcal{B}_i} \sum_{u=1}^U c_{u,i} (r_{u,i} - \hat{r}_{u,i})^2 \quad (33)$$

²Note that an alternative common choice consists in considering batches of users instead of items. However, in a content-aware framework, considering batches of items is more straightforward to train both the collaborative filtering part and the content extractor part, since the latter operates on items (and not on users).

where $\mathcal{B}_i \subset \{1, \dots, I\}$. Even though this approach imposes to use a smaller batch size, since playcounts for all users are considered for each sample $i \in \mathcal{B}_i$, it yields a significantly better performance for all tested methods. Consequently, we present the results obtained using this sampling strategy.

3.5 Relation to other works

The proposed framework encompasses and generalizes several collaborative filtering-based models from the literature. We summarize these in Table 2 along with their main characteristics, that is:

- whether the model is trained using a two-stage or a joint training approach;
- whether the training algorithm hybridizes two forms of updates or a single gradient descent;
- whether the model uses a shallow or a deep interaction model;
- whether it is content-aware, thus suitable for cold-start recommendation, or not.

The NCACF-Relaxed model is notably an extension of the NCF approach (He et al., 2017) with an additional branch to account for some content information, which makes it suitable for cold-start recommendation. In particular, when the interaction model is based on a multiplication of the embeddings and uses a single layer, NCACF reduces to the variant termed generalized matrix factorization (*cf.* (2)).

When the interaction model in NCACF reduces to a dot product, the corresponding models (MF-Hybrid and MF-Uni) also encompass some previous proposals. In particular, MF-Hybrid-Relaxed is similar to (Wang et al., 2015; Li and She, 2017), as these works use a one-stage approach for learning the model which is based on a hybrid algorithm (using both ALS and GD). Note however that our proposal is slightly more flexible since it allows for using an arbitrary number of GD epochs in between ALS updates (*cf.* Algorithm 1). If we further simplify the training procedure of MF-Hybrid by only conducting ALS updates (in a content-free framework) and then learning the deep network ϕ_θ , then it reduces to the deep content-based (DCB) recommendation technique (Oord et al., 2013) presented in Section 2.2.

On the other hand, MF-Uni-Strict is somewhat equivalent to the deep content-user embedding (DCUE) model introduced in (Lee et al., 2018). The model (Zheng et al., 2017) is also similar to MF-Uni-Strict when the interaction model is replaced by a factorization machine. These approaches are however not tested in a specific cold-start setting, even though they are suitable for this task, but rather consider a traditional collaborative filtering application. Lee et al. (2018) actually observe that DCUE yields worse results than a simple WMF-based model. This observation is reminiscent of prior works such as (Liang et al., 2015) where it is observed that warm-start recommendation usually does not benefit from adding extra content information. Therefore, we will test this approach in a cold-start scenario, which appears more appropriate for highlighting the potential of this content-aware model.

We also propose MF-Uni-Relaxed, which is a more flexible approach than its strict counterpart. Besides, still according to (Liang et al., 2015), it is expected to outperform the strict variant (Lee et al., 2018) in a warm-start collaborative filtering. MF-Uni-Relaxed shares some similarity with the content-boosted collaborative filtering network (CCCFNet) presented in (Lian et al., 2017). It uses a shallow content extractor and exploits two item embeddings (a collaborative filtering one and another one extracted from side-content). Like DCUE, CCFNet has only been evaluated in a warm-start recommendation scenario.

Overall, our unifying framework not only generalizes several methods from the literature, but also allows for making them suitable for cold-start recommendation, for which their potential was partly still left to explore.

4 Experimental protocol

In this section, we present the experimental protocol to assess the potential of the proposed model for music recommendation. Even though our framework is primarily intended to perform cold-start recommendation, we also consider a more traditional warm-start scenario, in order to show the potential of exploiting side content information in such a case. In the spirit of reproducible research, we provide the code related to these experiments.³ Our method is implemented using the PyTorch framework,⁴ and all the computations have been performed using an NVIDIA Tesla T4 GPU with 15 GB RAM.

4.1 Research questions

Our experiments aim at answering the following research questions:

1. Is a joint training approach preferable over a two-stage training for content-aware models?
2. How does a unified training algorithm perform compared to a method that hybridizes two forms of updates?
3. What is the impact of the interaction model (shallow vs. deep) onto performance for both warm- and cold-start recommendation?

³<https://github.com/magronp/ncacf>

⁴<https://pytorch.org/>

Table 2: Overview of the proposed models, related approaches in the literature, and main characteristics.

Model	Related work	Joint learning	GD-only	Deep interaction	Content-aware
NCF	(He et al., 2017)	-	✓	✓	✗
DCB					
Relaxed	(Oord et al., 2013)	✗	✗	✗	✓
Strict	(Oord et al., 2013)	✗	✗	✗	✓
MF-Hybrid					
Relaxed	(Wang et al., 2015)	✓	✗	✗	✓
Strict	-	✓	✗	✗	✓
MF-Uni					
Relaxed	(Lian et al., 2017)	✓	✓	✗	✓
Strict	(Lee et al., 2018)	✓	✓	✗	✓
NCACF					
Relaxed	-	✓	✓	✓	✓
Strict	-	✓	✓	✓	✓

4. Which training loss, or equivalently, which model variant (relaxed or strict) performs the best?
5. Does a content extractor improves performance even in the warm-start setting?
6. How do content-aware baselines perform for cold-start recommendation (generally, these have been tested only for warm-start recommendation)?

Questions 1, 2, and 3, will be addressed in Sections 5.2, 5.3, and 5.4, respectively, while Question 4 will be addressed throughout these sections. Questions 5 and 6 will be treated in Section 5.1.

4.2 Playcount data

As implicit feedback data, we use the Taste Profile dataset⁵ which is part of the Million Song Dataset (Bertin-Mahieux et al., 2011). It provides listening counts of 1 million users and 380,000 songs. After removing duplicates, we kept the songs whose acoustic content features were available (see Section 4.3). The playcount data is binarized by retaining values of seven or higher as implicit feedback (that is, $\tau = 7$ in (3)), since lower values yield feedback that are commonly considered as non-informative (Tran et al., 2019). In accordance with (3), other values are set at 0. As in (Liang et al., 2015; Tran et al., 2019), in order to keep the computational burden low, we retain the top songs and users (sorted by playcounts) and we remove inactive users and items: that is, we only keep users who listened to at least 20 songs, and songs which have been listened to by at least 50 users. The resulting dataset has a density of about 0.2 %, which is a relatively high level of sparsity when compared to other datasets on which similar methods are tested (Chen et al., 2019).

To train and evaluate the models, we resort to a 10-fold cross-validation strategy (Flexer, 2006). In the cold-start setting, we hold 20 % of the songs out as validation subset. We then randomly partition the remaining songs into 10 equally sized subsets, thus 90 % and 10 % of these songs serve for training and testing, respectively. The playcounts corresponding to the training subset are used for learning the various models. The held-out validation subset is used for tuning the hyperparameters (see Section 4.4), and the remaining test set is used for comparing the models. Evaluation on the validation and test sets corresponds to the cold-start scenario, since the models have not been trained on the playcounts corresponding to the songs they contain. In order to reduce the computational burden, hyperparameters are only tuned on the first training/testing split. The corresponding optimal values are then used for all 9 remaining splits when testing.

For warm-start recommendation, we use a similar strategy, except we split the non-zero playcounts instead of the songs: this ensures that all songs in the validation and test sets also have some non-zero interaction in the training set (otherwise, warm-start recommendation would not be possible). The resulting datasets' properties are detailed in Table 3.

4.3 Acoustic content features

Several previous works (Oord et al., 2013; Lee et al., 2018) rely on extracting acoustic features from the raw audio waveforms directly. Due to copyright restriction, the audio files corresponding to the songs in the Million Song Dataset cannot be distributed. Therefore, these prior works have relied on the 7digital platform,⁶ which allows to download short audio samples for small evaluations and prototyping. However, these audio snippets are not easily

⁵<http://millionsongdataset.com/tasteprofile/>

⁶<https://us.7digital.com/>

Table 3: Dataset statistics for the first split (interactions denote the number of non-zero playcounts).

	Total	Warm-start			Cold-start		
		Training	Validation	Test	Training	Validation	Test
Users	17,572	17,572	17,572	17,572	17,572	17,572	17,572
Songs	13,532	13,532	13,532	13,532	9,743	2,706	1,083
Interactions	509,177	366,607	101,835	40,735	355,891	105,897	47,389

obtained due to access restrictions of the 7digital-API. Besides, these are limited to 30s-long snippets, which makes it difficult to extract features that encompass the full temporal dynamic of the songs.⁷ An alternative option consists in using the pre-extracted features provided with the Million Song Dataset and computed using the Echo Nest platform. While a subset of these features is easily available for prototyping, the full dataset is however difficult to access, as it is stored on an Amazon snapshot which access requires a non-free account. Finally, we could not access the Echo Nest API to recompute these features.⁸

In the spirit of reproducible research, we aimed at using easily accessible acoustic features. Therefore, we used the statistical spectrum descriptors (SSDs) (Lidy and Rauber, 2005), since these features have shown good performance in several music information retrieval tasks such as genre recognition (Schindler and Rauber, 2012; Schindler et al., 2012), and are freely available online.⁹ The SSDs are sets of statistical moments extracted from the sonogram of each song, a mid-level time-frequency representation that reflects the human loudness sensation. To compute the SSDs, Schindler and Rauber (2012) first obtained the raw audio waveform of each song using the 7digital platform, from which they obtained a time-frequency representation by applying a short-time Fourier transform. The resulting frequency channels are then grouped into 24 psycho-acoustically motivated critical bands, accounting for several masking effects. Finally, seven statistical moments (mean, median, variance, skewness, kurtosis, min, and max) are computed in each critical band to account for the temporal dynamic of each song. This results in a set of low-level acoustic features for each song $\mathbf{x}_i \in \mathbb{R}^L$ with $L = 168$, which we scale to have 0 mean and unit variance.

Note that alternative pre-calculated low-level acoustic features are available, such as MFCC-based statistical moments computed with the Essentia toolkit (Bogdanov et al., 2013).¹⁰ However, these are computed from 30s-long excerpts, and stacked into a feature vector of length $L = 40$, while SSDs are vectors of length $L = 168$ computed from the whole songs. Therefore, we resort to using the SSDs, since they capture more comprehensively the acoustic content of the songs.

4.4 Training details

Following a previous music recommendation work using the same dataset (Tran et al., 2019), we set the user and item embeddings dimension at $K = 128$, whether these are computed using ALS or integrated within the network. The user/item embeddings are initialized (in all models) with random values drawn from a centered normal distribution with a standard deviation of 10^{-2} . The deep content feature extractor and deep interaction model parameters θ and γ are initialized using the default Le Cun’s initialization scheme (LeCun et al., 2012), except for the output layer of the interaction model in NCACF (and the NCF baseline), whose weights are all initialized with ones. GD is performed using the Adam algorithm (Kingma and Ba, 2015) with a learning rate of 10^{-4} and a batch size of 128. The hyperparameters λ_W and λ_H are tuned on the validation set and chosen to maximize the NDCG metric (see Section 4.5). The number of epochs is determined on the validation set, with a maximum value of 150. In order to save some computational time, NCACF uses the same hyperparameters as MF-Uni and is pretrained using no deep interaction model nor activation function (which corresponds to MF-Uni when the interaction model is a multiplication of the embeddings) for a maximum number of 100 epochs. The subsequent NCACF models are then trained for a maximum of 100 more epochs.

4.5 Evaluation metric

We use the NDCG (Wang et al., 2013) as a measure of the overall quality of the recommendation. For each user u , we compute a ranked list of items in the evaluation (i.e., validation or test) set based on the predicted preferences $\hat{\mathbf{r}}_u$. We then compute the relevance of this list with respect to the ground truth preferences, that is, the observed user/item interaction in the evaluation set: $\text{rel}_{u,i} = 1$ if the item i is in the listening history of user u (in the evaluation set) and 0 otherwise. In order to favor recommendations that place the evaluation items high

⁷For these reasons, we did not consider alternative datasets such as KKBox (<https://www.kaggle.com/c/kkbox-music-recommendation-challenge/data>), for which only the listening history is directly available.

⁸The Echo Nest developer API (<http://developer.echonest.com/>) was not accessible at the time of conducting this research.

⁹<http://www.ifs.tuwien.ac.at/mir/msd/download.html>

¹⁰<https://zenodo.org/record/3258042#.YpYacTnPiso>

in the list, we apply a discounted weight to the relevance, which yields the discounted cumulative gain (DCG):

$$\text{DCG}_u = \sum_{i=1}^{I'} \frac{\text{rel}_{u,i}}{\log_2(i+1)}, \quad (34)$$

where I' denotes the length of the list of ranked items. To obtain a metric that accounts for all items in the evaluation set, I' should be equal to the number of songs in this set. However, this considerably increases the computational load to evaluate the methods. Consequently, a common approach consists in considering the truncated list of top- I' items instead (Liang et al., 2016). Following similar work conducted on this dataset (Tran et al., 2019), we used $I' = 50$ in our experiments, which is reasonable for a music recommendation application, and significantly reduces the computational burden. The normalized version of the DCG is then obtained as follows:

$$\text{NDCG}_u = \frac{\text{DCG}_u}{\text{IDCG}_u}, \quad (35)$$

where IDCG is the ideal DCG, which corresponds to the DCG of a perfectly ranked list. Finally, these scores are averaged over users to yield an overall recommendation performance. The resulting NDCG ranges from 0 to 1 (higher is better), and will be expressed in % in our experiments for readability.

4.6 Baselines

We compare our proposed NCACF model with the following methods:

- **WMF** (Hu et al., 2008). This basic matrix factorization-based recommendation method does not leverage any side-content information. It is described in Section 2.1.
- **DCB** (Oord et al., 2013). This two-stage deep content-based (DCB) recommendation technique uses a linear interaction model, and it is described in Section 2.2.
- **CDL** (Wang et al., 2015). The collaborative deep learning framework leverages side-content information and relies on a linear interaction model. Unlike DCB, the whole model is trained jointly in a single stage by combining ALS updates and GD iterations, as described in Section 2.2.
- **DCUE** (Lee et al., 2018). This model is similar to CDL, but it is trained using a unified GD-based algorithm.
- **CCCFNet** (Lian et al., 2017). This model is similar to DCUE and is also trained using a unified GD algorithm, but it uses two different item embeddings (extracted from the collaborative filtering and side-content, respectively), thus sharing more similarity with our MF-Uni-Relaxed.
- **NCF** (He et al., 2017). This method exploits a deep interaction model and reaches state-of-the-art warm-start recommendation performance. It is described in Section 2.3.

Note that even though several baseline methods use alternative deep architectures for content extraction (e.g., convolutional layers for DCUE or auto-encoders for CDL), we implement them using the same MLP described in Section 3.1 for a fair comparison. The link between these methods and our implementation as variants of the NCACF model is described in Section 3.5.

5 Results and discussion

5.1 Overall recommendation results

Let us examine the recommendation performance of the baselines and the proposed NCACF model for a warm- and cold-start recommendation tasks. The results presented in Table 4 correspond to the best performing method in each category for a fair comparison.

Firstly, by comparing WMF and content-aware approaches using a shallow interaction model, we observe that leveraging side content information does not significantly improve performance in the warm-start setting (except for CCCFNet, for which a slight performance improvement can be noticed), but mostly allow to address the cold-start problem. This is in line with several previous works such as (Liang et al., 2015; Magron and Févotte, 2021) that suggested that accounting for extra content information was not beneficial in this setting. The slightly different behavior of CCCFNet might be explained by the fact that in this strict approach, acoustic features are used to directly predict the item embedding, thus it exploits the available content information in a way that is more optimal for the recommendation task at hand. Nonetheless, a different trend is observed when comparing NCF and NCACF (whose deep interaction models are similar) in the warm-start setting, since the latter significantly outperforms its content-free counterpart. This outlines the potential of content-aware approaches even in traditional collaborative filtering settings, provided a deep content feature extractor is used.

Let us now compare the various content-aware approaches using a shallow interaction model. DCB is outperformed by the other approaches in all configurations, except for DCUE in the warm-start setting. Overall,

Table 4: Overall performance of the compared methods on the test set (mean NDCG in % $\pm$ standard deviation over the 10 folds).

	Warm-start	Cold-start
WMF	31.5 ± 0.2	–
DCB	30.0 ± 1.8	23.7 ± 1.7
CDL	31.3 ± 0.2	25.3 ± 0.9
DCUE	27.3 ± 0.2	24.7 ± 1.0
CCCFNet	31.7 ± 0.1	23.9 ± 1.2
NCF	41.1 ± 0.4	–
NCACF	43.6 ± 0.5	26.9 ± 4.0

Table 5: Performance of the DCB and MF-Hybrid approaches (NDCG in %). Note that in the warm-start setting, DCB is only presented in its strict variant since the relaxed one simply reduces to WMF.

		Warm-start		Cold-start	
		Relaxed	Strict	Relaxed	Strict
DCB			21.0	17.5	15.0
MF-Hybrid	$N_{\text{gd}} = 1$	28.2	18.2	19.4	15.2
	$N_{\text{gd}} = 2$	28.1	13.8	19.3	13.8
	$N_{\text{gd}} = 5$	28.0	16.5	19.3	11.3
	$N_{\text{gd}} = 10$	28.0	15.9	19.2	11.3


this demonstrates the advantage of resorting to a joint learning strategy rather than a two-stage approach. Interestingly, while CCFNet yields the best results in the warm-start setting among these approaches, this is not the case in the cold-start setting, where it is outperformed by CDL. The main difference between these methods being the training algorithm (GD only vs. ALS+GD), leveraging closed-form updates should be preferred to a unified algorithm in the cold-start setting (when at least part of the model is tractable). Besides, note that these two methods (CDL and CCFNet) are relaxed approaches, since the item embedding is regularized by the content network (but not directly predicted by it). This emphasizes the benefit of using such a flexible approach, as it will be confirmed in the following experiments. Finally, we remark that DCUE yields poor results in the warm-setting, where it is outperformed by a basic WMF, a conclusion that is reminiscent of the original DCUE paper (Lee et al., 2018). However, the results in Table 4 also reveal that DCUE outperforms DCB in the cold-start setting: this finding outlines that the true potential of such an approach lies within the cold-start setting, which complements the experiments conducted in (Lee et al., 2018) where only warm-start recommendation was considered.

A comparison between NCF and the other baselines in the warm-start setting reveals the potential of using a deep interaction model. Indeed, NCF yields a large 8.1 % NDCG improvement over the best performing approach using a shallow interaction model (CCFNet), which confirms the findings of previous studies (He et al., 2017; Chen et al., 2019). As mentioned above, exploiting content information (i.e., using NCACF) allows to boost the performance even further in this setting. More interestingly, this approach allows to alleviate the cold-start problem faced by NCF: NCACF yields the best cold-start performance, outperforming the previous best performing approach (CDL) by 1.6 % in NDCG. This highlights the potential of using a deep user/item interaction model over a conventional matrix factorization-based technique, where both the interaction model and the content extractor are jointly trained for cold-start music recommendation.


5.2 Learning strategies for a linear interaction model

The goal of this experiment is to assess the potential of a joint training strategy over a two-stage approach for the variants based on a linear interaction model. To that end, we compare MF-Hybrid and DCB, which correspond to these two strategies.

First, let us select the optimal hyperparameters for the MF-Hybrid algorithm (a similar experiment is conducted for optimizing DCB). The results in the warm-start scenario are presented in Fig. 4. It can be noted that the performance of MH-Hybrid-Relaxed does not strongly depend on these parameters (except for very large values of λ_H , where the performance decreases). The strict variant exhibits a less stable behavior overall and also yields worse results than its relaxed counterpart. A similar trend can be observed in the cold-start setting, whose results are presented in Fig. 5. More specifically, MF-Hybrid-Relaxed exhibits a better performance and a smoother (quasi-monotonic) behavior for large values of λ_W and λ_H . This emphasizes the importance of a (relatively large) regularization on the user preference factor $\mathbf{W}$ in the cold-start setting.




(a) Relaxed




(b) Strict

Figure 4: NDCG on the validation set for the MF-Hybrid algorithms (relaxed variant on the top, strict variant on the bottom) in the warm-start setting, for several values of the hyperparameters.



(a) Relaxed



(b) Strict

Figure 5: NDCG on the validation set for the MF-Hybrid algorithms (relaxed variant on the top, strict variant on the bottom) in the cold-start setting, for several values of the hyperparameters.




Figure 6: NDCG on the validation set for the MF-Uni algorithms (relaxed variant on the top, strict variant on the bottom) in the warm-start setting, for several values of the hyperparameters.

Let us now investigate the impact of the parameter N_{gd} on the performance of MF-Hybrid algorithms. Indeed, as can be read in Algorithms 1 and 2, these approaches consist in alternating ALS updates for estimating $\mathbf{W}$ (and potentially $\mathbf{H}$ for the relaxed variant) and GD updates for estimating the deep content extractor’s parameters. As such, the number of GD epochs performed in each iteration in between ALS updates is expected to have an impact on the overall performance. We test the MF-Hybrid algorithms with $N_{\text{gd}} = 1, 2, 5$, and 10 (the total number of epochs being the same), and we present the results in Table 5.

For the relaxed variant, we observe that in both the warm- and cold-start scenarios, choosing a value of N_{gd} between 1 and 10 yields overall similar results. However, a different behavior is observed for the strict variants, where the performance drops when increasing N_{gd} . Since the total number of epochs is fixed for a fair comparison, increasing N_{gd} implies decreasing the total number of ALS updates, which might explain the decreasing performance in this case. The relaxed variants are therefore more promising since they exhibit less sensitivity to the number of GD updates. Overall, updating the deep content feature extractor with a single epoch after updating the embeddings appears to be the optimal choice, as done e.g., in the CDL baseline, which corresponds to MF-Hybrid-Relaxed when $N_{\text{gd}} = 1$.

Finally, we also report in Table 5 the performance of the DCB two-stage approach in its strict and relaxed variants, which completes the experimental results from Section 5.1. We observe that the MF-Hybrid approaches overall outperform this baseline. In particular, in the cold-start scenario the performance improvement is more significant for the relaxed variant (+1.9 %) than for its strict counterpart (+0.2 %). This confirms the potential of jointly estimating the user/item factors and the deep content extractor, rather than relying on a two-stage approach. We also note that the relaxed variant outperforms its strict counterpart (for both DCB and MF-Hybrid approaches), which highlights the interest of using a more flexible system, especially in the cold-start recommendation scenario.

5.3 Learning algorithms for the MF models

In this experiment, we compare the learning strategies of the MF models, that is, whether they are trained using a combination of ALS and GD (MF-Hybrid) or a single GD algorithm (MF-Uni).

We first tune the hyperparameters for MF-Uni on the validation set and present the results in Fig. 6 and 7 for the warm- and cold-start scenarios, respectively. Interestingly, we observe that unlike MF-Hybrid, the performance in both variants (and in both scenarios) does not strongly depend on these values (except when very large values for λ_W and λ_H are used, in which case the performance tends to decrease). This might explain why several previous studies incorporating the factors within the network did not specifically consider a regularization on the user preference matrix (He et al., 2017; Lee et al., 2018).

The performance of MF-Uni on the validation set is reported in Table 6, as well as the performance of MF-Hybrid with $N_{\text{gd}} = 1$. We observe that in the cold-start scenario, the strict variant of MF-Uni significantly




Figure 7: NDCG on the validation set for the MF-Uni algorithms (relaxed variant on the top, strict variant on the bottom) in the cold-start setting, for several values of the hyperparameters.

Table 6: Comparison between MF-Hybrid and MF-Uni.

		Warm-start		Cold-start	
		NDCG (%)	Time (s)	NDCG (%)	Time (s)
Relaxed	MF-Hybrid	28.2	41	19.4	296
	MF-Uni	29.0	1525	18.4	1334
Strict	MF-Hybrid	18.2	296	15.2	1048
	MF-Uni	23.4	1334	19.0	1334

outperforms MF-Hybrid, but a different behavior is observed for the relaxed variant, where MF-Uni is outperformed by MF-Hybrid: this reveals the importance of choosing the learning algorithm in accordance with the underlying model. In the warm-start setting however, different conclusions are drawn: MF-Uni outperforms MF-Hybrid in both variants, but the best performing variants are the relaxed ones. This complements the results obtained in (Lee et al., 2018), where DCUE (which is similar to MF-Uni-Strict) yielded a limited performance for a warm-start recommendation task: this model might have actually benefited from a more flexible formulation (MF-Uni-Relaxed), or might have revealed more potential for a cold-start recommendation task, where it is competitive with MF-Hybrid-Relaxed.

We also present in Table 6 the total computational time for these methods. While MF-Uni-Relaxed allows for a slight improvement over MF-Hybrid-Relaxed in the warm-start setting, it is significantly more computationally demanding. The fastest method overall is MF-Hybrid-Relaxed: indeed, training the network in this approach only involves the target attributes $\mathbf{H}$, as opposed to its strict counterpart and to the MF-Uni approach, which involve the set of all binarized playcounts $\mathbf{R}$. Therefore, this experiment shows that when (part of) the model is tractable, leveraging closed-form updates in a flexible formulation is more interesting than treating all the factors as embedding layers in a deep network trained with GD, in terms of both computational time and performance. Incorporating user/item factors in the network then becomes interesting when the interaction model no longer allows for deriving closed-form updates, as will be shown in the next experiment.

5.4 Impact of the interaction model in NCACF

In this experiment, we study the impact of the interaction model on the performance of NCACF. To save some computational time, we use the hyperparameters values obtained in the previous experiment. We consider two possible embedding combinations based on multiplication or concatenation, and a variable number of hidden




Figure 8: NDCG on the validation set for the NCACF model in warm-start (left) and cold-start (right) settings, for an interaction model based on multiplication (top) or concatenation (bottom) of the embeddings, for a variable amount of hidden layers Q .

layers Q in the interaction network. The results are presented in Fig. 8.

First, we remark that overall the relaxed variants of NCACF benefit from using a deep interaction model, since we observe a performance increase (in most cases up to $Q = 4$) when compared to a shallow interaction model ($Q = 0$) (except for model using embedding concatenation in the cold-start scenario). We also observe that in most cases, the relaxed variants significantly outperform their strict counterparts, which seems more prone to overfitting when Q increases. Nonetheless, in the warm-start scenario, the performance of NCACF-Strict (for both embedding multiplication and concatenation) remains superior to that of MF-Hybrid-Strict, which confirms the potential of fully learning a strict model with a single GD algorithm, provided a deep interaction model is used.

Interestingly, we remark that embedding multiplication outperforms concatenation by a large margin for all NCACF variants. Even though in the warm-setting embedding concatenation with $Q = 1$ outperforms the baseline using a shallow interaction model, the improvement is more important when using embedding multiplication. This might be explained by the ability of the (non-linear) concatenation technique to learn complex interaction patterns between user and items for which some shared feedback is available, which in turns does not properly generalize to unseen items, where a more simple multiplicative model yields better performance. Overall, this result shows that several architectural design choices employed in “pure” (i.e., warm-start) collaborative filtering techniques cannot be systematically exploited in the more challenging cold-start scenario, or when a deep content extractor is used. This somehow tempers conclusions from the literature, where embedding concatenation has been shown to perform better than multiplication in several studies addressing warm-start collaborative filtering (He et al., 2017; Chen et al., 2019): while we observed that for NCF embedding multiplication and concatenation yield overall similar results, this does not generalize to its content-aware counterpart, even for warm-start recommendation.

This conclusion is reminiscent of recent work such as (Rendle et al., 2020), where it is shown that a carefully tuned matrix factorization technique outperforms an NCF model using embeddings concatenation. The authors in (Rendle et al., 2020) also outline that a “leave-one-out strategy”, which is commonly employed for evaluating recommender systems (He et al., 2017; Chen et al., 2019), is not appropriate for drawing general conclusions, notably when a different train/test splitting strategy is used, which is the case in our experiment as described in Section 4.2. Further investigation is then required to fully identify to what extent a concatenation of the

embeddings is appropriate for addressing the cold-start problem.

6 Conclusion

In this work, we introduced neural content-aware collaborative filtering, a unifying framework for cold-start recommendation which encompasses many methods from the literature. We proposed several variants of this model in order to study the impact of the learning strategy (two-stage vs. joint learning), training algorithm (hybrid vs. unified), and interaction model (shallow vs. deep). In particular, the proposed NCACF method fully leverages deep learning for modeling user/item interactions and extracting content information, and reaches state-of-the-art results for both warm- and cold-start recommendation on a large scale and highly sparse music dataset.

In future work, we will investigate alternative generative models that are better suited for the considered implicit feedback data. For instance, compound Poisson models (Gopalan et al., 2014a; Gouvert et al., 2019) are appropriate for directly modeling the raw (non-binarized) playcounts and accounting for their over-dispersed nature. Alternatively, Bernoulli models (Bingham et al., 2009) are appropriate for modeling binarized playcounts, and yield alternative losses such as the binary cross entropy, which have shown superior performance than the quadratic loss (Chen et al., 2019). We will also explore more refined sampling strategy (Chen et al., 2017; Tran et al., 2019) in order to reduce the computational time while keeping the performance high. Besides, our work can be extended to handle other modalities in order to fully exploit the available content such as artists biographies (Oramas et al., 2017) or musical tags (Liang et al., 2015; Zuo et al., 2016), but also contextual data such as culture (Zangerle et al., 2020) or location (Gillhofer and Schedl, 2015). This work can also be combined with knowledge-enhanced recommenders (Wang et al., 2019, 2020; Li et al., 2020) in order to improve the interpretability of the item embeddings. Finally, alternative architectures could be exploited, such as convolutional networks that directly extract content features from the raw audio data (Lee et al., 2018) in an end-to-end fashion.

References

- Basbug ME, Engelhardt BE (2016) Hierarchical compound Poisson factorization. In: Proceedings of the 33rd International Conference on Machine Learning (ICML), pp 1795–1803
- Bertin-Mahieux T, Ellis DP, Whitman B, Lamere B (2011) The million song dataset. In: Proceedings of the 12th International Society for Music Information Retrieval Conference (ISMIR), pp 591–596, DOI 10.7916/D8NZ8J07
- Bingham E, Kabán A, Fortelius M (2009) The aspect Bernoulli model: Multiple causes of presences and absences. *Pattern Analysis & Applications* 12(1):55–78, DOI 10.1007/s10044-007-0096-4
- Bogdanov D, Wack N, Gómez E, Gulati S, Herrera P, Mayor O, Roma G, Salamon J, Zapata J, Serra X (2013) ESSENTIA: An audio analysis library for music information retrieval. In: Proceedings of the 14th International Conference on Music Information Retrieval (ISMIR), pp 319–326
- Chen T, Sun Y, Shi Y, Hong L (2017) On sampling strategies for neural network-based collaborative filtering. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’17), p 767–776, DOI 10.1145/3097983.3098202
- Chen W, Cai F, Chen H, Rijke MD (2019) Joint neural collaborative filtering for recommender systems. *ACM Transactions on Information Systems* 37(39):1–30, DOI 10.1145/3343117
- Cheng HT, Koc L, Harmsen J, Shaked T, Chandra T, Aradhye H, Anderson G, Corrado G, Chai W, Ispir M, Anil R, Haque Z, Hong L, Jain V, Liu X, Shah H (2016) Wide & deep learning for recommender systems. In: Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, p 7–10, DOI 10.1145/2988450.2988454
- Cheng Z, Shen J (2016) On effective location-aware music recommendation. *ACM Transactions on Information Systems* 34(13), DOI 10.1145/2846092
- Fang Y, Si L (2011) Matrix co-factorization for recommendation with rich side information and implicit feedback. In: Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec ’11), p 65–69, DOI 10.1145/2039320.2039330
- Ferwerda B, Yang E, Schedl M, Tkalcic M (2015) Personality traits predict music taxonomy preferences. In: Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA ’15), pp 2241–2246, DOI 10.1145/2702613.2732754
- Flexer A (2006) Statistical evaluation of music information retrieval experiments. *Journal of New Music Research* 35(2):113–120, DOI 10.1080/09298210600834946

- Gillhofer M, Schedl M (2015) Iron Maiden while jogging, Debussy for dinner? an analysis of music listening behavior in context. In: Proceedings of the 21st International conference on MultiMedia Modeling (MMM 2015), p 380–391, DOI 10.1007/978-3-319-14442-9_44
- Glorot X, Bordes A, Bengio Y (2011) Deep sparse rectifier neural networks. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, vol 15, pp 315–323
- Gopalan PK, Charlin L, Blei D (2014a) Content-based recommendations with Poisson factorization. In: Proceedings of the 27th International Conference on Neural Information Processing Systems (NIPS’14), p 3176–3184
- Gopalan PK, Hofman JM, Blei D (2014b) Scalable recommendation with hierarchical Poisson factorization. In: Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI’15), p 326–335
- Gouvert O, Oberlin T, Févotte C (2018) Matrix co-factorization for cold-start recommendation. In: Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR), p 792–798
- Gouvert O, Oberlin T, Févotte C (2019) Recommendation from raw data with adaptive compound Poisson factorization. In: Proceedings of the International Conference on Machine Learning (ICML)
- He K, Zhang X, Ren S, Sun J (2016a) Deep residual learning for image recognition. In: Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp 770–778, DOI 10.1109/CVPR.2016.90
- He X, Zhang H, Kan MY, Chua TS (2016b) Fast matrix factorization for online recommendation with implicit feedback. In: Proceedings of the 39th International ACM Conference on Research and Development in Information Retrieval (SIGIR ’16), p 549–558, DOI 10.1145/2911451.2911489
- He X, Liao L, Zhang H, Nie L, Hu X, Chua TS (2017) Neural collaborative filtering. In: Proceedings of the 26th International Conference on World Wide Web (WWW ’17), pp 173–182, DOI 10.1145/3038912.3052569
- He X, Du L, Wang X, Tian F, Tang J, Chua TS (2018) Outer product-based neural collaborative filtering. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18), pp 2227–2233
- Hsieh CK, Yang L, Cui Y, Lin TY, Belongie S, Estrin D (2017) Collaborative metric learning. In: Proceedings of the 26th International Conference on World Wide Web (WWW ’17), pp 193–201, DOI 10.1145/3038912.3052639
- Hu Y, Koren Y, Volinsky C (2008) Collaborative filtering for implicit feedback datasets. In: Proceedings of the 2008 Eighth IEEE International Conference on Data Mining (ICDM ’08), pp 263–272, DOI 10.1109/ICDM.2008.22
- Huan H, Wei Z, Liang L, Yang L (2017) Collaborative filtering recommendation model based on convolutional denoising auto encoder. In: Proceedings of the 12th Chinese Conference on Computer Supported Cooperative Work and Social Computing (ChineseCSCW ’17), p 64–71, DOI 10.1145/3127404.3127420
- Jeunen O, Van Balen J, Goethals B (2020) Closed-form models for collaborative filtering with side-information. In: Proceedings of the 14th ACM Conference on Recommender Systems (RecSys ’20), p 651–656, DOI <https://doi.org/10.1145/3383313.3418480>
- Kim T, Lee J, Nam J (2018) Sample-level CNN architectures for music auto-tagging using raw waveforms. In: Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp 366–370
- Kingma DP, Ba JL (2015) Adam: a method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (ICLR)
- Koren Y, Bell R, Volinsky C (2009) Matrix factorization techniques for recommender systems. *Computer* 42(8):30–37, DOI 10.1109/MC.2009.263
- Laplanche A (2014) Improving music recommender systems: What can we learn from research on music tastes? In: Proceedings of the 15th International Society for Music Information Retrieval Conference (ISMIR), pp 451–456
- LeCun YA, Bottou L, Orr GB, Müller KR (2012) Efficient BackProp, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 9–48. DOI 10.1007/978-3-642-35289-8_3
- Lee J, Lee K, Park J, Park J, Nam J (2018) Deep content-user embedding model for music recommendation. URL <https://arxiv.org/pdf/1807.06786>
- Li S, Kawale J, Fu Y (2015) Deep collaborative filtering via marginalized denoising auto-encoder. In: Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM ’15), p 811–820, DOI 10.1145/2806416.2806527

- Li X, She J (2017) Collaborative variational autoencoder for recommender systems. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17), p 305–314, DOI 10.1145/3097983.3098077
- Li Z, Xu Q, Jiang Y, Cao X, Huang Q (2020) Quaternion-based knowledge graph network for recommendation. In: Proceedings of the 28th ACM International Conference on Multimedia, MM '20, p 880–888, DOI 10.1145/3394171.3413992
- Lian J, Zhang F, Xie X, Sun G (2017) CCCFNet: A content-boosted collaborative filtering neural network for cross domain recommender systems. In: Proceedings of the 26th International Conference on World Wide Web Companion, p 817–818, DOI 10.1145/3041021.3054207
- Liang D, Zhan M, Ellis DP (2015) Content-aware collaborative music recommendation using pre-trained neural networks. In: Proceedings of the 16th International Society for Music Information Retrieval Conference (ISMIR)
- Liang D, Charlin L, McInerney J, Blei DM (2016) Modeling user exposure in recommendation. In: Proceedings of the International World Wide Web Conference (WWW), pp 951–961, DOI 10.1145/2872427.2883090
- Liang D, Krishnan RG, Hoffman MD, Jebara T (2018) Variational autoencoders for collaborative filtering. In: Proceedings of the 2018 World Wide Web Conference, WWW '18, p 689–698, DOI 10.1145/3178876.3186150
- Lidy T, Rauber A (2005) Evaluation of feature extractors and psycho-acoustic transformations for music genre classification. In: Proceedings of the 6th International Society for Music Information Retrieval Conference (ISMIR)
- Liu X, Ouyang Y, Rong W, Xiong Z (2015) Item category aware conditional restricted boltzmann machine based recommendation. In: Proceedings, Part II, of the 22nd International Conference on Neural Information Processing - Volume 9490, ICONIP 2015, p 609–616, DOI 10.1007/978-3-319-26535-3_69
- Magron P, Févotte C (2021) Leveraging the structure of musical preference in content-aware music recommendation. In: <https://arxiv.org/abs/2010.10276>
- Marlin B, Zemel RS (2004) The multiple multiplicative factor model for collaborative filtering. In: Proceedings of the twenty-first international conference on Machine learning Proceedings of the International Conference on Machine Learning (ICML '04), DOI 10.1145/1015330.1015437
- Oord Avd, Dieleman S, Schrauwen B (2013) Deep content-based music recommendation. In: Proceedings of the 26th International Conference on Neural Information Processing Systems, NIPS'13, p 2643–2651
- Oramas S, Nieto O, Sordo M, Serra X (2017) A deep multimodal approach for cold-start music recommendation. In: Proceedings of the 2nd Workshop on Deep Learning for Recommender Systems (DLRS 2017), p 32–37, DOI 10.1145/3125486.3125492
- Pichl M, Zangerle E (2021) User models for multi-context-aware music recommendation. *Multimedia Tools and Applications* 80:22509–22531, DOI <https://doi.org/10.1007/s11042-020-09890-7>
- Pons J, Nieto O, Prockup M, Schmidt E, Ehmann A, Serra X (2018) End-to-end learning for music audio tagging at scale. In: Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR), pp 637–644
- Rendle S, Freudenthaler C, Gantner Z, Schmidt-Thieme L (2009) BPR: Bayesian personalized ranking from implicit feedback. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09), p 452–461
- Rendle S, Krichene W, Zhang L, Anderson J (2020) Neural collaborative filtering vs. matrix factorization revisited. In: Proceedings of the Fourteenth ACM Conference on Recommender Systems (RecSys '20), p 240–248, DOI 10.1145/3383313.3412488
- Salakhutdinov R, Mnih A (2007) Probabilistic matrix factorization. In: Proceedings of the 20th International Conference on Neural Information Processing Systems (NIPS'07), p 1257–1264
- Salakhutdinov R, Mnih A (2008) Bayesian probabilistic matrix factorization using markov chain monte carlo. In: Proceedings of the 25th international conference on Machine learning (ICML '08), p 880–887, DOI 10.1145/1390156.1390267
- Salakhutdinov R, Mnih A, Hinton G (2007) Restricted Boltzmann machines for collaborative filtering. In: Proceedings of the International Conference on Machine Learning (ICML), p 791–798, DOI 10.1145/1273496.1273596
- Schedl M, Knees P, McFee B, Bogdanov D, Kaminskas M (2015) Music recommender systems. In: *Recommender Systems Handbook*, Springer, pp 453–492

- Schedl M, Knees P, Gouyon F (2017) New paths in music recommender systems research. In: Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17), p 392–393, DOI 10.1145/3109859.3109934
- Schedl M, Zamani H, Chen CW, Deldjoo Y, Elahi M (2018) Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval* 7(2):95–116, DOI 10.1007/s13735-018-0154-2
- Schein AI, Popescul A, Ungar LH, Pennock DM (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, SIGIR '02, p 253–260, DOI 10.1145/564376.564421
- Schindler A, Rauber A (2012) Capturing the temporal domain in Echonest features for improved classification effectiveness. In: Proceedings of the International Workshop on Adaptive Multimedia Retrieval (AMR), pp 214–227, DOI 10.1007/978-3-319-12093-5_13
- Schindler A, Mayer R, Rauber A (2012) Facilitating comprehensive benchmarking experiments on the million song dataset. In: Proceedings of the 13th International Society for Music Information Retrieval Conference (ISMIR), pp 469–474
- Soleymani M, Aljanaki A, Wiering F, Veltkamp RC (2015) Content-based music recommendation using underlying music preference structure. In: Proceedings IEEE International Conference on Multimedia and Expo (ICME), DOI 10.1109/ICME.2015.7177504
- Tran VA, Hennequin R, Royo-Letelier J, Moussallam M (2019) Improving collaborative metric learning with efficient negative sampling. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19, p 1201–1204, DOI 10.1145/3331184.3331337
- Wang C, Blei DM (2011) Collaborative topic modeling for recommending scientific articles. In: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '11), pp 448–456, DOI doi.org/10.1145/2020408.2020480
- Wang H, Wang N, Yeung DY (2015) Collaborative deep learning for recommender systems. In: Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '15), p 1235–1244, DOI 10.1145/2783258.2783273
- Wang H, Shi X, Yeung DY (2016) Collaborative recurrent autoencoder: Recommend while learning to fill in the blanks. In: Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16), p 415–423
- Wang X, Wang Y (2014) Improving content-based and hybrid music recommendation using deep learning. In: Proceedings of the 22nd ACM international conference on Multimedia (MM '14), pp 627–636, DOI 10.1145/2647868.2654940
- Wang X, He X, Cao Y, Liu M, Chua TS (2019) Kgat: Knowledge graph attention network for recommendation. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '19)
- Wang Y, Wang L, Yuanzhi L, He D, Liu TY (2013) A theoretical analysis of NDCG type ranking measures. In: Proceedings of the 26th Conference on Learning Theory (COLT), pp 25–54
- Wang Z, Lin G, Tan H, Chen Q, Liu X (2020) CKAN: Collaborative Knowledge-Aware Attentive Network for Recommender Systems. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, p 219–228, DOI 10.1145/3397271.3401141
- Wu Y, DuBois C, Zheng AX, Ester M (2016) Collaborative denoising auto-encoders for top-N recommender systems. In: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM '16), p 153–162, DOI 10.1145/2835776.2835837
- Xue F, He X, Wang X, Xu J, Liu K, Hong R (2019) Deep item-based collaborative filtering for top-N recommendation. *ACM Transactions on Information Systems* 37(3), DOI 10.1145/3314578
- Xue HJ, Dai X, Zhang J, Huang S, Chen J (2017) Deep matrix factorization models for recommender systems. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17), pp 3203–3209
- Yoshii K, Goto M, Komatani K, Ogata T, Okuno HG (2006) Hybrid collaborative and content-based music recommendation using probabilistic model with latent user preferences. In: Proceedings of the 7th International Society for Music Information Retrieval Conference (ISMIR)

- Zangerle E, Pichl M, Schedl M (2020) User models for culture-aware music recommendation: Fusing acoustic and cultural cues. *Transactions of the International Society for Music Information Retrieval (TISMIR)* 3(1):1–16, DOI 10.5334/tismir.37/
- Zhang S, Yao L, Sun A, Tay Y (2019) Deep learning based recommender system: A survey and new perspectives. *ACM Compututing Surveys* 52(1), DOI 10.1145/3285029
- Zheng L, Noroozi V, Yu PS (2017) Joint deep modeling of users and items using reviews for recommendation. In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*, p 425–434, DOI 10.1145/3018661.3018665
- Zuo Y, Zeng J, Gong M, Jiao L (2016) Tag-aware recommender systems based on deep neural networks. *Neuro-computing* 204:51 – 60, DOI 10.1016/j.neucom.2015.10.134