

Personalized TV Recommendation: Fusing User Behavior and Preferences

ANONYMOUS AUTHOR(S)*

In this paper, we propose a two-stage ranking approach for recommending linear TV programs. The proposed approach first leverages user viewing patterns regarding time and TV channels to identify potential candidates for recommendation and then further leverages user preferences to rank these candidates given textual information about programs. To evaluate the method, we conduct empirical studies on a real-world TV dataset, the results of which demonstrate the superior performance of our model in terms of both recommendation accuracy and time efficiency.

ACM Reference Format:

Anonymous Author(s). 2018. Personalized TV Recommendation: Fusing User Behavior and Preferences. *J. ACM* 37, 4, Article 111 (August 2018), 8 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Linear TV programs play crucial roles in our daily lives. With the quickly increasing number of TV channels and programs, it is important to develop effective recommender systems for TV users. Although the development of recommender systems has been stimulated by the rapid growth of information on the Internet, and many algorithms have been successfully applied to various online services (e.g., music and video streaming services) [3, 8, 13], little has been done for personalized TV recommendation (TV Rec) in the literature. Most well-developed recommendation algorithms are not applicable for such a recommendation problem due to the following two key challenges of TV Rec: (1) Complete-item cold start: Unlike video on demand (VOD), new TV programs are released on a daily basis (although some drama or movies are replayed, they usually have different titles or descriptions);¹ (2) Context awareness: user viewing behavior for TV programs strongly depends on their conditions (e.g., time and mood); for instance, watching news during dinner but preferring sports in the morning.

To address the first challenge, some studies adopt content-based approaches combined with collaborative filtering (CF) for TV Rec [1, 6, 7, 11, 15]. However, these approaches do not consider the second key characteristic—context awareness—in TV Rec, for which another line of work focuses mainly on characterizing users’ time-aware preferences [2, 10, 12, 14, 15]. Although, these studies model users’ time-aware preferences regarding channels and program genres, they do not precisely reflect users’ viewing preferences regarding program content. This is due to the fact that users’ access to channels also depends on their viewing habits or location, and that genres are merely coarse-grained information about programs and thus provide little information about program content. Moreover, [9] further accounts user moods but such user data is difficult to obtain and even harder to measure.

To address the above two challenges within a unified framework, we propose a two-stage ranking approach for TV Rec which consists of two components: one to model viewing behavior and the other for viewing preferences. Specifically,

¹Another practical challenge is that the programs that share common content do not share an identical ID, which rules out directly adopting collaborative filtering or matrix factorization in real-world scenarios.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

viewing behavior refers to users’ viewing patterns regarding time and TV channels, whereas viewing preferences refers to preferences regarding the content of TV programs. For the former, we adopt a finer granularity in terms of time than previous work (e.g., days×hours in [2, 12]), whereas for the latter, we leverage textual information about programs to better model user viewing preferences. Moreover, inspired by the capabilities and limitations of the two components, we propose fusing them with a simple yet effective two-stage ranking algorithm that locates potential candidates based on the first component and then further ranks them based on the second component. Also note that in the literature, this is the first work to formally define the problem of TV Rec and provide a unified approach to capture both user viewing behavior and preferences. Empirical results on a real-world TV dataset demonstrate its effectiveness in recommendation; at the same time, this approach is advantageous and practical for real-world applications due to its time-efficient and parameter-free design.

2 METHODOLOGY

2.1 Problem Formulation

Personalized TV recommendation (TV Rec) is the task of recommending yet-to-be-released TV programs to a group of users. To properly formulate the problem and our proposed method, we first define three terms: 1) weekly time slot, 2) interaction tensor, and 3) program meta information required for TV Rec. With these definitions, we formalize personalized TV Rec as a top- k recommendation problem given user-implicit feedback.

Definition 2.1 (Weekly time slot). A weekly time interval can be equally divided into n weekly time slots, each of which is denoted as $w_i = (t_i, t_{i+1}]$, where t_i (t_{i+1}) denotes the beginning time (the end time, respectively) of the i -th time slot. Together, all of the time slots compose set $W = \{w_i | 1 \leq i \leq n\}$. Thus, any given timestamp $s \in S$ can be projected onto a weekly time slot $w_{\mathcal{T}(s)} \in W$ by function $\mathcal{T}(\cdot) : S \rightarrow \{1, \dots, n\}$, where S denotes a set of arbitrary timestamps.

For example, when we divide a week into 168 time slots (i.e., one hour for each time slot), we have $W = \{w_1 = [\text{Mon } 00:00, \text{Mon } 01:00), \dots, w_{168} = [\text{Sun } 23:00, \text{Mon } 00:00)\}$, in which the specific timestamp “May 11, 2020, 05:30 (Mon)” belongs to the 6th time slot, w_6 . Note that a given time span $[s, e]$ can also be projected onto a set of time slots $\{w_j | \mathcal{T}(s) \leq j \leq \mathcal{T}(e)\}$. Also note that in our later empirical studies, we adopt a finer granularity in terms of time (i.e., 15 minutes as the length of the time slot) than prior art.

Definition 2.2 (Interaction tensor). Let U , I , and C denote the sets of users, TV programs, and TV channels, respectively. An interaction tensor, denoted as $\mathcal{A} = (a_{u,i,w,c}) \in \mathbb{R}^{|U| \times |I| \times |W| \times |C|}$, represents user-item associations through a certain channel within a certain weekly time slot, where $a_{u,i,w,c}$ denotes the weight of the association. Note that the tensor is binary for implicit feedback; that is, if user $u \in U$ views program $i \in I$ played in channel $c \in C$ within time slot $w \in W$, $a_{u,i,w,c} = 1$; otherwise, $a_{u,i,w,c} = 0$.

Definition 2.3 (Program meta information). Given a set of TV programs I , meta information for each $i \in I$ records that program i is broadcast by channel $\text{CH}(i) \in C$ at the time interval $[s_i, e_i]$ with the content information $\text{CNT}(i)$, where $\text{CH}(\cdot)$ and $\text{CNT}(\cdot)$ are the projection functions respectively mapping program i to its channel and its textual information (e.g., title, artists, and abstract).

PROBLEM 1. Top- k TV Recommendation from Implicit Feedback. Let I_{train} and I_{test} denote the sets of TV programs broadcast in the past (training data) and in the future (test data), respectively; note that for the problem of TV Rec, $I_{\text{train}} \cap I_{\text{test}} = \emptyset$. Given a historical interaction tensor $\mathcal{A}_{\text{train}} = (a_{u,i,w,c}) \in \mathbb{R}^{|U| \times |I_{\text{train}}| \times |W| \times |C|}$, for each user $u \in U$, we

identify the top- k programs from the set of yet-to-be-released (new) programs I_{test} by leveraging the information from $\mathcal{A}_{\text{train}}$ and meta information of $I_{\text{train}} \cup I_{\text{test}}$.

2.2 Proposed Method

With a TV recommender system, we seek to leverage historical viewing logs and content information of programs to infer two user characteristics: (1) behavior and (2) preferences, which are addressed in Sections 2.2.1 and 2.2.2, respectively. We then propose a simple yet effective two-stage ranking method in Section 2.2.3 that takes into account both user characteristics, thereby fusing user viewing habits and preferences into the modeling process.

2.2.1 Viewing behavior. Here, we define the so-called *viewing behavior* of users based on the following observations. As suggested by [12], most TV users exhibit predictable viewing behavior strongly connected to weekly time slots and TV channels. Intuitively, users prefer to watch TV during their leisure time, which heavily depends on their work and lifestyle. In addition, users tend to switch between a limited number of channels even though they have a large number to choose from. Thus a user's TV viewing behavior can be defined as the probability distribution of watching TV on a given channel at a given time.

Given a historical user-item interaction tensor $\mathcal{A}_{\text{train}} = (a_{u,i,w,c}) \in \mathbb{R}^{|U| \times |I_{\text{train}}| \times |W| \times |C|}$, we extract each user u 's viewing behavior by computing his or her viewing probability distribution over weekly time slots W and TV channels C . Formally speaking, we represent each u 's viewing behavior as a probability distribution matrix, $\mathcal{B}^u = (b_{w,c}^u) \in \mathbb{R}^{|W| \times |C|}$, where each element $b_{w,c}^u$ is defined as

$$b_{w,c}^u = \left(\sum_{i,w,c} a_{u,i,w,c} \mathbb{1}_{\{w=w\}} \mathbb{1}_{\{c=c\}} \right) / \left(\sum_{i,w,c} a_{u,i,w,c} \right). \quad (1)$$

Additionally, in order to recommend yet-to-be-released TV programs for users based on their viewing behavior, we construct the matrix $\mathcal{B}^{i'} = (b_{w,c}^{i'}) \in \mathbb{R}^{|W| \times |C|}$ for each new item $i' \in I$ using the meta information defined in Definition 2.3, where $b_{w,c}^{i'} = \mathbb{1}_{\{w \in \{w_j | \mathcal{T}(s_{i'}) \leq j \leq \mathcal{T}(e_{i'})\}\}} \cdot \mathbb{1}_{\{\text{CH}(i')=c\}}$. Recall that $[s_{i'}, e_{i'}]$ denotes the time interval during which program i' is broadcast. Finally, we compute the matching score between u and i' given viewing behavior as

$$s_{u,i'}^b = \text{MAX} \left(\mathcal{B}^u \odot \mathcal{B}^{i'} \right) \text{ and } (w, c) = \text{IdxMax} \left(\mathcal{B}^u \odot \mathcal{B}^{i'} \right), \quad (2)$$

where \odot denotes element-wise multiplication between two matrices, $\text{MAX}(\cdot)$ is the function to extract the maximum element in a matrix, and $\text{IdxMax}(\cdot)$ locates the indices of the maximum element.² Note that $s_{u,i'}^b$ is the estimated probability that user u views item i' given his or her historical viewing behavior.

2.2.2 Viewing Preferences. In contrast to the aforementioned user behavior, a user's preferences are usually associated with the content of his or her preferred items. We formally define a user's *viewing preferences* as the program contents he or she prefers to watch, which we represent in the proposed method using the textual information of programs. Note that as with a typical TV Rec scenario, all candidate items in I_{test} for recommendation are new, which is the same as the complete cold-start problem in typical recommender systems. Such a problem is commonly addressed using content-based approaches [4, 15]; likewise, we here use textual item information to locate new items for recommendation.

For each program $i \in I_{\text{train}}$, we map its content information to a d -dimensional embedding h_i using a text encoder \mathcal{E} :

$$h_i = \mathcal{E}(\text{CNT}(i)) \in \mathbb{R}^d. \quad (3)$$

²In practice, there is no need to conduct the element-wise multiplication to get $s_{u,i'}^b$; instead, for each i' , $s_{u,i'}^b$ is the maximum in the set $\{b_{w,c}^u | w \in \{w_j | \mathcal{T}(s_{i'}) \leq j \leq \mathcal{T}(e_{i'})\} \wedge c = \text{CH}(i')\}$.

In order to map user \mathbf{u} 's preferences to the same embedding space, we gather all the programs associated with \mathbf{u} in the training data, after which we compute the average pooling over their embeddings to obtain \mathbf{u} 's viewing preferences $h_{\mathbf{u}}$ as

$$h_{\mathbf{u}} = \frac{\sum_{i \in I_{\text{train}}^{\mathbf{u}}} h_i}{|I_{\text{train}}^{\mathbf{u}}|} \in \mathbb{R}^d, \quad (4)$$

where $I_{\text{train}}^{\mathbf{u}} = \{i \mid i \in I_{\text{train}} \wedge \exists w \in W, c \in C a_{\mathbf{u},i,w,c} = 1\}$. Similarly, for each item $i' \in I_{\text{test}}$, we project its content information using the same text encoder \mathcal{E} from Eq. (3). Finally, the matching score for \mathbf{u} and i' in terms of viewing preferences is computed as

$$s_{\mathbf{u},i'}^p = \langle h_{\mathbf{u}}, h_{i'} \rangle, \quad (5)$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product of two vectors.

In addition, for TV Rec, it is common that multiple users (i.e., family members) share the same account, under which these users may have different viewing preferences and watch TV at different weekly time slots. For example, whereas children enjoy watching cartoons after school, parents prefer to watch news or dramas after work. We address this by further tailoring the viewing preferences of an ‘‘account’’ to time-aware preferences; that is, for each account $\mathbf{u} \in U$ and each time slot $\mathbf{w} \in W$, we have

$$h_{\mathbf{u},\mathbf{w}} = \frac{\sum_{i \in I_{\text{train}}^{\mathbf{u},\mathbf{w}}} h_i}{|I_{\text{train}}^{\mathbf{u},\mathbf{w}}|} \in \mathbb{R}^d, \quad (6)$$

where $I_{\text{train}}^{\mathbf{u},\mathbf{w}} = \{i \mid i \in I_{\text{train}} \wedge \exists c \in C a_{\mathbf{u},i,\mathbf{w},c} = 1\}$. With these fine-grained viewing preferences, the score of user \mathbf{u} for item i' becomes

$$s_{\mathbf{u},i'}^p = \langle h_{\mathbf{u},\mathbf{w}_j}, h_{i'} \rangle, \quad (7)$$

where $\mathbf{w}_j \in W$ denotes the time slot in which item i' begins playing; i.e., $j = \mathcal{T}(s_{i'})$.

2.2.3 Two-stage Ranking. In this section, we propose a two-stage ranking approach that leverages the above two features—user viewing behavior and user viewing preferences—for TV Rec. Before describing the proposed approach, we make observations and lay out the motivation of our design based on the limitations of each feature as follows.

- **Viewing behavior:** In practice, there are usually multiple programs broadcast on the same channel at the same time slot; in this case, these programs are given the same matching score for a user in terms of his or her viewing behavior. Thus, recommendation that is based solely on user viewing behavior chooses all the programs from a certain channel and time slot.³ However, in a real-world scenario, it is unlikely that a user at a given time slot watches more than one TV program, especially for short time slots;⁴ in this case recommending multiple programs from the same channel at a given time slot could lead to poor recommendation quality.
- **Viewing preferences:** Although user preferences are useful for recommendation, recommending linear TV programs based solely thereon usually results in low accuracy. For example, if an office worker enjoys watching action movies during the weekend, it is unreasonable to recommend action movies at midnight during weekdays.

Based on the above characteristics and limitations, we propose two-stage ranking to leverage the two features for TV Rec, as detailed in Algorithm 1. Briefly speaking, for each user \mathbf{u} , we propose first ranking the program set I_{test} according to viewing behavior ($s_{\mathbf{u},i'}^b$) (lines 2–6); then, at the second stage (lines 7–15), for those programs broadcast on the same channel at the same weekly time slot, we choose only one program among them according to the user's viewing

³When multiple programs have the same score, we assign a higher rank to programs with earlier starting times.

⁴In the experiments, we adopted 15 minutes as our time slot interval, an optimal setting for using only viewing behavior for recommendation; even in this case, each time slot nevertheless contains 1.5 programs on average.

preferences ($s_{u,i'}^p$). Note that we put the model for viewing behavior at the first stage as previous studies indicate that the viewing behavior usually dominates the recommendation performance [12], which is also consistent with the finding in our later experiments. This approach boasts two advantages: 1) it is parameter-free, and 2) it is computationally efficient as only a limited number of preference matching scores $s_{u,i'}^b$ are computed at the second stage. Thus, the computational cost of the proposed two-stage ranking method is only slightly higher than for recommendation based solely on viewing behavior; this is also discussed in later experiments.

Algorithm 1: Two-stage Ranking

Input: $\mathcal{A}_{\text{train}}, I_{\text{train}}, I_{\text{test}}, k, \mathbf{u}$
Output: \hat{I}_{test}^u (set consisting of recommended programs in I_{test} for user \mathbf{u})

- 1 $\mathcal{S}^b \leftarrow []; \mathcal{S}^p \leftarrow []; \hat{I}_{\text{test}}^u \leftarrow []$
- 2 Construct \mathcal{B}^u with Eq. (1)
- 3 **for each** i' **in** I_{test} **do**
- 4 Compute $s_{u,i'}^b$ and (w, c) with Eq. (2)
- 5 $\mathcal{S}^b.\text{append}((i', (w, c), s_{u,i'}^b))$
- 6 Sort \mathcal{S}^b in ascending order according to $s_{u,i'}^b$
- 7 **while** $(|\hat{I}_{\text{test}}^u| < k)$ **do**
- 8 $(i', (w, c), s_{u,i'}^b) \leftarrow \mathcal{S}^b.\text{pop}()$
- 9 Compute $h_{w_j, u}$ (or h_u), $h_{i'}$ and $s_{u,i'}^p$ with Eqs. (3)–(7)
- 10 **if** $\mathcal{S}^p \neq \emptyset$ **and** $(w, c) \neq (w_0, c_0)$ **then**
- 11 Sort \mathcal{S}^p in ascending order according to $s_{u,i'}^p$
- 12 $\hat{I}_{\text{test}}^u.\text{append}(\mathcal{S}^p.\text{pop}())$
- 13 $\mathcal{S}^p \leftarrow []$
- 14 $(w_0, c_0) \leftarrow (w, c)$
- 15 $\mathcal{S}^p.\text{append}((i', s_{u,i'}^p))$
- 16 **return** \hat{I}_{test}^u

3 EXPERIMENT

3.1 Dataset and Preprocessing

We collected user viewing logs, denoted as D_{raw} , from a set of set-top boxes providing linear television service to end users in Japan from Jan 1, 2019 to June 1, 2019. This period comprises a total of 42,301 unique users and 875,550 distinct programs (denoted as I_{raw}), where each user was anonymized using a hashed ID. Each log records a channel-switching event for a user, denoted as $d = (u, i, c, t, \Delta t)$, indicating that user u switched to channel c broadcasting program i at UTC timestamp t . Above, Δt is the interval between channel-switching events, which can be considered as the duration of the user’s viewing of the program. Note that each program was broadcast only once on a channel in the linear TV system. In addition, each program $i \in I_{\text{raw}}$ was associated with its meta information (see Definition 2.3).

Given these data logs D_{raw} and TV programs I_{raw} , we first removed viewing logs whose duration was less than Δt_θ (e.g., 15 minutes in the experiments) to filter out logs where users were just flipping channels rather than watching a program. Formally, we constructed the preprocessed data logs $D = \{d = (u, i, c, t, \Delta t) | d \in D_{\text{raw}} \wedge \Delta t \geq \Delta t_\theta\}$. We then generated training and testing sets by splitting the processed data logs D based on a timestamp t_{split} and extracting the logs of period $T_{\text{train}} = [t_{\text{split}} - \Delta t_{\text{train}}, t_{\text{split}})$ for training (denoted as D_{train}) and $T_{\text{test}} = [t_{\text{split}}, t_{\text{split}} + \Delta t_{\text{test}})$ for testing (D_{test}); thus $I_{\text{train}} = \{i | i \in I_{\text{raw}}, \mathbf{s}_i \in T_{\text{train}}\}$ and $I_{\text{test}} = \{i | i \in I_{\text{raw}}, \mathbf{s}_i \in T_{\text{test}}\}$. In our experiments, we constructed four datasets with different values for t_{split} and set $\Delta t_{\text{train}}, \Delta t_{\text{test}}$ to 90 and 7 days, respectively. Table 1

Dataset	t_{split}	$ D_{\text{train}} $	$ I_{\text{train}} $	$ C $	$ U $	$ I_{\text{test}} $	$\overline{ I_{\text{test}}^u }$
1	APR. 01, 2019	37,859,993	257,370	173	34,392	31,556	53.45
2	APR. 08, 2019	38,212,364	259,514	174	34,129	32,504	55.47
3	APR. 15, 2019	38,335,769	260,466	174	33,803	32,773	55.31
4	APR. 22, 2019	38,415,448	261,212	177	33,817	33,811	55.24

Table 1. Data statistics

contains the dataset statistics. With user logs in D_{train} , the interaction tensor is $\mathcal{A}_{\text{train}} = (a_{\mathbf{u},i,w,c}) \in \mathbb{R}^{|U| \times |I_{\text{train}}| \times |W| \times |C|}$, where $a_{\mathbf{u},i,w,c} = \sum_{(u,i,c,t,\Delta t) \in D_{\text{train}}} \mathbb{1}\{(u,i,w_T(t),c) = (\mathbf{u},i,w,c)\}$. Here we consider only user sets U appearing at least once both in D_{train} and D_{test} . The length of each weekly time slot $w_i \in W$ was set to 15 minutes by setting n to 672. For validation, we adopted user-implicit feedback extracted from I_{test} ; that is, for each user $\mathbf{u} \in U$, we constructed program set $I_{\text{test}}^{\mathbf{u}} = \{i \mid i \in I_{\text{test}} \wedge (\mathbf{u}, i, c, t, \Delta t) \in D_{\text{test}}\}$ as our ground truth.

3.2 Baselines and Experimental Setup

We first built two baselines based on viewing behavior and viewing preferences, the user characteristics introduced in Sections 2.2.1 and 2.2.2, respectively. Note that for viewing preferences, we tokenized the textual information of each program using MeCab,⁵ after which we used the term frequency-inverse document frequency (tf-idf) vectorizer as the text encoder (see $\mathcal{E}(\cdot)$ in Eq. (3)) to represent items in $I_{\text{train}} \cup I_{\text{test}}$.

In addition, we compared the proposed two-stage ranking approach with a ranking fusion method that combines the recommendations from the above two baselines using reciprocal rank fusion (RRF) [5]. In information retrieval (IR), RRF is a simple but effective method for combining document rankings from multiple IR systems. Formally speaking, given a set of items I_{test} and a set of ranking functions \mathcal{K} , where each $\kappa \in \mathcal{K}$ is a function mapping item $i \in I_{\text{test}}$ to its ranking $\kappa(i)$, the fusion score for each item i is computed as $s_{\text{RRF}}(i) = \sum_{\kappa \in \mathcal{K}} \frac{1}{\kappa(i) + \eta}$, where η is a hyperparameter to reduce the impact of high-ranking items from any of the systems. With the two ranking functions based on viewing behavior and preferences (denoted as κ_b and κ_p , respectively), we have $s_{\text{RRF}}(i) = \frac{1}{\kappa_b(i) + \eta} + \frac{1}{\kappa_p(i) + \eta}$. Another baseline is an RRF variant with an additional hyperparameter ξ to control the impact of two ranking systems, $s_{\text{RRF}}^{\xi}(i) = \frac{\xi}{\kappa_b(i) + \eta} + \frac{1 - \xi}{\kappa_p(i) + \eta}$.

We use the following metrics to evaluate our models: (1) nDCG, (2) precision, and (3) recall. For each user $\mathbf{u} \in U$, we recommend $k = 30$ programs among I_{test} and evaluate model performance with cut-offs $N \in \{10, 20, 30\}$. To fine-tune the hyperparameters for the RRF fusion methods (denoted as RRF and RRF $^{\xi}$), we randomly selected 10% of the users in Dataset 1 as the development set and searched η and ξ in the range of $\{1, 2, \dots, 100\}$ and $\{0, 0.1, \dots, 1\}$, respectively, for the best performance in terms of Recall@30. Additionally, to examine the efficiency of each model, we evaluated each model’s CPU time cost for inference (seconds/user).⁶ For models using viewing preferences (including fusion methods), we computed and indexed $h_{\mathbf{u},w}$ (or $h_{\mathbf{u}}$) and $h_{i'}$ in advance; thus, for each user at the inference stage, the computation cost is mainly associated with the dot product between $h_{\mathbf{u},w}$ (or $h_{\mathbf{u}}$) and $h_{i'}$ (for all programs $i' \in I_{\text{test}}$). In modeling the viewing behavior, the time cost results are primarily due to the construction of matrix $\mathcal{B}^{\mathbf{u}}$ and the calculation of $s_{\mathbf{u},i'}^b$.

3.3 Quantitative Results

Table 2 compares model performance in terms of the aforementioned metrics and inference time. In the table, the best result for each column is in boldface; † denotes statistical significance at $p < 0.05$ (paired t -test over four datasets) with

⁵<https://taku910.github.io/mecab/>

⁶As the inference time is measured on a per-user basis, the number of threads does not impact the measurement.

		N = 10			N = 20			N = 30			
	Time-aware	nDCG	Prec.	Recall	nDCG	Prec.	Recall	nDCG	Prec.	Recall	Time
Behavior		35.25	33.79	12.26	34.68	30.42	18.39	34.25	27.97	22.91	† 0.27
Preferences	✓	13.79	12.91	4.61	13.96	12.13	7.63	14.11	11.45	9.98	1.45
		43.82	38.27	12.89	38.87	30.30	17.97	36.41	26.01	21.59	1.45
RRF	✓	45.99	40.64	13.15	41.02	32.58	18.72	38.25	27.82	22.43	1.46
		45.44	39.93	13.69	41.78	33.53	19.66	39.80	29.60	23.83	1.45
Fusion RRF ^ξ	✓	47.79	41.90	13.93	43.35	34.65	19.86	41.13	30.53	24.11	1.46
		46.32	40.92	13.61	42.61	34.45	19.23	40.54	30.44	23.27	0.30
Two-stage	✓	† 48.92	† 43.28	14.12	† 44.90	† 36.41	19.98	† 42.64	† 32.13	24.20	0.31

Table 2. Recommendation performance. ✓ denotes methods using time-aware user preferences, and † denotes statistical significance at $p < 0.05$.

respect to all other methods, and ✓ indicates methods using time-aware user preferences (i.e., $h_{\mathbf{u},\mathbf{w}}$ in Eq. (6)) as opposed to global preferences (i.e., $h_{\mathbf{u}}$ in Eq. (4)).

First, the comparison between the methods using only behavior or preferences (denoted as Behavior and Preferences, respectively, in the table and hereafter) is strong evidence that in the TV Rec scenario, user viewing behavior dominates recommendation performance, which underscores the importance of putting the model for viewing behavior at the first stage of the proposed two-stage ranking approach. In addition, note that the inference time cost of Behavior is five times less than that of Preferences. On the other hand, as demonstrated in the table, fusing the two user characteristics significantly boosts ranking performance. Specifically, RRF outperforms Behavior in terms of nDCG and Precision by over 7% in the low cut-off regions (i.e., $N = \{10, 20\}$). Tuning the impact of Behavior and Preferences (the second row of RRF^ξ with $\xi = 0.6$) further improves overall ranking performance in terms of nDCG and precision by over 10% and recall by over 5%. However, both RRF and RRF^ξ include exhaustive dot product computation over all programs in I_{test} , resulting in a time cost per user approximately equal to that of Preferences.

Table 2 shows that the proposed two-stage ranking consistently outperforms other (fusion) methods in terms of efficiency and the three evaluation metrics. Specifically, the method significantly surpasses the strongest baseline RRF^ξ by over 2% in terms of nDCG and precision when modeling user preferences both globally and in a time-dependent fashion; also note that time-aware preferences better capture user viewing preferences and thus yield better performance. Most importantly, from an efficiency perspective, the time cost of the two-stage ranking shown in the table is much lower than that of the two fusion methods and is approximate to that of Behavior, because in our method, only a limited number of preference matching scores involving the dot product operation are computed at the second stage. Combining such efficiency and the fact that our method is parameter-free, we conclude that the proposed method is much more practical than RRF-based methods.

4 CONCLUSION

We propose a two-stage ranking approach to fuse two user characteristics—viewing behavior and viewing preferences—in a unified manner for TV Rec. The empirical results on a real-world TV dataset show that our proposed approach consistently outperforms other baseline methods; more importantly, our two-stage ranking approach is parameter-free and efficient at inference, making it applicable and practical to real-world TV Rec scenarios.

REFERENCES

- [1] Kamal Ali and Wijnand Van Stam. 2004. TiVo: Making Show Recommendations Using a Distributed Collaborative Filtering Architecture. In *Proc. of SIGKDD*. 394–401.
- [2] Liliana Ardissono, Cristina Gena, Pietro Torasso, Fabio Bellifemine, Angelo Difino, and Barbara Negro. 2004. User Modeling and Recommendation Techniques for Personalized Electronic Program Guides. In *Personalized Digital Television*. 3–26.
- [3] Chih-Ming Chen, Chuan-Ju Wang, Ming-Feng Tsai, and Yi-Hsuan Yang. 2019. Collaborative Similarity Embedding for Recommender Systems. In *Proc. of WWW*. 2637–2643.
- [4] Szu-Yu Chou, Yi-Hsuan Yang, Jyh-Shing Roger Jang, and Yu-Ching Lin. 2016. Addressing Cold Start for Next-Song Recommendation. In *Proc. of RecSys*. 115–118.
- [5] Gordon V. Cormack, Charles L A Clarke, and Stefan Buettcher. 2009. Reciprocal Rank Fusion Outperforms Condorcet and Individual Rank Learning Methods. In *Proc. of SIGIR*. 758–759.
- [6] Paul Cotter and Barry Smyth. 2000. Ptv: Intelligent Personalised TV Guides. In *Proc. AAAI/IAAI*. 957–964.
- [7] Yolanda Blanco Fernández, José J Pazos Arias, Martín López Nores, Alberto Gil Solla, and Manuel Ramos Cabrer. 2006. AVATAR: An Improved Solution For Personalized TV Based on Semantic Inference. *IEEE Transactions on Consumer Electronics* 52, 1 (2006), 223–231.
- [8] Carlos A. Gomez-Uribe and Neil Hunt. 2016. The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Transaction Management Information System* (Dec. 2016), Article No. 19.
- [9] Shang H. Hsu, Ming-Hui Wen, Hsin-Chieh Lin, Chun-Chia Lee, and Chia-Hoang Lee. 2007. AIMED- A Personalized TV Recommendation System. In *Proc. of EuroITV*. 166–174.
- [10] Noo-Ri Kim, Sungtak Oh, and Jee-Hyong Lee. 2018. A Television Recommender System Learning a User’s Time-Aware Watching Patterns Using Quadratic Programming. *Applied Sciences* 8, 8 (2018), 1323.
- [11] Barry Smyth and Paul Cotter. 1999. Surfing the Digital Wave. In *Proc. of ICCBR*. 561–571.
- [12] Roberto Turrin, Andrea Condorelli, Paolo Cremonesi, and Roberto Pagano. 2014. Time-based TV Programs Prediction. In *Proc. of RecSys Workshops for Television and Online Video*.
- [13] Jheng-Hong Yang, Chih-Ming Chen, Chuan-Ju Wang, and Ming-Feng Tsai. 2018. HOP-Rec: High-Order Proximity for Implicit Recommendation. In *Proc. of RecSys*. 140–144.
- [14] Chenguang Yu, Hao Ding, Houwei Cao, Yong Liu, and Can Yang. 2017. Follow Me: Personalized IPTV Channel Switching Guide. In *Proc. of MMSys*. 147–157.
- [15] David Zibriczky, Balázs Hidasi, Zoltán Petres, and Domonkos Tikk. 2012. Personalized Recommendation of Linear Content on Interactive TV Platforms: Beating the Cold Start and Noisy Implicit User Feedback. In *Proc. of UMAP Workshops*.