

Personalized Multimedia Item and Key Frame Recommendation

Le Wu¹, Lei Chen¹, Yonghui Yang¹, Richang Hong¹,
Yong Ge², Xing Xie³, Meng Wang¹

¹Hefei University of Technology

²The University of Arizona

³Microsoft Research

{lew.ustc, chenlei.hfut, yyh.hfut, hongrc.hfut}@gmail.com, yongge@email.arizona.edu,
xingx@microsoft.com, eric.mengwang@gmail.com

Abstract

When recommending or advertising items to users, an emerging trend is to present each multimedia item with a *key frame* image (e.g., the poster of a movie). As each multimedia item can be represented as multiple fine-grained visual images (e.g., related images of the movie), personalized key frame recommendation is necessary in these applications to attract users' unique visual preferences. However, previous personalized key frame recommendation models relied on users' fine-grained image behavior of multimedia items (e.g., user-image interaction behavior), which is often not available in real scenarios. In this paper, we study the general problem of joint multimedia item and key frame recommendation in the absence of the fine-grained user-image behavior. We argue that the key challenge of this problem lies in discovering users' visual profiles for key frame recommendation, as most recommendation models would fail without any users' fine-grained image behavior. To tackle this challenge, we leverage users' item behavior by projecting users (items) in two latent spaces: a collaborative latent space and a visual latent space. We further design a model to discern both the collaborative and visual dimensions of users, and model how users make decisive item preferences from these two spaces. As a result, the learned user visual profiles could be directly applied for key frame recommendation. Finally, experimental results on a real-world dataset clearly show the effectiveness of our proposed model on the two recommendation tasks.

1 Introduction

Living in a digital world with overwhelming information, the visual based content, e.g., pictures, and images, is usually the most eye-catching for users and convey specific views to users [Zhang *et al.*, 2018; Yin *et al.*, 2018;

Chen *et al.*, 2019]. Therefore, when recommending or advertising multimedia items to users, an emerging trend is to present each multimedia item with a display image, which we call the *key frame* in this paper. E.g., as shown in Fig. 1, the movie recommendation page usually displays each movie with a poster to attract users' attention. Similarly, for (short) video recommendation, the title cover is directly presented to users to quickly spot the visual content. Besides, image-based advertising promotes the advertising item with an attractive image, with 80% of marketers use visual assets in the social media marketing [Stelzner, 2018].

As a key frame distills the visual essence of a multimedia item, key frame extraction and recommendation for multimedia items is widely investigated in the past [Mundur *et al.*, 2006; Zhang *et al.*, 2013; Lee *et al.*, 2011]. Some researchers focused on how to summarize representative video content from videos, or apply image retrieval models from text descriptions for universal key frame extraction [Mundur *et al.*, 2006; Wan *et al.*, 2014]. In the real-world, users' visual preferences are not the same but vary from person to person [Matz *et al.*, 2017; Wu *et al.*, 2019]. By collecting users' behaviors to the images of the multimedia items, many recommendation models could be applied for personalized key frame recommendation [Adomavicius and Tuzhilin, 2005; Rendle *et al.*, 2009; Covington *et al.*, 2016]. Recently, researchers have made one of the first few attempts to design a computational KFR model that is tailored for personalized key frame recommendation in videos [Chen *et al.*, 2017b]. By exploring the time-synchronized comment behavior from video sharing platforms, the authors proposed to encode both users' interests and the visual and textual features of the frame in a unified multi-modal space to enhance key frame recommendation performance.

Despite the preliminary attempts for personalized key frame recommendation, we argue that these models not well designed due to the following two reasons. First, as the key frame is a visual essence of a multimedia item, recommending a key frame is always associated with the corresponding recommended multimedia item. Therefore, is it possible to design a model that simultaneously recommends both items and personalized key

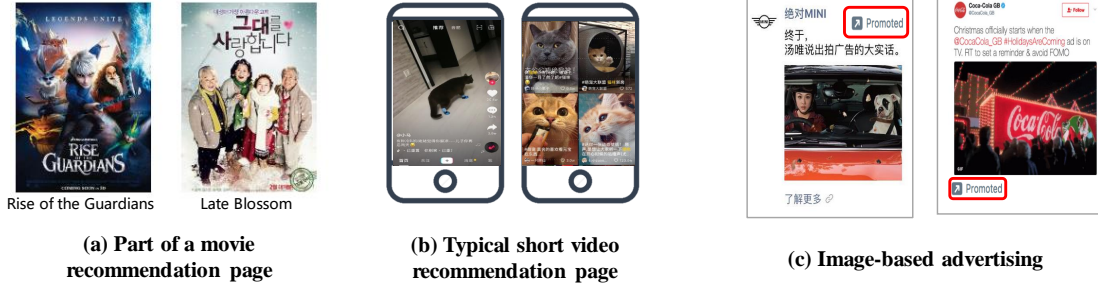


Figure 1: Three application scenarios with the key frame presentation of a multimedia item.

frames? Second, most of the current key frame recommendation models relied on the fine-grained frame-level user preference in the modeling process, i.e., users' behaviors to the frames of the multimedia item, which is not available for most platforms. Specifically, without any users' frame behavior, the classical collaborative filtering based models fail as these models rely on the user-frame behavior for recommendation [Rendle *et al.*, 2009; Koren *et al.*, 2009]. The content based models also could not work as these models need users' actions to learn the visual dimensions of users [Hu *et al.*, 2018; Lei *et al.*, 2016]. Therefore, the problem of how to build the user visual profile for key frame recommendation when the fine-grained user-image behavior data is not available remains challenging.

In this paper, we study the general problem of personalized multimedia item and key frame recommendation without fine-grained user-image interaction data. The key idea of our proposed model is to distinguish each user's latent collaborative preference and the visual preference from her multimedia item interaction history, such that each user's visual dimensions could be transferred for visual based key frame recommendation. We design a *Joint Item and key Frame Recommendation (JIFR)* model to discern both the collaborative and visual dimensions of users, and model how users make item preference decisions from these two spaces. Finally, extensive experimental results on a dataset clearly show the effectiveness of our proposed model for the two personalized recommendation tasks.

2 Problem Definition

In a multimedia item recommender system, there are a userset \mathcal{U} ($|\mathcal{U}| = M$), and an itemset \mathcal{V} ($|\mathcal{V}| = N$). Each multimedia item v is composed of many *frames*, where each frame is an image that shows a part of the visual content of this item. Without confusion, we would use the term of frame and image interchangeably in this paper. E.g., in a movie recommender system, each movie is a multimedia item. There are many images that could describe the visual content of this movie, e.g., the official posters in different countries, the trailer posters, and the frames contained in this video. For video recommendation, it is hard to directly analyze the video content from frame to frame, a natural practice is to extract several typical frames that summarize the content of this video [Mundur *et al.*, 2006]. Besides, for image-

based advertising, for each advertising content with text descriptions, many related advertising images (frames) could be retrieved with text based image retrieval techniques [Wan *et al.*, 2014].

Therefore, besides users and items, all the frames of the multimedia items in the itemset compose a frameset \mathcal{C} ($|\mathcal{C}| = L$). The relationships between items and frames are represented as a *frame-item correlation matrix* $\mathbf{S} \in \mathbb{R}^{L \times N}$. We use $S_v = [j | s_{jv} = 1]$ to denote the frames of a multimedia item v . For each item v , the *key frame* is the display image of this multimedia item when it is presented or recommended to users. Therefore, the key frame belongs to v 's frameset S_v . Besides, in the multimedia system, users usually show implicit feedbacks (e.g., watching a movie, clicking an advertisement) to items, which could be represented as a *user-item rating matrix* $\mathbf{R} \in \mathbb{R}^{M \times N}$. In this matrix, r_{ai} equals 1 if a shows preferences for item i , otherwise it equals 0.

Definition 1 [Multimedia Item and Key Frame Recommendation] In a multimedia recommender system, there are three kinds of entities: a userset \mathcal{U} , an itemset \mathcal{V} , and a frameset \mathcal{C} . The item multimedia content is represented as a frame-item correlation matrix \mathbf{S} . Users show item preference with user-item implicit rating matrix \mathbf{R} . For visual based multimedia recommendation, our goal is two-fold: 1) Multimedia Item Recommendation: Predict each user a 's unknown preference \hat{r}_{ai} to multimedia item i ; 2) Key Frame Recommendation: For user a and the recommended multimedia item i , predict her unknown fine-grained preference to multimedia content \hat{l}_{ak} , where k is a multimedia image content of i ($s_{ki} = 1$).

3 The Proposed Model

In this section, we introduce our proposed JIFR model for multimedia item and key frame recommendation. We start with the overall intuition, followed by the detailed model architecture and the model training process. At the end of this section, we analyze the proposed model.

With the absence of the fine-grained user behavior data, it is natural to leverage the user-item rating matrix \mathbf{R} to learn each user's profile for key frame recommendation. Therefore, in each user's decisive process for multimedia items, we adopt a hybrid recommendation model that projects users and multimedia items into two spaces: a latent space to characterize the collaborative

behavior, and a visual space that shows the visual dimensions that influences users' decisions. Let $\mathbf{U} \in \mathbb{R}^{d1 \times M}$ and $\mathbf{V} \in \mathbb{R}^{d1 \times N}$ denote the free user and item latent vectors in the collaborative space, and $\mathbf{W} \in \mathbb{R}^{d2 \times M}$ and $\mathbf{X} \in \mathbb{R}^{d2 \times N}$ are the visual representations of users and items in the visual dimensions. For the visual dimension construction, as each multimedia item i is composed of multiple images, its visual representation \mathbf{x}_i is summarized from the related visual embeddings of the corresponding frame set S_i as:

$$\mathbf{x}_i = \sum_{k \in S_i} \frac{\mathbf{Pc}_k}{|S_i|}, \quad (1)$$

where \mathbf{c}_k is the visual representation of image k . Due to the success of convolutional neural networks, similar as many visual modeling approaches, we use the last fully connected layer in VGG-19 to represent the visual content of each image k as $\mathbf{c}_k \in \mathbb{R}^{4096}$ [Simonyan and Zisserman, 2015; He and McAuley, 2016]. \mathbf{Pc}_k transforms the original item visual content representation from 4096 dimensions into a low latent visual space, which is usually less than 100 dimensions.

Given the multimedia item representation, each user's predicted preference could be seen as a hybrid preference decision process that combines the collaborative filtering preference and the visual content preference as:

$$\hat{r}_{ai} = \mathbf{u}_a^T \mathbf{v}_i + \mathbf{w}_a^T \mathbf{x}_i, \quad (2)$$

where \mathbf{w}_a is the visual embedding of user a from the user visual matrix \mathbf{W} . In fact, by summarizing the item visual content from its related frames, the above equation is similar to the VBPR model that uncovers the visual and latent dimensions of users [He and McAuley, 2016].

With the implicit feedbacks of the rating matrix \mathbf{R} , Bayesian Personalized Ranking (BPR) is widely used for modeling the pair-wise based optimization function [Rendle *et al.*, 2009]:

$$\min \mathcal{L}_R = \sum_{a=1}^M \sum_{(i,j) \in D_a^R} \sigma(\hat{r}_{ai} - \hat{r}_{aj}) + \lambda_1 \|\Theta_1\|_F^2, \quad (3)$$

where $\Theta_1 = [\mathbf{U}, \mathbf{V}, \mathbf{W}]$ is the parameter set, and $\sigma(x)$ is a sigmoid function that transforms the output into range (0, 1). The training data for user a is $D_a^R = \{(i, j) | i \in R_a \wedge j \in V - R_a\}$, where R_a denotes the set of implicit positive feedbacks of a (i.e., $r_{ai}=1$), and $j \in V - R_a$ is an unobserved feedback.

In key frame decision process, each key frame image presentation of the current item is the foremost visual impression to influence and persuade users. By borrowing the learned user visual representation matrix \mathbf{W} from the user-item interaction behavior (Eq.(3)), each user's visual preference for image k is predicted as:

$$\hat{l}_{ak} = \mathbf{w}_a^T (\mathbf{Pc}_k), \quad (4)$$

where \mathbf{w}_a is the visual latent embedding of user a learned from user-item interaction behavior. Please note that, the predicted \hat{l}_{ak} is only used in the test data for the key frame recommendation without any training process.

Under the above approach, for each user, by optimizing the user-item based loss function (Eq.(3)), we could

align users and images in the visual space without any user-image interaction data for multimedia item and key frame recommendation. However, we argue that the above approach is not well designed for the proposed problem due to the overlook of the *content indecisiveness* and *rating indecisiveness* in the modeling process. The content indecisiveness is correlated to the visual representation of each item (Eq.(1)), which refers to which images are more important to reflect the visual content of the multimedia item are unknown. E.g., some frames in the movie convey more visual semantics than other frames that are not informative. Simply using an average operation that summarizes the item visual representation (Eq.(1)) would neglects the visual highlights of the item semantics. Besides, the rating indecisiveness appears in each user-item preference decision process as shown in Eq.(2), which refers to the implicitness of whether to concentrate more on the collaborative part or the visual item part for the preference decision process. For example, sometimes a user chooses a movie since this movie is visually stunning, even though this movie does not follow her historical watching histories. Therefore, the recommendation performance is limited by the assumption that the hybrid preference is equally contributed by the collaborative and visual content based models as shown in Eq.(2).

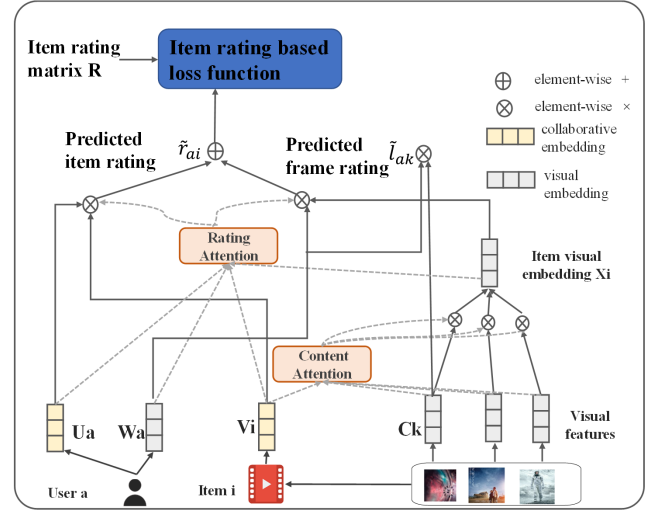


Figure 2: The framework of our proposed JIFR model.

3.1 The Proposed JIFR Model

In this part, we illustrate our proposed *Joint Item and Key Frame Recommendation* (JIFR) model, with the architecture is show in Fig. 2. The key idea of JIFR are two carefully designed attention networks for dealing with the content indecisiveness and rating indecisiveness. Specifically, to tackle the content indecisiveness, the visual content attention attentively learns the visual representation \mathbf{x}_i of each item. By taking the user-item predicted collaborative rating $\mathbf{u}_a^T \mathbf{v}_i$, and the visual content based rating $\mathbf{w}_a^T \mathbf{x}_i$, the rating attention module learns to attentively combine these two kinds of predictions for rating indecisiveness problem.

Visual Content Attention. For each multimedia item i , the goal of the visual attention is to select the frames from the frameset S_i that are representative for item visual representation. Therefore, instead of simply averaging all the related images' visual embeddings as the item visual dimension (Eq.(1)), we model the item visual embedding as:

$$\mathbf{x}_i = \sum_{j=1}^L \alpha_{ij} s_{ji} \mathbf{P} \mathbf{c}_j, \quad (5)$$

where $s_{ji} = 1$ if image j belongs to the image element of multimedia item i . α_{ij} denotes the attentive weight of image j for multimedia item i . The larger the value of α_{ij} , the more the current visual frame contributes to the item visual content representation.

Since α_{ij} is not explicitly given, we model the attentive weight as a three-layered attention neural network:

$$\alpha_{ij} = \mathbf{w}^1 \times f(\mathbf{W}^1[\mathbf{v}_i, \mathbf{W}^0 \mathbf{c}_j]), \quad (6)$$

where $\Theta_c = [\mathbf{w}^1, \mathbf{W}^1, \mathbf{W}^0]$ is the parameter set in this three layered attention network, and $f(x)$ is a non-linear activation function. Specifically, \mathbf{W}^0 is a dimension reduction matrix that transforms the original visual embeddings (i.e., $\mathbf{c}_j \in \mathbb{R}^{4096}$) in a low dimensional visual space.

Then, the final attentive upload history score α_{ji} is obtained by normalizing the above attention scores as:

$$\alpha_{ij} = \frac{\exp(\alpha_{ij})}{\sum_{k=1}^L s_{ki} \exp(\alpha_{ik})}. \quad (7)$$

Hybrid Rating Attention. The hybrid rating attention part models the importance of the collaborative preference and the content based preference for users' final decision making as:

$$\hat{r}_{ai} = \beta_{ai1} \mathbf{u}_a^T \mathbf{v}_i + \beta_{ai2} \mathbf{W}_a^T \mathbf{x}_i, \quad (8)$$

where the first part models the collaborative predicted rating, and the second part denotes the user's visual preference for items. β_{ai1} and β_{ai2} are the weights that balance these two effects for the user's final preference.

As the underlying reasons for users to balance these two aspects are unobservable, we model the hybrid rating attention as:

$$\beta_{ai1} = \mathbf{w}^2 \times f(\mathbf{W}^2[\mathbf{u}_a, \mathbf{v}_i]), \quad \beta_{ai2} = \mathbf{w}^2 \times f(\mathbf{W}^2[\mathbf{w}_a, \mathbf{x}_i]). \quad (9)$$

Then, the final rating attention values β_{ai1} and β_{ai2} are also normalized as:

$$\beta_{ai1} = \frac{\exp(\beta_{ai1})}{\exp(\beta_{ai1}) + \exp(\beta_{ai2})} = 1 - \beta_{ai2}. \quad (10)$$

Model Learning and Prediction. With the two proposed attention networks, the optimization function is the same as Eq. (3). To optimize the objective function, we implement the model in TensorFlow [Abadi *et al.*, 2016] to train model parameters with mini-batch Adam, which is a stochastic gradient descent based optimization model with adaptive learning rates. For the user-item interaction behavior, we could only observe the

positive feedbacks with huge missing ratings. In practice, similar as many implicit feedback based optimization approaches, in each iteration, for each positive feedback, we randomly sample 10 missing feedbacks as pseudo negative feedbacks in the training process [Chen *et al.*, 2017a; Wu *et al.*, 2017]. As in each iteration, the pseudo negative feedbacks change with each missing record gives very weak negative signal.

After the model learning process, the multimedia recommendation could be directly computed as in Eq.(8). For each recommended item, as users and images are also learned in the visual dimensions, the key frame recommendation could be predicted as in Eq.(4).

3.2 Connections to related models.

VBPR [He and McAuley, 2016] extends the classical collaborative filtering model with the additional visual dimensions of users and items (Eq.(4)). By assigning the same weight for all of an items' frames as shown in Eq (1), and without any hybrid rating attention, our proposed item recommendation task degenerates to VBPR.


ACF [Chen *et al.*, 2017a] is proposed to combine each user's historical rated items and the item components with attentive modeling on top the CF model of SVD++ [Koren *et al.*, 2009]. ACF did not explicitly model users in the visual space, and could not be transferred for visual key frame recommendation.

VPOI [Wang *et al.*, 2017] is proposed for POI recommendation by leveraging the uploaded images of users at a particular POI. In VPOI, each item has a free embedding, and the relationship between images and POIs are used as side information in the regularization terms.


KFR [Chen *et al.*, 2017b] is the one of the first attempts for personalized key frame recommendation. As KFR relied on users' fine-grained interaction data with frames, it fails when the user-frame interaction data is not available. Besides, KFR is not designed for item recommendation at the same time.

Attention Mechanism is also closely related to our modeling techniques. Attention has been widely used in recommendation, such as the importance of historical items in CF models [He *et al.*, 2018; Chen *et al.*, 2017b], the helpfulness of review for recommendation [Chen *et al.*, 2018; Chen *et al.*, 2017b; Cheng *et al.*, 2018], the strength of social influence for social recommendation [Sun *et al.*, 2018]. In particular, ACCM is an attention based hybrid item recommendation model. It also fuses users and items in the collaborative space and the content space for recommendation [Shi *et al.*, 2018]. However, the content representation of users and items rely on the features of both users and items. As it is sometimes hard to collect user profiles, this model could not be applied when users do not have any features, including our proposed problem.

In summary, our proposed model differs greatly from these related models as we perform both joint personalized item and key frame recommendation at the same time. The application scenario has rarely been studied before. From the technical perspective, we carefully de-




(a) HR@K




(b) NDCG@K

Figure 3: Item recommendation performance (Better viewed in color).



(a) HR@K



(b) NDCG@K

Figure 4: Key frame recommendation performance.

sign two levels of attentions for dealing with content indecisiveness and rating indecisiveness, which is tailored to discern the visual profiles of users for joint item and key frame recommendation.

4 Experiments

We conduct experiments on a real-world dataset. To the best of our knowledge, there is no public available datasets with fine-grained user behavior data for evaluating the key frame recommendation performance. To this end, we crawl a large dataset from *Douban*¹, which is one of the most popular movie sharing websites in China. We crawl this dataset as for each movie, this platform allows users to show their preference to each frame of this movie by clicking the “Like” button just below each frame.

After data crawling, we pre-process the data to ensure each user and each item have at least 5 rating records. In data splitting process, we randomly select 70% user-movie ratings for training, 10% for validation and 20% for test. The pruned dataset has 16,166 users, 12,811 movies, 379,767 training movie ratings, 98,425 test movie ratings, and 4,760 test frame ratings. For each user-item record in the test data, if the user has rated the images of this multimedia item, the correlated user-image records are used for evaluating the key frame recommendation performance in the test data. Please note that, to make the proposed model general to the multimedia recommendation scenarios, the fine-grained image ratings in the training data is not used for model learning.

4.1 Overall Performance

We adopt two widely used evaluation metrics to measure the top-K ranking performance: the Hit Ra-

Model	Input		Task	
	Rating	Image	Item Rec	Frame Rec
BPR [Rendle <i>et al.</i> , 2009]	✓	×	✓	×
CDL [Lei <i>et al.</i> , 2016]	✓	✓	✓	✓
VBPR [He and McAuley, 2016]	✓	✓	✓	×
VPOI [Wang <i>et al.</i> , 2017]	✓	✓	✓	×
ACF [Chen <i>et al.</i> , 2017a]	✓	✓	✓	×
<i>JIFR_NA</i>	✓	✓	✓	✓
<i>JIFR</i>	✓	✓	✓	✓

Table 1: The characteristics of the models.

tio (HR@K) and Normalized Discounted Cumulative Gain (NDCG@K) [He and McAuley, 2016; Sun *et al.*, 2018; Chen *et al.*, 2017a]. In our proposed JIFR model, we choose the collaborative latent dimension size d_1 and the visual dimension size d_2 in the set $[16, 32, 64]$, and find when $d_1 = d_2 = 32$ reaches the best performance. The non-linear activation function $f(x)$ in the attention networks is set as the ReLU function. Besides, the regularization parameter is set in range $[0.0001, 0.001, 0.01, 0.1]$, and $\lambda_1 = 0.001$ reaches the best performance.

For better illustration, we summarize the details of the comparison models in Table 1, where the second and third column shows the input data of each model, and the last two columns show whether this model could be used for item recommendation and key frame recommendation. The last two rows are our proposed models, with JIFR_NA is a simplified version of our proposed model without any attention modeling.

Item Recommendation Performance. In the item recommendation evaluation process, as the multimedia item size is large, for each user, we randomly select 1000 unrated items as negative samples, and then mix them with the positive feedbacks to select the top-K ranking items. To avoid bias at each time, the process is repeated for 10 times and we report the average results. Fig. 3 shows the results with different top-K values. In this figure, some top-K results for CDL are not shown as the performance is lower than the smallest range values of the y-axis. This is because CDL is a visual content based model without any collaborative filtering modeling, while the collaborative signals are very important for enhancing recommendation performance. Please note that, for item recommendation, our simplified model JIFR_NA degenerates to VBPR without any attentive modeling. VPOI and VBPR improves over BPR by

¹www.douban.com




Figure 5: Case study of the key frame recommendation of *Interstellar* for a typical user *a*.

leveraging the visual data. ACF further improves the performance by learning the attentive weights in the user’s history. Our proposed JIFR model achieves the best performance by explicitly modeling users and items in both the latent space and the visual space with two attention networks. For the two metrics, the improvement of NDCG is larger than HR as NDCG considers the ranking position of the hit items.

Key Frame Recommendation Performance. For key frame recommendation, as the detailed user-image information is not available, the models that relied on the collaborative information fails, including *BPR*, *VBPR*, *KFR* [Chen *et al.*, 2017b]. We also show a simplified *RND* baseline that randomly selects a frame from the movie frames for evaluation. In the evaluation process, all the related frames of this movie is used as the candidate key frames. Please note that, the ranking list size K for key frame recommendation is very small as we could only recommend one key frame of each multimedia item, while the item recommendation list could be larger. As observed in Fig 4, all models improve over *RND*, showing the effectiveness of modeling user visual profile. Among all models, our proposed JIFR model shows the best performance, followed by the simplified JIFR_NA model. This clearly shows the effectiveness of discerning the visual profiles and the collaborative interests of users with attentive modeling techniques.

Visual Att	Rating Att	Item Rec@15		Frame Rec@3	
		HR	NDCG	HR	NDCG
AVG	AVG	/	/	/	/
AVG	ATT	2.25%	2.42%	4.20%	4.35%
ATT	AVG	4.02%	4.44%	8.07%	8.62%
ATT	ATT	5.49%	5.94%	9.70%	10.53%

Table 2: Improvement of attention modeling.

4.2 Detailed Model Analysis

Attention Analysis. Table 2 shows the performance improvement of different attention networks compared to the average setting, i.e., $\alpha_{ji} = \frac{1}{\sum_{k=1}^C s_{ki}}$ for content attention modeling, and $\beta_{ai1} = \beta_{ai2}$ for rating fusion. The ranking list value K is set as $K = 15$ for item recommendation and $K = 3$ for key frame recommendation. As can be observed from this table, both attention techniques improve the performance of the two recommendation tasks. On average, the visual attention improvement is larger than the rating attention. By combining these two attention networks, the two recommendation tasks achieve the best performance.

Frame Recommendation Case Study. Fig 5 shows the a case study of the recommended frames for

user *a* with the movie *Interstellar*. It is regarded as a sci-fi that describes a team explorers travel through a wormhole in space to ensure humanity’s survival. In the meantime, the love between the leading actor, and his daughter also touches the audiences. For ease of understanding, we list the training data of this user in the last column, with each movie is represented with an official poster. In the training data, the liked movies contain both sci-fi and love categories. Our proposed JIFR model could correctly recommend the key frame, which is different from the official poster of this movie. However, the remaining models fail. We guess a possible reason is that, as shown in the four column, most frames of this movie are correlated to sci-fi. As the comparison models could not distinguish the important of these frames, the love related visual frame is overwhelmed by the sci-fi visual frames. Our model tackles this problem by learning the attentive frame weights for item visual representation, and user visual representation from her historical movies.

5 Conclusions

In this paper, we studied the general problem of personalized multimedia item and key frame recommendation in the absence of fine-grained user behavior. We proposed a JIFR model to project both users and items into a latent collaborative space and a visual space. Two attention networks are designed to tackle the content indecisiveness and rating indecisiveness for better discerning the visual profiles of users. Finally, extensive experimental results on a real-world dataset clearly showed the effectiveness of our proposed model for the two recommendation tasks.

Acknowledgments

This work was supported in part by grants from the National Natural Science Foundation of China(Grant No. 61725203, 61722204, 61602147, 61732008, 61632007), the Anhui Provincial Natural Science Foundation(Grant No. 1708085QF155), and the Fundamental Research Funds for the Central Universities(Grant No. JZ2018HG TB0230).

References

[Abadi *et al.*, 2016] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.

- [Adomavicius and Tuzhilin, 2005] Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *TKDE*, 17(6):734–749, 2005.
- [Chen *et al.*, 2017a] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. In *SIGIR*, pages 335–344, 2017.
- [Chen *et al.*, 2017b] Xu Chen, Yongfeng Zhang, Qingyao Ai, Hongteng Xu, Junchi Yan, and Zheng Qin. Personalized key frame recommendation. In *SIGIR*, pages 315–324, 2017.
- [Chen *et al.*, 2018] Chong Chen, Min Zhang, Yiqun Liu, and Shaoping Ma. Neural attentional rating regression with review-level explanations. In *WWW*, pages 1583–1592, 2018.
- [Chen *et al.*, 2019] Lei Chen, Le Wu, Zhenzhen Hu, and Meng Wang. Quality-aware unpaired image-to-image translation. *TMM*, 2019.
- [Cheng *et al.*, 2018] Zhiyong Cheng, Ying Ding, Xiangnan He, Lei Zhu, Xuemeng Song, and Mohan S Kankanhalli. A³ncf: An adaptive aspect attention model for rating prediction. In *IJCAI*, pages 3748–3754, 2018.
- [Covington *et al.*, 2016] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Recsys*, pages 191–198, 2016.
- [He and McAuley, 2016] Ruining He and Julian McAuley. Vbpr: Visual bayesian personalized ranking from implicit feedback. In *AAAI*, pages 144–150, 2016.
- [He *et al.*, 2018] Xiangnan He, Zhenkui He, Jingkuan Song, Zhenguang Liu, Yu-Gang Jiang, and Tat-Seng Chua. Nais: Neural attentive item similarity model for recommendation. *TKDE*, 2018.
- [Hu *et al.*, 2018] Liang Hu, Songlei Jian, Longbing Cao, and Qingkui Chen. Interpretable recommendation via attraction modeling: Learning multilevel attractiveness over multimodal movie contents. In *IJCAI*, pages 3400–3406, 2018.
- [Koren *et al.*, 2009] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- [Lee *et al.*, 2011] Yong Jae Lee, Jaechul Kim, and Kristen Grauman. Key-segments for video object segmentation. In *CVPR*, pages 1995–2002, 2011.
- [Lei *et al.*, 2016] Chenyi Lei, Dong Liu, Weiping Li, Zheng-Jun Zha, and Houqiang Li. Comparative deep learning of hybrid representations for image recommendations. In *CVPR*, pages 2545–2553, 2016.
- [Matz *et al.*, 2017] Sandra C Matz, Michal Kosinski, Gideon Nave, and David J Stillwell. Psychological targeting as an effective approach to digital mass persuasion. *PNAS*, 114(48):12714–12719, 2017.
- [Mundur *et al.*, 2006] Padmavathi Mundur, Yong Rao, and Yelena Yesha. Keyframe-based video summarization using delaunay clustering. *International Journal on Digital Libraries*, 6(2):219–232, 2006.
- [Rendle *et al.*, 2009] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, pages 452–461, 2009.
- [Shi *et al.*, 2018] Shaoyun Shi, Min Zhang, Yiqun Liu, and Shaoping Ma. Attention-based adaptive model to unify warm and cold starts recommendation. In *CIKM*, pages 127–136, 2018.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [Stelzner, 2018] Michael Stelzner. 2018 social media marketing industry report. *Social Media Examiner*, pages 1–44, 2018.
- [Sun *et al.*, 2018] Peijie Sun, Le Wu, and Meng Wang. Attentive recurrent social recommendation. In *SIGIR*, pages 185–194, 2018.
- [Wan *et al.*, 2014] Ji Wan, Dayong Wang, Steven Chu Hong Hoi, Pengcheng Wu, Jianke Zhu, Yongdong Zhang, and Jintao Li. Deep learning for content-based image retrieval: A comprehensive study. In *MM*, pages 157–166, 2014.
- [Wang *et al.*, 2017] Suhang Wang, Yilin Wang, Jiliang Tang, Kai Shu, Suhas Ranganath, and Huan Liu. What your images reveal: Exploiting visual contents for point-of-interest recommendation. In *WWW*, pages 391–400, 2017.
- [Wu *et al.*, 2017] Le Wu, Yong Ge, Qi Liu, Enhong Chen, Richang Hong, Junping Du, and Meng Wang. Modeling the evolution of users preferences and social links in social networking services. *TKDE*, 29(6):1240–1253, 2017.
- [Wu *et al.*, 2019] Le Wu, Lei Chen, Richang Hong, Yanjie Fu, Xing Xie, and Meng Wang. A hierarchical attention model for social contextual image recommendation. *TKDE*, 2019.
- [Yin *et al.*, 2018] Yu Yin, Zhenya Huang, Enhong Chen, Qi Liu, Fuzheng Zhang, Xing Xie, and Guoping Hu. Transcribing content from structural images with spotlight mechanism. In *SIGKDD*, pages 2643–2652, 2018.
- [Zhang *et al.*, 2013] Dong Zhang, Omar Javed, and Mubarak Shah. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *CVPR*, pages 628–635, 2013.
- [Zhang *et al.*, 2018] Kun Zhang, Guangyi Lv, Le Wu, Enhong Chen, Qi Liu, Han Wu, and Fangzhao Wu.

Image-enhanced multi-level sentence representation net for natural language inference. In *ICDM*, pages 747–756, 2018.