



**HAL**  
open science

## Do We Really Need to Drop Items with Missing Modalities in Multimodal Recommendation?

Daniele Malitesta, Emanuele Rossi, Claudio Pomo, Tommaso Di Noia,  
Fragkiskos D Malliaros

► **To cite this version:**

Daniele Malitesta, Emanuele Rossi, Claudio Pomo, Tommaso Di Noia, Fragkiskos D Malliaros. Do We Really Need to Drop Items with Missing Modalities in Multimodal Recommendation?. CIKM 2024 - The 33rd ACM International Conference on Information and Knowledge Management, Oct 2024, Boise, ID, United States. pp.3943-3948, 10.1145/3627673.3679898 . hal-04847987

**HAL Id: hal-04847987**

<https://inria.hal.science/hal-04847987v1>

Submitted on 19 Dec 2024

**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.



Distributed under a Creative Commons Attribution 4.0 International License

# Do We Really Need to Drop Items with Missing Modalities in Multimodal Recommendation?

Daniele Malitesta\*  
Université Paris-Saclay  
CentraleSupélec, Inria,  
Centre for Visual Computing (CVN)  
Gif-sur-Yvette, France  
daniele.malitesta@centralesupelec.fr

Emanuele Rossi  
VantAI  
Barcelona, Spain  
emanuele.rossi1909@gmail.com

Claudio Pomo  
Politecnico di Bari  
Bari, Italy  
claudio.pomo@poliba.it

Tommaso Di Noia  
Politecnico di Bari  
Bari, Italy  
tommaso.dinoia@poliba.it

Fragkiskos D. Malliaros  
Université Paris-Saclay  
CentraleSupélec, Inria,  
Centre for Visual Computing (CVN)  
Gif-sur-Yvette, France  
fragkiskos.malliaros@centralesupelec.fr

## ABSTRACT

Generally, items with missing modalities are dropped in multimodal recommendation. However, with this work, we question this procedure, highlighting that it would further damage the pipeline of any multimodal recommender system. First, we show that the lack of (some) modalities is, in fact, a widely-diffused phenomenon in multimodal recommendation. Second, we propose a pipeline that imputes missing multimodal features in recommendation by leveraging traditional imputation strategies in machine learning. Then, given the graph structure of the recommendation data, we also propose three more effective imputation solutions that leverage the item-item co-purchase graph and the multimodal similarities of co-interacted items. Our method can be plugged into any multimodal RSs in the literature working as an untrained pre-processing phase, showing (through extensive experiments) that any data pre-filtering is not only unnecessary but also harmful to the performance.

## CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval; Personalization.**

## KEYWORDS

Multimodal Recommendation, Missing Modalities, Graph Neural Networks

### ACM Reference Format:

Daniele Malitesta, Emanuele Rossi, Claudio Pomo, Tommaso Di Noia, and Fragkiskos D. Malliaros. 2024. Do We Really Need to Drop Items with Missing Modalities in Multimodal Recommendation?. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3627673.3679898>

\*Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '24, October 21–25, 2024, Boise, ID, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0436-9/24/10.  
<https://doi.org/10.1145/3627673.3679898>

'24), October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3627673.3679898>

## 1 INTRODUCTION AND MOTIVATION

Despite the large success of recommender systems (RSs) algorithms built around the paradigm of collaborative filtering [9] (CF), they still face several challenges. Among those, it is acknowledged that RSs struggle when learning from very *sparse* user-item data [12], a widely common setting in real-world applications. In such contexts, the literature suggests to enhance the information carried by the user-item data through additional *side information*, such as *multimodal* data (e.g., images, descriptions, audio tracks) accompanying items in several domains (e.g., fashion [1, 4, 8], music [6, 35, 45], food [18, 34, 49], and micro-video [3, 5, 54] recommendation). Multimodal RSs [28] leverage high-level features extracted from multimodal items' data to inject/process them into the recommendation pipeline, outperforming CF recommendation [12, 52–54, 59, 61, 62].

While the multimodal recommendation pipeline hugely relies on the *quality* of multimodal data to extract meaningful items' features [7, 27, 28], in real-world scenarios, such data may be *noisy* or (even worse) **missing** [29]. Concretely, if we consider the e-commerce scenario, it is not rare to observe inaccessible URLs of product images and unavailable descriptions/reviews on the product webpage. Even assuming that this phenomenon only applies to real-world recommendation data, a simple analysis of a popular dataset for multimodal recommendation (the Amazon Reviews data [33]) reveals that this is also true in academic research (Table 1).

In machine learning, the literature recognizes a whole research line addressing *missing information* [10, 21, 31]. The same applies to RSs, where works have focused on missing user-item interactions [19, 20, 32, 40, 41, 43, 44, 48, 51, 56, 60] and (in very few cases) users' and items' metadata [22, 42]. However, even acknowledging the extensive literature on **missing multimodal information** in deep learning [14, 17, 25, 26, 30, 46, 57, 58], to date, very limited effort has been devoted to this issue in multimodal recommendation. The only work is [47], an **end-to-end** framework that uses modalities correlations and multimodal dropout during training, and a reconstruction autoencoder for the inference.

**Table 1: Datasets statistics, where we show items with missing visual (V) and textual (T) information, before and after they have been dropped (without and with the apex).**

Datasets	$\mathcal{U}$   /   $\mathcal{U}'$	$\mathcal{I}$   /   $\mathcal{I}'$	$\mathcal{R}$   /   $\mathcal{R}'$	Missing	
				V	T
Office	4,905 / 4,891	2,420 / 1,746	53,258 / 35,185	0	674
Music	5,541 / 5,349	3,568 / 2,453	64,706 / 51,516	2	1,114
Beauty	22,363 / 22,293	12,101 / 11,124	198,502 / 165,772	7	977

Nevertheless, in most cases, multimodal RSs pre-process the data by dropping items with missing multimodal information. This procedure inevitably leads to losing further data regarding users and user-item interactions in the system (Table 1), which is completely *counteractive* to the idea of enhancing the *sparse* user-item recommendation data through multimodal information. Hence, the question is: “**Do we really need to drop items with missing modalities in multimodal recommendation?**”. The short answer is **no**, as we empirically demonstrate in this work (Table 2), showing that the performance improvement of multimodal RSs (e.g., VBPR [12]) over their pure CF versions (e.g., BPRMF [38]) can even widen when, as a **pre-processing**, we impute the items’ missing modalities through **untrained** strategies instead of dropping them.

To this end, in this work, we provide the following contributions. (1) We formalize the understudied problem of *missing modalities* in multimodal recommendation. (2) We propose a pipeline that imputes the items’ missing multimodal features acting as an *untrained pre-processing* procedure. (3) Alongside *traditional* machine learning imputations, we propose three novel **graph-aware** imputation methods that leverage the user-item graph, and take inspiration from **node features propagation** [39] and **diffusion** [15] in graph learning. (4) Through extensive experiments, we prove that dropping items with missing modalities in multimodal recommendation is **not only unnecessary but also harmful**; we further support our findings and proposal through an ablation study and hyper-parameter sensitivity analysis to assess the goodness of the proposed graph-aware imputation methods.

## 2 METHODOLOGY

In this section, we present our proposed pipeline where missing multimodal features are imputed in an *untrained pre-processing* manner. First, we formalize the missing modalities setting as observed in the collected recommendation datasets. Then, we describe the imputation methods to recover missing modalities, categorized as *traditional* and our proposed *graph-aware* ones.

### 2.1 Missing modalities in recommendation

**Preliminaries.** In a recommendation system, let  $\mathcal{U}$  and  $\mathcal{I}$  be the sets of users and items, respectively. Then, we indicate with  $\mathbf{R} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$  the user-item interaction matrix, where  $\mathbf{R}_{ui} = 1$  if there exists a recorded implicit feedback of user  $u \in \mathcal{U}$  over item  $i \in \mathcal{I}$ , 0 otherwise. As in any latent factor-based recommendation approach [16], we denote with  $\mathbf{E}_u \in \mathbb{R}^d$  and  $\mathbf{E}_i \in \mathbb{R}^d$  the  $d$ -dimensional embeddings for user  $u \in \mathcal{U}$  and item  $i \in \mathcal{I}$ .

In a multimodal recommendation setting, items’ representation may be suitably enriched through their multimodal features. Thus,

we introduce  $\mathcal{M}$  as the set of modalities, and  $\mathbf{F} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{M}| \times c}$  as the multimodal feature tensor of all items, where  $\mathbf{F}_{im} \in \mathbb{R}^c$  is the feature of item  $i$  accounting for the  $m$  modality.

**Missing modalities.** As already observed in Table 1, the problem of missing modalities in multimodal recommendation implies the unavailability of **some of (or the whole) multimodal feature vectors** for a specific subset of items. For instance, in a double modalities setting (e.g., visual and textual) we might have that item  $i$  is missing either the visual or textual features or both; in every such case, we say item  $i$  has some missing modalities, and indicate the missing multimodal features as  $\mathbf{F}'_{im} \in \mathbb{R}^c$  for every  $m \in \mathcal{M}$  that is missing. Note that the considered setting implies the presence/missingness of the whole feature vector  $\mathbf{F}'_{im} \in \mathbb{R}^c$ .

### 2.2 Modalities imputation methods

**Traditional imputation.** The problem of missing data information in machine learning has been historically discussed and addressed for decades. Following this research line, we adapt three simple and common solutions to impute missing modalities: (i) ZEROS, where a vector of all zero entries replaces missing multimodal features; (ii) RANDOM, where a vector of all random entries replaces missing multimodal features; (iii) GLOBALMEAN, where missing multimodal features are replaced by the mean of available multimodal features.

**Graph-aware imputation.** Differently from other data types, recommendation data is a bipartite and undirected graph of user-item interactions. Thus, we leverage the topology of the user-item recommendation graph to propose more enhanced imputing strategies.

Specifically, we aim to impute missing multimodal features of an item through the existing (i.e., non-missing) multimodal features of the connected items in the **item-item co-purchase graph** at multiple hops. Indeed, our assumption (supported by several works [23, 24, 59, 61]) is that **co-interacted items are likely to present semantically-similar multimodal features**, aligning with the same principle behind collaborative filtering (users sharing similar preferences tend to interact with the same items). The item-item co-interaction matrix is obtained as  $\mathbf{R}^I = \mathbf{R}^T \mathbf{R}$ , where  $\mathbf{R}^I \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{I}|}$  and  $\mathbf{R}^I_{ij}$  is the number of users who interacted with both items  $i$  and  $j$ . We sparsify it through top- $k$  sparsification [59] obtaining  $\bar{\mathbf{R}}^I$  as the item-item sparsified matrix.

Then, inspired by recent works from graph learning [39], we propose to frame the imputation of missing modalities in multimodal recommendation as a process involving the **propagation of nodes multimodal features over the item-item graph**. Specifically, we aim to recover the nodes missing multimodal features through the propagation (at multiple hops) of the existing multimodal features of their neighborhood nodes in the item-item graph.

Let  $\mathbf{F}'_{im}$  be the missing multimodal features for item  $i$  and modality  $m$ , and  $\mathcal{N}_i = \{j \mid \bar{\mathbf{R}}^I_{ij} = 1\}$  be the set of neighborhood nodes of  $i$  in the item-item graph. By leveraging feature propagation from the nearest neighborhood nodes in the item-item graph (i.e., one hop), we extend the GLOBALMEAN imputation through the mean of multimodal features from the one-hop neighborhood (NEIGHMEAN):  $\mathbf{F}'_{im} = \frac{\sum_{j \in \mathcal{N}_i} \mathbf{F}_{jm}}{|\mathcal{N}_i|}$ ,  $\forall m \in \mathcal{M}$ . Since item  $i$  may be linked to items with missing modalities,  $\mathbf{F}_{jm}$  would not be known in advance. For this

reason, before running the propagation, we initialize all multimodal features of items with missing modalities to the zero vector.

As we leverage the item-item graph, we can also consider **multi-hops** relationships between items. Thus, our second graph-aware method **iteratively** performs feature propagation (MULTIHOP):

$$\mathbf{F}'_{im}(t) = \frac{\sum_{j \in \mathcal{N}_i} \mathbf{F}_{jm}^{(t-1)}}{\sqrt{|\mathcal{N}_i|} \sqrt{|\mathcal{N}_j|}}, \forall m \in \mathcal{M}, \forall t \in \{1, \dots, T\}.$$

Here,  $\mathbf{F}_{jm}^{(0)}$  is again initialized to the zero vector for all items with missing modalities. Then, we set  $\mathbf{F}'_{jm}(t-1) = \mathbf{F}_{jm}^{(0)}$  for items with existing modalities as, at each new iteration  $t \in \{1, \dots, T\}$ , we want to re-initialize the multimodal features to their original values [39].

Above, we perform the symmetric Laplacian normalization of the item-item co-interaction matrix to smooth the noisy item-item connections we initially built. While this operation only accounts for the 1-hop neighborhoods, the recent literature has demonstrated that **graph diffusion at multiple hops** [15] may improve graph learning; the idea is to leverage an enhanced representation of the interaction matrix, as personalized PageRank [36].

Hence, we propose a third graph-aware imputation that modifies the MULTIHOP strategy by incorporating **personalized PageRank** as a normalization of the item-item co-interactions matrix. Let  $\mathbf{B} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{I}|}$  be a matrix such that  $\mathbf{B}_{ij} = 1 - (1 - \alpha) / (\sqrt{|\mathcal{N}_i|} \sqrt{|\mathcal{N}_j|})$  if  $i = j$ ,  $\mathbf{B}_{ij} = -(1 - \alpha) / (\sqrt{|\mathcal{N}_i|} \sqrt{|\mathcal{N}_j|})$  otherwise. Then, we have (PERSPAGERANK):  $\mathbf{F}'_{im}(t) = \sum_{j \in \mathcal{N}_i} (\alpha \mathbf{B}^{-1})_{ij} \mathbf{F}_{jm}^{(t-1)}$ ,  $\forall m \in \mathcal{M}, \forall t \in \{1, \dots, T\}$ , where  $\alpha \in [0, 1]$  is the teleport probability of going from an item  $i$  to the 1-hop connected item  $j$  in the item-item graph.

### 3 EXPERIMENTS AND RESULTS

In this section, we report on the experiments and discuss the results. First, we describe the datasets, and baselines, and provide reproducibility details. Then, we answer four research questions: **RQ1.** What is the performance improvement of multimodal recommender systems over their pure collaborative filtering versions in both the dropped and imputed dataset settings? **RQ2.** What is the impact of missing modalities imputation on recent multimodal RSs? **RQ3.** How does each imputation strategy contribute to the performance? **RQ4.** What is the effect of top- $k$  sparsification and the number of propagation hops on the performance?

#### 3.1 Experimental settings

**Datasets.** We select three categories of the popular Amazon Reviews dataset [33], namely: Office Products (**Products**), Digital Music (**Music**), and **Beauty** (Table 1). For each of them, we consider two distinct versions: (i) **dropped**, where we remove items with at least one missing modality; (ii) **imputed**, where we impute the items' missing multimodal features. From the available items' multimodal data (i.e., product images and textual descriptions), we extract [2] visual [11] and textual embeddings [37].

**Baselines.** We select CF and multimodal models: BPRMF [38], NGCF [50], LightGCN [13], SGL [55], VBPR [12], NGCF-M, FREEDOM [61], and BM3 [62]. We introduce NGCF-M as the multimodal version of NGCF where multimodal features are injected in the same manner as VBPR for BPRMF. For the imputation, we use three

*classical machine learning* methods (i.e., ZEROS, RANDOM, GLOBALMEAN), and propose three novel *graph-aware* approaches (i.e., NEIGHMEAN, MULTIHOP, and PERSPAGERANK).

**Reproducibility.** We split the datasets into train and test sets (80%/20%) and run the experiments with [27]. We explore the learning rate in  $\{0.0001, 0.0005, 0.001, 0.005, 0.01\}$  and the regularization term in  $[1e-5, 1e-2]$ , leaving the other parameters to the best values, and fix the batch size at 1024 and the epochs at 200. Then, we explore the top- $k$  sparsification in  $[10, 20, \dots, 100]$  and the propagation hops  $T$  in  $[1, 2, \dots, 20]$ . We use the Recall@20 and the nDCG@20 as test recommendation metrics and retain 10% of the train set as validation, with Recall@20 as a validation metric. The code is at: <https://github.com/sisinflab/Graph-Missing-Modalities>.

#### 3.2 Results and discussion

**Multimodal improvement over CF (RQ1).** Table 2 displays the performance improvement of VBPR and NGCF-M over their CF versions (i.e., BPRMF and NGCF) in the **dropped** and **imputed** recommendation settings. Note that we decide to consider VBPR and NGCF-M for the current analysis as, differently from the other multimodal baselines we will use later, they differ from their respective CF versions for the **sole** injection of items' multimodal features. Thus, we can precisely assess the impact of multimodal features imputation and avoid the risk that other factors (e.g., denoising procedures as in FREEDOM [61]) could inadvertently influence the results. Moreover, since the dropped and imputed versions of each dataset may largely differ (refer again to Table 1), we are mainly interested in the **performance improvement** between CF and multimodal RSs; we decide to report also the performance for all models/datasets pairs only for the sake of completeness.

In the vast majority of cases, we observe that the multimodal version outperforms the CF one. Nevertheless, while such improvement may sometimes not occur in the dropped setting (see Music), the imputed setting **consistently** shows an improvement from CF to multimodal RSs. Additionally, the imputed performance improvement is always quite **higher in magnitude** than the dropped one (refer to the asterisks in the table). We assume that the imputed setting permits (i) retaining important information in the recommendation data (i.e., users, items, and interactions) and (ii) recovering missing multimodal features in a high-quality manner.

**Impact on recent multimodal RSs (RQ2).** We test if the imputation of missing modalities in **recent** multimodal RSs can lead to performance trends that are aligned with the related literature. To this aim, we extend the set of RSs with LightGCN, SGL (pure CF), and FREEDOM, BM3 (multimodal). Table 3 displays the recommendation results in the **imputed** setting. Overall, we notice two trends aligned with the literature. On the one hand, more recent approaches (e.g., LightGCN and SGL on the CF side, FREEDOM and BM3 on the multimodal side) can often provide higher performance. On the other hand, the adoption of refined approaches leveraging multimodality (e.g., FREEDOM and BM3) can frequently lead to improved performance over pure CF (e.g., see Office and Beauty).

**Ablation study (RQ3).** Table 4 shows the performance of FREEDOM and BM3 on Music and Beauty, highlighting the various imputation methods. Overall, we note that our proposed **graph-aware** imputation methods (i.e., NEIGHMEAN, MULTIHOP, PERSPAGERANK)

**Table 2: Performance improvement of VBPR and NGCF-M over BPRMF and NGCF in the dropped and imputed settings.**

Models	Dropped						Imputed (Ours)					
	Office		Music		Beauty		Office		Music		Beauty	
	Recall	nDCG	Recall	nDCG	Recall	nDCG	Recall	nDCG	Recall	nDCG	Recall	nDCG
BPRMF	7.74	3.65	26.50	14.24	9.81	5.00	8.51	4.55	24.59	13.69	9.43	4.85
VBPR	9.99	4.82	25.22	13.71	11.12	5.86	10.09	5.51	26.25	14.58	10.93	5.74
Improvement (%)	<b>+29.07%</b>	<b>+32.05%</b>	<b>-4.83%</b>	<b>-3.72%</b>	<b>+13.35%</b>	<b>+17.20%</b>	<b>+18.57%</b>	<b>+21.10%</b>	<b>+6.75%*</b>	<b>+6.50%*</b>	<b>+15.91%*</b>	<b>+18.35%*</b>
NGCF	9.90	4.49	24.39	12.61	8.64	4.28	8.18	3.91	24.12	13.06	8.92	4.56
NGCF-M	13.10	6.36	24.12	12.73	9.68	4.94	11.65	5.85	25.18	13.64	11.00	5.76
Improvement (%)	<b>+31.32%</b>	<b>+41.65%</b>	<b>-1.11%</b>	<b>+0.95%</b>	<b>+12.04%</b>	<b>+15.42%</b>	<b>+42.42%*</b>	<b>+49.62%*</b>	<b>+4.39%*</b>	<b>+4.44%*</b>	<b>+23.32%*</b>	<b>+26.32%*</b>

\* The performance improvement is higher in our imputed setting than in the dropped one.

**Table 3: Performance in the imputed setting. Boldface and underline are the best and second-best values.**

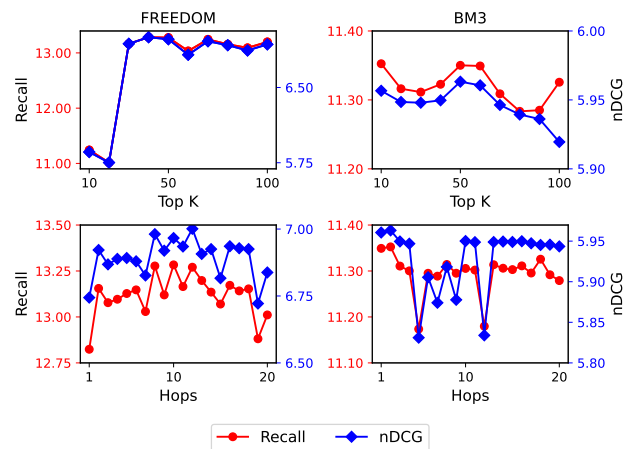
Models	Office		Music		Beauty	
	Recall	nDCG	Recall	nDCG	Recall	nDCG
BPRMF	8.51	4.55	24.59	13.69	9.43	4.85
NGCF	8.18	3.91	24.12	13.06	8.92	4.56
LightGCN	11.55	6.12	26.61	14.49	11.61	6.02
SGL	10.28	5.81	<b>27.59</b>	<b>15.86</b>	<u>12.20</u>	<u>6.72</u>
VBPR	10.09	5.51	26.25	14.58	10.93	5.74
NGCF-M	<u>11.65</u>	5.85	25.18	13.64	11.00	5.76
FREEDOM	<b>12.79</b>	<b>6.76</b>	<u>27.18</u>	<u>14.91</u>	<b>13.28</b>	<b>7.00</b>
BM3	11.28	<u>6.15</u>	25.11	13.55	11.49	6.04

**Table 4: Ablation study with varying imputation methods.**

Models	Imputation	Music		Beauty	
		Recall	nDCG	Recall	nDCG
FREEDOM	+ZEROS	25.60	13.78	10.99	5.76
	+RANDOM	26.81	14.46	13.00	6.73
	+GLOBALMEAN	26.62	14.61	13.02	6.80
	+NEIGHMEAN	26.17	14.35	13.23	6.88
	+MULTIHOP	<u>27.09</u>	<b>14.91</b>	<u>13.26</u>	<u>6.96</u>
	+PERSPAGERANK	<b>27.18</b>	<u>14.88</u>	<b>13.28</b>	<b>7.00</b>
BM3	+ZEROS	24.64	13.00	11.31	5.96
	+RANDOM	24.29	12.93	11.35	5.88
	+GLOBALMEAN	24.27	12.88	11.35	5.91
	+NEIGHMEAN	24.46	12.95	<u>11.48</u>	<u>6.01</u>
	+MULTIHOP	<u>25.00</u>	<b>13.55</b>	<b>11.49</b>	<b>6.04</b>
	+PERSPAGERANK	<b>25.11</b>	<u>13.44</u>	11.35	5.96

always **outperform** traditional machine learning imputation (i.e., ZEROS, RANDOM, and GLOBALMEAN); as expected, we see MULTIHOP and PERSPAGERANK steadily settle as the superior solutions. Thus, we can effectively leverage the **graph structure** of the user-item interaction data to adopt more tailored and refined graph-aware imputation methods that exploit item-item co-interactions and semantic similarities, especially at multiple distance hops.

**Hyper-parameter sensitivity (RQ4).** Finally, we assess the impact of top- $k$  sparsification and propagation layers for the best

**Figure 1: Impact of top- $k$  sparsification (top) and propagation hops (bottom) on the Beauty dataset for PERSPAGERANK.**

**graph-aware** imputation strategy, PERSPAGERANK. Figure 1 shows the performance changes when varying the top- $k$  sparsification (top) and the propagation layers (bottom) on Beauty. Regarding top- $k$  sparsification, we obtain the best results on FREEDOM for  $k \geq 30$ , while results on BM3 are better around  $k = 50$ . Regarding the propagation layers, the outcomes suggest that on FREEDOM we reach the highest metric values around  $T = 10$ , while on BM3 the trend is almost stable for  $T \geq 10$ . Conclusively, **high top- $k$  sparsifications** and many propagation hops can improve the results on both FREEDOM and BM3, as they remove **noisy** item-item interactions and leverage **multiple hops** in the item-item graph.

## 4 CONCLUSION AND FUTURE WORK

Unlike the common practice of dropping items with missing modalities in multimodal recommendation, in this work, we proposed an untrained pre-processing pipeline to impute missing multimodal features through traditional methods and novel graph-aware strategies. First, we showed that imputing can preserve (or even improve) the performance gap between multimodal and pure collaborative recommender systems. Then, we demonstrated the superior efficacy of our proposed graph-aware imputation methods. We plan to

further extend the experimental setting to additional datasets and techniques, as well as design new imputation solutions to integrate end-to-end in the recommendation pipeline [47].

## ACKNOWLEDGMENTS

Daniele Malitesta and Fragkiskos D. Malliaros acknowledge the support by ANR (French National Research Agency) under the JCJC project GraphIA (ANR-20-CE23-0009-01). Claudio Pomo and Tommaso Di Noia acknowledge the partial support by CT\_FINCONS\_III, OVS Fashion Retail Reloaded, LUTECH DIGITALE 4.0, Secure Safe Apulia, IDENTITA, REACH-XY. The authors acknowledge Inria and the CINECA award under the ISCRA initiative for the availability of high-performance computing resources.

## REFERENCES

- [1] Vito Walter Anelli, Yashar Deldjoo, Tommaso Di Noia, Eugenio Di Sciascio, Antonio Ferrara, Daniele Malitesta, and Claudio Pomo. 2022. Reshaping Graph Recommendation with Edge Graph Collaborative Filtering and Customer Reviews. In *DLASR@CIKM (CEUR Workshop Proceedings, Vol. 3317)*. CEUR-WS.org.
- [2] Matteo Attimonelli, Danilo Danese, Daniele Malitesta, Claudio Pomo, Giuseppe Gassi, and Tommaso Di Noia. 2024. Ducho 2.0: Towards a More Up-to-Date Unified Framework for the Extraction of Multimodal Features in Recommendation. In *WWW (Companion Volume)*. ACM, 1075–1078.
- [3] Desheng Cai, Shengsheng Qian, Quan Fang, and Changsheng Xu. 2022. Heterogeneous Hierarchical Feature Aggregation Network for Personalized Micro-Video Recommendation. *IEEE Trans. Multim.* 24 (2022), 805–818.
- [4] Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, and Binqiang Zhao. 2019. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. In *KDD*. ACM.
- [5] Xusong Chen, Dong Liu, Zhiwei Xiong, and Zheng-Jun Zha. 2021. Learning and Fusing Multiple User Interest Representations for Micro-Video and Movie Recommendations. *IEEE Trans. Multim.* 23 (2021), 484–496.
- [6] Zhiyong Cheng, Jialie Shen, and Steven C. H. Hoi. 2016. On Effective Personalized Music Retrieval by Exploring Online User Behaviors. In *SIGIR*. ACM, 125–134.
- [7] Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. 2021. A Study on the Relative Importance of Convolutional Neural Networks in Visually-Aware Recommender Systems. In *CVPR Workshops*. Computer Vision Foundation / IEEE, 3961–3967.
- [8] Yashar Deldjoo, Tommaso Di Noia, Daniele Malitesta, and Felice Antonio Merra. 2022. Leveraging Content-Style Item Representation for Visual Recommendation. In *ECIR (2) (Lecture Notes in Computer Science, Vol. 13186)*. Springer, 84–92.
- [9] Michael D. Ekstrand, John Riedl, and Joseph A. Konstan. 2011. Collaborative Filtering Recommender Systems. *Found. Trends Hum. Comput. Interact.* 4, 2 (2011), 175–243.
- [10] Tlameo Emmanuel, Thabiso M. Maupong, Dimane Mpoeleng, Thabo Semong, Banyatsang Mphago, and Oteng Tabona. 2021. A survey on missing data in machine learning. *J. Big Data* 8, 1 (2021), 140.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. IEEE Computer Society, 770–778.
- [12] Ruining He and Julian J. McAuley. 2016. VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. In *AAAI*. AAAI Press, 144–150.
- [13] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*. ACM, 639–648.
- [14] Natasha Jaques, Sara Taylor, Akane Sano, and Rosalind W. Picard. 2017. Multimodal autoencoder: A deep learning approach to filling in missing sensor data and enabling better mood prediction. In *ACII*. IEEE Computer Society, 202–208.
- [15] Johannes Klicpera, Stefan Weissenberger, and Stephan Günnemann. 2019. Diffusion Improves Graph Learning. In *NeurIPS*. 13333–13345.
- [16] Yehuda Koren, Robert M. Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37.
- [17] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. 2023. Multimodal Prompting with Missing Modalities for Visual Recognition. In *CVPR*. IEEE, 14943–14952.
- [18] Zhenfeng Lei, Anwar Ul Haq, Adnan Zeb, Md Suzauddola, and Defu Zhang. 2021. Is the suggested food your desired?: Multi-modal recipe recommendation with demand-based knowledge graph. *Expert Syst. Appl.* 186 (2021), 115708.
- [19] Haoxuan Li, Chunyuan Zheng, and Peng Wu. 2023. StableDR: Stabilized Doubly Robust Learning for Recommendation on Data Missing Not at Random. In *ICLR*. OpenReview.net.
- [20] Daryl Lim, Julian J. McAuley, and Gert R. G. Lanckriet. 2015. Top-N Recommendation with Missing Implicit Feedback. In *RecSys*. ACM, 309–312.
- [21] Wei-Chao Lin and Chih-Fong Tsai. 2020. Missing value imputation: a review and analysis of the literature (2006-2017). *Artif. Intell. Rev.* 53, 2 (2020), 1487–1509.
- [22] Fan Liu, Zhiyong Cheng, Lei Zhu, Chenghao Liu, and Liqiang Nie. 2022. An Attribute-Aware Attentive GCN Model for Attribute Missing in Recommendation. *IEEE Trans. Knowl. Data Eng.* 34, 9 (2022), 4077–4088.
- [23] Yong Liu, Susen Yang, Chenyi Lei, Guoxin Wang, Haihong Tang, Juyong Zhang, Aixun Sun, and Chunyan Miao. 2021. Pre-training Graph Transformer with Multimodal Side Information for Recommendation. In *ACM Multimedia*. ACM, 2853–2861.
- [24] Zhuang Liu, Yunpu Ma, Matthias Schubert, Yuanxin Ouyang, and Zhang Xiong. 2022. Multi-Modal Contrastive Pre-training for Recommendation. In *ICMR*. ACM, 99–108.
- [25] Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. 2022. Are Multimodal Transformers Robust to Missing Modality?. In *CVPR*. IEEE, 18156–18165.
- [26] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. 2021. SMIL: Multimodal Learning with Severely Missing Modality. In *AAAI*. AAAI Press, 2302–2310.
- [27] Daniele Malitesta, Giandomenico Cornacchia, Claudio Pomo, Felice Antonio Merra, Tommaso Di Noia, and Eugenio Di Sciascio. 2024. Formalizing Multimedia Recommendation through Multimodal Deep Learning. *ACM Trans. Recomm. Syst.* (apr 2024). <https://doi.org/10.1145/3662738> Just Accepted.
- [28] Daniele Malitesta, Giuseppe Gassi, Claudio Pomo, and Tommaso Di Noia. 2023. Ducho: A Unified Framework for the Extraction of Multimodal Features in Recommendation. In *ACM Multimedia*. ACM, 9668–9671.
- [29] Daniele Malitesta, Emanuele Rossi, Claudio Pomo, Fragkiskos D. Malliaros, and Tommaso Di Noia. 2024. Dealing with Missing Modalities in Multimodal Recommendation: a Feature Propagation-based Approach. *CoRR* abs/2403.19841 (2024).
- [30] Manuel J. Marín-Jiménez, Francisco M. Castro, Rubén Delgado-Escañó, Vicky Kalogeiton, and Nicolás Guil. 2021. UGaitNet: Multimodal Gait Recognition With Missing Input Modalities. *IEEE Trans. Inf. Forensics Secur.* 16 (2021), 5452–5462.
- [31] Benjamin M. Marlin. 2008. *Missing Data Problems in Machine Learning*. Ph.D. Dissertation. University of Toronto, Canada.
- [32] Benjamin M. Marlin, Richard S. Zemel, Sam T. Roweis, and Malcolm Slaney. 2011. Recommender Systems, Missing Data and Statistical Model Estimation. In *IJCAI/IJCAI/AAAI*, 2686–2691.
- [33] Julian J. McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-Based Recommendations on Styles and Substitutes. In *SIGIR*. ACM, 43–52.
- [34] Weiqing Min, Shuqiang Jiang, and Ramesh C. Jain. 2020. Food Recommendation: Framework, Existing Solutions, and Challenges. *IEEE Trans. Multim.* 22, 10 (2020), 2659–2671.
- [35] Sergio Oramas, Oriol Nieto, Mohamed Sordo, and Xavier Serra. 2017. A Deep Multimodal Approach for Cold-start Music Recommendation. In *DLRS@RecSys*. ACM, 32–37.
- [36] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report 1999-66. Stanford InfoLab.
- [37] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *EMNLP/IJCNLP (1)*. Association for Computational Linguistics, 3980–3990.
- [38] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI*.
- [39] Emanuele Rossi, Henry Kenlay, Maria I. Gorinova, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. 2022. On the Unreasonable Effectiveness of Feature Propagation in Learning on Graphs With Missing Node Features. In *LoG (Proceedings of Machine Learning Research, Vol. 198)*. PMLR, 11.
- [40] DONALD B. RUBIN. 1976. Inference and missing data. *Biometrika* 63, 3 (12 1976), 581–592. <https://doi.org/10.1093/biomet/63.3.581> arXiv:<https://academic.oup.com/biomet/article-pdf/63/3/581/756166/63-3-581.pdf>
- [41] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. 2020. Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. In *WSDM*. ACM, 501–509.
- [42] Shaoyun Shi, Min Zhang, Xinxing Yu, Yongfeng Zhang, Bin Hao, Yiqun Liu, and Shaoping Ma. 2019. Adaptive Feature Sampling for Recommendation with Missing Content Feature Values. In *CIKM*. ACM, 1451–1460.
- [43] Harald Steck. 2010. Training and testing of recommender systems on data missing not at random. In *KDD*. ACM, 713–722.
- [44] William E. Strawderman. 1989. Statistical Analysis with Missing Data (Roderick J. A. Little and Donald B. Rubin). *SIAM Rev.* 31, 2 (1989), 348–349.
- [45] Kunal Vaswani, Yudhik Agrawal, and Vinoo Alluri. 2021. Multimodal Fusion Based Attentive Networks for Sequential Music Recommendation. In *BigMM*. IEEE, 25–32.
- [46] Johannes Wagner, Elisabeth André, Florian Lingens, and Jonghwa Kim. 2011. Exploring Fusion Methods for Multimodal Emotion Recognition with Missing Data. *IEEE Trans. Affect. Comput.* 2, 4 (2011), 206–218.

- [47] Cheng Wang, Mathias Niepert, and Hui Li. 2018. LRMM: Learning to Recommend with Missing Modalities. In *EMNLP*. Association for Computational Linguistics, 3360–3370.
- [48] Menghan Wang, Mingming Gong, Xiaolin Zheng, and Kun Zhang. 2018. Modeling Dynamic Missingness of Implicit Feedback for Recommendation. In *NeurIPS*. 6670–6679.
- [49] Wenjie Wang, Ling-Yu Duan, Hao Jiang, Peiguang Jing, Xuemeng Song, and Liqiang Nie. 2021. Market2Dish: Health-aware Food Recommendation. *ACM Trans. Multim. Comput. Commun. Appl.* 17, 1 (2021), 33:1–33:19.
- [50] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *SIGIR*. ACM, 165–174.
- [51] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. 2019. Doubly Robust Joint Learning for Recommendation on Data Missing Not at Random. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 6638–6647.
- [52] Wei Wei, Chao Huang, Lianghao Xia, and Chuxu Zhang. 2023. Multi-Modal Self-Supervised Learning for Recommendation. In *WWW*. ACM, 790–800.
- [53] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, and Tat-Seng Chua. 2020. Graph-Refined Convolutional Network for Multimedia Recommendation with Implicit Feedback. In *ACM Multimedia*. ACM, 3541–3549.
- [54] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal Graph Convolution Network for Personalized Recommendation of Micro-video. In *ACM Multimedia*. ACM, 1437–1445.
- [55] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised Graph Learning for Recommendation. In *SIGIR*. ACM, 726–735.
- [56] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge J. Belongie, and Deborah Estrin. 2018. Unbiased offline recommender evaluation for missing-not-at-random implicit feedback. In *RecSys*. ACM, 279–287.
- [57] Jiandian Zeng, Tianyi Liu, and Jiantao Zhou. 2022. Tag-assisted Multimodal Sentiment Analysis under Uncertain Missing Modalities. In *SIGIR*. ACM, 1545–1554.
- [58] Chaohu Zhang, Xu Chu, Liantao Ma, Yinghao Zhu, Yasha Wang, Jiantao Wang, and Junfeng Zhao. 2022. M3Care: Learning with Missing Modalities in Multimodal Healthcare Data. In *KDD*. ACM, 2418–2428.
- [59] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, Shu Wu, Shuhui Wang, and Liang Wang. 2021. Mining Latent Structures for Multimedia Recommendation. In *ACM Multimedia*. ACM, 3872–3880.
- [60] Xiaolin Zheng, Menghan Wang, Renjun Xu, Jianmeng Li, and Yan Wang. 2022. Modeling Dynamic Missingness of Implicit Feedback for Sequential Recommendation. *IEEE Trans. Knowl. Data Eng.* 34, 1 (2022), 405–418.
- [61] Xin Zhou and Zhiqi Shen. 2023. A Tale of Two Graphs: Freezing and Denoising Graph Structures for Multimodal Recommendation. In *ACM Multimedia*. ACM, 935–943.
- [62] Xin Zhou, Hongyu Zhou, Yong Liu, Zhiwei Zeng, Chunyan Miao, Pengwei Wang, Yuan You, and Feijun Jiang. 2023. Bootstrap Latent Representations for Multimodal Recommendation. In *WWW*. ACM, 845–854.