Characterizing Impression-Aware Recommender Systems

Fernando B. Pérez Maurera[1,2,*], Maurizio Ferrari Dacrema[1], Pablo Castells[3] and Paolo Cremonesi[1]

[1]Politecnico di Milano, Milan, Italy
[2]ContentWise, Milan, Italy
[3]Universidad Autónoma de Madrid, Madrid, Spain

Abstract

Impression-aware recommender systems (IARS) are a type of recommenders that learn user preferences using their interactions and the recommendations (also known as impressions) shown to users. The community's interest in this type of recommenders has steadily increased in recent years. To aid in characterizing this type of recommenders, we propose a theoretical framework to define IARS and classify the recommenders present in the state-of-the-art. We start this work by defining core concepts related to this type of recommenders, such as impressions and user feedback. Based on this theoretical framework, we identify and define three properties and three taxonomies that characterize IARS. Lastly, we undergo a systematic literature review where we discover and select papers belonging to the state-of-the-art. Our review analyzes papers under the properties and taxonomies we propose; we highlight the most and least common properties and taxonomies used in the literature, their relations, and their evolution over time, among others.

Keywords

Recommender Systems, Impression, Slate, Exposure, Taxonomy

1. Introduction

A recommender system (RS) is a collection of software tools that, in conjunction, assist users in discovering meaningful items of interest. To achieve their goal, RS learn users' preferences by creating a model of the user. They create such a model by gathering different types of data from varied sources and extracting relations between users, between items, and among users and items.

Many types of RS exist in the literature, each tailored to specific tasks, data sources, or domains. In this work, we study one type of RS, which we term impression-aware recommender system (IARS). IARS learn user preferences by leveraging user *interactions* and *impressions*. Interactions are the actions users perform over items of a recommender system, such as ratings, purchases, or media watching. Instead, Impressions are the items the system recommends to users, i.e., the items presented to the user on-screen. Impressions are not exclusive to recommender systems; any system that selects a limited amount of items to show to the user can be considered a system that generates impressions, e.g., editorial selections or search systems also generate impressions.

In Pérez Maurera et al. [1], we present an significant extension of this work, where we address and discuss in further detail several topics that we present in this work. Nevertheless, this work presents some specifics that were not included in our previous works. For instance, in our characterization of IARS, we identify the inputs, outputs, and computational tasks of IARS. The detailed contributions of this work are:

1. We define those recommender systems that leverage impressions to learn users' preferences: impression-aware recommender system (IARS).
2. We identify different properties that IARS share. Furthermore, we propose several categories within those identified properties.
3. We propose a classification system for papers describing IARS. Such taxonomies inspect each paper from different perspectives and provide a comprehensive paper overview.
4. We classify and discuss the current state-of-the-art on IARS. In our discussion, we provide a comprehensive view of past works.

2. Related Works

Some previous works have provided initial contributions to the description and formalization of IARS. The majority of papers describing IARS, e.g., [2, 3, 4, 5], mention impressions are a type of negative user feedback. In this

Workshop on Learning and Evaluating Recommendations with Impressions (LERI) @ RecSys 2023, September 18-22 2023, Singapore

*Corresponding author.
fernandobenjamin.perez@polimi.it (F. B. Pérez Maurera); maurizio.ferrari@polimi.it (M. Ferrari Dacrema); pablo.castells@uam.es (P. Castells); paolo.cremonesi@polimi.it (P. Cremonesi)
0000-0001-6578-7404 (F. B. Pérez Maurera); 0000-0001-7103-2788 (M. Ferrari Dacrema); 0000-0003-0668-6317 (P. Castells); 0000-0002-1253-8081 (P. Cremonesi)

work, we do not assume impressions represent either a positive or negative type of user feedback; instead, we identify which signals impressions may carry and classify papers accordingly.

Several papers address topics related to the characterization of IARS. For instance, Pérez Maurera et al. [6] provides an overview of IARS while describing five categories of this type of recommenders. Compared to their paper, we integrate their proposed categories into a single taxonomy and provide additional novel taxonomies that analyze other aspects of IARS. Moreover, we analyze how papers can be classified among all the taxonomies, present how these taxonomies evolve over time, and discuss how they are associated to others. These discussions are not included in the extended version of this work Pérez Maurera et al. [1].

Other papers have pursued different directions for IARS, particularly evaluating this type of recommenders while considering the recent calls for sound evaluation strategies and better replicability measures in the recommender systems community. For instance, Zhao et al. [7] studied the user preference for those items in impressions which received no interactions. Their results suggest that impressions are not strictly a type of negative user feedback but are complex signals dependent on the context and time of the recommendation. Pérez Maurera et al. [8, 9] performed two evaluation studies of some IARS from the literature. The results of their experiments suggest that using impressions increases the accuracy and beyond-accuracy metrics of some existing recommenders. Those papers emphasize the need for further experiments and the evaluation of various recommenders in future works. This work differs from those papers mainly in evaluating recommenders, as we do not perform an evaluation study of IARS. Instead, we cover the topic of the evaluation of IARS by highlighting the existing recommendation tasks and public datasets.

Overall, this work aims to propose a theoretical framework to define IARS and describe this type of recommenders through their properties and taxonomies.

3. Impression-Aware Recommender Systems

Previously, we briefly presented the concept of an impression as a collection of items shown to the user. This section provides a broader definition of impressions and other subjects needed to formalize our taxonomies of this type of recommenders.

Impression An impression is an ordered sequence of potentially relevant items to the user whose contents are arranged on the user's screen. The source of an impression, i.e., the entity that generates it, does not necessarily need to be a RS. For instance, a search engine or an editorial system also generate impressions, i.e., impressions may be the results of a search query or a selection of items the editors of the system want to promote.

Interaction An interaction is the implicit or explicit action of the user over an item where this item is inside an impression shown on-screen. Examples of implicit interactions are clicks or purchases, while examples of explicit actions are ratings. Traditionally, implicit interactions represent the users' favorable preference, i.e., users interact with items they consider align with their tastes.

User Feedback User feedback refers to the implicit or explicit preferences of the user to recommended items, i.e., to impressions. When a user receives an impression on their screen, it is assumed they scan the impression and decide to interact with certain items while others remain without interactions. We catalog those items that received an interaction as **interacted impressions** and those that did not as **non-interacted impressions**.

Impression-Aware Recommender System (IARS) An IARS is a type of recommender system that leverages impressions and interactions to learn users' preferences toward items in the catalog. Inside a RS, one module called **recommendation model** is in charge of learning the user preferences.

Inputs The input of a recommendation model is, at minimum, a user, an item identifier, and users' profiles. The users' profiles are a collection of all the interactions the users performed with the system and the impressions the system showed to those users. Depending on the design of a recommendation model, the input may change; commonly, the input may include additional information about users and items, e.g., user demographics or item attributes.

Output The output of any recommendation model is a real number called the predicted relevance score. Generally, this score tells the preference of a given user over an item in the catalog. Despite this typical case, the relevance score can also portray different meanings, e.g., it may reflect the expected click-through rate or the probability of purchase.

Features Features refer to statistical properties of impressions that recommendation models compute and utilize to learn user preferences. For instance, one common feature in the literature is the number of times a given item has been impressed by a user. This feature is commonly called the number of impressions.

Computational Tasks Based on the definition of IARS, recommendation models, their inputs, outputs, and features, the more evident computational task for a IARS is computing the relevance score. However, the recommender computes more than the relevance score. Depending on the recommender's design, it may compute several features or needs to interpret interactions and impressions in specific manners. For instance, a recommender may be designed to predict the next song in a playlist; thus, it needs to order impressions and interactions by the date and time they occurred.

3.1. Related Recommenders

Based on the definition of IARS shown in the previous section, a recommender system is considered impression-aware when it learns user preferences using impressions and interactions. Impressions are a data type that complements interactions and do not pose any restrictions or constraints to the underlying recommender. In other words, a recommender of a different type may incorporate impressions without becoming exclusively IARS. For instance, a sequential-aware recommender learns user preferences from ordered sequences of interactions [10]. Sequential-aware recommenders may also become impression-aware if they learn from the sequence of interactions and impressions i.e., they incorporate into their input sequence the impressions generated by the system.

Regarding seemingly similar types of recommenders, Context-Aware Recommender Systems (CARS) seem comparable to IARS in their definition and may pose as equivalent. A CARS learns user preferences using interactions and *contextual features* [11]. In this sense, impressions can be seen as contextual features of interactions. However, these two types of recommenders are different because the inputs of the recommenders are different. Mainly, when producing a relevance score (as defined in the previous section), part of the input of a CARS is the current context; IARS cannot receive an impression at this stage because it generates the impression after it computes all the relevance scores.

3.2. Evaluation of IARS

This section presents specific properties for the evaluation of IARS. In particular, we present the typical recommendation tasks for IARS and list the thirteen public datasets with impressions available for the recommender systems community. Other relevant aspects for the evaluation of IARS; for instance, evaluation methodologies and challenges, have already been discussed in other works [8, 6, 9].

Recommendation Tasks The community in RS has primarily focused on the task of **top-N recommendations**, which generates a personalized collection of N items to a user called a recommendation list or impression. Another task, which has increased in popularity, is called **re-ranking**. Under that task, the recommender receives an impression holding N items. Then the recommender produces a permutation of such impression [12].

Public Datasets To evaluate IARS, the community has access to thirteen datasets with impressions from different recommendation domains. The distribution of the datasets by recommendation domain is: three datasets in news [13, 14, 15, 16, 17], three in online advertisement [18, 19],[1] three in media [20, 21], two in fashion [22], and two in e-commerce [23, 24, 25].

The information in those datasets varies greatly. Some contain the entire impression list and the interactions each item received by the users, while others only hold the number of interactions or impressions of each user-item pair. Pérez Maurera et al. [9] provide a thorough description of three datasets. Another paper Pérez Maurera et al. [6] shortly describe other types of datasets in the literature, e.g., those not publicly available to the community.

4. Properties

This section depicts the properties shared among IARS. In this context, a property is a characteristic of the RS and how such characteristic is involved with impressions. Notably, we study three properties: which type of impressions IARS use, how IARS deem impressions in terms of the users' preferences, and what kind of recommendation model uses an IARS.

4.1. Impressions Type

This property refers to the type of information used by an IARS when learning user preferences. We identify two types of impressions based on the information available within impressions:

- **Contextual:** the recommender has access to the user, a possibly interacted item, and the impression holding such an item. In other words, with a contextual impression, the recommender knows every impression shown to users and their feedback on every impressed item, i.e., the feedback indicates whether the user interacted with the impressed item.
- **Global:** the recommender has access to users' feedback on impressed items, i.e., interacted or

[1] https://www.kaggle.com/competitions/kddcup2012-track2

non-interacted impressions, but does not have access to the contents of the impressions. In other words, the recommender knows whether a user interacted with an item but does not know to which impression such an item belongs.

4.2. Impressions Signal

In traditional recommender systems research, it is usually considered that interacted items convey positive signals of users' preferences, i.e., users show their implicit acceptance of an item by interacting with it. Analogously, the non-interacted items are assumed to represent negative signals of user preferences. This scenario is the so-called the *missing as negatives* assumption [26].

As stated in Section 3, impressions may receive two types of user feedback: interacted impressions and non-interacted impressions. As commonly agreed by the community, we also assume that interacted impressions, i.e., interactions, represent positive signals. However, we do not assign a negative signal to non-interacted impressions beforehand as the literature in IARS does not converge on a single signal for non-interacted impressions. Moreover, we identify three signals to non-interacted impressions:

- **Positive:** meaning users prefer non-interacted impressions; however, they decided not to interact with them when shown.
- **Negative:** meaning users dislike non-interacted impressions. Under the traditional *missing as negatives* assumption, non-interacted impressions are deemed as negative signals.
- **Neutral:** meaning users do not have a positive or a negative preference for non-interacted impressions.

4.3. Recommender Type

As explained in Section 3, recommendation models are the module of an IARS that learns user preferences and computes the relevance score for any given user-item pair. Despite this, the design of the recommendation model depends on the recommendation task at hand, i.e., top-N recommendations, or re-ranking of impressions. As such, we identify three types of recommenders:

- **End-to-end:** a type of recommendation model used in top-N recommendations. The model receives the user preferences and generates an impression containing N elements.
- **Plug-in:** a type of recommendation model used in top-N recommendations. The model receives the user preferences and the predicted relevance scores created by another recommender. The model generates an impression containing N elements by transforming those relevance scores.
- **Re-ranking:** a type of recommendation model exclusively used in re-ranking tasks. This type of recommender not only receives the users' profiles as input but also receives an impression. It generates a permutation of the input impression.

5. Taxonomies

In this section, we present one of the main contributions of this work: three taxonomies for IARS. This section thoroughly defines each taxonomy, indicating their similarities and differences and their main categories. These taxonomies are different classification systems for recommenders in the literature, which we inspect under three perspectives: model-centric, data-centric, and signal-centric.

5.1. Model-Centric Taxonomy

In the model-centric taxonomy, we analyze the recommendation model a specific paper uses, i.e., we examine the design of the recommendation model and its learning paradigm. Hence, this taxonomy does not classify papers based on how they use impressions nor how they deem the user preferences to impressions. Instead, the taxonomy focuses on classifying the recommendation model of an IARS. Under this taxonomy, we propose five categories of IARS:

- **Heuristics:** contains all papers which describe a recommendation model using rules, associations, or ad-hoc approaches to model user preferences. For example, Buchbinder et al. [27] does not recommend an item after a certain number of impressions.
- **Statistical:** contains all papers which describe a recommendation model using probability or statistical models to model user preferences. For example, Zhang et al. [28] uses logistic regression.
- **Machine learning:** contains all papers which describe a recommendation model using *shallow* machine learning approaches to model user preferences. For example, Liu et al. [29] uses gradient boosting decision trees.
- **Deep learning:** contains all papers which describe a recommendation model using *deep* machine learning approaches to model user preferences. For example, Covington et al. [30] uses multilayer perceptrons.

- **Reinforcement learning:** contains all papers which describe a recommendation model using *Markov decision problems* to model the recommendation problem. For example, Gruson et al. [31] uses multi-armed bandits.

We propose the previous categories based on the types of recommendation models currently existing in the literature. However, the taxonomy may be expanded to cover recommenders with additional learning paradigms.

5.2. Data-Centric Taxonomy

In the data-centric taxonomy, we analyze how a recommendation model uses impressions to learn users' preferences. In other words, we analyze the input of the recommendation model. Under this taxonomy, we propose three categories of IARS:

- **Learn:** contains all papers which describe a recommendation model where its input is an impression in any of its forms, e.g., the recommendation list shown to the user or a single impressed item. For example, Ma et al. [32] learn to classify items into two classes: interacted and non-interacted impressions.

- **Sample:** contains all papers which describe a recommendation model where its input is a collection (e.g., a set, a sequence, or a vector) containing items sampled from interactions, impressions, or both. For example, Pérez Maurera et al. [20] sample impressions as negative items and interactions as positive items when training a BPR-optimized [33] matrix factorization recommender.

- **Features:** contains all papers which describe a recommendation model where its input is one or several features computed from impressions. For example, Gong et al. [4] compute the amount of time a user has watched an impressed item.

Unlike the previous taxonomy, in the data-centric taxonomy, one paper may belong to two categories simultaneously. This relaxed property is allowed to avoid the creation of categories that combine two of the previous categories. For example, Aharon et al. [34] propose a recommender that learns from impressions and computes frequency features; therefore, we classify the paper as *learn* and *features* under the data-centric taxonomy.

5.3. Signal-Centric Taxonomy

This taxonomy analyzes how a paper treats non-interacted impressions in terms of how relevant impressions are to the users. We specifically study non-interacted impressions as their counterparts, i.e., interactions, are already deemed as positive user feedback in the literature under the *missing as negatives* [26] assumption. Under this taxonomy, we propose two categories of IARS:

- **Assume:** contains all papers describing a recommendation model that assumes users' specific preference toward non-interacted impressions. For example, Xi et al. [2] assume non-interacted impressions represent implicit negative feedback.

- **Learn:** contains all papers describing a recommendation model that learns users' preference toward non-interacted impressions. For example, Deffayet et al. [35] use a variational autoencoder to learn users' preferences for impressions and their items using the user feedback on items in the impression.

6. Classification of the State-of-the-Art

In this section, we classify relevant papers in the literature, which we deem as state-of-the-art, under the properties and taxonomies of IARS discussed in Section 4 and Section 5, respectively. Before the classification, we identify the state-of-the-art in IARS by conducting a systematic literature exploration, describing the discovery process and selection criteria for papers. Overall, in two taxonomies, we find that works are distributed almost uniformly amongst the proposed categories, while in the remaining taxonomy and the three properties, most papers favor one category over the rest. We also find that almost all taxonomies and properties show large statistically significant associations against others. In contrast, only one taxonomy does not have statistically significant associations with the rest of the properties and taxonomies. Table 1 shows the distribution of papers according to each taxonomy and property, and Table 2 shows the association between properties and taxonomies.

6.1. Paper Selection Criteria

In this work, we consider a paper to belong to the state-of-the-art of IARS when the paper meets these conditions:

1. The paper is peer-reviewed.

2. The paper is published in a conference or a journal.

3. The paper is published in a high-level venue.

4. The paper describes or evaluates a IARS.

Table 1
Classification of the state-of-the-art according to the taxonomies and properties of IARS proposed and defined in this work. **Count** is the number of papers belonging to a given taxonomy or property. **Percentage** is the percentage of papers belonging to the classification inside a taxonomy or property.

Classification	Category	Papers References	Count	Percentage
Model-centric taxonomy	heuristic	[36, 29, 37, 27, 38]	5	13.9%
	statistical	[39, 28, 40, 41]	4	11.1%
	machine learning	[20, 34, 29, 42]	4	11.1%
	deep learning	[43, 2, 44, 3, 4, 32, 45, 22, 46, 47, 48, 49, 30]	13	36.1%
	reinforcement learning	[35, 5, 50, 51, 52, 31, 53, 54, 17]	9	25.0%
	not described	[7]	1	2.8%
Data-centric taxonomy	features & learn	[45, 22, 49, 34, 42, 30, 37, 38]	8	22.2%
	features	[4, 50, 48, 52, 36, 29, 39, 40, 27, 41]	10	27.8%
	learn	[43, 2, 44, 35, 3, 5, 32, 46, 51, 31, 53, 54, 29, 28, 17]	15	41.6%
	sample	[47, 20]	2	5.6%
	not described	[7]	1	2.8%
Signal-centric taxonomy	assume	[2, 3, 4, 5, 32, 22, 46, 47, 51, 20, 49, 52, 31, 53, 34, 54, 36, 29, 29, 28, 42, 17, 30, 27]	24	66.7%
	learn	[44, 35, 45, 7, 39, 40, 37, 38, 41]	9	25.0%
	not described	[43, 50, 48]	3	8.3%
Impressions type	contextual	[43, 2, 44, 35, 4, 5, 47, 51, 49, 52, 7, 42]	12	33.3%
	global	[3, 32, 45, 22, 46, 50, 48, 20, 31, 53, 34, 54, 36, 29, 29, 39, 28, 40, 17, 30, 37, 27, 38, 41]	24	66.7%
Impressions signal	negative	[2, 3, 4, 5, 32, 22, 46, 47, 51, 20, 52, 31, 53, 34, 54, 36, 29, 29, 28, 42, 17, 30, 37, 27, 38]	25	69.4%
	neutral	[43, 44, 35, 45, 50, 48, 39, 40, 41]	9	25.0%
	not described	[49, 7]	2	5.6%
Recommender type	end-to-end	[43, 35, 3, 5, 32, 45, 22, 46, 47, 50, 48, 51, 20, 52, 31, 53, 34, 54, 29, 29, 28, 40, 42, 17, 27, 41]	26	72.2%
	plug-in	[39, 37, 38]	3	8.3%
	re-ranking	[2, 44, 4, 49, 36, 30]	6	16.7%
	not described	[7]	1	2.8%

The first condition excludes pre-prints. The second condition excludes posters, demos, extended abstracts, workshops, and short papers. The third condition excludes conference venues classified with a ranking of "B" or lower according to the CORE Rank 2021. It also excludes journals classified with a ranking of "Q2" or lower according to the Scimago 2021 ranking in computer science.[2] Conversely, the third condition only accepts "A" or "A*" conferences or "Q1" journals. Lastly, condition 4 ensures that discovered papers are relevant to the IARS community. The condition excludes papers with keywords related to impressions without describing an IARS.

We searched through five popular academic search engines to discover, review, analyze, and select those papers deemed relevant to this work. We queried the ACM Digital Library, IEEE Xplore, SpringerLink, ScienceDirect, and Google Scholar academic search engines with the query "recommender system AND (impression OR exposure OR slate OR past recommendation OR previous recommendation)". This query matches the keyword recommender system and keywords related to IARS.

6.2. Recommendation Models

We identify a learning paradigm shift regarding the types of published recommenders. Notably, early works (before 2018) primarily published recommenders using heuristic, probabilistic, or shallow machine learning; moreover, only one paper uses reinforcement learning. We do not identify any recommendation model used to a greater extent than others amongst reviewed papers on those four groups.

When analyzing more recent papers, i.e., those published after 2018, we see that most papers describe either recommenders using deep learning or reinforcement learning. From those papers using deep learning architectures, four papers [30, 45, 22, 3] use multilayer

[2]CORE ranking: https://portal.core.edu.au/conf-ranks/ and Scimago ranking: https://www.scimagojr.com/journalrank.php?area=1700

perceptrons, two [44, 43] use the encoder-decoder architecture, three [48, 46, 47] use the two-tower framework, and four [49, 32, 2, 4] use multi-gate mixture of experts. Regarding those papers using reinforcement learning, two papers [54, 31] use multi-armed bandits, four [50, 51, 53, 35] use deep learning architectures adapted to the reinforcement learning paradigm, and two [5, 52] use other reinforcement learning strategies.

6.3. Data-Centric Taxonomy

Regarding the data-centric taxonomy, we see three major trends: papers only learn using impressions, only use features, or both. Moreover, only two papers sample from impressions; those papers also use impressions as negative signals.

Regarding features computed from impressions, the most common feature is the number of impressions a given user received for a given item. Other features used in a few papers include the average position of an item in an impression, the number of days between two impressions with the same item, average watch time of impressed items, among others.

6.4. Signal-Centric Taxonomy

On the signal-centric taxonomy, most papers assume a specific signal for non-interacted impressions; from those, the assumed signal is negative, i.e., the papers deem non-interacted impressions as negative items to user preferences.

Nine papers learn the user preferences for non-interacted impressions; seven assume non-interacted impressions are neutral signals, while two assume negative signals. For the former, they allow the recommendation models to learn such a preference. For the latter, they design their recommenders to learn how much the user dislikes non-interacted impressions.

6.5. Types of Impressions

In Section 4, we presented two types of impressions: global and contextual. Their difference is that the former does not associate interactions and impressions, while the latter provides such association.

Most papers, approximately a 66 % of them, use global impressions. Moreover, those papers use one particular type of global impressions: impressions as user-item-label triplets. Under this type of impressions, the recommender's input is a triplet containing a user identifier, an item identifier, and a binary label that indicates whether the item is an interacted or non-interacted impression.

Despite most papers using a particular type of impressions, Table 2 shows no statistically significant association between the type of impressions and the remaining properties or taxonomies.

6.6. Signals in Impressions

Close to 69 % of papers in the state-of-the-art follow the missing as negatives assumption, i.e., where non-interacted impressions are seen as negative signals of users' preferences.

Despite many papers following the missing as negatives assumption, some papers in the literature do not consider impressions as negative signals; instead, they dissect the signals within non-interacted impressions. In particular, Zhao et al. [7] performed a user study where they surveyed participants to express their preference toward non-interacted impressions. The results of such a study suggest that users are mostly *unaware* of all impressed items, i.e., users only scan part of the impression list. Moreover, the paper shows that only 5.8 % of non-impressed items were disliked by users. The results of those papers suggest that considering non-interacted impressions as *negative signals* may be an oversimplification of users' preferences.

6.7. Types of Recommenders

In Section 3, we defined three types of recommenders: end-to-end, plug-in, and re-ranking. Their main difference is how they generate an impression. Table 1 shows that most papers, approximately 72 %, describe end-to-end recommenders. Those papers describe a recommender that uses impressions and generates the final impression shown to the user without further processing by another recommender.

As shown in Table 2, the type of recommender has a statistically significant association with the taxonomies we propose: the model-centric, data-centric, and signal-centric taxonomies. Mainly, plug-in recommenders only appear in papers describing heuristic or statistical approaches. On the data-centric taxonomy, re-ranking recommenders do not sample from impressions. Lastly, on the signal-centric taxonomy, plug-in recommenders exclusively learn the signal from non-interacted impressions, while the majority of re-ranking ones assume a specific signal.

7. Conclusions

In this work, we provide a theoretical framework for the definition, characterization, and classification of impression-aware recommender system, i.e., recommenders learning users' preferences from impressions and interactions. In our first discussions, we provide the definitions of core concepts to IARS, such as **impression**,

Table 2
Association at 5% significance between taxonomies and properties of reviewed papers computed using Cramér's V [55]. NSS indicate non-statistically significant association. Large indicate a statistically significant and *large* association according to Cohen [56]. Diagonal omitted to avoid the redundant association of a taxonomy or property with itself.

	Model-Centric taxonomy	Data-Centric taxonomy	Signal-Centric Taxonomy	Impressions type	Impressions signal	Recommender type
Model-Centric taxonomy		Large	NSS	NSS	NSS	Large
Data-Centric taxonomy	Large		NSS	NSS	NSS	Large
Signal-Centric taxonomy	NSS	NSS		NSS	Large	Large
Impressions type	NSS	NSS	NSS		NSS	NSS
Impressions signal	NSS	NSS	Large	NSS		NSS
Recommender type	Large	Large	Large	NSS	NSS	

interactions, **user feedback**, among others. Also, we compare IARS to other types of recommenders, where we identified that IARS are a distinct type of recommenders despite their resemblance to others.

We present one of our main contributions to this work: the characterization of IARS in terms of their **properties** and **taxonomies**. For the former, we identify three properties that cover how recommenders consider user preferences to impressions, the kind of impressions they use, and how the recommender generates future impressions. For the latter, we propose three classifications, covering different aspects of IARS, i.e., the learning paradigm of the recommender, how the recommender uses impressions, and whether the recommender assumes or learns users' preference toward impressions.

Lastly, we classify papers belonging to the state-of-the-art under our proposed properties and taxonomies. We select relevant papers by applying specific selection criteria, focusing on papers published in high-level conferences and journals. In our classification, we discuss the general trends of the state-of-the-art under each property or taxonomy; at the same time, we indicate how they relate. Our study reveals a strong association between the signal of impressions and papers assuming or learning a specific signal in impressions. Furthermore, most papers assume a negative signal to non-interacted impressions, following the traditional *missing as negatives* assumption. However, as indicated by Zhao et al. [7], this assumption may be erroneous. Some of our previous works [57, 8, 9] align to both assumptions, i.e., in some cases impressions represent negative signals, while in others it does not. Such results call for future works that address this topic in particular.

References

[1] F. B. Pérez Maurera, M. Ferrari Dacrema, P. Castells, P. Cremonesi, Impression-aware recommender systems, CoRR abs/2308.07857 (2023). doi:10.48550/arXiv.2308.07857.

[2] Y. Xi, J. Lin, W. Liu, X. Dai, W. Zhang, R. Zhang, R. Tang, Y. Yu, A bird's-eye view of reranking: From list level to page level, in: T. Chua, H. W. Lauw, L. Si, E. Terzi, P. Tsaparas (Eds.), Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM 2023, Singapore, 27 February 2023 - 3 March 2023, ACM, 2023, pp. 1075–1083. doi:10.1145/3539597.3570399.

[3] Y. Zhang, Z. Chan, S. Xu, W. Bian, S. Han, H. Deng, B. Zheng, KEEP: an industrial pre-training framework for online recommendation via knowledge extraction and plugging, in: M. A. Hasan, L. Xiong (Eds.), Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17-21, 2022, ACM, 2022, pp. 3684–3693. doi:10.1145/3511808.3557106.

[4] X. Gong, Q. Feng, Y. Zhang, J. Qin, W. Ding, B. Li, P. Jiang, K. Gai, Real-time short video recommendation on mobile devices, in: M. A. Hasan, L. Xiong (Eds.), Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17-21, 2022, ACM, 2022, pp. 3103–3112. doi:10.1145/3511808.3557065.

[5] R. Xie, S. Zhang, R. Wang, F. Xia, L. Lin, A peep into the future: Adversarial future encoding in rec-

ommendation, in: K. S. Candan, H. Liu, L. Akoglu, X. L. Dong, J. Tang (Eds.), WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022, ACM, 2022, pp. 1177–1185. doi:10.1145/3488560.3498476.

[6] F. B. Pérez Maurera, M. Ferrari Dacrema, P. Castells, P. Cremonesi, Impressions in recommender systems: Present and future, in: Proceedings of the 13th Italian Information Retrieval Workshop 2023, Pisa, Italy, June 08-09, 2023, volume 3488 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 97–104. URL: https://ceur-ws.org/Vol-3448/paper-23.pdf.

[7] Q. Zhao, M. C. Willemsen, G. Adomavicius, F. M. Harper, J. A. Konstan, Interpreting user inaction in recommender systems, in: S. Pera, M. D. Ekstrand, X. Amatriain, J. O'Donovan (Eds.), Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018, ACM, 2018, pp. 40–48. doi:10.1145/3240323.3240366.

[8] F. B. Pérez Maurera, M. Ferrari Dacrema, P. Cremonesi, Replication of recommender systems with impressions, in: G. Pasi, P. Cremonesi, S. Orlando, M. Zanker, D. Massimo, G. Turati (Eds.), Proceedings of the 12th Italian Information Retrieval Workshop 2022, Milan, Italy, June 29-30, 2022, volume 3177 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022. URL: http://ceur-ws.org/Vol-3177/paper20.pdf.

[9] F. B. Pérez Maurera, M. Ferrari Dacrema, P. Cremonesi, Towards the evaluation of recommender systems with impressions, in: J. Golbeck, F. M. Harper, V. Murdock, M. D. Ekstrand, B. Shapira, J. Basilico, K. T. Lundgaard, E. Oldridge (Eds.), RecSys '22: Sixteenth ACM Conference on Recommender Systems, Seattle, WA, USA, September 18 - 23, 2022, ACM, 2022, pp. 610–615. doi:10.1145/3523227.3551483.

[10] M. Quadrana, P. Cremonesi, D. Jannach, Sequence-aware recommender systems, ACM Comput. Surv. 51 (2018) 66:1–66:36. doi:10.1145/3190616.

[11] G. Adomavicius, A. Tuzhilin, Context-aware recommender systems, in: F. Ricci, L. Rokach, B. Shapira (Eds.), Recommender Systems Handbook, Springer, 2015, pp. 191–226. doi:10.1007/978-1-4899-7637-6_6.

[12] W. Liu, Y. Xi, J. Qin, F. Sun, B. Chen, W. Zhang, R. Zhang, R. Tang, Neural re-ranking in multi-stage recommender systems: A review, in: L. D. Raedt (Ed.), Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022, ijcai.org, 2022, pp. 5512–5520. doi:10.24963/ijcai.2022/771.

[13] F. Wu, Y. Qiao, J. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu, M. Zhou, MIND: A large-scale dataset for news recommendation, in: D. Jurafsky, J. Chai, N. Schluter, J. R. Tetreault (Eds.), Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020, Association for Computational Linguistics, 2020, pp. 3597–3606. doi:10.18653/v1/2020.acl-main.331.

[14] L. Li, W. Chu, J. Langford, X. Wang, Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms, in: I. King, W. Nejdl, H. Li (Eds.), Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM 2011, Hong Kong, China, February 9-12, 2011, ACM, 2011, pp. 297–306. doi:10.1145/1935826.1935878.

[15] W. Chu, S. Park, T. Beaupre, N. Motgi, A. Phadke, S. Chakraborty, J. Zachariah, A case study of behavior-driven conjoint analysis on yahoo!: front page today module, in: J. F. E. IV, F. Fogelman-Soulié, P. A. Flach, M. J. Zaki (Eds.), Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July 1, 2009, ACM, 2009, pp. 1097–1104. doi:10.1145/1557019.1557138.

[16] C. Gentile, S. Li, G. Zappella, Online clustering of bandits, in: Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014, volume 32 of *JMLR Workshop and Conference Proceedings*, JMLR.org, 2014, pp. 757–765. URL: http://proceedings.mlr.press/v32/gentile14.html.

[17] S. Li, A. Karatzoglou, C. Gentile, Collaborative filtering bandits, in: R. Perego, F. Sebastiani, J. A. Aslam, I. Ruthven, J. Zobel (Eds.), Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016, ACM, 2016, pp. 539–548. doi:10.1145/2911451.2911548.

[18] S. Sidana, C. Laclau, M. Amini, Learning to recommend diverse items over implicit feedback on PANDOR, in: S. Pera, M. D. Ekstrand, X. Amatriain, J. O'Donovan (Eds.), Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018, ACM, 2018, pp. 427–431. doi:10.1145/3240323.3240400.

[19] Q. Shen, H. Wen, W. Tao, J. Zhang, F. Lv, Z. Chen, Z. Li, Deep interest highlight network for click-through rate prediction in trigger-induced recommendation, in: F. Laforest, R. Troncy, E. Simperl, D. Agarwal, A. Gionis, I. Herman, L. Médini (Eds.), WWW '22: The ACM Web Conference 2022, Virtual

Event, Lyon, France, April 25 - 29, 2022, ACM, 2022, pp. 422–430. doi:10.1145/3485447.3511970.

[20] F. B. Pérez Maurera, M. Ferrari Dacrema, L. Saule, M. Scriminaci, P. Cremonesi, Contentwise impressions: An industrial dataset with impressions included, in: M. d'Aquin, S. Dietze, C. Hauff, E. Curry, P. Cudré-Mauroux (Eds.), CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020, ACM, 2020, pp. 3093–3100. doi:10.1145/3340531.3412774.

[21] J. Wang, W. Ma, J. Li, H. Lu, M. Zhang, B. Li, Y. Liu, P. Jiang, S. Ma, Make fairness more fair: Fair item utility estimation and exposure re-distribution, in: A. Zhang, H. Rangwala (Eds.), KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022, ACM, 2022, pp. 1868–1877. doi:10.1145/3534678.3539354.

[22] C. Zhu, P. Du, W. Zhang, Y. Yu, Y. Cao, Combo-fashion: Fashion clothes matching CTR prediction with item history, in: A. Zhang, H. Rangwala (Eds.), KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022, ACM, 2022, pp. 4621–4629. doi:10.1145/3534678.3539101.

[23] S. Eide, D. S. Leslie, A. Frigessi, J. Rishaug, H. Jenssen, S. Verrewaere, Finn.no slates dataset: A new sequential dataset logging interactions, all viewed items and click responses/no-click for recommender systems research, in: H. J. C. Pampín, M. A. Larson, M. C. Willemsen, J. A. Konstan, J. J. McAuley, J. Garcia-Gathright, B. Huurnink, E. Oldridge (Eds.), RecSys '21: Fifteenth ACM Conference on Recommender Systems, Amsterdam, The Netherlands, 27 September 2021 - 1 October 2021, ACM, 2021, pp. 556–558. doi:10.1145/3460231.3474607.

[24] S. Eide, D. S. Leslie, A. Frigessi, Dynamic slate recommendation with gated recurrent units and thompson sampling, CoRR abs/2104.15046 (2021). URL: https://arxiv.org/abs/2104.15046.

[25] X. Ma, L. Zhao, G. Huang, Z. Wang, Z. Hu, X. Zhu, K. Gai, Entire space multi-task model: An effective approach for estimating post-click conversion rate, in: K. Collins-Thompson, Q. Mei, B. D. Davison, Y. Liu, E. Yilmaz (Eds.), The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018, ACM, 2018, pp. 1137–1140. doi:10.1145/3209978.3210104.

[26] B. Pradel, N. Usunier, P. Gallinari, Ranking with non-random missing ratings: influence of popularity and positivity on evaluation metrics, in: P. Cunningham, N. J. Hurley, I. Guy, S. S. Anand (Eds.), Sixth ACM Conference on Recommender Systems, RecSys '12, Dublin, Ireland, September 9-13, 2012, ACM, 2012, pp. 147–154. doi:10.1145/2365952.2365982.

[27] N. Buchbinder, M. Feldman, A. Ghosh, J. Naor, Frequency capping in online advertising, J. Sched. 17 (2014) 385–398. doi:10.1007/s10951-014-0367-z.

[28] X. Zhang, Y. Zhou, Y. Ma, B. Chen, L. Zhang, D. Agarwal, Glmix: Generalized linear mixed models for large-scale response prediction, in: B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, R. Rastogi (Eds.), Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016, ACM, 2016, pp. 363–372. doi:10.1145/2939672.2939684.

[29] D. C. Liu, S. K. Rogers, R. Shiau, D. Kislyuk, K. C. Ma, Z. Zhong, J. Liu, Y. Jing, Related pins at pinterest: The evolution of a real-world recommender system, in: R. Barrett, R. Cummings, E. Agichtein, E. Gabrilovich (Eds.), Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017, ACM, 2017, pp. 583–592. doi:10.1145/3041021.3054202.

[30] P. Covington, J. Adams, E. Sargin, Deep neural networks for youtube recommendations, in: S. Sen, W. Geyer, J. Freyne, P. Castells (Eds.), Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15-19, 2016, ACM, 2016, pp. 191–198. doi:10.1145/2959100.2959190.

[31] A. Gruson, P. Chandar, C. Charbuillet, J. McInerney, S. Hansen, D. Tardieu, B. Carterette, Offline evaluation to make decisions about playlist recommendation algorithms, in: J. S. Culpepper, A. Moffat, P. N. Bennett, K. Lerman (Eds.), Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019, ACM, 2019, pp. 420–428. doi:10.1145/3289600.3291027.

[32] N. Ma, M. Ispir, Y. Li, Y. Yang, Z. Chen, D. Z. Cheng, L. Nie, K. Barman, An online multi-task learning framework for google feed ads auction models, in: A. Zhang, H. Rangwala (Eds.), KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022, ACM, 2022, pp. 3477–3485. doi:10.1145/3534678.3539055.

[33] S. Rendle, C. Freudenthaler, Z. Gantner, L. Schmidt-Thieme, BPR: bayesian personalized ranking from implicit feedback, in: J. A. Bilmes, A. Y. Ng (Eds.), UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21,

2009, AUAI Press, 2009, pp. 452–461. URL: https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=1630&proceeding_id=25.

[34] M. Aharon, Y. Kaplan, R. Levy, O. Somekh, A. Blanc, N. Eshel, A. Shahar, A. Singer, A. Zlotnik, Soft frequency capping for improved ad click prediction in yahoo gemini native, in: W. Zhu, D. Tao, X. Cheng, P. Cui, E. A. Rundensteiner, D. Carmel, Q. He, J. X. Yu (Eds.), Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019, ACM, 2019, pp. 2793–2801. doi:10.1145/3357384.3357801.

[35] R. Deffayet, T. Thonet, J. Renders, M. de Rijke, Generative slate recommendation with reinforcement learning, in: T. Chua, H. W. Lauw, L. Si, E. Terzi, P. Tsaparas (Eds.), Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM 2023, Singapore, 27 February 2023 - 3 March 2023, ACM, 2023, pp. 580–588. doi:10.1145/3539597.3570412.

[36] Q. Zhao, G. Adomavicius, F. M. Harper, M. C. Willemsen, J. A. Konstan, Toward better interactions in recommender systems: Cycling and serpentining approaches for top-n item lists, in: C. P. Lee, S. E. Poltrock, L. Barkhuus, M. Borges, W. A. Kellogg (Eds.), Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW 2017, Portland, OR, USA, February 25 - March 1, 2017, ACM, 2017, pp. 1444–1453. doi:10.1145/2998181.2998211.

[37] P. Lee, L. V. S. Lakshmanan, M. Tiwari, S. Shah, Modeling impression discounting in large-scale recommender systems, in: S. A. Macskassy, C. Perlich, J. Leskovec, W. Wang, R. Ghani (Eds.), The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014, ACM, 2014, pp. 1837–1846. doi:10.1145/2623330.2623356.

[38] D. Agarwal, B. Chen, R. Gupta, J. Hartman, Q. He, A. Iyer, S. Kolar, Y. Ma, P. Shivaswamy, A. Singh, L. Zhang, Activity ranking in linkedin feed, in: S. A. Macskassy, C. Perlich, J. Leskovec, W. Wang, R. Ghani (Eds.), The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014, ACM, 2014, pp. 1603–1612. doi:10.1145/2623330.2623362.

[39] F. Borisyuk, L. Zhang, K. Kenthapadi, Lijar: A system for job application redistribution towards efficient career marketplace, in: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017, ACM, 2017, pp. 1397–1406. doi:10.1145/3097983.3098028.

[40] C. Wu, C. V. Alvino, A. J. Smola, J. Basilico, Using navigation to improve recommendations in real-time, in: S. Sen, W. Geyer, J. Freyne, P. Castells (Eds.), Proceedings of the 10th ACM Conference on Recommender Systems, Boston, MA, USA, September 15-19, 2016, ACM, 2016, pp. 341–348. doi:10.1145/2959100.2959174.

[41] D. Agarwal, B. Chen, P. Elango, Spatio-temporal models for estimating click-through rate, in: Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009, ACM, 2009, pp. 21–30. doi:10.1145/1526709.1526713.

[42] H. Ma, X. Liu, Z. Shen, User fatigue in online news recommendation, in: J. Bourdeau, J. Hendler, R. Nkambou, I. Horrocks, B. Y. Zhao (Eds.), Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016, ACM, 2016, pp. 1363–1372. doi:10.1145/2872427.2874813.

[43] X. Xin, J. Yang, H. Wang, J. Ma, P. Ren, H. Luo, X. Shi, Z. Chen, Z. Ren, On the user behavior leakage from recommender system exposure, ACM Trans. Inf. Syst. 41 (2023). doi:10.1145/3568954.

[44] Y. Ren, X. Han, X. Zhao, S. Zhang, Y. Zhang, Slate-aware ranking for recommendation, in: T. Chua, H. W. Lauw, L. Si, E. Terzi, P. Tsaparas (Eds.), Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM 2023, Singapore, 27 February 2023 - 3 March 2023, ACM, 2023, pp. 499–507. doi:10.1145/3539597.3570380.

[45] R. Zhan, C. Pei, Q. Su, J. Wen, X. Wang, G. Mu, D. Zheng, P. Jiang, K. Gai, Deconfounding duration bias in watch-time prediction for video recommendation, in: A. Zhang, H. Rangwala (Eds.), KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022, ACM, 2022, pp. 4472–4481. doi:10.1145/3534678.3539092.

[46] H. Wang, T. Chang, T. Liu, J. Huang, Z. Chen, C. Yu, R. Li, W. Chu, ESCM2: entire space counterfactual multi-task model for post-click conversion rate estimation, in: E. Amigó, P. Castells, J. Gonzalo, B. Carterette, J. S. Culpepper, G. Kazai (Eds.), SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022, ACM, 2022, pp. 363–372. doi:10.1145/3477495.3531972.

[47] S. Gong, K. Q. Zhu, Positive, negative and neutral: Modeling implicit feedback in session-based news recommendation, in: E. Amigó, P. Castells, J. Gonzalo, B. Carterette, J. S. Culpepper, G. Kazai (Eds.), SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in

Information Retrieval, Madrid, Spain, July 11 - 15, 2022, ACM, 2022, pp. 1185–1195. doi:10.1145/3477495.3532040.

[48] Z. Chen, J. Wu, C. Li, J. Chen, R. Xiao, B. Zhao, Co-training disentangled domain adaptation network for leveraging popularity bias in recommenders, in: E. Amigó, P. Castells, J. Gonzalo, B. Carterette, J. S. Culpepper, G. Kazai (Eds.), SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022, ACM, 2022, pp. 60–69. doi:10.1145/3477495.3531952.

[49] Z. Zhao, L. Hong, L. Wei, J. Chen, A. Nath, S. Andrews, A. Kumthekar, M. Sathiamoorthy, X. Yi, E. H. Chi, Recommending what video to watch next: a multitask ranking system, in: T. Bogers, A. Said, P. Brusilovsky, D. Tikk (Eds.), Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019, ACM, 2019, pp. 43–51. doi:10.1145/3298689.3346997.

[50] Y. Ge, X. Zhao, L. Yu, S. Paul, D. Hu, C. Hsieh, Y. Zhang, Toward pareto efficient fairness-utility trade-off in recommendation through reinforcement learning, in: K. S. Candan, H. Liu, L. Akoglu, X. L. Dong, J. Tang (Eds.), WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022, ACM, 2022, pp. 316–324. doi:10.1145/3488560.3498487.

[51] R. Xie, S. Zhang, R. Wang, F. Xia, L. Lin, Explore, filter and distill: Distilled reinforcement learning in recommendation, in: G. Demartini, G. Zuccon, J. S. Culpepper, Z. Huang, H. Tong (Eds.), CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, ACM, 2021, pp. 4243–4252. doi:10.1145/3459637.3481917.

[52] C. Pei, X. Yang, Q. Cui, X. Lin, F. Sun, P. Jiang, W. Ou, Y. Zhang, Value-aware recommendation based on reinforcement profit maximization, in: L. Liu, R. W. White, A. Mantrach, F. Silvestri, J. J. McAuley, R. Baeza-Yates, L. Zia (Eds.), The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019, ACM, 2019, pp. 3123–3129. doi:10.1145/3308558.3313404.

[53] M. Chen, A. Beutel, P. Covington, S. Jain, F. Belletti, E. H. Chi, Top-k off-policy correction for a REINFORCE recommender system, in: J. S. Culpepper, A. Moffat, P. N. Bennett, K. Lerman (Eds.), Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019, ACM, 2019, pp. 456–464. doi:10.1145/3289600.3290999.

[54] J. McInerney, B. Lacker, S. Hansen, K. Higley, H. Bouchard, A. Gruson, R. Mehrotra, Explore, exploit, and explain: personalizing explainable recommendations with bandits, in: S. Pera, M. D. Ekstrand, X. Amatriain, J. O'Donovan (Eds.), Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018, ACM, 2018, pp. 31–39. doi:10.1145/3240323.3240354.

[55] H. Cramér, Mathematical methods of statistics, Princeton university press, 1946.

[56] J. Cohen, Statistical power analysis for the behavioral sciences, 2 ed., Lawrence Erlbaum Associates, 1988.

[57] F. B. Pérez Maurera, M. Ferrari Dacrema, P. Castells, P. Cremonesi, Incorporating impressions to graph-based recommenders, in: LERI 2023: Proceedings of the 1st Workshop on Learning and Evaluating Recommendations with Impressions, Singapore, September 19, 2023, CEUR Workshop Proceedings, CEUR-WS.org, 2023.