Expert Systems With Applications 249 (2024) 123642

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Model-based approaches to profit-aware recommendation

Alvise De Biasio [a,b], Dietmar Jannach [c,*], Nicolò Navarin [a]

[a] Department of Mathematics, University of Padova, Via Trieste 63, Padova, 35131 PD, Italy
[b] R&D Department, estilos srl, Via Ca' Marcello 67/D, Venezia, 30172 VE, Italy
[c] Department of AI & Cybersecurity, University of Klagenfurt, Universitätsstraße 65–67, 9020 Klagenfurt, Austria

ARTICLE INFO

Keywords:
Recommender systems
Economic recommendations
Business value
Profit awareness
Multi-stakeholder systems

ABSTRACT

Recommender systems are traditionally optimized to facilitate content discovery for consumers by ranking items based on predicted relevance. As such, these systems often do not consider the varying profitability of items for service providers. Since the purpose of recommender systems is usually to create value for both consumers and providers, we hypothesize that integrating *profit awareness* into recommender systems, considering both consumer relevance and provider profitability, can enhance recommendation outcomes from a provider point of view. In this study, we design and evaluate novel *modeling* approaches for different families of collaborative filtering algorithms to overcome the existing limitations of conventional reranking methods. Specifically, we show how to embed the business value perspective directly into the loss functions during model training. Through empirical evaluations on three datasets, we show that our proposed models effectively generate recommendations that balance profitability and relevance. Overall, our analyses indicate that these models offer a promising alternative to traditional reranking approaches, particularly because they exhibit improved efficiency during prediction.

1. Introduction

Recommender systems (*RSs*) (Jannach, Zanker, Felfernig, & Friedrich, 2010) are currently widely used in modern online services, including e-commerce shops, media streaming platforms, and social media sites. Commonly, the primary purpose of such systems is to support users in locating relevant content in situations of information overload (Bollen, Knijnenburg, Willemsen, & Graus, 2010). In such settings, the central task of the underlying algorithms is thus to predict the relevance of individual content items for a given user, which is then used to create a relevance-ranked list of recommendations (Zhao, Lin, Feng, Wang, & Wen, 2023).

Traditionally, research on recommender systems has focused on developing new machine learning models to improve the accuracy of predicting relevant items (Gunawardana, Shani, & Yogev, 2022). A common assumption in the literature is that focusing recommendations on the most relevant items for consumers, i.e., on consumer value, will directly or indirectly create value for the provider of the recommendations as well (Chen, Wu, & Yoon, 2004). A direct effect for organizations could, for example, be increased sales numbers on an e-commerce site when consumers discover more relevant items through recommendations (Lee & Hosanagar, 2014; Panniello, Gorgoglione and Tuzhilin, 2016). An indirect effect could be increased engagement and thus higher customer retention, that results from the use of a recommender system (Cavenaghi et al., 2022; Cooke, Sujan, Sujan, & Weitz, 2002; Gomez-Uribe & Hunt, 2016; Gorgoglione, Panniello, & Tuzhilin, 2011).

In recent years, researchers have gained awareness that optimizing a system exclusively for consumer value may be suboptimal or insufficient (De Biasio, Monaro, Oneto, Ballan and Navarin, 2023; Montagna, De Biasio, Navarin, Aiolli, et al., 2023; Panniello, Hill and Gorgoglione, 2016). In fact, there are various scenarios in which the objectives of different stakeholders should be considered in parallel (Basu, 2021; Ren & Zhang, 2021; Zhan et al., 2021). A characterizing feature of such *multistakeholder* recommendation (Abdollahpouri et al., 2020; Abdollahpouri & Burke, 2022; Burke, Abdollahpouri, Malthouse, Thai, & Zhang, 2019) problem settings is that the often competing objectives of the different stakeholders must be balanced to achieve sustainable success through the service.

Profit-aware approaches are special examples of multistakeholder recommender systems (Jannach & Adomavicius, 2017). These approaches operate under the realistic assumption that not every recommendable item is equally profitable for the service provider (De Biasio, Montagna, Aiolli and Navarin, 2023), e.g., an e-commerce marketplace, in the case of a purchase. A service provider might, therefore, be interested in generating recommendations that balance relevance for

* Corresponding author.
E-mail addresses: alvise.debiasio@phd.unipd.it (A. De Biasio), dietmar.jannach@aau.at (D. Jannach), nnavarin@math.unipd.it (N. Navarin).

https://doi.org/10.1016/j.eswa.2024.123642
Received 6 September 2023; Received in revised form 14 February 2024; Accepted 8 March 2024
Available online 15 March 2024

the consumer and profitability for the provider (Jannach & Adomavicius, 2017). Finding the right balance, however, is crucial in such a situation. An overly strong focus on profitability may result in frequent recommendations of irrelevant items (Chen et al., 2007; Chen, Hsu, Chen, & Hsu, 2008), which may lead to reduced consumer trust over time (Basu, 2021; Ghanem et al., 2022; Hosanagar, 2008). However, not considering monetary aspects may lead to missed opportunities for increasing profits (Basu, 2021; Krasnodebski & Dines, 2016; Panniello, Hill et al., 2016). If, for example, one recommendable item is assumed to be equally relevant for a consumer as a similar alternative, it may be beneficial for the provider to recommend the item that results in a higher profit (Zhang, Chen, Wang, & Si, 2017).

Several profit-aware, or more generally, value-aware recommendation approaches have been proposed in recent years; see De Biasio, Montagna et al. (2023), De Biasio, Navarin and Jannach (2023) and Jannach and Jugovac (2019) for related surveys. Many of these approaches rely on *reranking* techniques (Chen et al., 2007; Chen et al., 2008; Das, Mathieu, & Ricketts, 2009; Ghanem et al., 2022; Kompan, Gaspar, Macina, Cimerman, & Bielikova, 2022; Malthouse, Vakeel, Hessary, Burke, & Fudurić, 2019; Wang & Wu, 2009, 2012), where a given baseline recommendation list, which is optimized for consumer relevance, is postprocessed to promote items with greater profitability. Commonly, certain guardrails are implemented in the reranking process to avoid items with little consumer relevance appearing in the highest places on the reranked lists (Das et al., 2009; Wang & Wu, 2009, 2012). A general advantage of such postprocessing techniques is that any recommendation model can be used to generate a relevance-optimized baseline list (Adomavicius & YoungOk, 2012). Furthermore, on high-traffic e-commerce sites, we assume that postprocessing every single recommendation list may easily lead to significant computational overhead (Yang, Xu, Jones, & Samatova, 2017). Moreover, another assumption is that the effectiveness of the reranking process may be limited when the guardrails are set too narrow (Ghanem et al., 2022).

In this work, we hypothesize that *model-based* (or in-processing) (De Biasio, Montagna et al., 2023) approaches for building profit-aware recommender systems (Jannach & Adomavicius, 2017) are particularly beneficial for overcoming the limitations of conventional reranking techniques. In such model-based approaches (Cai & Zhu, 2019; Li et al., 2021; Piton, Blanchard, & Guillet, 2011; Qu, Zhu, Liu, Liu, & Xiong, 2014), the task of balancing the competing goals of consumer relevance and provider profitability (Ghanem et al., 2022; Pei et al., 2019) is embedded directly in the learning process. Specifically, we propose profit-aware loss functions for three important families of collaborative filtering techniques: matrix factorization (Koren & Bell, 2021; Koren, Bell, & Volinsky, 2009), learning-to-rank (Rendle, 2022; Rendle, Freudenthaler, Gantner, & Schmidt-Thieme, 2009), and neural models (He et al., 2017). Moreover, we consider a profit-aware variant of the model-free nearest neighbors recommendation approach (Nikolakopoulos, Ning, Desrosiers, & Karypis, 2022) that was recently proposed in a different context in Cai and Zhu (2019). Experiments on three real-world e-commerce datasets (Microsoft Corporation, 1998; Ni, Li, & McAuley, 2019; Zhang, Zhao, & LeCun, 2015) reveal that model-based approaches effectively balance consumer and provider values. Moreover, we compare our models with recent postprocessing techniques (Chen et al., 2008; Ghanem et al., 2022; Jannach & Adomavicius, 2017). This additional comparison reveals that model-based approaches can be favorable alternatives to existing reranking approaches because they obtain comparable or better performance in recommending more profitable yet relevant items to users (with respect to the baselines) but with lower prediction times. Because of this greater efficiency, in-processing models might, therefore, be preferred in practical cases where postprocessing methods might be inapplicable, e.g., considering large-scale production systems with millions of active users and catalog items.

The contributions of this paper can be summarized as follows:

- We studied how to optimize the profitability of recommendations by integrating the objective functions of four different families of state-of-the-art recommender systems widely used in industry, i.e., nearest neighbors, matrix factorization, learning-to-rank, and neural models.
- We compared our in-processing models with several of the most commonly used reranking approaches on three real-world datasets, demonstrating how the proposed models may represent more efficient alternatives that achieve comparable or better performance.

The remainder of the paper is organized as follows. In Section 2, we provide additional background and discuss previous related works. The technical approaches for embedding profit awareness for different families of collaborative filtering techniques are presented in Section 3. Section 4 describes the experimental setting of our research and the outcomes of the evaluation. The paper ends with a discussion of the findings in Section 5 and an outlook on future works in Section 6. Finally, Section 7 concludes the article.

2. Background and related work

In this section, we first review the different families of recommendation algorithms that are relevant to our work (Section 2.1) and subsequently discuss earlier work in the area of profit-aware and value-aware recommender systems (Section 2.2).

2.1. Recommendation algorithms

Countless recommendation algorithms have been proposed since the early 1990s (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994; Resnick & Varian, 1997). Traditionally, we distinguish between collaborative filtering algorithms, content-based approaches, and hybrid systems (Adomavicius & Tuzhilin, 2005; Jannach et al., 2010). *Collaborative filtering* (*CF*) (Su & Khoshgoftaar, 2009) systems base their recommendations solely on knowledge about the past behavior of a community of users, e.g., their previous purchases on an online shop or the feedback they provide on items on a media streaming site. CF systems are currently widely used in industry (Amatriain & Basilico, 2016; Gomez-Uribe & Hunt, 2016) and are the basis for our work on profit-aware recommendations.

Content-based (*CB*) systems (Lops, Gemmis, & Semeraro, 2011) have roots in *information retrieval* (*IR*) (Bellogín & Said, 2019; Kobayashi & Takeda, 2000) and rely on (meta-)information about the available items and the past content preferences of individual users. Such systems also have practical applications, particularly when no large user community exists, which is required to build a collaborative system.

Since both CF and CB systems can have limitations (Adomavicius & Tuzhilin, 2005), various methods of combining these approaches in a *hybrid system* have been proposed (Burke, 2007). The most common hybrid systems today are likely collaborative systems that use additional knowledge about items as *side information*; see (Sun et al., 2019; Villegas, Sánchez, Díaz-Cely, & Tamura, 2018) for two comprehensive surveys. Profit-aware and value-aware recommender systems (De Biasio, Montagna et al., 2023) also fall into this category. These systems commonly aim to leverage the power of collaborative filtering (Su & Khoshgoftaar, 2009), but they additionally account for item prices or profits (Jannach & Adomavicius, 2017) when creating the final recommendation lists to be presented to users.

The following main algorithmic approaches can be identified within the family of collaborative filtering systems.

- *Nearest neighbors* techniques were used in the earliest recommendation systems both in academia and industry (Linden, Smith, & York, 2003; Nikolakopoulos et al., 2022; Resnick et al., 1994). While such approaches are not model-based and conceptually simple, they can often achieve competitive results in terms of prediction accuracy, at least for small datasets (Anelli, Bellogín, Noia, Jannach, & Pomo, 2022).

- *Matrix factorization* approaches were initially explored in the late 1990s (Billsus & Pazzani, 1998), and they became state-of-the-art (Koren, 2008; Koren & Bell, 2021; Koren et al., 2009) in the context of the Netflix Prize competition (Bennett, Lanning, et al., 2007). Despite the recent wave of modern deep learning algorithms (He et al., 2017), matrix factorization methods are still relevant today (Rendle, Krichene, Zhang, & Anderson, 2020), as they often perform well in purely collaborative filtering approaches (Rendle, Krichene, Zhang, & Koren, 2022).
- *Learning-to-rank* techniques such as *Bayesian Personalized Ranking* (Rendle, 2022; Rendle et al., 2009) became popular around the 2010s when the community increasingly started to move away from focusing on the rating prediction problem and started to target the *top-k* recommendation problem more directly (Deshpande & Karypis, 2004; Gunawardana et al., 2022; Zhao et al., 2023). Therefore, the goal of such approaches is not to make relevance predictions on an absolute scale but rather to find an optimal ranking of the recommended items.
- The use of *neural networks* (deep learning techniques) for collaborative filtering dates back at least to 2007 (Salakhutdinov, Mnih, & Hinton, 2007). Today, deep learning techniques such as *Neural Collaborative Filtering* (He et al., 2017) are considered to be state-of-the-art techniques. One main advantage of such systems in practice is that various types of side information can be easily integrated into such networks (Steck et al., 2021).

In our present work, we show in Section 3 how each of the discussed model-based approaches can be made *profit-aware*.

2.2. Value-aware recommendations

There are several types of recommendation algorithms in the literature (Adomavicius & Tuzhilin, 2005; Jannach et al., 2010), each designed primarily to serve certain user needs. However, it has recently emerged that, especially in certain application domains (e.g., in e-commerce) (Schafer, Konstan, & Riedl, 1999, 2001; Xiao & Benbasat, 2007), it may be necessary to consider not only the customers' goals (Adomavicius & Tuzhilin, 2005; Amatriain & Basilico, 2016; Gomez-Uribe & Hunt, 2016) but also the potentially competing objectives of other stakeholders involved, such as organizations (e.g., Amazon) and suppliers (e.g., Amazon's merchants) (Abdollahpouri et al., 2020; Abdollahpouri & Burke, 2022; Burke et al., 2019; Jannach & Adomavicius, 2016). In particular, a certain subset of such multistakeholder systems, called *value-aware recommender systems* (*VARSs*) (De Biasio, Montagna et al., 2023), aims to more directly target certain organizational purposes to optimize one or more of the following business value categories (Jannach & Jugovac, 2019):

- the number of *user clicks*, often measured through the *click-through rate*;
- the degree of *consumer adoption* of the system, often measured via the *conversion rate*;
- the *revenue* generated by the *sales* of the products and services sold by the firm;
- the *sales distribution*, e.g., shifting the items being recommended toward the most profitable;
- the degree of *user engagement* with the platform as an indicator of consumer satisfaction.

The specific type of value optimized may depend on various variables (Jannach & Jugovac, 2019), often related to particular organizational business strategies (Chen, Chou, & Kauffman, 2009; Herder, 2019; Hoffman & Novak, 2005; Resnick & Varian, 1997). In some cases, it may be advantageous for the organization to maximize the conversion rate of recommendations (Goyal & Lakshmanan, 2012; Karlsson & Nilsson, 2013; Ren, Kauffman, & King, 2019), e.g., to increase the number of consumers on the platform. In other circumstances, it may be helpful to optimize user engagement (Gomez-Uribe & Hunt, 2016; Maslowska, Malthouse, & Hollebeek, 2022; Zou et al., 2019), e.g., to retain acquired consumers and guarantee stable cash flow levels. However, companies often want to optimize a specific type of business value if doing so will contribute to increasing corporate profitability (De Biasio, Montagna et al., 2023; Jannach & Adomavicius, 2017; Jannach & Jugovac, 2019), either in the short (Chen et al., 2007; Chen et al., 2008) or long term (Guo et al., 2021; He, Liu, Zhao, Liu, & Tang, 2022; Pei et al., 2019; Zhang, Zhao et al., 2022). In the following, we mainly focus on short-term profit optimization considering the particular class of profit-aware recommender systems (Jannach & Adomavicius, 2017).

2.2.1. Profit-aware recommender systems

Profit-aware recommender systems (*PARSs*) (Jannach & Adomavicius, 2017) have recently emerged in the literature as tools to directly optimizing profit using different techniques. It is possible to divide the underlying approaches into two main macrocategories (De Biasio, Montagna et al., 2023) called *postprocessing* and *in-processing* methods, depending on when the profit optimization step is applied. Postprocessing methods can be applied to the outputs of any recommendation algorithm, treating the algorithm as a black box. In-processing methods, in contrast, can be used to optimize profits directly during learning.

Postprocessing profit-aware approaches. Of the two types of approaches, postprocessing is the currently the most widely employed in the literature (Azaria et al., 2013; Chen et al., 2007; Chen et al., 2008; Das et al., 2009; Demirezen & Kumar, 2016; Ghanem et al., 2022; Kompan et al., 2022; Malthouse et al., 2019; Seymen, Sachs, & Malthouse, 2022; Wang & Wu, 2009, 2012; Zhang et al., 2017). Although postprocessing algorithms require performing reranking operations at prediction time, which typically involve some computational overhead, these approaches are conceptually simple and can be applied to all backbone recommenders (Adomavicius & YoungOk, 2012). For example, in several early studies (Chen et al., 2007; Chen et al., 2008), it was proposed that the predicted scores of a user-based nearest neighbors collaborative filtering algorithm (Nikolakopoulos et al., 2022) should be weighted by item profitability to optimize the average expected profit. However, as acknowledged many times in the literature, e.g., in De Biasio, Montagna et al. (2023) and Jannach and Adomavicius (2017), although it may be possible to provide recommendations of higher business value using this approach, a company may risk losing its consumers if the system recommends only profitable items that are not of interest to the users. Hence, other studies (Das et al., 2009; Ghanem et al., 2022; Jannach & Adomavicius, 2017; Kompan et al., 2022; Malthouse et al., 2019; Wang & Wu, 2009, 2012) often employ constrained or multiobjective variations of the previous algorithm exploiting additional hyperparameters to properly balance consumers' interests with organizational profit (*profitability/relevance tradeoff*) (Ghanem et al., 2022; Jannach & Adomavicius, 2017). For example, in some studies (Das et al., 2009; Jannach & Adomavicius, 2017; Wang & Wu, 2009, 2012), a static threshold on the predicted scores is used to consider only the most potentially relevant items in the reranking phase. In other works, the predicted scores and the profit are balanced using a regularizer (Azaria et al., 2013; Ghanem et al., 2022) that may give more or less weight to the recommendation provider profit or consumer utility. Moreover, in some instances, reranking algorithms are used to target specific application contexts (Demirezen & Kumar, 2016; Malthouse et al., 2019; Seymen et al., 2022; Zhang et al., 2017) by embracing broader perspectives. For example, some studies have focused on advertising (Malthouse et al., 2019; Zhang et al., 2017) by considering the perspective of suppliers who pay a certain fee to see their products advertised on the platform. Other studies have addressed inventory management issues (Demirezen & Kumar, 2016; Seymen et al., 2022) to recommend profitable items while avoiding out-of-stock or perishability issues related to consumer goods.

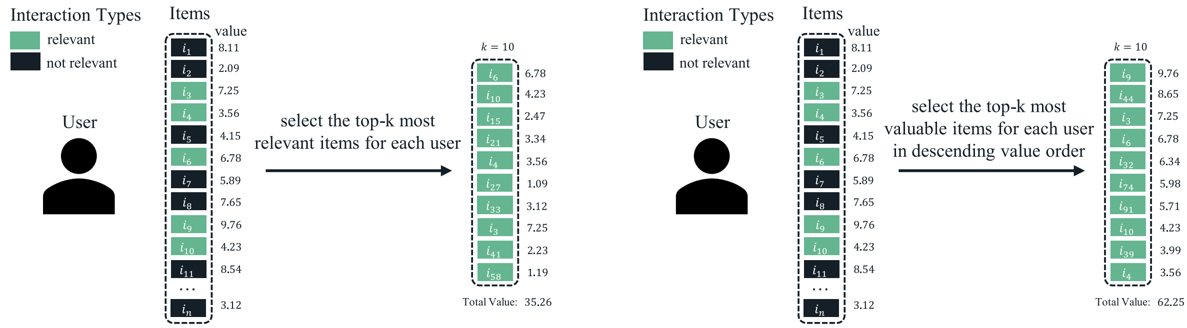

(a) Top-k recommendation problem. (b) Top-k value maximization problem.

Fig. 1. Comparison of the top-k recommendation and top-k value maximization problems. (a) Considering an implicit feedback setting, with $x_{u,i} = 1$ if an item is relevant to a user and $x_{u,i} = 0$ if not, the top-k recommendation problem aims to determine the top-k most relevant items for each user. To target this problem, a traditional RS may first rank the items by predicted relevance and then select the top-k as described in Section 3.1. (b) Considering for example a generic business value $v_i \in [0, 10]$ associated with each item, the top-k value maximization problem aims to find the top-k most valuable yet relevant items for each user in descending value order. To target this problem, a profit-aware RS may exploit in-processing or postprocessing approaches described in Sections 3.1 and 3.2.

In-processing profit-aware approaches. In addition to postprocessing approaches, in-processing approaches (Akoglu & Faloutsos, 2010; Cai & Zhu, 2019; Concha-Carrasco, Vega-Rodríguez, & Pérez, 2023; Guo et al., 2021; He et al., 2022; Li, Fang, Bai, & Sheng, 2017; Li et al., 2021; Nemati & Khademolhosseini, 2020; Piton et al., 2011; Qu et al., 2014; Wang & Su, 2002; Wang, Zhou, & Han, 2002) are sometimes employed in the literature. In-processing approaches can be used to incorporate profit awareness during learning, thus avoiding additional reranking operations and possible postprocessing overhead at prediction time. However, this branch of the literature is highly scattered (De Biasio, Montagna et al., 2023), and most of the methods are tailored to certain algorithm families, application domains, or business contexts. For example, earlier studies (Piton et al., 2011; Wang & Su, 2002; Wang et al., 2002) proposed algorithms to incorporate profit into association rules mining techniques (Cai, Fu, Cheng, & Kwong, 1998; Hipp, Güntzer, & Nakhaeizadeh, 2000). However, unlike modern RSs based on collaborative filtering algorithms (Su & Khoshgoftaar, 2009), association rules (Hipp et al., 2000) are not personalized, i.e., the same items are recommended to different users, and these approaches may face challenges when the number of recommendable items is large (e.g., as in e-commerce) (Wang et al., 2002). In other studies, graph-based (Akoglu & Faloutsos, 2010; Li et al., 2017; Qu et al., 2014), reinforcement learning (Li et al., 2021), and evolutionary algorithms (Concha-Carrasco et al., 2023; Nemati & Khademolhosseini, 2020) have been used to optimize short-term profit. However, most of these methods are proposed in isolated contexts to target specific application domains, such as taxi drivers (Qu et al., 2014), social networks (Akoglu & Faloutsos, 2010; Li et al., 2017), insurance (Li et al., 2021) or telecommunications (Dookeram, Hosein, & Hosein, 2022), and these algorithms are not suitable for use in different contexts without major adaptations. For example, significant efforts are required to employ a method originally developed to optimize the profit of driving routes for taxi drivers in e-commerce settings. Other methods are rather niche and can be inefficient (Concha-Carrasco et al., 2023; Nemati & Khademolhosseini, 2020); thus, they are unlikely to be useful in industrial applications, e.g., because large-scale catalogs would require considerable computational power. Finally, although a profit-aware nearest neighbors algorithm suitable for use in other settings is found in one case (Cai & Zhu, 2019), the study focused on security aspects, investigating how to protect such an RS from malicious user attacks.

Hence, although in-processing methods may, in principle, bring benefit over postprocessing approaches in certain cases, at least considering the computational efficiency aspect, not much of the literature has studied the pros and cons of embedding profit optimization into the major families of collaborative filtering algorithms (De Biasio, Montagna et al., 2023). Focusing on this topic, in the following, we study profit-aware in-processing methods that can be applied to various application domains to optimize profitability at training time. In this way, unlike current postprocessing methods, no additional costs are required to perform additional reranking operations at prediction time. In particular, we propose profit-aware loss functions for three important algorithm families, i.e., matrix factorization, learning-to-rank, and neural algorithms. Moreover, we also consider a profit-aware nearest neighbors algorithm recently proposed in a study focused on security issues of RSs (Cai & Zhu, 2019) as an additional use case. In our evaluation, we compare these four in-processing methods with more popular postprocessing approaches to determine which specific approaches are more favorable than others.

3. Profit-aware recommendation algorithms

In this paper, we target an integrated problem formulation, which we call the *top-k value maximization problem*. As Fig. 1 shows, the top-k value maximization problem, which differs from the well-known top-k recommendation problem (Rendle, 2022; Zhao et al., 2023), aims to identify the most valuable, yet relevant items for each user (i.e., that the user may consider for future purchases) in descending value order.[1] The problem can be formally characterized as follows by focusing on the optimization of *short-term profit* as a particular category of business value.[2]

In Table 1 we introduce the main notation used in the paper. Let $\mathcal{U} = \{u_1, \dots, u_m\}$ be a set of m users and $\mathcal{I} = \{i_1, \dots, i_n\}$ a set of n items. For the sake of simplicity, let $\mathcal{I}_u^+$ be a set of items user u has interacted with (e.g., purchases) and $\mathcal{I}_u^-$ be a set of items the user has not interacted with. In the following, focusing on an implicit feedback setting (Rendle et al., 2009) that is widely employed in practical applications (e.g., e-commerce), we consider a recommendation algorithm that learns a *scoring function* $\mathbf{X} \rightarrow \hat{\mathbf{X}}$ to predict a matrix $\hat{\mathbf{X}} \in \{x \in \mathbb{R} : 0 \leq x \leq 1\}^{m \times n}$ from a binary interaction matrix $\mathbf{X} \in \{0, 1\}^{m \times n}$. In these matrices, $x_{u,i}$ is a single user-item interaction of a given user u for a certain item i, and $\hat{x}_{u,i}$ is a *predicted score* from the RS algorithm.

[1] Like other settings in the RS literature (Ghanem et al., 2022), the top-k value maximization problem considers that consumer attention is often limited and that purchase probability likely decreases according to the position of recommended items (*position bias*).

[2] The study focuses on the optimization of short-term profit, but the top-k value maximization problem can be extended to include possible long-term value optimization perspectives in future work (De Biasio, Montagna et al., 2023).

Table 1
Main notation.

Notation	Definition
u	User
i	Item
v_i	Item's profit
m	Number of overall users
n	Number of overall items
k	Number of items to recommend
$\mathcal{U} = \{u_1, \dots, u_m\}$	Set of users
$\mathcal{I} = \{i_1, \dots, i_n\}$	Set of items
$\mathcal{I}_u^+$	Set of items user u has interacted with
$\mathcal{I}_u^-$	Set of items user u has not interacted with
$\mathbf{X} \in \{0,1\}^{m\times n}$	User–item interaction matrix
$x_{u,i} \in \{0,1\}$	User–item feedback
$\mathbf{X} \rightarrow \hat{\mathbf{X}}$	Scoring function
$\hat{\mathbf{X}} \in \{x \in \mathbb{R} : 0 \le x \le 1\}^{m\times n}$	Matrix of prediction scores
$\hat{x}_{u,i} \in [0,1]$	User-item predicted interest
$\mathcal{Y}_{u,k}$	Recommendations list

In typical circumstances (Jannach et al., 2010; Rendle, 2022), where an RS is designed to recommend the most relevant items (*top-k recommendation problem*), $\hat{x}_{u,i}$ represents the *expected interest* (or utility) of a given item for a certain user. In those cases, an *ordered list* $\mathcal{Y}_{u,k}$ of k items for each user can be obtained:

$$\underset{\mathcal{Y}_{u,k}}{\operatorname{argmax}} \sum_{i\in\mathcal{Y}_{u,k}} \hat{x}_{u,i} \tag{1}$$

by choosing the items with the highest scores. However, the most relevant items may not necessarily be the most profitable ones. Therefore, an RS algorithm based on Eq. (1), although it may optimize the relevance of recommendations for the end users, does not guarantee the optimization of the profit for the firm. Using the above notation, in the next sections, we explore how profit-aware postprocessing and in-processing algorithms (Section 2.2.1) (De Biasio, Montagna et al., 2023; Jannach & Adomavicius, 2017) can be used to properly address the top-k value maximization problem by considering the item's profit v_i in their objective function.

3.1. Postprocessing profit-aware algorithms for addressing the top-k value maximization problem

By analyzing the literature on profit-aware recommendations (Section 2.2.1) (De Biasio, Montagna et al., 2023; Jannach & Adomavicius, 2017), we can observe that reranking approaches are viable methods for addressing the top-k value maximization problem. For example, these methods may weigh the consumer's expected interest (Chen et al., 2007; Chen et al., 2008):

$$\underset{\mathcal{Y}_{u,k}}{\operatorname{argmax}} \sum_{i\in\mathcal{Y}_{u,k}} \hat{x}_{u,i} \cdot v_i \tag{2}$$

with the items' profit v_i to rank the items of higher value for the company in the highest positions. Some algorithms, e.g., as in Jannach and Adomavicius (2017), may also exploit constrained variations of this approach:

$$\underset{\mathcal{Y}_{u,k}}{\operatorname{argmax}} \sum_{i\in\mathcal{Y}_{u,k}} \hat{x}_{u,i} \cdot v_i \quad \text{s.t. } \hat{x}_{u,i} \ge \beta \tag{3}$$

to consider only those items having predicted scores above a certain threshold $\beta \in [0,1]$, as these may be the ones most interesting for consumers. Similarly, other algorithms, e.g., that in Ghanem et al. (2022), may consider balancing business and consumer interests by exploiting an additional regularizer $\gamma \in [0,1]$:

$$\underset{\mathcal{Y}_{u,k}}{\operatorname{argmax}} \sum_{i\in\mathcal{Y}_{u,k}} \gamma \cdot \hat{x}_{u,i} + (1-\gamma)\cdot v_i \tag{4}$$

Considering the above equations, when β is very small, Eq. (3) falls back to Eq. (2), whereas when $\gamma = 1$, Eq. (4) falls back to the base case in Eq. (1)

3.2. In-processing profit-aware algorithms for addressing the top-k value maximization problem

To study how in-processing algorithms can be used to address our problem more appropriately, we focus on four collaborative filtering algorithms. Specifically, we first describe how to embed profit awareness into the neighbors selection procedure of a *User-Based Nearest Neighbors* algorithm (Nikolakopoulos et al., 2022) by referring to a recently proposed paper (Cai & Zhu, 2019) (Section 3.2.1). Then, we propose profit-aware loss functions for three well-known model-based algorithms: *Matrix Factorization* (Koren et al., 2009), *Neural Collaborative Filtering* (He et al., 2017), and *Bayesian Personalized Ranking* (Rendle et al., 2009) (Sections 3.2.2, 3.2.3).

3.2.1. Profit-aware nearest neighbors adaptations

User-Based Collaborative Filtering (*UCF*) (Nikolakopoulos et al., 2022; Resnick et al., 1994) is a well-known nearest neighbors algorithm that has been successfully applied in various application domains (Su & Khoshgoftaar, 2009). Considering an explicit feedback context (e.g., where $x_{u,i}$ is a rating in the range $[0,5]$), the algorithm calculates the predicted score $\hat{x}_{u,i}$ of an item i that user u has never interacted with based on a weighted sum of similarities:

$$\hat{x}_{u,i} = \frac{\sum_{v\in\mathcal{S}(u,i)} \mathrm{sim}(u,v)\cdot(x_{v,i} - \bar{x}_u)}{\sum_{v\in\mathcal{S}(u,i)} |\mathrm{sim}(u,v)|} \tag{5}$$

between users belonging to a given neighborhood $\mathcal{S}(u,i)$. Commonly, correlation criteria (e.g., Pearson correlation) are applied to the user-item interaction matrix to determine the similarity $\mathrm{sim}(u,v)$ between user pairs. In addition, since each user can rate items subjectively (thus having a different scale), the individual user's average rating $\bar{x}_u$ is subtracted from the given rating $x_{v,i}$. By selecting the neighbors most similar to each user, as follows:

$$\underset{\mathcal{Y},\mathcal{S}}{\operatorname{argmax}} \frac{\sum_{i\in\mathcal{Y}_{u,k}} \sum_{v\in\mathcal{S}(u,i)} \mathrm{sim}(u,v)\cdot(x_{v,i} - \bar{x}_u)}{\sum_{v\in\mathcal{S}(u,i)} |\mathrm{sim}(u,v)|} \tag{6}$$

it is thus possible to generate a list $\mathcal{Y}_{u,k}$ of k recommendations for each user. However, although the algorithm can be used to determine the most potentially interesting items for each user, it may not optimize the profit for the business.

Instead, as noted in a recent study focused on a different context (Cai & Zhu, 2019), by extending Eq. (6) and selecting a set of similar but more profitable neighbors $\mathcal{P}(u,i)$, as follows,

$$\underset{\mathcal{Y},\mathcal{P}}{\operatorname{argmax}} \frac{\sum_{i\in\mathcal{Y}_{u,k}} v_i \sum_{v\in\mathcal{P}(u,i)} \mathrm{sim}(u,v)\cdot(x_{v,i} - \bar{x}_u)}{\sum_{v\in\mathcal{P}(u,i)} |\mathrm{sim}(u,v)|} \tag{7}$$

it may be possible to increase profitability while still keeping the relevance of recommendations high. Intuitively, it is possible to determine the set of profitable neighbors for each user by selecting those with higher similarity-weighted cumulative profits, where the cumulative profit can be calculated considering the k most profitable items from each neighbor's purchase history.

In this paper, we propose adapting Eq. (7) for the fundamental class of implicit filtering-based CF algorithms employed in various practical applications:

$$\underset{\mathcal{Y},\mathcal{P}}{\operatorname{argmax}} \sum_{i\in\mathcal{Y}_{u,k}} v_i \sum_{v\in\mathcal{P}(u,i)} \mathrm{sim}(u,v) \tag{8}$$

In particular, we removed the mean centering and normalization operations from Eq. (7). Considering each user-item interaction $x_{u,i} \in \{0,1\}$, the k items included in the list $\mathcal{Y}_{u,k}$ may not be the optimal ones. In fact, in an implicit feedback setting, the user-item interaction $x_{v,i}$ of neighbor v and the average rating $\bar{x}_u$ of user u would equal one. Hence, the nominator in Eq. (7) would always be equal to zero.

3.2.2. Profit-aware matrix factorization and neural collaborative filtering extensions

Matrix Factorization (*MF*) (Koren, 2008; Koren & Bell, 2021; Koren et al., 2009) is a well-known latent factors model for recommendation. The algorithm aims to estimate the expected interest of user u in item i:

$$\hat{x}_{u,i} = \mathbf{q}_i^\intercal \mathbf{p}_u \tag{9}$$

through the dot product between l-dimensional embeddings. The user and item embeddings, $\mathbf{p}_u \in \mathbb{R}^l$ and $\mathbf{q}_i \in \mathbb{R}^l$, are traditionally learned through a dimensionality reduction algorithm applied to the user-item interaction matrix.

The model can handle both explicit and implicit feedback, albeit with some adaptation. If the feedback is implicit (i.e., $x_{u,i} \in \{0,1\}$), as in this paper, the learning algorithm typically optimizes a *binary cross-entropy loss* function:

$$\mathcal{L} = -\sum_{(u,i)\in\mathcal{K}} x_{u,i} \log \hat{x}_{u,i} + (1 - x_{u,i}) \log(1 - \hat{x}_{u,i}) \tag{10}$$

where $\mathcal{K} = \{(u,i) : i \in \mathcal{I}_u^+ \cup \mathcal{I}_u^-\}$ is the set of known interactions for each user. The algorithm is trained to recommend the most relevant items for each user by optimizing the loss function above.

Adopting the underlying principles of the profit-aware in-processing approaches presented in Section 2.2, in this paper, we propose extending the loss function of MF as follows to optimize profitability and relevance:

$$\mathcal{L} = -\sum_{(u,i)\in\mathcal{K}} v_i \cdot x_{u,i} \log \hat{x}_{u,i} + (1 - x_{u,i}) \log(1 - \hat{x}_{u,i}) \tag{11}$$

where v_i is the profit of the item. In this way, the algorithm can give more weight to higher-profit items in the learning process. In addition, other adaptations can be proposed by considering, for example, explicit feedback scenarios (e.g., $x_{u,i} \in [0,5]$). In these cases, the widely employed *squared loss* function can be also weighted as follows:

$$\mathcal{L} = -\sum_{(u,i)\in\mathcal{K}} v_i \cdot (x_{u,i} - \hat{x}_{u,i})^2 \tag{12}$$

to optimize overall profitability of the item.

The two proposed profit-aware loss functions in Eqs. (11) and (12) can also be used to consider business value in the *Neural Collaborative Filtering* (*NCF*) (He et al., 2017) model, a deep learning variant of matrix factorization (Koren et al., 2009) that uses the same loss functions but replaces the user-item dot product with a multilayer perceptron to learn any arbitrary pattern from the data. Similar extensions can be proposed considering other neural recommendation algorithms.

3.2.3. Profit-aware Bayesian personalized ranking adaptations

Bayesian Personalized Ranking (*BPR*) (Rendle, 2022; Rendle et al., 2009) is a state-of-the-art optimization framework applicable to various algorithms to generate recommendations in implicit feedback settings. Typically, BPR is applied on top of matrix factorization (Koren et al., 2009) by exploiting a *pairwise loss*[3] function that approximates the *area under the ROC curve* (*AUC*) ranking statistic:

$$\mathcal{L} = -\sum_{(u,i,j)\in D} \ln \sigma(\hat{x}_{u,i} - \hat{x}_{u,j}), \tag{13}$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is a sigmoid function and $D = \{(u,i,j) : i \in \mathcal{I}_u^+ \wedge j \in \mathcal{I}_u^-\}$ is a set of pairwise training examples consisting of pairs of positive and negative items for each user. The term $\sigma(\hat{x}_{u,i} - \hat{x}_{u,j})$ represents the probability that user u prefers item i over item j. Since matrix factorization is the primary underlying model of BPR, $\hat{x}_{u,i}$ and $\hat{x}_{u,j}$ generally represent the predicted scores for a positive item $i \in \mathcal{I}_u^+$ and a negative item $j \in \mathcal{I}_u^-$, respectively. By minimizing the loss function $\mathcal{L}$, the scores of positive items become higher than those of negative items.

[3] In Eq. (13), the regularization term has been omitted for clarity.

In this way, the algorithm can be trained to recommend the items most relevant to the user.

Inspired by the principles behind profit-aware in-processing approaches presented in Section 2.2, in this paper, we propose modifying BPR's objective function as follows to optimize profitability and relevance.

$$\mathcal{L} = -\sum_{(u,i,j)\in D} v_i \cdot \ln \sigma(\hat{x}_{u,i} - \hat{x}_{u,j}) \tag{14}$$

by weighting the probability $\sigma(\hat{x}_{u,i} - \hat{x}_{u,j})$ of user u preferring positive item i over negative item j by the profit v_i of the positive item. In this way, the algorithm can give more weight to higher-profit items relevant to the user, thus guiding the overall learning process. In addition, we also note that similarly to profit-aware reranking methods, there can be alternative variations of the loss function.

$$\mathcal{L} = -\alpha \sum_{(u,i,j)\in D} \ln \sigma(\hat{x}_{u,i} - \hat{x}_{u,j}) - (1-\alpha) \sum_{(u,i)\in D} v_i \tag{15}$$

thus weighting the profit of positive items v_i and the probability $\sigma(\hat{x}_{u,i} - \hat{x}_{u,j})$ according to a regularization parameter $\alpha \in [0,1]$ to balance consumer utility and business value during the optimization process at learning time.

4. Experiments

In this section, we evaluate the effectiveness of our proposed profit-aware in-processing algorithms by comparing them with the most commonly used postprocessing approaches in terms of the business value of their recommendations and their computational efficiency. We first describe the experimental setup (Section 4.1) and subsequently discuss the results (Section 4.2).

The experiments are executed on an on-premise server running Ubuntu 20.04 OS equipped with 12 vCPUs, 32 GB RAM, and 2 NVIDIA GeForce RTX 2080 Ti GPUs based on the CUDA® 11.6 architecture. The code[4] is developed in Python 3.9.15, is based on TensorFlow 2.11.0, and extends LibRecommender 0.10.2.[5]

4.1. Experimental setup

Below, we discuss the datasets, evaluation metrics, the compared algorithms, and the hyperparameter tuning process.

4.1.1. Data preparation

We chose three real-world datasets to run the experiments on to ensure a comprehensive evaluation. Each dataset comes from a different application domain and correspondingly has certain distinctive characteristics:

- **Amazon**[6] (Ni et al., 2019): This dataset contains product reviews and corresponding metadata (e.g., price, brand) from Amazon.com. Each review is associated with a rating on a $[1,5]$ scale. The data are organized into different categories (e.g., Books, Fashion, Electronics). Like in another study (Zheng, Gao, He, Jin, & Li, 2021), since the number of categories is large, we limited our analysis by selecting only the *Tools and Home Improvement* category. In addition, in accordance with many real-world business cases, e.g., Heien (1980) and Sammut-Bonnici and Channon (2015), we assume that *markup pricing* is used, associating the item price with a proportional profit, i.e., a profit equal to 20% of the price.

[4] https://github.com/estilos-lab/mba2par.

[5] https://github.com/massquantity/LibRecommender.

[6] https://nijianmo.github.io/amazon/index.html.

Table 2
The number of users, items, interactions, and the corresponding density of datasets used for experiments (Microsoft Corporation, 1998; Ni et al., 2019; Zhang et al., 2015) after the data preparation phase.

Dataset	# Users	# Items	# Interactions	Density
Amazon	$7.039 \cdot 10^3$	$56.365 \cdot 10^3$	$182.379 \cdot 10^3$	0.046%
Foodmart	$4.115 \cdot 10^3$	$1.559 \cdot 10^3$	$212.547 \cdot 10^3$	3.313%
Yelp	$1.959 \cdot 10^3$	$9.392 \cdot 10^3$	$58.065 \cdot 10^3$	0.316%

- **Foodmart**[7] (Microsoft Corporation, 1998): This dataset contains a sample of sales transactions from various consumers of a supermarket chain, a dataset that is usually exploited in Microsoft SQL Server as a test sample (Chen et al., 2007; Melomed, Gorbach, Berger, & Bateman, 2006). Each product belongs to a different category (e.g., Food, Drink, Non-Consumable). Since the dataset is not very large, we consider all the categories for the experiments. Moreover, given that each transaction includes the product price and its corresponding cost for the firm, similarly to the methods of other studies (Chen et al., 2007; Chen et al., 2008), we use this information to calculate the profit of each item, i.e., subtracting the item's cost from the price.
- **Yelp**[8] (Zhang et al., 2015): This dataset contains the user reviews of various real-world businesses organized into different categories (e.g., Shopping, Automotive, Medical). Like in the case of Amazon, each review is associated with a rating on a $[1, 5]$ scale. As in two other studies (Zheng et al., 2021; Zheng, Gao, He, Li, & Jin, 2020), we consider the *Restaurants* category, where the price bucket of each item is indicated using a different number of dollar symbols (from $ to $$$$). For this dataset, we associate this price bucket indicator with a proportional economic value, considering a hypothetical case in which the item's profit is difficult to estimate a priori with certainty, e.g., due to highly variable costs.

Before performing the experiments, some preliminary data preparation was conducted. In accordance with the objectives of the top-k value maximization problem described in Section 3, we prepare the various datasets for an implicit feedback recommendation task (Rendle, 2022) as follows: every purchase transaction in *Foodmart* is considered a positive user-item interaction; every review with a rating greater than or equal to four in *Amazon* and *Yelp* is considered a positive user-item interaction. All users who did not positively interact with at least 20 items are excluded, as done in the well-known MovieLens 20M dataset (Harper & Konstan, 2016), because we are not examining cold-start situations. However, no cold-start item is excluded because unpopular items are often those associated with the highest business value (Ghanem et al., 2022). Instead, we exclude all items with null, zero, or negative economic value. Although in real circumstances, negligible or negative profits may occur, this may occur only as a result of specific business strategies (Amatriain & Basilico, 2016; Gomez-Uribe & Hunt, 2016); e.g., unprofitable popular items may be used as loss leaders to stimulate the purchase of complementary higher-margin niche items (Garfinkel, Gopal, Pathak, Venkatesan, & Yin, 2007). In the following, we do not assume any of those cases apply for the datasets considered.

The statistics of the datasets after the data preparation phase are shown in Table 2. As shown in the table, *Amazon* is the least dense dataset with the largest number of items, *Foodmart* is the most dense dataset with the highest number of interactions, and *Yelp* has an intermediate density and the lowest number of interactions. Popularity and profit histograms of the datasets with fitted gamma distributions are shown in Fig. 2. Considering popularity, both *Amazon* and *Yelp* show a long-tail distribution. Instead, *Foodmart* exhibits a normal popularity distribution. Considering profit, *Amazon* has a long-tail distribution where the profit of most items is very low and very few items are highly profitable. In contrast, *Foodmart* has a distribution similar to that of a normal distribution, with most of the profit generated by the central bins. Finally, *Yelp* shows a left-skewed distribution with the majority of items having medium-low profits. By analyzing Pearson's correlation between popularity and profit, we find that for Amazon ($corr = -0.03486$) and Foodmart ($corr = 0.00720$) there is no correlation; however, for Yelp ($corr = 0.20893$) there is a weak positive correlation. This fact could have an impact on the experimental results since, generally, RSs tend to recommend the most popular items more frequently (*popularity bias*) (Abdollahpouri, 2019).

4.1.2. Evaluation metrics

To evaluate the performance of profit-aware algorithms according to the goals of the top-k value maximization problem defined in Section 3, we select two metrics that can be used to measure different aspects of recommendations. Using the *Normalized Discounted Cumulative Value* ($NDCV@k$) metric, we aim to assess the ability of the algorithms to place the most profitable items actually purchased by each user in the highest positions of the ranking. In addition, using the more widely-known *Normalized Discounted Cumulative Gain* ($NDCG@k$), we aim to measure how any increase in profitability might adversely affect the relevance of recommendations for consumers. Given that $NDCG@k$ and $NDCV@k$ measure partially competing aspects (i.e., consumer vs. business value), we expect that optimizing one metric will result in a reduction in the other. Below, we explain the underlying rationale of these metrics.

The *Normalized Discounted Cumulative Gain* is a widely used metric that can be used to measure the consumer relevance of recommendations. Let $rel^{\mathcal{Y}}_{u,j}$ be a relevance variable (Järvelin & Kekäläinen, 2017) that indicates whether the item recommended at position j in the ordered ranking $\mathcal{Y}_{u,k}$ of k items is relevant or not for user u. In an implicit feedback setting (Rendle, 2022), each item's relevance corresponds to its ground truth information $x_{u,i} \in \{0, 1\}$, e.g., assuming $x_{u,i} = 1$ if the user actually purchased the item, and $x_{u,i} = 0$ if not. Correspondingly, the *Normalized Discounted Cumulative Gain* at position k (Gunawardana et al., 2022; Zhao et al., 2023)

$$NDCG@k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{\sum_{j=1}^{k} \frac{rel^{\mathcal{Y}}_{u,j}}{log_2(j+1)}}{IDCG_u@k} \tag{16}$$

is defined as an inverse log reward over all the ranking positions with relevant items among the top-k recommended ones. In the equation, $IDCG_u@k$ is usually referred to as the *Ideal Discounted Cumulative Gain* obtained by sorting all the items relevant to the user in descending order. Hence, $NDCG@k$ measures how precise an RS algorithm is in recommending the most relevant items actually purchased by each user in the highest-ranking positions.

The *Normalized Discounted Cumulative Value* is a metric that combines consumer relevance and organizational value. The idea underlying this metric is taken from a recent paper that measures performance using the *Price-Based* $NDCG@k$ (Louca, Bhattacharya, Hu, & Hong, 2019), i.e., a variant of the previously defined $NDCG@k$ where the gain is given by the item's price. In our context, instead of explicitly considering the price, we consider a *generic business value* (e.g., short-term profit) (Jannach & Adomavicius, 2017; Jannach & Jugovac, 2019) that the company may aim to optimize in accordance with the purposes of value-aware RSs (De Biasio, Montagna et al., 2023). Hence, considering v_j as the value an organization obtains if an item recommended at position j is purchased by a user,[9] we define the *Normalized Discounted*

[7] https://github.com/julianhyde/foodmart-data-json.
[8] https://yelp.com/dataset.
[9] For the sake of notation, v_j is used to indicate the value of the recommended item at position j, but typically the value depends only on the item and not on its ranking position.

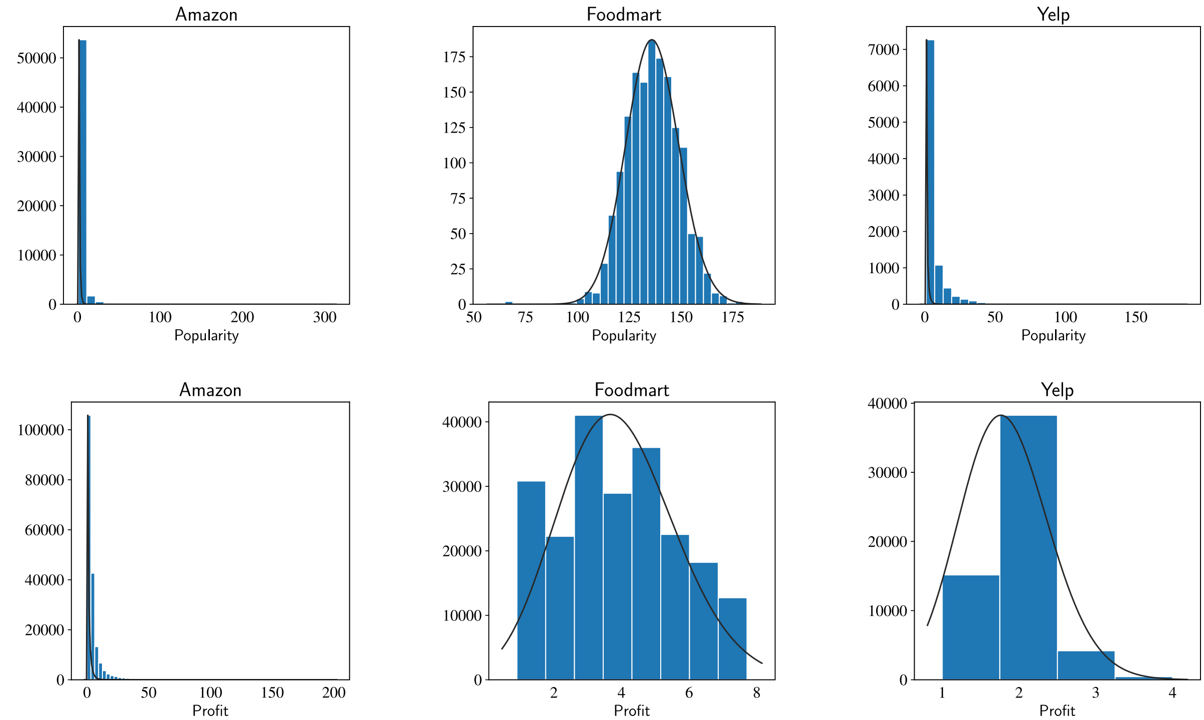

Fig. 2. Popularity and profit histograms (in blue) with best-fit gamma distributions (in black) of datasets used for experiments (Microsoft Corporation, 1998; Ni et al., 2019; Zhang et al., 2015) after the data preparation phase.

Cumulative Value at position k

$$NDCV@k = \frac{1}{|\mathcal{U}|}\sum_{u\in\mathcal{U}} \frac{\sum_{j=1}^{k} \frac{rel_{u,j}^{\mathcal{Y}} \cdot v_j}{log_2(j+1)}}{IDCV_u@k} \tag{17}$$

as an inverse value-based log reward over all positions with valuable yet relevant items among the top-k recommended ones. In the equation, similar to $IDCG_u@k$, we refer to $IDCV_u@k$ as the *Ideal Discounted Cumulative Value* obtained by sorting all the items relevant to the user in descending value order. Therefore, $NDCV@k$ can be used to measure how precise an RS algorithm is in recommending the most valuable yet relevant items actually purchased by each user in the highest-ranking positions.

Note that although other metrics from the literature on profit-aware RSs can also be used to measure business value (e.g., overall profit or expected profit) (Cai & Zhu, 2019; Jannach & Adomavicius, 2017), these metrics are not *rank-aware*; i.e., they do not consider the items' ranking positions for evaluation purposes. Therefore, these metrics are limited for use in measuring the ability of an RS to recommend the most valuable items in descending profit order for each user, as mandated by the top-k value maximization problem. In addition, unlike in other papers on profit-aware RSs (Ghanem et al., 2022; Jannach & Adomavicius, 2017), we do not unrealistically assume that the user must always buy one item among the recommended items (*guaranteed purchase*). Instead, we rely on the actual consumer purchasing history for performance evaluation.

4.1.3. Compared algorithms

Various algorithms are compared in the experiments. In particular, we select representative profit-aware in-processing algorithms belonging to the main classes described in Section 3.2, namely:

- **Value Neighbor Selection** (**VNS**): a UCF variant that selects the most profitable neighbors to generate recommendations as defined in Eq. (8).
- **Value Matrix Factorization** (**VMF**): an MF variant we propose in this paper that exploits the profit-aware cross-entropy loss defined in Eq. (11).
- **Value Neural Collaborative Filtering** (**VNCF**): an NCF variant we propose in this paper that exploits the profit-aware cross-entropy loss defined in Eq. (11).
- **Value Bayesian Personalized Ranking** (**VBPR**): a BPR variant we propose in this paper that exploits the profit-aware pairwise loss defined in Eq. (14).

Moreover, to denote the profit-aware postprocessing algorithms presented in Section 3.1, we refer to:

- **Hybrid Perspective Recommender System** (**HPRS**) (Chen et al., 2008): a profit-aware reranking algorithm that recommends the top-k items with the highest profit-weighted predicted scores, as defined in Eq. (2).
- **Constrained Profit Ranking** (**CPR**) (Jannach & Adomavicius, 2017): a constrained variant of HPRS that generates recommendations considering only items with an expected interest above a certain threshold $\beta \in [0,1]$, as in Eq. (3).
- **Multi-Objective Profit Ranking** (**MOPR**) (Ghanem et al., 2022): a multiobjective variant of HPRS that balances consumer and organizational interests with a regularizer $\gamma \in [0,1]$, as in Eq. (4).

4.1.4. Hyperparameter tuning

The hyperparameter tuning procedure proceeds as follows. The users in each dataset are split into training, validation, and test sets (60%/20%/20%), ensuring that users in one set do not appear in any other set. Four items are kept as known positive interactions for each validation and test set user to avoid cold-start situations. The remaining positive interaction items are used as the only relevant items for evaluating performance. For each model, a grid search is performed by optimizing the $NDCV$ on the validation set to find the best hyperparameters. All the models are trained for a maximum of 1000 epochs using early stopping with a *patience* of 10 epochs. Experiments are performed considering a different number of recommended items $k \in \{10, 20\}$. The results are averaged across three random splits of users using different seeds. In the following experiments, we report the mean and the standard deviation over the different runs.

The following hyperparameter ranges are explored in the grid search. In particular, for UCF and VNS, the number of neighbors is selected from $\{1, 2, 3, 5, 8, 10, 25, 50\}$. Regarding MF and BPR and their profit-aware variants VMF and VBPR, the embedding sizes are selected from $\{32, 64, 128, 256\}$, and learning rates from $\{10^{-3}, 10^{-4}\}$ are explored while fixing the batch size at 128. As for NCF and VNCF, embedding sizes are selected from $\{16, 32, 64, 128, 256\}$ and the learning rates are selected from $\{10^{-3}, 10^{-4}\}$. Batch sizes from $\{64, 128, 256\}$ are explored while setting the multilayer perceptron hidden units as suggested in the original paper (He et al., 2017) to $\{2 \cdot h, h, \frac{h}{2}\}$, where h is the embedding size. In addition, concerning postprocessing approaches, CPR's threshold β in Eq. (3) is varied in the range $\{0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$ and MOPR's regularizer γ in Eq. (4) is varied in the range $\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$.

4.1.5. Assumptions and limitations of experiments

In this section, we summarize the assumptions and limitations of the experiments that were described primarily in Sections 4.1.1, 4.1.2, and 4.1.4 The assumptions underlying the experiments (AEs) are reported below:

- **AE1**: The profit of each item is calculated differently for each dataset. For Amazon, we assumed, as in real-world business cases (Heien, 1980; Sammut-Bonnici & Channon, 2015), that *markup pricing* is proportional to profit by setting the latter to 20% of the item's price. For Foodmart, similarly to other studies (Chen et al., 2007; Chen et al., 2008), we calculated each item's profit by subtracting its price from its cost. Finally, for Yelp, we assumed that profit was an integer in the range $\{1, 2, 3, 4\}$ based on the number of dollar symbols in each item's price bucket (from $ to $$$$).
- **AE2**: We assumed that any review with a rating greater than or equal to four was a positive user-item interaction to place Amazon and Yelp in an implicit recommendation setting, as also done in other studies in the literature (Rendle, 2022)
- **AE3**: We assumed that the history of consumer interactions collected in the datasets is representative of algorithm performance evaluation, as in almost all previous RS literature. Specifically, we did not consider cases where there are items of potential interest to the user with which he or she has never interacted with; these items are therefore not present in the ground truth.

The possible limitations of the experiments (LEs) are as follows:

- **LE1**: Only users who had positively interacted with at least 20 items were considered in the experiments to avoid considering cold-start recommendation settings, which was also done in the well-known MovieLens 20M dataset (Harper & Konstan, 2016). For the same reasons, four items were kept as known positive interactions for each validation and test set user during model training.
- **LE2**: All items with null, zero, or negative business value were excluded from the experiments; that is, we did not consider particular business strategies that exploit unprofitable popular items as loss leaders to stimulate the purchase of complementary products (Garfinkel et al., 2007).
- **LE3**: We focused our analyses on Amazon's *Tools and Home Improvement* and Yelp's *Restaurants* categories to avoid considering a large number of product categories in the experiments, as also done in previous studies (Zheng et al., 2021, 2020).

4.2. Results

In this section, we discuss the results of the experiments. We first analyze the performance of our proposed in-processing profit-aware algorithms by comparing them with relevance-based baselines. Then, we compare the results of the most widely known postprocessing approaches with those of the baselines and our proposed in-processing algorithms.

4.2.1. Results for profit-aware in-processing methods

In Table 3, we report the results obtained by applying four in-processing methods (i.e., VNS, VMF, VNCF, and VBPR), one from the literature (i.e., VNS), and three proposed in this paper (i.e., VMF, VNCF, and VBPR). Each method is related to an underlying baseline recommendation model. Three real-world datasets with different characteristics are considered in the experiments. The number of recommended items is varied according to two widely used settings in the literature (i.e., $k \in \{10, 20\}$). Given that the four underlying models (i.e., UCF, NCF, MF, and BPR) are widely used in industry, and that one model may be preferred over another by a firm for various reasons (e.g., explainability, cost of ownership), with this experiment, we are not aiming to identify one model superior to all others in terms of performance, e.g.; a firm might want to optimize the business value of a specific model that is already in production. Instead, we aim to demonstrate how any baseline model can be adapted to optimize profitability by exploiting the methodologies discussed in Section 3.

All three proposed model-based methods (i.e., VMF, VNCF, and VBPR) successfully improved $NDCV@k$ over the baselines for all the datasets and the number of recommended items considered. The VNS algorithm proposed earlier in Cai and Zhu (2019) also proved effective, except for the Amazon dataset. This indicates that compared to the baselines, profit-aware algorithms can generally balance consumer relevance and profit by recommending higher profit items that are still relevant to the users. As expected (see Section 4.1.2), an increase in $NDCV@k$ almost always results in a corresponding decrease in $NDCG@k$ because profit-aware algorithms give more weight to the company's interests in the learning process. Similar to the findings of other studies in the literature (Chen et al., 2007; Chen et al., 2008; Ghanem et al., 2022; Jannach & Adomavicius, 2017), profit-aware algorithms are able to include items of higher profit in the top-k recommendations list, but this occurs at the expense of a more or less significant loss of relevance.

In addition, from a computational point of view, the overall prediction times of the model-based algorithms are comparable to those of the various baselines. Only the VNS model has a slightly prediction time because it performs more computational operations than does the UCF baseline. This indicates that in-processing methods are efficient overall, and no particular computational overhead (e.g., that incurred due to reranking operations) is required at prediction time to optimize profit. Moreover, our proposed in-processing algorithms do not lead to additional hyperparameters to tune compared to those of the baseline models.

4.2.2. Results for the profit-aware postprocessing methods

In Table 3, we also report the results obtained by applying three postprocessing methods (i.e., HPRS, CPR, and MOPR) on top of the four baseline recommendation models.

As shown by $NDCV@k$, the results are mixed. For instance, considering the Foodmart dataset, an improvement over the baselines is exhibited, but not always. These results are different from from those of our modeling approaches, which were always proven to be effective. For example, the Amazon and Yelp datasets seem particularly challenging for postprocessing algorithms, e.g., considering results for MF and BPR.

Moreover, considering each postprocessing algorithm individually, we also found different behaviors. For example, for HPRS, the algorithm almost never succeeds in improving $NDCV@k$ over the baseline (except for Foodmart and a few other cases). In contrast, the CPR and MOPR algorithms often succeed in improving performance. However, unlike HPRS and our proposed modeling methods, these postprocessing methods require additional hyperparameters to tune the weights of users' and organizational interests in the reranking process.

Finally, observing the overall prediction time, we note that postprocessing methods take longer to generate recommendations than the baseline models. For example, considering the Amazon dataset and the

Table 3
Results (i.e., NDCV, NDCG, and overall prediction time in seconds) of different profit-aware in-processing (i.e., VNS, VNCF, VMF, VBPR) and postprocessing algorithms (i.e., HPRS, CPR, MOPR) compared to their actual baseline recommendation models (i.e., UCF, NCF, MF, BPR) for different datasets (i.e., Amazon, Foodmart, Yelp), obtained by varying the number of recommended items (i.e., $k \in \{10, 20\}$). *The prediction time is not reported for different cutoff lengths k because the time needed to compute a recommendation list, which is the focus here, is independent of how many items are used to compute a certain metric.

			$NDCV@10$		$NDCG@10$		$NDCV@20$		$NDCG@20$		$Pred.\ Time$ (s)*	
Dataset	Model	Algorithm	mean	std	mean	std	mean	std	mean	std	mean	std
Amazon	UCF	Base	0.0326	0.0062	0.1802	0.0087	0.0325	0.0041	0.1920	0.0077	3.53	0.09
		VNS	0.0139	0.0021	0.0273	0.0032	0.0180	0.0022	0.0357	0.0042	5.45	0.01
		HPRS	0.0207	0.0058	0.0400	0.0071	0.0246	0.0023	0.0507	0.0079	7.18	1.63
		CPR	0.0133	0.0013	0.1059	0.0070	0.0130	0.0020	0.1060	0.0070	7.25	0.38
		MOPR	**0.0345**	0.0078	0.1693	0.0074	**0.0360**	0.0045	0.1762	0.0060	8.12	0.03
	NCF	Base	0.0060	0.0013	0.0336	0.0143	0.0090	0.0050	0.0466	0.0287	6.47	4.04
		VNCF	**0.0150**	0.0045	0.0249	0.0055	**0.0237**	0.0069	0.0339	0.0107	9.53	7.09
		HPRS	0.0021	0.0009	0.0013	0.0007	0.0042	0.0016	0.0020	0.0009	40.15	2.52
		CPR	0.0029	0.0011	0.0058	0.0039	0.0052	0.0052	0.0048	0.0056	30.60	3.31
		MOPR	0.0057	0.0021	0.0263	0.0165	0.0094	0.0045	0.0431	0.0306	36.24	2.57
	MF	Base	0.0141	0.0020	0.0869	0.0104	0.0217	0.0021	0.1050	0.0115	4.79	0.05
		VMF	**0.0205**	0.0005	0.0880	0.0107	**0.0309**	0.0004	0.1061	0.0088	3.71	0.91
		HPRS	0.0029	0.0022	0.0018	0.0011	0.0036	0.0035	0.0017	0.0012	31.14	0.53
		CPR	0.0164	0.0029	0.0322	0.0124	0.0269	0.0051	0.0287	0.0054	20.54	0.67
		MOPR	0.0101	0.0009	0.0723	0.0147	0.0171	0.0014	0.0849	0.0087	31.38	0.95
	BPR	Base	0.0260	0.0037	0.1291	0.0034	0.0375	0.0032	0.1464	0.0031	4.94	0.28
		VBPR	**0.0327**	0.0007	0.1223	0.0023	**0.0428**	0.0068	0.1289	0.0174	5.02	0.30
		HPRS	0.0035	0.0011	0.0021	0.0007	0.0051	0.0020	0.0024	0.0007	31.10	0.51
		CPR	0.0210	0.0032	0.0318	0.0031	0.0329	0.0053	0.0485	0.0033	20.62	0.19
		MOPR	0.0181	0.0025	0.0740	0.0051	0.0277	0.0025	0.0932	0.0064	30.81	0.66
Foodmart	UCF	Base	0.0202	0.0013	0.1223	0.0088	0.0267	0.0020	0.1670	0.0080	3.20	0.11
		VNS	0.0213	0.0016	0.0781	0.0072	0.0276	0.0007	0.1102	0.0025	7.63	4.00
		HPRS	0.0299	0.0028	0.1160	0.0098	0.0392	0.0023	0.1560	0.0057	7.32	0.02
		CPR	0.0272	0.0009	0.1175	0.0089	0.0319	0.0022	0.1601	0.0099	7.30	0.06
		MOPR	**0.0303**	0.0031	0.1101	0.0109	**0.0389**	0.0028	0.1482	0.0079	7.36	0.04
	NCF	Base	0.0099	0.0008	0.0730	0.0030	0.0145	0.0014	0.1099	0.0049	1.89	0.33
		VNCF	0.0191	0.0039	0.0750	0.0136	0.0249	0.0030	0.1033	0.0096	1.81	0.38
		HPRS	0.0212	0.0013	0.0760	0.0044	**0.0280**	0.0026	0.1049	0.0097	4.21	0.12
		CPR	0.0075	0.0116	0.0278	0.0392	0.0278	0.0025	0.1059	0.0076	4.22	0.22
		MOPR	**0.0213**	0.0007	0.0757	0.0031	0.0269	0.0024	0.1029	0.0095	4.17	0.22
	MF	Base	0.0112	0.0021	0.0768	0.0134	0.0153	0.0015	0.1105	0.0130	0.14	0.01
		VMF	0.0195	0.0002	0.0800	0.0026	0.0261	0.0002	0.1166	0.0038	0.15	0.01
		HPRS	**0.0205**	0.0003	0.0735	0.0022	**0.0284**	0.0010	0.1069	0.0055	2.92	0.05
		CPR	0.0196	0.0005	0.0697	0.0022	0.0275	0.0002	0.1029	0.0023	2.88	0.04
		MOPR	0.0196	0.0019	0.0722	0.0083	0.0284	0.0012	0.1076	0.0048	2.92	0.09
	BPR	Base	0.0164	0.0006	0.1044	0.0015	0.0223	0.0009	0.1461	0.0043	0.14	0.00
		VBPR	0.0215	0.0016	0.0828	0.0036	0.0273	0.0018	0.1125	0.0058	0.11	0.04
		HPRS	**0.0260**	0.0004	0.0921	0.0010	0.0352	0.0016	0.1302	0.0041	2.96	0.07
		CPR	0.0249	0.0021	0.0912	0.0029	**0.0354**	0.0011	0.1311	0.0036	2.87	0.10
		MOPR	0.0255	0.0010	0.0894	0.0060	0.0351	0.0016	0.1295	0.0049	2.87	0.05
Yelp	UCF	Base	0.1725	0.0045	0.4552	0.0065	0.2025	0.0042	0.4773	0.0045	0.89	0.10
		VNS	0.1755	0.0051	0.4153	0.0130	0.2064	0.0063	0.4444	0.0090	1.41	0.07
		HPRS	**0.1802**	0.0056	0.4180	0.0044	**0.2133**	0.0044	0.4450	0.0041	1.83	0.07
		CPR	0.1532	0.0056	0.3989	0.0068	0.1616	0.0066	0.4171	0.0068	1.79	0.04
		MOPR	0.1785	0.0026	0.4478	0.0007	0.2101	0.0045	0.4637	0.0082	1.83	0.03
	NCF	Base	0.0578	0.0278	0.1947	0.0951	0.0699	0.0339	0.2212	0.1018	0.91	0.15
		VNCF	**0.0657**	0.0267	0.2043	0.0679	0.0789	0.0310	0.2379	0.0685	0.79	0.17
		HPRS	0.0367	0.0111	0.0963	0.0229	0.0386	0.0157	0.1079	0.0343	1.81	0.11
		CPR	0.0248	0.0089	0.1033	0.0121	0.0091	0.0079	0.0575	0.0518	1.44	0.37
		MOPR	0.0650	0.0343	0.1918	0.0933	**0.0821**	0.0318	0.2186	0.0938	1.76	0.03
	MF	Base	0.1239	0.0043	0.3540	0.0030	0.1451	0.0051	0.3857	0.0047	0.20	0.01
		VMF	**0.1378**	0.0047	0.3792	0.0084	**0.1622**	0.0044	0.4115	0.0051	0.17	0.00
		HPRS	0.0555	0.0040	0.1412	0.0057	0.0691	0.0023	0.1740	0.0045	1.35	0.05
		CPR	0.0849	0.0048	0.2437	0.0158	0.0905	0.0050	0.2655	0.0046	1.01	0.05
		MOPR	0.1027	0.0058	0.2744	0.0096	0.1217	0.0014	0.3143	0.0053	1.35	0.03
	BPR	Base	0.1578	0.0052	0.4234	0.0125	0.1810	0.0031	0.4500	0.0125	0.17	0.02
		VBPR	**0.1600**	0.0048	0.4126	0.0206	**0.1902**	0.0015	0.4428	0.0144	0.18	0.02
		HPRS	0.0740	0.0094	0.1870	0.0140	0.0921	0.0034	0.2202	0.0091	1.34	0.04
		CPR	0.0764	0.0091	0.2270	0.0134	0.0677	0.0085	0.2401	0.0083	0.83	0.02
		MOPR	0.1435	0.0069	0.3969	0.0123	0.1716	0.0039	0.4153	0.0260	1.35	0.10

underlying BPR model, the baseline and VBPR in-processing algorithms take about five seconds to generate predictions, while the HPRS and MOPR postprocessing methods take approximately 30 s[10] This behavior is expected because, unlike our proposed model-based algorithms, postprocessing methods need to perform a subsequent reranking step that may incur significant computational overhead after generating recommendations. This limitation is fundamental to consider because, in practical cases, postprocessing methods could be inapplicable for large-scale production systems with millions of active users and large item catalogs.

5. Discussion

As mentioned in the previous sections, both in-processing and postprocessing methods can theoretically be used for generating profit-aware recommendations. However, these methods have rarely been compared in the profit-aware literature and may have some peculiarities that make them more suitable for use in some contexts than in others. For example, literature (De Biasio, Montagna et al., 2023; Jannach & Adomavicius, 2017) has shown that postprocessing methods are flexible and can be implemented on top of various recommender systems. Moreover, although in-processing methods are typically tailored to specific RS families, they are potentially more efficient since they avoid reranking overhead. Below, we discuss the performance and computational aspects of both profit-aware postprocessing and in-processing approaches.

5.1. Performance aspects of profit-aware algorithms

The experiments presented in this paper demonstrate that our three proposed model-based algorithms (i.e., VMF, VNCF, and VBPR) successfully improved $NDCV@k$ in all the considered cases. Adapting the VNS algorithm (Cai & Zhu, 2019) for the implicit feedback setting also proved effective, except for the Amazon dataset. This may depend on the particular characteristics of this dataset (see Section 4.1.1). In particular, the Amazon dataset is very sparse and exhibits long-tailed distributions of both popularity and profits. Moreover, given that there is also no correlation between popularity and profit, by selecting the most profitable neighbors to generate recommendations instead of those most similar to the current user (see Section 3.2.1), much relevance is lost, thus negatively impacting $NDCV@k$.

The postprocessing methods, were also effective, but not under all circumstances. In particular, the Amazon and Yelp datasets proved especially challenging. This behavior may occur because postprocessing methods exploit heuristic criteria to rerank recommendations from an underlying model (see Section 3.1). In the case of Amazon and Yelp datasets, for example, the sparsity is high (see Table 2), and since postprocessing algorithms perform reranking operations on the entire spectrum of items, this may negatively affect the subsequent quality of recommendations, including in the final ranking items that are highly profitable but not relevant to users.

5.2. Computational aspects of profit-aware algorithms

From a computational point of view, unlike postprocessing methods, in-processing algorithms do not incur any computational overhead at prediction time. In practice, postprocessing methods could have major limitations in many commercial applications because the high prediction times of these methods could be prohibitive in large-scale production systems with millions of active users and large item catalogs. Instead, in-processing algorithms may be preferable for reducing computational resources or when it is necessary to instantly provide recommendations to users.

Moreover, considering the actual implementation of the methods, the CPR and MOPR postprocessing algorithms use an additional hyperparameter to balance consumer utility and provider profits, thus requiring more time to train than the various in-processing algorithms. In fact, in many cases, especially for the Amazon dataset, the HPRS postprocessing algorithm, which does not require additional hyperparameters, failed to achieve a higher $NDCV@k$ performance than the baseline.

6. Limitations and future research

In this paper, we addressed what we called the top-k value maximization problem by comparing in-processing and postprocessing approaches that we used to build profit-aware recommender systems (Section 3). A variety of extensions of our work are possible in future work. Below, we discuss several possible future algorithm adaptations and comparative analyses.

6.1. Possible algorithm extensions for future studies

We have identified several research directions for possible future algorithm extensions. First, in this paper, we limited ourselves to incorporating profitability aspects through in-processing methods into major RS algorithmic classes, such as nearest neighbors, matrix factorization, learning-to-rank, and neural algorithms. In future works, we might consider embedding profit awareness in other algorithmic classes (e.g., based on linear models, graph neural networks, or association rule mining techniques) (Ning & Karypis, 2011; Wang et al., 2002; Wu, Sun, Zhang, Xie, & Cui, 2022) or in other algorithms belonging to the same class (e.g., neural algorithms) (Cheng et al., 2016; Guo, Tang, Ye, Li, & He, 2017).

In addition, we limited the focus of our research to model-based approaches (i.e., algorithms based on MF, BPR and NCF). In the future, it might be interesting to study in more detail how to improve the performance of the VNS algorithm (Cai & Zhu, 2019). For example, by modifying the neighbors selection criterion with an additional hyperparameter, we may adjust the number of profitable neighbors with that of similar neighbors. In this way, the algorithm may be able to increase the profitability of recommendations without losing too much relevance, thus enabling it to perform well even on the Amazon dataset.

Moreover, we limited the scope of our work to the implicit feedback setting. Therefore, although it is possible to extend the proposed algorithms in various ways, in the experiments, we compared only profit-aware in-processing algorithms exploiting the loss functions in Eqs. (8), (11) and (14). Hence, we have left the comparison of algorithms designed to handle explicit feedback, possibly by exploiting the losses in Eqs. (7) and (12), for future work. Moreover, comparisons of other variants of in-processing profit-aware algorithms that may use additional hyperparameters, such as those in Eq. (15), could be incorporated in future experiments.

Furthermore, considering the current experiments, we limited the comparison of the in-processing methods we designed with three postprocessing algorithms we identified in the literature. In the future, it could be interesting to supplement experiments by comparing preprocessing methods that, although not found in the literature, may be possible alternatives for generating profit-aware recommendations, e.g., completely ruling out unprofitable items with a static threshold before the training phase. Moreover, considering computational aspects, the analysis highlighted several limitations of current postprocessing approaches. Hence, a future research direction may be to design more effective or efficient postprocessing algorithms that could reduce computational overhead by applying reranking only to the most potentially

[10] Note that the prediction times of the various postprocessing algorithms (i.e., HPRS, CPR, and MOPR) are comparable given the same dataset and underlying model. Only CPR yields slightly shorter prediction times because it performs reranking not on the entire item spectrum but on a subset of items with predicted scores above a certain threshold (see Eq. (3)).

relevant items, thus avoiding considering the complete spectrum of items. Additionally, studying the possibility of combining pre, in-, and postprocessing approaches to achieve better results could be an interesting future research direction.

6.2. Possible future comparative studies

We identified the following research directions for possible future comparative analyses. First, in this study, we focused on optimizing short-term profit as a particular business value category using collaborative filtering algorithms that are widely used in practice. Given this objective, we leave for future work the study of algorithms that may consider temporal dynamics to optimize long-term business value (e.g., based on reinforcement learning) (Afsar, Crump, & Far, 2022; Guo et al., 2021; He et al., 2022; Iwata, Saito, & Yamada, 2008; Ji, Qin, Han, & Yang, 2021; Pei et al., 2019; Theocharous, Thomas, & Ghavamzadeh, 2015; Zhang, Zhao et al., 2022; Zou et al., 2019). We also leave for future work the possible study of niche methods (Akoglu & Faloutsos, 2010; Concha-Carrasco et al., 2023; Nemati & Khademolhosseini, 2020) or applications (e.g., considering the taxi driver domain) (Li et al., 2021; Qu et al., 2014). An interesting future research direction could also be to compare promotional approaches (Jannach & Adomavicius, 2017) that may increase profitability by incentivizing impulsive purchasing behaviors (e.g., dynamic pricing or bundling methods) (Adelnia Najafabadi, Shekarchizadeh, Nabiollahi, Khani, & Rastgari, 2022; Ettl, Harsha, Papush, & Perakis, 2020; Garfinkel, Gopal, Pathak, & Yin, 2008; Ghoshal, Mookerjee, & Sarkar, 2021; Sun et al., 2022; Zhao, Zhang, Friedman, & Tan, 2015). Furthermore, algorithms leveraging consumer-oriented strategies (Jannach & Adomavicius, 2017) that may, in turn, bring greater profit to the company (e.g., price sensitivity or economic utility modeling approaches) (Ge et al., 2019; Greenstein-Messica & Rokach, 2018; Huang, Ding, Hu, Jiang, & Li, 2021; Umberto, 2015; Wang & Zhang, 2011; Zhang, Xu et al., 2022; Zheng et al., 2021), can be studied in the future.

In addition, in our work, the $NDCV@k$ metric was used mainly to evaluate the ability of a recommendation algorithm to place the most profitable yet relevant items to the users in descending profit order in the ranking. However, the metric gives equal weight to relevance and profitability during evaluation. In the present study, considering an implicit feedback setting, we normalized both profitability and relevance before calculating the metric. Nevertheless, in possible future studies that may consider an explicit feedback setting (e.g., where the relevance range can be $[0, 5]$), the width of the profitability and relevance ranges before normalization may impact the final results. Correspondingly, a future research direction might be to consider these factors to investigate how to evaluate profit-aware recommendation algorithms offline, which may also involve the design of additional metrics.

Finally, three datasets with different characteristics were selected to evaluate the algorithms. However, the distributions of popularity and profitability and the correlation between these two factors may impact the final results. Since more popular items are generally more relevant to users, they are more likely to bring an increase in profitability if they are also positively correlated with profits. Thus, a future research direction might be to study the relationship between popularity and profitability (and related cold-start aspects) in more depth to understand the contributions of both factors into the final performance of algorithms.

7. Conclusion

In various practical contexts, such as electronic commerce, companies strive to leverage recommender systems to enhance business value through increasing profits, conversion rates, or customer retention. In this work, we explore the use of novel modeling approaches across key collaborative filtering families. These model-based strategies, proposed as alternatives to prevalent reranking methodologies, have been proven to be consistently effective at generating more profitable yet relevant recommendations in a computationally efficient way. Our findings suggest that model-based approaches present potential for overcoming the limitations of today's prevalent reranking techniques. Based on the effectiveness of our model-based methods, we anticipate that this study will promote additional research for improving recommender systems and their impact on business outcomes.

CRediT authorship contribution statement

Alvise De Biasio: Conceptualization, Methodology, Software, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Dietmar Jannach:** Conceptualization, Methodology, Validation, Resources, Writing – original draft, Writing – review & editing, Supervision. **Nicolò Navarin:** Conceptualization, Methodology, Validation, Resources, Writing – original draft, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

A link to a public repository containing code and data is provided in the article.

Acknowledgment

This work was partially funded by estilos srl.

References

Abdollahpouri, H. (2019). Popularity bias in ranking and recommendation. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 529–530).

Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., et al. (2020). Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction*, *30*, 127–158. http://dx.doi.org/10.1007/s11257-019-09256-1.

Abdollahpouri, H., & Burke, R. (2022). Multistakeholder recommender systems. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender Systems Handbook* (pp. 647–677). http://dx.doi.org/10.1007/978-1-0716-2197-4_17.

Adelnia Najafabadi, H., Shekarchizadeh, A., Nabiollahi, A., Khani, N., & Rastgari, H. (2022). Dynamic pricing for information goods using revenue management and recommender systems. *Journal of Revenue and Pricing Management*, *21*, 153–163. http://dx.doi.org/10.1057/s41272-020-00276-w.

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, *17*, 734–749. http://dx.doi.org/10.1109/TKDE.2005.99.

Adomavicius, G., & YoungOk, K. (2012). Improving aggregate recommendation diversity using ranking-based techniques. *IEEE Transactions on Knowledge and Data Engineering*, *24*, 896–911. http://dx.doi.org/10.1109/TKDE.2011.15.

Afsar, M. M., Crump, T., & Far, B. (2022). Reinforcement learning based recommender systems: A survey. *ACM Computing Surveys*, *55*, 1–38. http://dx.doi.org/10.1145/3543846.

Akoglu, L., & Faloutsos, C. (2010). ValuePick: Towards a value-oriented dual-goal recommender system. In *Proceedings of the IEEE International Conference on Data Mining* (pp. 1151–1158). http://dx.doi.org/10.1109/ICDMW.2010.68.

Amatriain, X., & Basilico, J. (2016). Past, present, and future of recommender systems: An industry perspective. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 211–214). http://dx.doi.org/10.1145/2959100.2959144.

Anelli, V. W., Bellogín, A., Noia, T. D., Jannach, D., & Pomo, C. (2022). Top-n recommendation algorithms: A quest for the state-of-the-art. In *Proceedings of the ACM Conference on User Modeling, Adaptation and Personalization* (pp. 121–131). http://dx.doi.org/10.1145/3503252.3531292.

Azaria, A., Hassidim, A., Kraus, S., Eshkol, A., Weintraub, O., & Netanely, I. (2013). Movie recommender system for profit maximization. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 121–128). http://dx.doi.org/10.1145/2507157.2507162.

Basu, S. (2021). Personalized product recommendations and firm performance. *Electronic Commerce Research and Applications*, *48*, Article 101074. http://dx.doi.org/10.1016/j.elerap.2021.101074.

Bellogín, A., & Said, A. (2019). Information retrieval and recommender systems. In A. Said, & V. Torra (Eds.), *Data Science in Practice, Vol. 46* (pp. 79–96). http://dx.doi.org/10.1007/978-3-319-97556-6_5.

Bennett, J., Lanning, S., et al. (2007). The netflix prize. In *Proceedings of the KDD Cup* (p. 35).

Billsus, D., & Pazzani, M. J. (1998). Learning collaborative information filters. In *Proceedings of the International Conference on Machine Learning* (pp. 46–54).

Bollen, D., Knijnenburg, B. P., Willemsen, M. C., & Graus, M. (2010). Understanding choice overload in recommender systems. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 63–70). http://dx.doi.org/10.1145/1864708.1864724.

Burke, R. (2007). Hybrid web recommender systems. In *The Adaptive Web* (pp. 377–408). Springer.

Burke, R., Abdollahpouri, H., Malthouse, E. C., Thai, K., & Zhang, Y. (2019). Recommendation in multistakeholder environments. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 566–567).

Cai, C., Fu, A., Cheng, C., & Kwong, W. (1998). Mining association rules with weighted items. In *Proceedings of the International Symposium on Database Engineering and Applications* (pp. 68–77). http://dx.doi.org/10.1109/IDEAS.1998.694360.

Cai, Y., & Zhu, D. (2019). Trustworthy and profit: A new value-based neighbor selection method in recommender systems under shilling attacks. *Decision Support Systems*, *124*, Article 113112. http://dx.doi.org/10.1016/j.dss.2019.113112.

Cavenaghi, E., Camaione, L., Minasi, P., Sottocornola, G., Stella, F., & Zanker, M. (2022). An online experiment of a price-based re-rank algorithm. In *Proceedings of the Workshop on Recommenders in Tourism*.

Chen, M., Chen, L., Hsu, F., Hsu, Y., & Chou, H. (2007). HPRS: A profitability based recommender system. In *Proceedings of the IEEE International Conference on Industrial Engineering and Engineering Management* (pp. 219–223). http://dx.doi.org/10.1109/IEEM.2007.4419183.

Chen, P., Chou, Y., & Kauffman, R. (2009). Community-based recommender systems: Analyzing business models from a systems operator's perspective. In *Proceedings of the Hawaii International Conference on System Sciences* (pp. 1–10). http://dx.doi.org/10.1109/HICSS.2009.117.

Chen, L., Hsu, F., Chen, M., & Hsu, Y. (2008). Developing recommender systems with the consideration of product profitability for sellers. *Information Sciences*, *178*, 1032–1048. http://dx.doi.org/10.1016/j.ins.2007.09.027.

Chen, P. Y., Wu, S. y., & Yoon, J. (2004). The impact of online recommendations and consumer feedback on sales. In *Proceedings of the International Conference on Information Systems* (pp. 711–724).

Cheng, H. T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., et al. (2016). Wide & deep learning for recommender systems. In *Proceedings of the Workshop on Deep Learning for Recommender Systems, co-located with ACM RecSys '16* (pp. 7–10).

Concha-Carrasco, J. A., Vega-Rodríguez, M. A., & Pérez, C. J. (2023). A multi-objective artificial bee colony approach for profit-aware recommender systems. *Information Sciences*, *625*, 476–488. http://dx.doi.org/10.1016/j.ins.2023.01.050.

Cooke, A. D., Sujan, H., Sujan, M., & Weitz, B. A. (2002). Marketing the unfamiliar: the role of context and item-specific information in electronic agent recommendations. *Journal of Marketing Research*, *39*, 488–497. http://dx.doi.org/10.1509/jmkr.39.4.488.1912.

Das, A., Mathieu, C., & Ricketts, D. (2009). Maximizing profit using recommender systems. http://dx.doi.org/10.48550/arXiv.0908.3633.

De Biasio, A., Monaro, M., Oneto, L., Ballan, L., & Navarin, N. (2023). On the problem of recommendation for sensitive users and influential items: Simultaneously maintaining interest and diversity. *Knowledge-Based Systems*, *275*, Article 110699. http://dx.doi.org/10.1016/j.knosys.2023.110699.

De Biasio, A., Montagna, A., Aiolli, F., & Navarin, N. (2023). A systematic review of value-aware recommender systems. *Expert Systems with Applications*, *226*, Article 120131. http://dx.doi.org/10.1016/j.eswa.2023.120131.

De Biasio, A., Navarin, N., & Jannach, D. (2023). Economic recommender systems - a systematic review. *Electronic Commerce Research and Applications*, *63*, Article 101352.

Demirezen, E. M., & Kumar, S. (2016). Optimization of recommender systems based on inventory. *Production and Operations Management*, *25*, 593–608. http://dx.doi.org/10.1111/poms.12540.

Deshpande, M., & Karypis, G. (2004). Item-based top-n recommendation algorithms. *ACM Transactions on Information Systems (TOIS)*, *22*, 143–177.

Dookeram, N., Hosein, Z., & Hosein, P. (2022). A recommender system for the upselling of telecommunications products. In *Proceedings of the International Conference on Advanced Communication Technology* (pp. 66–72). http://dx.doi.org/10.23919/ICACT53585.2022.9728818.

Ettl, M., Harsha, P., Papush, A., & Perakis, G. (2020). A data-driven approach to personalized bundle pricing and recommendation. *Manufacturing and Service Operations Management*, *22*, 461–480. http://dx.doi.org/10.1287/msom.2018.0756.

Garfinkel, R., Gopal, R., Pathak, B., Venkatesan, R., & Yin, F. (2007). Empirical analysis of the business value of recommender systems. *Social Science Research Network*, http://dx.doi.org/10.2139/SSRN.958770.

Garfinkel, R., Gopal, R., Pathak, B., & Yin, F. (2008). Shopbot 2.0: Integrating recommendations and promotions with comparison shopping. *Decision Support Systems*, *46*, 61–69. http://dx.doi.org/10.1016/j.dss.2008.05.006.

Ge, Y., Xu, S., Liu, S., Geng, S., Fu, Z., & Zhang, Y. (2019). Maximizing marginal utility per dollar for economic recommendation. In *Proceedings of the International Conference on World Wide Web* (pp. 2757–2763). http://dx.doi.org/10.1145/3308558.3313725.

Ghanem, N., Leitner, S., & Jannach, D. (2022). Balancing consumer and business value of recommender systems: A simulation-based analysis. *Electronic Commerce Research and Applications*, *55*, Article 101195. http://dx.doi.org/10.1016/j.elerap.2022.101195.

Ghoshal, A., Mookerjee, V. S., & Sarkar, S. (2021). Recommendations and cross-selling: Pricing strategies when personalizing firms cross-sell. *Journal of Management Information Systems*, *38*, 430–456. http://dx.doi.org/10.1080/07421222.2021.1912930.

Gomez-Uribe, C. A., & Hunt, N. (2016). The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, *6*, 1–19. http://dx.doi.org/10.1145/2843948.

Gorgoglione, M., Panniello, U., & Tuzhilin, A. (2011). The effect of context-aware recommendations on customer purchasing behavior and trust. In *Proceedings of ACM Conference on Recommender Systems* (pp. 85–92). http://dx.doi.org/10.1145/2043932.2043951.

Goyal, A., & Lakshmanan, L. V. (2012). RecMax: exploiting recommender systems for fun and profit. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1294–1302). http://dx.doi.org/10.1145/2339530.2339731.

Greenstein-Messica, A., & Rokach, L. (2018). Personal price aware multi-seller recommender system: Evidence from ebay. *Knowledge-Based Systems*, *150*, 14–26. http://dx.doi.org/10.1016/j.knosys.2018.02.026.

Gunawardana, A., Shani, G., & Yogev, S. (2022). Evaluating recommender systems. In L. Ricci, & B. Shapira (Eds.), *Recommender Systems Handbook* (pp. 547–601). http://dx.doi.org/10.1007/978-1-0716-2197-4_15.

Guo, L., Jin, J., Zhang, H., Zheng, Z., Yang, Z., Xing, Z., et al. (2021). We know what you want: An advertising strategy recommender system for online advertising. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 2919–2927). http://dx.doi.org/10.1145/3447548.3467175.

Guo, H., Tang, R., Ye, Y., Li, Z., & He, X. (2017). Deepfm: a factorization-machine based neural network for CTR prediction. In *Proceedings of the International Joint Conference on Artificial Intelligence* (pp. 1725–1731).

Harper, F. M., & Konstan, J. A. (2016). The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, *5*, 1–19. http://dx.doi.org/10.1145/2827872.

He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. S. (2017). Neural collaborative filtering. In *Proceedings of the International Conference on World Wide Web* (pp. 173–182).

He, P., Liu, H., Zhao, X., Liu, H., & Tang, J. (2022). PROPN: Personalized probabilistic strategic parameter optimization in recommendations. In *Proceedings of the ACM International Conference on Information and Knowledge Management* (pp. 3152–3161). http://dx.doi.org/10.1145/3511808.3557130.

Heien, D. M. (1980). Markup pricing in a dynamic model of the food industry. *American Journal of Agricultural Economics*, *62*, 10–18.

Herder, E. (2019). The need for identifying ways to monetize personalization and recommendation. In *Adjunct Proceedings of the ACM Conference on User Modeling, Adaptation and Personalization* (pp. 291–294). http://dx.doi.org/10.1145/3314183.3323844.

Hipp, J., Güntzer, U., & Nakhaeizadeh, G. (2000). Algorithms for association rule mining – a general survey and comparison. *ACM SIGKDD Explorations Newsletter*, *2*, 58–64. http://dx.doi.org/10.1145/360402.360421.

Hoffman, D. L., & Novak, T. P. (2005). A conceptual framework for considering web-based business models and potential revenue streams. *International Journal of Marketing Education*, *1*, 7–34.

Hosanagar, R. (2008). Recommended for you: the impact of profit incentives on the relevance of online recommendations. In *Proceedings of the International Conference on Information Systems* (p. 31).

Huang, C., Ding, Y., Hu, W., Jiang, Y., & Li, Y. (2021). Cost-based attraction recommendation for tour operators under stochastic demand. *Omega*, *102*, Article 102314. http://dx.doi.org/10.1016/j.omega.2020.102314.

Iwata, T., Saito, K., & Yamada, T. (2008). Recommendation method for improving customer lifetime value. *IEEE Transactions on Knowledge and Data Engineering*, *20*, 1254–1263. http://dx.doi.org/10.1109/TKDE.2008.55.

Jannach, D., & Adomavicius, G. (2016). Recommendations with a purpose. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 7–10). http://dx.doi.org/10.1145/2959100.2959186.

Jannach, D., & Adomavicius, G. (2017). Price and profit awareness in recommender systems. In *Proceedings of the Workshop on Value-Aware and Multistakeholder Recommendation*.

Jannach, D., & Jugovac, M. (2019). Measuring the business value of recommender systems. *ACM Trans. Manage. Inf. Syst.*, *10*, 1–23. http://dx.doi.org/10.1145/3370082.

Jannach, D., Zanker, M., Felfernig, A., & Friedrich, G. (2010). *Recommender Systems: An Introduction*. New York, NY, USA: Cambridge University Press.

Järvelin, K., & Kekäläinen, J. (2017). IR evaluation methods for retrieving highly relevant documents. In *ACM SIGIR Forum* (pp. 243–250). New York, NY, USA: ACM.

Ji, L., Qin, Q., Han, B., & Yang, H. (2021). Reinforcement learning to optimize lifetime value in cold-start recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management* (pp. 782–791). http://dx.doi.org/10.1145/3459637.3482292.

Karlsson, R., & Nilsson, B. J. (2013). Using maximum coverage to optimize recommendation systems in e-commerce. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 265–272). http://dx.doi.org/10.1145/2507157.2507169.

Kobayashi, M., & Takeda, K. (2000). Information retrieval on the web. *ACM Computing Surveys*, (32), http://dx.doi.org/10.1145/358923.358934.

Kompan, M., Gaspar, P., Macina, J., Cimerman, M., & Bielikova, M. (2022). Exploring customer price preference and product profit role in recommender systems. *IEEE Intelligent Systems*, *37*, 89–98. http://dx.doi.org/10.1109/MIS.2021.3092768.

Koren, Y. (2008). Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 426–434).

Koren, Y., & Bell, R. (2021). Advances in collaborative filtering. In *Recommender Systems Handbook* (pp. 91–142).

Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, *42*, 30–37. http://dx.doi.org/10.1109/MC.2009.263.

Krasnodebski, J., & Dines, J. (2016). Considering supplier relations and monetization in designing recommendation systems. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 381–382). http://dx.doi.org/10.1145/2959100.2959124.

Lee, D., & Hosanagar, K. (2014). Impact of recommender systems on sales volume and diversity. In *Proceedings of the International Conference on Information Systems* (p. 15).

Li, Z., Fang, X., Bai, X., & Sheng, O. R. L. (2017). Utility-based link recommendation for online social networks. *Management Science*, *63*, 1938–1952. http://dx.doi.org/10.1287/mnsc.2016.2446.

Li, Y., Zhang, Y., Gan, L., Hong, G., Zhou, Z., & Li, Q. (2021). RevMan: Revenue-aware multi-task online insurance recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*, 303–310.

Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, *7*, 76–80. http://dx.doi.org/10.1109/MIC.2003.1167344.

Lops, P., Gemmis, M., & Semeraro, G. (2011). Content-based recommender systems: State of the art and trends. In *Recommender Systems Handbook* (pp. 73–105). Springer.

Louca, R., Bhattacharya, M., Hu, D., & Hong, L. (2019). Joint optimization of profit and relevance for recommendation systems in E-commerce. In *Proceedings of the Workshop on Recommendation in Multistakeholder Environments*.

Malthouse, E. C., Vakeel, K. A., Hessary, Y. K., Burke, R., & Fudurić, M. (2019). A multistakeholder recommender systems algorithm for allocating sponsored recommendations. In *Proceedings of the Workshop on Recommendation in Multistakeholder Environments, co-located with ACM RecSys '19*.

Maslowska, E., Malthouse, E. C., & Hollebeek, L. D. (2022). The role of recommender systems in fostering consumers' long-term platform engagement. *Journal of Service Management*, *33*, 721–732. http://dx.doi.org/10.1108/JOSM-12-2021-0487.

Melomed, E., Gorbach, I., Berger, A., & Bateman, P. (2006). *Microsoft SQL Server 2005 Analysis Services*. Sams Publishing.

Microsoft Corporation (1998). Example database FoodMart of microsoft analysis services. URL: https://github.com/julianhyde/foodmart-data-hsqldb.

Montagna, A., De Biasio, A., Navarin, N., Aiolli, F., et al. (2023). Graph-based explainable recommendation systems: Are we rigorously evaluating explanations? In *CEUR Workshop Proceedings, CEUR-WS*.

Nemati, Y., & Khademolhosseini, H. (2020). Devising a profit-aware recommender system using multi-objective GA. *Journal of Advances in Computer Research, 10*, 9–120.

Ni, J., Li, J., & McAuley, J. (2019). Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing* (pp. 188–197).

Nikolakopoulos, A. N., Ning, X., Desrosiers, C., & Karypis, G. (2022). Trust your neighbors: A comprehensive survey of neighborhood-based methods for recommender systems. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender Systems Handbook* (pp. 39–89). http://dx.doi.org/10.1007/978-1-0716-2197-4_2.

Ning, X., & Karypis, G. (2011). SLIM: Sparse linear methods for top-n recommender systems. In *Proceedings of the IEEE International Conference on Data Mining* (pp. 497–506). IEEE.

Panniello, U., Gorgoglione, M., & Tuzhilin, A. (2016). In carss we trust: How context-aware recommendations affect customers' trust and other business performance measures of recommender systems. *Information Systems Research*, *27*, 182–196. http://dx.doi.org/10.1287/isre.2015.0610.

Panniello, U., Hill, S., & Gorgoglione, M. (2016). The impact of profit incentives on the relevance of online recommendations. *Electronic Commerce Research and Applications*, *20*, 87–104. http://dx.doi.org/10.1016/j.elerap.2016.10.003.

Pei, C., Yang, X., Cui, Q., Lin, X., Sun, F., Jiang, P., et al. (2019). Value-aware recommendation based on reinforcement profit maximization. In *Proceedings of the International Conference on World Wide Web* (pp. 3123–3129).

Piton, T., Blanchard, J., & Guillet, F. (2011). CAPRE: A new methodology for product recommendation based on customer actionability and profitability. In *Proceedings of the IEEE International Conference on Data Mining Workshops* (pp. 466–473). http://dx.doi.org/10.1109/ICDMW.2011.42.

Qu, M., Zhu, H., Liu, J., Liu, G., & Xiong, H. (2014). A cost-effective recommender system for taxi drivers. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 45–54). http://dx.doi.org/10.1145/2623330.2623668.

Ren, J., Kauffman, R. J., & King, D. (2019). Two-sided value-based music artist recommendation in streaming music services. In *Proceedings of the Hawaii International Conference on System Sciences*.

Ren, J., & Zhang, Y. (2021). Is all-win possible? A novel smart contracts-based music recommendation system. In *Proceedings of the International Conference on Information Communication and Management* (pp. 18–23). http://dx.doi.org/10.1145/3484399.3484402.

Rendle, S. (2022). Item recommendation from implicit feedback. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender Systems Handbook* (pp. 143–171). http://dx.doi.org/10.1007/978-1-0716-2197-4_4.

Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2009). BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence* (pp. 452–461).

Rendle, S., Krichene, W., Zhang, L., & Anderson, J. (2020). Neural collaborative filtering vs. matrix factorization revisited. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 240–248). http://dx.doi.org/10.1145/3383313.3412488.

Rendle, S., Krichene, W., Zhang, L., & Koren, Y. (2022). Revisiting the performance of iALS on item recommendation benchmarks. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 427–435). http://dx.doi.org/10.1145/3523227.3548486.

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 175–186).

Resnick, P., & Varian, H. R. (1997). Recommender systems. *Communications of the ACM*, 40.

Salakhutdinov, R., Mnih, A., & Hinton, G. (2007). Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the International Conference on Machine Learning* (pp. 791–798). http://dx.doi.org/10.1145/1273496.1273596.

Sammut-Bonnici, T., & Channon, D. F. (2015). Pricing strategy. *Wiley Encyclopedia of Management*, 1–3.

Schafer, J. B., Konstan, J., & Riedl, J. (1999). Recommender systems in E-commerce. In *Proceedings of the ACM Conference on Electronic Commerce* (pp. 158–166). http://dx.doi.org/10.1145/336992.337035.

Schafer, J. B., Konstan, J. A., & Riedl, J. (2001). E-commerce recommendation applications. *Data Mining and Knowledge Discovery*, *5*, 115–153. http://dx.doi.org/10.1023/A:1009804230409.

Seymen, S., Sachs, A. L., & Malthouse, E. C. (2022). Making smart recommendations for perishable and stockout products. In *Proceedings of the Workshop on Multi-Objective Recommender Systems, co-located with ACM RecSys 2022*.

Steck, H., Baltrunas, L., Elahi, E., Liang, D., Raimond, Y., & Basilico, J. (2021). Deep learning for recommender systems: A netflix case study. *AI Magazine*, *42*, 7–18. http://dx.doi.org/10.1609/aimag.v42i3.18140.

Su, X., & Khoshgoftaar, T. M. (2009). A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009. http://dx.doi.org/10.1155/2009/421425.

Sun, Z., Guo, Q., Yang, J., Fang, H., Guo, G., Zhang, J., et al. (2019). Research commentary on recommendations with side information: A survey and research directions. *Electronic Commerce Research and Applications*, 37. http://dx.doi.org/10.1016/j.elerap.2019.100879.

Sun, Z., Yang, J., Feng, K., Fang, H., Qu, X., & Ong, Y. S. (2022). Revisiting bundle recommendation: Datasets, tasks, challenges and opportunities for intent-aware product bundling. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2900–2911). http://dx.doi.org/10.1145/3477495.3531904.

Theocharous, G., Thomas, P. S., & Ghavamzadeh, M. (2015). Personalized ad recommendation systems for life-time value optimization with guarantees. In *Proceedings of the International Conference on World Wide Web* (p. 7). http://dx.doi.org/10.1145/2740908.2741998.

Umberto, P. (2015). Developing a price-sensitive recommender system to improve accuracy and business performance of ecommerce applications. *International Journal of Electronic Commerce Studies*, *6*, 1–18. http://dx.doi.org/10.7903/ijecs.1348.

Villegas, N. M., Sánchez, C., Díaz-Cely, J., & Tamura, G. (2018). Characterizing context-aware recommender systems: A systematic literature review. *Knowledge-Based Systems*, *140*, 173–200. http://dx.doi.org/10.1016/j.knosys.2017.11.003.

Wang, K., & Su, M. Y. T. (2002). Item selection by hub-authority profit ranking. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 652–657). http://dx.doi.org/10.1145/775047.775144.

Wang, H. F., & Wu, C. T. (2009). A mathematical model for product selection strategies in a recommender system. *Expert Systems with Applications*, *36*, 7299–7308. http://dx.doi.org/10.1016/j.eswa.2008.09.006.

Wang, H. F., & Wu, C. T. (2012). A strategy-oriented operation module for recommender systems in E-commerce. *Computers & Operations Research*, *39*, 1837–1849. http://dx.doi.org/10.1016/j.cor.2010.03.011.

Wang, J., & Zhang, Y. (2011). Utilizing marginal net utility for recommendation in e-commerce. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 1003–1012). http://dx.doi.org/10.1145/2009916.2010050.

Wang, K., Zhou, S., & Han, J. (2002). Profit mining: From patterns to actions. In G. Goos, J. Hartmanis, J. van Leeuwen, C. S. Jensen, S. Šaltenis, K. G. Jeffery, J. Pokorny, E. Bertino, K. Böhn, & M. Jarke (Eds.), *Advances in Database Technology, Vol. 2287* (pp. 70–87). http://dx.doi.org/10.1007/3-540-45876-X_7.

Wu, S., Sun, F., Zhang, W., Xie, X., & Cui, B. (2022). Graph neural networks in recommender systems: a survey. *ACM Computing Surveys*, *55*, 1–37.

Xiao, B., & Benbasat, I. (2007). E-commerce product recommendation agents: Use, characteristics, and impact. *MIS Quarterly*, 137–209. http://dx.doi.org/10.2307/25148784.

Yang, R., Xu, M., Jones, P., & Samatova, N. (2017). Real time utility-based recommendation for revenue optimization via an adaptive online top-k high utility itemsets mining model. In *Proceedings of the International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery* (pp. 1859–1866). http://dx.doi.org/10.1109/FSKD.2017.8393050.

Zhan, R., Christakopoulou, K., Le, Y., Ooi, J., Mladenov, M., Beutel, A., et al. (2021). Towards content provider aware recommender systems: A simulation study on the interplay between user and provider utilities. In *Proceedings of the Web Conference* (pp. 3872–3883). http://dx.doi.org/10.1145/3442381.3449889.

Zhang, Z., Chen, N., Wang, J., & Si, L. (2017). SMART: Sponsored mobile app recommendation by balancing app downloads and appstore profit. In *Proceedings of the IEEE International Conference on Big Data* (pp. 1600–1609). http://dx.doi.org/10.1109/BigData.2017.8258094.

Zhang, X., Xu, B., Yang, L., Li, C., Ma, F., Liu, H., et al. (2022). Price DOES matter!: Modeling price and interest preferences in session-based recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 1684–1693). http://dx.doi.org/10.1145/3477495.3532043.

Zhang, X., Zhao, J., & LeCun, Y. (2015). Character-level convolutional networks for text classification. *Advances in Neural Information Processing Systems*, 28.

Zhang, H., Zhao, P., Xian, X., Sheng, V. S., Hao, Y., & Cui, Z. (2022). Click is not equal to purchase: Multi-task reinforcement learning for multi-behavior recommendation. In R. Chbeir, H. Huang, F. Silvestri, Y. Manolopoulos, & Y. Zhang (Eds.), *Proceedings of the Conference on Web Information Systems Engineering, Vol. 13724* (pp. 443–459). http://dx.doi.org/10.1007/978-3-031-20891-1_32.

Zhao, W. X., Lin, Z., Feng, Z., Wang, P., & Wen, J. R. (2023). A revisiting study of appropriate offline evaluation for top-N recommendation algorithms. *ACM Transactions on Information Systems*, *41*, 1–41. http://dx.doi.org/10.1145/3545796.

Zhao, Q., Zhang, Y., Friedman, D., & Tan, F. (2015). E-commerce recommendation with personalized promotion. In *Proceedings of the ACM Conference on Recommender Systems* (pp. 219–226). http://dx.doi.org/10.1145/2792838.2800178.

Zheng, Y., Gao, C., He, X., Jin, D., & Li, Y. (2021). Incorporating price into recommendation with graph convolutional networks. *IEEE Transactions on Knowledge and Data Engineering*, *35*, 1609–1623. http://dx.doi.org/10.1109/TKDE.2021.3091160.

Zheng, Y., Gao, C., He, X., Li, Y., & Jin, D. (2020). Price-aware recommendation with graph convolutional networks. In *Proceedings of the IEEE International Conference on Data Engineering* (pp. 133–144). http://dx.doi.org/10.1109/ICDE48307.2020.00019.

Zou, L., Xia, L., Ding, Z., Song, J., Liu, W., & Yin, D. (2019). Reinforcement learning to optimize long-term user engagement in recommender systems. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 2810–2818). http://dx.doi.org/10.1145/3292500.3330668.