

The impact of consumer preferences on the accuracy of collaborative filtering recommender systems

Sebastian Köhler¹  · Thomas Wöhner¹ · Ralf Peters¹

Received: 11 March 2016 / Accepted: 5 August 2016 / Published online: 26 August 2016
© Institute of Applied Informatics at University of Leipzig 2016

Abstract Despite the omnipresent use of recommender systems in electronic markets, previous research has not analyzed how consumer preferences affect the accuracy of recommender systems. Markets, however, are characterized by a certain structure of consumers' preferences. Consequently, it is not known in which markets recommender systems perform well. In this paper, we introduce a microeconomic model that allows a systematical analysis of different structures of consumers' preferences. We develop a model-specific metric to measure the recommendation accuracy. We employ our model in a simulation to evaluate the impact of the structure of the consumers' preferences on the accuracy of a popular collaborative filtering algorithm. Our study shows that recommendation accuracy is significantly affected by the similarity and number of consumer types and the distribution of consumers. The investigation reveals that in certain markets even random product recommendations outperform the collaborative filtering algorithm.

Keywords Recommender system · Consumer preferences · Recommendation accuracy · Collaborative filtering

JEL Classification M2 Business economics

Responsible Editor: Rainer Alt

✉ Sebastian Köhler
sebastian.koehler@wiwi.uni-halle.de

¹ Universitätsring 3, 06108 Halle (Saale), Germany

Introduction

Recommender systems are widely used today in electronic markets (Schafer et al. 1999), including [Amazon.com](#) (Linden et al. 2003), [MovieLens](#) (Miller et al. 2003), [Netflix](#) (Bennett and Lanning 2007), and [YouTube](#) (Davidson et al. 2010). With the wide variety of products in e-commerce, recommender systems aim at reducing consumers' search costs (Resnick and Varian 1997) and facilitating the discovery of preferred products. Consumers express their preferences (valuation of different products) while shopping online by giving implicit and explicit feedback, e.g., searching for products or rating products already bought (Jawaheer et al. 2014; Ricci et al. 2011). Recommender systems collect and use this buying behavior information to infer consumers' preferences and recommend products automatically.

Markets are characterized by a particular *structure of consumers' preferences*, which means that consumers' preferences are similar in some markets but different in others (Levin et al. 2003). The term *structure of consumers' preferences* describes the model of the consumers' preferences in a given market. In this paper, the structure of the consumers' preferences is represented by the following features: the number of consumer types (consumers with identical preferences), the similarity (commonly preferred products) of consumer types and the distribution of consumers to the various types.

Collaborative filtering recommender systems suggest products that are preferred by consumers with similar preferences (Ekstrand 2010; Jannach et al. 2011). Consequently, the product recommendations generated by the system depend on the underlying structure of the consumers' preferences (Adomavicius et al. 2011; Felfernig et al. 2007; Hotelling 1929; Tirole 1988). Despite the omnipresent use of recommender systems in electronic markets, there has been no research analyzing the impact of consumer preferences on the

accuracy of those systems. In particular, it is unknown which features of the structure of the consumers' preferences influence the recommendation accuracy in which way. Furthermore, the question arises whether particular structures of consumers' preferences make recommender systems ineffective. Consequently, online retailers do not know in which markets recommender systems perform well. There is no model to systematically investigate the impact of consumer preferences on recommendation accuracy.

This paper introduces a microeconomic model that enables the evaluation of recommender systems with regard to different structures of consumers' preferences. Consumer preference modeling is based on Hotelling's linear city model (Hotelling 1929; Tirole 1988). We also develop a model-specific metric to measure recommendation accuracy. Using that model, we conducted a simulation to study how the structure of consumers' preferences affects recommendation accuracy using an established collaborative filtering recommender system (Desrosiers and Karypis 2011; Jannach et al. 2011). In particular, we analyze three features that represent the structure of consumers' preferences in a market: the similarity and number of consumer types as well as the distribution of consumers to the various consumer types.

Our study follows a microeconomic research paradigm. We do not investigate existing markets; instead, we examine the general effects of different structures of consumers' preferences on the accuracy of recommender systems. Therefore, we keep our model as simple as possible to focus on the question whether and to what extent the analyzed features influence the recommendation accuracy. Our study finds that the similarity of consumer types has the largest effect. Recommendation accuracy is high only with a high or low similarity of consumer types. Moreover, our study finds that the number of consumer types facilitates a high recommendation accuracy. As a final result, our investigation shows that the distribution of consumers to the consumer types affects the recommendation accuracy only if a single consumer type comprises the majority of consumers.

The paper is structured as follows: Section 2 presents the related work and points out the research gap addressed in this study. Section 3 explains our modeling of consumer preferences and the collaborative filtering recommender system. In Section 4 we illustrate the simulation procedure and our model-specific quality metric. Subsequently, we describe the parameterization of our model, explicate the simulation scenarios, and analyze and discuss the results. Finally, Section 5 summarizes the paper and points out further research directions.

State of the art and related work

The most important recommendation approaches in literature and practice are content-based and collaborative filtering

(Adomavicius and Tuzhilin 2005). Content-based filtering systems recommend products that are similar to those products consumers liked in the past (Jannach et al. 2011; Lops et al. 2011). Thus, the recommender system needs information about the similarity of items. A collaborative filtering system recommends products for a given consumer that were bought by other consumers with similar buying behavior (Ekstrand 2010; Jannach et al. 2011). In contrast to content-based filtering, the similarity of items is derived implicitly. Collaborative filtering is the most widely used recommendation approach in e-commerce (Jannach et al. 2011). A well-known example for collaborative filtering is the recommender system of online retailer Amazon.com (Linden et al. 2003).

Collaborative filtering recommender systems collect and use consumers' product ratings to generate product recommendations. In literature, collaborative filtering systems are categorized into neighborhood-based and model-based methods (Adomavicius and Tuzhilin 2005; Desrosiers and Karypis 2011). Model-based recommendation approaches employ Machine Learning methods to train complex predictive models and calculate individual purchase probabilities (Desrosiers and Karypis 2011; Su and Khoshgoftaar 2009). In contrast, the rules for generating product recommendations are defined manually with neighborhood-based recommendation approaches (Desrosiers and Karypis 2011). Recommendation generation is realized using item-based or user-based techniques. Item-based filtering uses consumers' product ratings to measure the similarity of products and recommend products that are similar to the products already purchased by a consumer (Jannach et al. 2011; Sarwar et al. 2001; Su and Khoshgoftaar 2009). In comparison, user-based filtering uses the product ratings to quantify consumers' similarities and recommends products that similar consumers purchased (Ekstrand 2010; Jannach et al. 2011). In our paper, we employ the popular user-based collaborative filtering (Desrosiers and Karypis 2011).

These approaches can be implemented differently, by selecting different similarity metrics, for example. Numerous research papers study the configuration and possible extensions of recommender systems to improve prediction accuracy (Adomavicius and Tuzhilin 2005; Herlocker et al. 1999; Herlocker et al. 2004). However, the developed artefacts are usually evaluated for a given application scenario using real data.¹ Therefore, the findings refer to the given data set and the underlying structure of consumers' preferences.

Also, there is a growing amount of research investigating the effects of recommender systems. Several publications in this research branch address the diversity of sales and, as a consequence, the impact of recommender systems on the long tail (Fleder and Hosanagar 2007; Hinz and Eckert 2010; Zhou et al. 2010).

¹ For example, the MovieLens data set is often used for evaluation (Herlocker et al. 2004).

However, researchers have not investigated the effect of the structure of consumers' preferences on recommendation accuracy. Nevertheless, a few studies on recommender systems also deal with consumer preferences. Some publications that focus on the effects of recommender systems on sales diversity apply simple models of consumer preferences (Hervas-Drane 2015; Hinz and Eckert 2010; Wu et al. 2011). Furthermore, recommender systems that follow the multi-attribute utility theory (MAUT) use queries to explicitly collect information about consumers' preferences (Pfeiffer and Scholz 2013; Scholz et al. 2015). Finally, several publications address the inference of preferences from the consumers' buying behavior (Gemmis et al. 2011; Karatzoglou and Weimer 2011; Rashid et al. 2002). However, there is no model that allows a systematic analysis of different structures of consumers' preferences and their effects on recommendation accuracy.

Model

This section describes the modeling of consumer preferences and explains the collaborative filtering algorithm for generating recommendations.

Modeling of consumer preferences

As a typical scenario for recommender systems we assume a horizontal product differentiation (Bergemann and Ozmen 2006; Hervas-Drane 2015; Hinz and Eckert 2010). Consumer preferences for such a scenario are modelled in microeconomics Hotelling's linear city model (Hotelling 1929; Tirole 1988).

According to this model, we consider a *market* where ng different products of a certain product category are offered. The products $i \in \{1, \dots, ng\}$ differ in a single characteristic g_i . The products are located at the positions $G = (g_1, \dots, g_{ng})$ within a one-dimensional differentiation spectrum. The differentiation characteristic is normalized for simplification so the position of a given product i is defined by $g_i \in [0, 1]$. A chocolate bar that is differentiated according to the cocoa content may serve as an example to illustrate horizontal product differentiation. The chocolate bar i at position $g_i = 0$ (minimum cocoa content) corresponds to white chocolate, whereas the product at $g_i = 1$ is dark chocolate.

Consumers $u \in \{1, \dots, nc\}$ differ according to their preferences. *Consumer preferences* describe consumers' valuations for the products available in the market. The position of the most preferred product of consumer u within the differentiation spectrum is defined by $c_u \in [0, 1]$. For this product, consumers have their maximum willingness to pay. In our model, we assume an identical maximum willingness to pay v^{\max} for all consumers. If product i at position g_i does not

correspond to the most preferred product of a consumer u at position c_u , the willingness to pay is reduced with increasing difference $|c_u - g_i|$. Following Hotelling's linear city (Hotelling 1929), there is a linear dependency between this difference and the willingness to pay. Thus, the willingness to pay $v(u, i)$ of consumer u for the product i is calculated as follows:

$$v(u, i) = v^{\max} - \tau * |c_u - g_i| \quad (1)$$

Parameter τ is a weighting coefficient that specifies how consumers assess the distance to their most preferred product. In our model, we assume that τ is identical for all consumers. The consumer surplus $CS(u, i)$ of consumer u from buying product i at price p is defined as:

$$CS(u, i) = v(u, i) - p \quad (2)$$

The consumer's willingness to pay and the consumer surplus decreases as the distance between the most preferred and the products offered grows. The *preference spectrum* of a given consumer u defines the entire set of products leading to a positive consumer surplus $CS(u, i) > 0$. It contains all products that would be purchased by the consumer u . Figure 1 uses a hypothetical example to illustrate the model of consumer preferences.

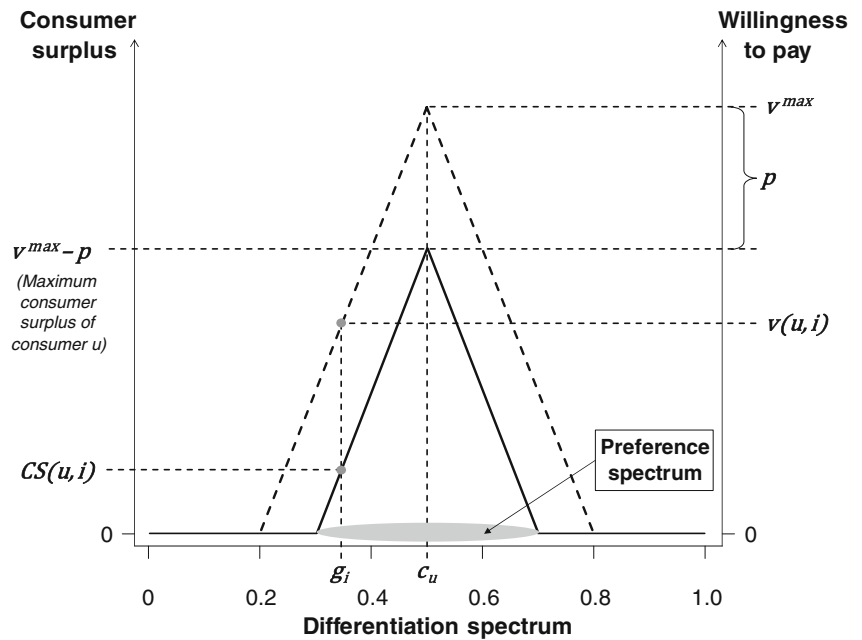
Furthermore, our model considers different consumer types. Each member of a given consumer type $x \in \{1, \dots, nt\}$ has exactly the same consumer preference. The position of the most preferred product of consumer type x is denoted by t_x . Parameter h_x defines the number of consumers for a given type x such that the condition $\sum_x h_x = nc$ holds.

As described above, the objective of this paper is to study how consumers' preferences affect the accuracy of a collaborative filtering recommender system. According to our model, this structure of the consumers' preferences is defined by the vector $T = (t_1, \dots, t_{nt})$, which represents the positions of the consumer types' most preferred products, and the vector $H = (h_1, \dots, h_{nt})$, which describes the size of the various consumer types. Moreover, the weighting coefficient τ , product price p , and maximum willingness to pay v^{\max} determine the length of the preference spectrums and the similarity of consumer types (commonly preferred products). These parameters are also part of the definition of the structure of consumers' preferences.

Modeling the collaborative filtering recommender system

Our analysis used the popular user-based collaborative filtering system. The implementation of the recommender system follows Desrosiers and Karypis (2011). As mentioned above, collaborative filtering is based on consumers' product ratings. Given a vector of product ratings $R_u = (r_{u,1}, \dots, r_{u,ng})$ for each consumer u and a set of ng products, the generation of

Fig. 1 Model of consumer preferences



product recommendations proceeds as follows: first, consumers’ similarity is determined. We use the Pearson product-moment correlation coefficient (Desrosiers and Karypis 2011; Jannach et al. 2011) to measure the similarity $sim(u, w)$ between consumer u and consumer w . The rating vectors R_u and R_w are required for this calculation. Only the ratings of the products that have already been purchased and rated by both consumers are considered. Consequently, for generating recommendations, only those consumers are taken into account that have common ratings with the consumer concerned. This group of consumers is referred to as neighbors. The recommendation value $rec(u, i)$ for a product i and consumer u is calculated as follows (Desrosiers and Karypis 2011):

$$rec(u, i) = \frac{\sum_w sim(u, w) * r_{w,i}}{\sum_w |sim(u, w)|} \quad \forall \text{consumers } w \in 1, \dots, nc \text{ applying to: } u \neq w \text{ and } \exists sim(u, w) \text{ and } r_{w,i} > 0 \quad (3)$$

The recommendation value $rec(u, i)$ is a weighted average of consumer ratings $r_{w,i}$, while the similarity $sim(u, w)$ is used as a weighting factor. Finally, the product with the highest recommendation value that has not already been purchased by consumer u is recommended.

Simulation and results

In this section, we examine how the structure of consumers’ preferences affects the accuracy of recommender systems using a simulation. First, we explain the simulation procedure and our model-specific quality metric and subsequently

describe the parameterization used. Finally, we introduce the simulation scenarios and discuss the results.

Simulation procedure and evaluation

Recommender systems are typically used for products such as books and movies, e.g., on Amazon.com (Linden et al. 2003) and MovieLens (Miller et al. 2003). In such markets, consumers purchase various products in a given product category. To analyze such a scenario with multiple purchases, we employed a round-based simulation implemented in Java. The consumers are processed in order of their index in each round $l \in \{1, \dots, nl\}$. Each consumer gets a product recommendation in each round according to the algorithm described in the previous section. The product will be purchased if there is a consumer surplus $CS(u, i) > 0$. Subsequent to a purchase, a product rating correspondent to the achieved consumer surplus is generated. In detail, the rating $r_{u,i} = CS(u, i)$ will be created if consumer u purchases product i . We assume a durable good that can be purchased only once per consumer.

The rating vectors R_u are initialized with zero for each consumer u (i.e. $r_{u,i} = 0, \forall$ consumers u and products i), since there are no purchases at the beginning of the simulation. With user-based collaborative filtering, the recommendation generation is only possible if at least one neighbor (consumer with common ratings) exists for a given consumer. Thus, there is a cold-start problem at the time of initialization (Jannach et al. 2011). Our simulation overcomes this issue by recommending a randomly chosen product when the collaborative filtering could not calculate any product recommendations. Additionally, recommendation generation is not possible if

all neighbors have purchased exactly the same products. In this case, the recommended product is also chosen by chance.

We developed the efficiency E as a model-specific metric to measure recommendation accuracy. The maximum consumer surplus CS^{\max} after nl simulation rounds that is attained in case of perfect recommendations is used as point of reference. In this case, in each round, consumers receive the product with the highest consumer surplus that has not been purchased by the respective consumer. The products are recommended in descending order of the achievable consumer surplus. Based on this, the efficiency E denotes the quotient of all achieved consumer surpluses within the simulation and the total maximum consumer surplus. The calculation is described in formula (4), where $CS_k(u)$ denotes the achieved consumer surplus of consumer u in simulation round k .

$$E = \frac{\sum_{u=1}^{nc} \sum_{k=1}^{nl} CS_k(u)}{nc * CS^{\max}} \quad (4)$$

The efficiency can be any value on the interval $[0, 1]$ and refers to the percentage of the maximum consumer surplus that is achieved based on the recommendations of the recommender system. If the consumers always get perfect recommendations according to their preferences and therefore achieve the maximum consumer surplus, the efficiency amounts to 1. Otherwise, the efficiency is 0 if only products outside of the preference spectrum are recommended.

Parameterization

The impact of the structure of consumers' preferences on recommendation accuracy of collaborative filtering recommender systems is investigated based on three different scenarios. Each scenario considers a market setting that comprises $ng = 200$ products and $nc = 200$ consumers. The products $1, \dots, ng$ are uniformly distributed within the differentiation spectrum on the interval $[0, 1]$. The price of each product amounts to $p = 1$ €. The consumers have an identical maximum willingness to pay $v^{\max} = 2$ €. The distance between an offered product and the consumer's most preferred product is valued with $\tau = 4$ €. The willingness to pay $v(u, i)$ of a consumer u varies on the interval $[0, 2]$ depending on the position of an offered product i . All the products i with a distance $|c_u - g_i| \leq (v^{\max} - p)/\tau$ to the most preferred product of consumer u are located within the preference spectrum. The consumers' most preferred products are positioned on the interval $c_u \in [0.25, 0.75]$ in our simulation, such that the preference spectrum of each consumer has exactly the same length of $2 * (v^{\max} - p)/\tau = 0.5$.²

² If the most preferred product is positioned at the edge of the differentiation spectrum, the preference spectrum would be smaller. For example, $c_u = 0.1$ would lead to a length of the preference spectrum of 0.35.

Once again, the aim of this study is not the investigation of existing markets. The parameters p , v^{\max} and τ are chosen in such a way that the preference spectrum comprises a set of 100 products. This leads to a probability of 50 % that a random product recommendation is within the preference spectrum of a given consumer. Consequently, the cold-start problem is overcome in a few rounds according to our initialization strategy. Furthermore, the chosen parameterization allows an adequate variation of consumer types (100 different consumer types are possible). The number of simulation rounds is fixed to $nl = 100$, so that each consumer buys all the products within his or her preference spectrum in the case of optimal recommendations. Table 1 summarizes the important parameters of our model and the parameterization used in the simulation.

We investigate three simulation scenarios. Each scenario addresses a particular feature of the structure of consumers' preferences. The simulation scenarios are chosen in such a way that the effect of each feature can be examined separately. In detail, we focus on the following three features:

- Similarity of consumer types (Scenario 1)
- Number of consumer types (Scenario 2)
- Distribution of consumers to the consumer types (Scenario 3)

The particular feature is systematically varied in each scenario in order to explore the impact on recommendation accuracy. We carry out 30 simulation runs for each feature configuration to compensate for potential biases caused by the random bootstrapping. The averaged results are used for evaluation and interpretation. The number of replications was determined experimentally. There was no significant change in the averaged results with more than 30 replications.

Similarity of consumer types (Scenario 1)

The similarity of consumer types corresponds to the conformity in the tastes of the different consumer types. It is represented by the intersection of the preference spectrums of the consumer types. In other words, the similarity refers to the percentage of products that are commonly preferred by different consumer types. Scenario 1 investigates the impact of the similarity on the example of $nt = 2$ consumer types with 100 consumers each. Given the parameterization for p , v^{\max} and τ introduced above, the similarity is modelled by the distance $|t_1 - t_2|$ between the most preferred products of the two consumer types.

For a practical explanation of the similarity of consumer types, we again employ the chocolate example. We assume that chocolate bars vary in the cocoa content from 0 g ($g_i = 0$) to 20 g ($g_i = 1$). Given our parameterization of $p = 1$ €, $v^{\max} = 2$ € and $\tau = 4$ €, the preference spectrum of each consumer has a length of 0.5 (10 g cocoa content), meaning that a difference of 5 g cocoa content to the

Table 1 Parameters and parameterization

Parameter	Explanation	Parameterization
g_i	Position of product i within the differentiation spectrum	$g_i \in [0, 1]$
c_u	Position of the most preferred product of consumer u	$c_u \in [0.25, 0.75]$
v^{\max}	Maximum willingness to pay	$v^{\max} = 2 \text{ €}$
τ	Weighting coefficient that specifies how consumers value the distance $ c_u - g_i $ to their most preferred product	$\tau = 4 \text{ €}$
$v(u, i)$	Willingness to pay of consumer u for product i	$v(u, i) \in [0, 2]$
p	Product price	$p = 1 \text{ €}$
CS^{\max}	Maximum consumer surplus after nl simulation rounds	$CS^{\max} = 50$
$CS(u, i)$	Consumer surplus of consumer u for product i	$CS(u, i) \in [0, 1]$
$r_{u,i}$	Rating of consumer u for product i	$r_{u,i} \in [0, 1]$
$sim(u, w)$	Similarity (Pearson product-moment correlation coefficient) between consumer u and consumer w	$sim(u, w) \in [0, 1]$
$rec(u, i)$	Recommendation value for product i and consumer u	$rec(u, i) \in [0, 1]$
E	Efficiency of the recommender system	$E \in [0, 1]$
t_x	Position of the most preferred product of consumer type x	$t_x \in [0.25, 0.75]$
ng	Number of products	$ng = 200$
nc	Number of consumers	$nc = 200$
nl	Number of simulation rounds	$nl = 100$
nt	Number of consumer types	Depends on the scenario
h_x	Number of consumers of consumer type x	Depends on the scenario

most preferred chocolate bar is accepted. If the consumer types are positioned at $t_1 = 0.375$ (7.5 g cocoa content) and $t_2 = 0.625$ (12.5 g cocoa content), for example, the intersection of the two preference spectrums is 50 %. In this case, chocolate bars with a cocoa content from 7.5 g to 12.5 g are accepted by both consumer types.

In order to analyze how the similarity affects the accuracy of collaborative filtering, we systematically vary the intersection of the two consumer types in our simulation. We increase the intersection from 0 % ($t_1 = 0.25$ and $t_2 = 0.75$, i.e. $|t_1 - t_2| = 0.5$; distinct preference spectrums of the two consumer types) to 100 % ($t_1 = t_2 = 0.5$, i.e. $|t_1 - t_2| = 0$; identical consumer preferences) in steps of 4 %. Figure 2 illustrates the intersection of preference spectrums on the example of an intersection of 0 % and 60 %. The grey-colored area refers to the intersection.

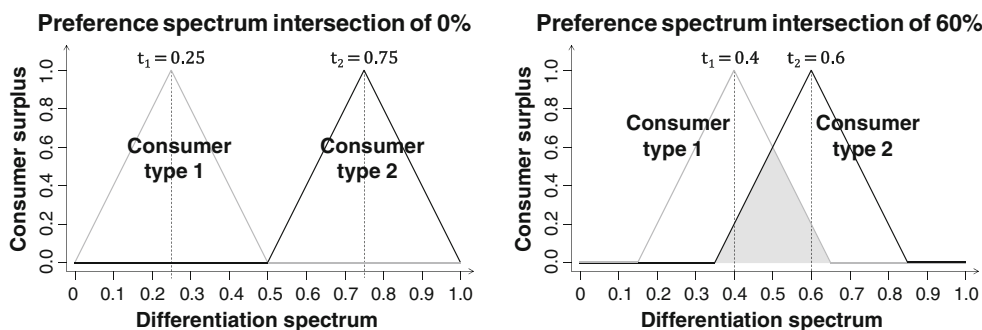
Figure 3 depicts the efficiency of the recommender system as a function of the similarity of both consumer types. Our analysis shows that the similarity has a considerable impact on the recommendation accuracy of the collaborative filtering system. The efficiency amounts to nearly 100 %³ in the case of disjoint preference spectrums (intersection of 0 %). The efficiency decreases with growing similarity until the minimum efficiency of 10 % is achieved at an intersection of 50 %. If the similarity further increases (intersection between 50 % and 100 %), the efficiency will grow up to nearly 100 % for identical consumer types.

³ The minor gap to an efficiency of 100 % (perfect recommendations) is caused by the random bootstrapping.

Irrelevant product recommendations ($CS \leq 0$) are possible if a given consumer has a higher correlation to another consumer type than to its own. Even in this case, the recommendation is irrelevant only if the recommended product, purchased by the other consumer type, is located outside of the given consumer's preference spectrum. The U-shaped curve in Fig. 3 is explained by two competing effects. Generally, the size of the intersection refers to the number of commonly and non-commonly preferred products of consumers of different consumer types. As the intersection and number of commonly preferred products grow, the probability increases that consumers of different types buy the same products. As a first effect, a growing intersection increases the probability that a higher correlation to the wrong consumer type is calculated and unsuitable products are recommended ($CS \leq 0$), reducing recommendation accuracy. As a second effect, the set of products that are relevant for both consumer types ($CS > 0$) grows as the similarity increases. Therefore, an increasing intersection reduces the probability that a product purchased by another consumer type is located outside of the preference spectrum of the given consumer. This effect leads to an increasing recommendation accuracy as the similarity grows. The first effect is dominant for intersection rates ≤ 50 %, while the second effect is dominant for intersection rates > 50 %. Consequently, the increasing similarity of consumer types causes in total a U-shaped curve of efficiency.

Another interesting finding is the comparison of the efficiency between the collaborative filtering system and random product recommendations. The recommendation accuracy

Fig. 2 Intersection of the preference spectrums of two consumer types of 0 % and 60 %



with randomly chosen products is independent from the similarity of consumer types and amounts to 39.1 % (dashed line in Fig. 3) for the given parameterization. Our study finds that the efficiency of collaborative filtering is partially lower compared to random recommendations. This shows that in some markets collaborative filtering performs worse than random recommendations.

Number of consumer types (Scenario 2)

Scenario 2 examines how the number of consumer types nt affects recommendation accuracy. We vary the number of consumer types and therefore the number of consumers per consumer type in each configuration of this scenario. We consider all the settings that allow an equal distribution of the consumers to the consumer types so that each type x comprises $h_x = 200/nt$ consumers. Table 2 shows the configurations we investigate. The two consumer types 1 and nt at the edge of the differentiation spectrum are always located at the positions $t_1 = 0.25$ and $t_{nt} = 0.75$. Any other types are positioned equidistantly between them. Consequently, the differentiation spectrum is completely covered by the consumers and the preference spectrum of each consumer comprises 100 products. Therefore, the consumer types are positioned closer and the intersection of their preference spectrums increases as the number of types grows.

The efficiency of our collaborative filtering recommender system as a function of the number of consumer types is

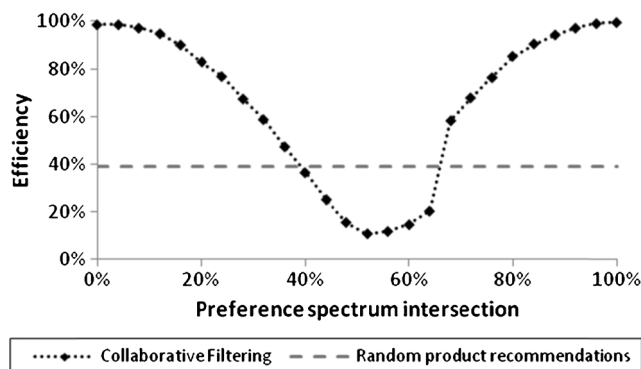


Fig. 3 Efficiency as a function of the similarity of consumer types

plotted in Fig. 4. The analysis demonstrates that recommendation accuracy is also affected by the number of consumer types. The curve shows that the efficiency decreases at first as the number of consumer types grows. With $nt = 5$ consumer types ($h_x = 40$ consumers per type), we have the minimum efficiency of 17 % in scenario 2. Subsequently, the efficiency increases as the number of consumer types grows.

Two competing effects can be observed again: on one hand, there are more common purchases by consumers of different types as the number of consumer types grows. This first effect increases the probability of a higher correlation to the wrong consumer type and unsuitable product recommendations. On the other hand, the growing number of products commonly preferred by consumers of adjacent types increases the probability of relevant product recommendations. This second effect is dominant for more than five consumer types in the scenario shown and results in an increasing efficiency curve. In comparison to the similarity of consumer types, the efficiency curve as a function of the number of consumer types is continuously rising for nearly all configurations.

Our analysis in Fig. 4 points out that in each configuration the collaborative filtering recommender system has a lower efficiency than random product recommendations. The low efficiency is caused by the disadvantageous relation between commonly and non-commonly preferred products of different consumer types. Figure 5 illustrates the efficiency as a function of the number of consumer types when the two consumer types 1 and nt at the edge of the differentiation spectrum are located at the positions $t_1 = 0.375$ and $t_{nt} = 0.625$. Hence, the consumer types are positioned more closely than in the configuration considered in Fig. 4. The efficiency curve shows the same trend and increases as the number of consumer types grows. However, the efficiency is on a higher level in general, so the collaborative filtering algorithm performs better than random product recommendations.

Distribution of consumers to consumer types (Scenario 3)

The distribution of consumers to consumer types describes the fractions of consumers that belong to the different consumer types. According to our model, the distribution of consumers

Table 2 Configurations to be investigated in simulation scenario 2

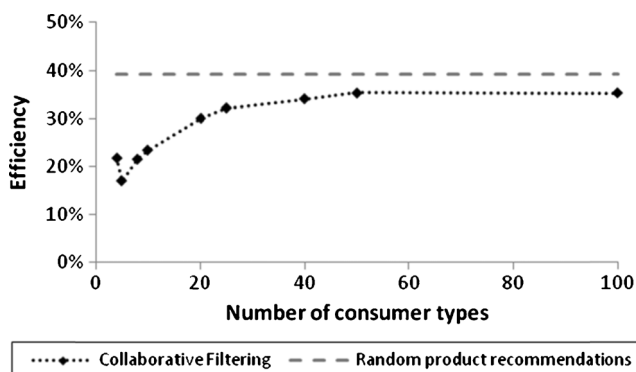
Number of consumer types nt	4	5	8	10	20	25	40	50	100
Number of consumers per type h_x	50	40	25	20	10	8	5	4	2

is represented by the vector $H = (h_1, \dots, h_m)$. In Scenario 1 and 2 we assumed an equal distribution of consumers to the consumer types. However, some existing markets such as the music industry are characterized by the *Long Tail* (Anderson 2006). That means that the majority of consumers belong to one consumer type that prefers blockbuster products. Each of the remaining consumer types prefers niche products and consists of a small fraction of consumers. Scenario 3 investigates how such a concentration of consumers affects the recommendation accuracy.

The scenario is based on $nt=2$ consumer types. Starting from an equal distribution of the consumers ($h_1 = h_2 = 100$), the consumers of type 2 are redistributed to type 1 stepwise (increasing consumer concentration). The redistribution is realized in steps of 5 % until consumer type 1 comprises 95 % of the consumers ($h_1 = 190$ and $h_2 = 10$). We additionally consider different intersections of the preference spectrums of the two consumer types (25 %, 50 %, and 75 %) to analyze dependencies between the distribution of consumers and the similarity of consumer types.

The three curves in Fig. 6 show that the efficiency of the recommender system is barely affected if consumer type 1 comprises between 50 % and 70 % of the consumers ($100 \leq h_1 \leq 140$). A variance analysis confirms that the distribution of consumers within this range is statistically insignificant independent of the intersection rate ($\alpha > 0.1$). The concentration of consumers causes a moderate effect if more than 70 % of the consumers belong to one consumer type ($h_1 > 140$). The recommendation accuracy increases with a growing concentration of consumers for intersection rates of 25 % and 50 %. In contrast, an intersection of 75 % leads to an opposed curve shape since the efficiency decreases as the concentration of consumers grows.

Our study shows that the distribution of consumers to consumer types affects the recommendation accuracy only for a

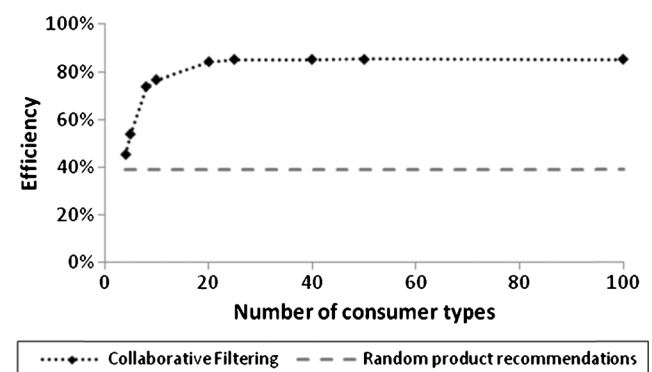
**Fig. 4** Efficiency as a function of the number of consumer types for $t_1 = 0.25$ and $t_m = 0.75$

high concentration of consumers. If the intersection rate is moderate (25 % and 50 %), an increasing concentration of consumers improves the efficiency since the probability that consumers are assigned to the correct consumer type grows. If the intersection rate is high (75 %), a higher correlation to the wrong consumer type is more likely as the concentration of consumers grows, so the number of unsuitable product recommendations increases. Similar to scenario 2, the overall level of efficiency strongly depends on the intersection rate. Therefore, the recommender system performs better than random product recommendations for intersection rates of 25 % and 75 % and worse for an intersection rate of 50 %.

Discussion

In contrast to our model based on Hotelling's linear city (Hotelling 1929), existing products typically differ in more than one characteristic, and consumers cannot be classified to specific consumer types precisely in real markets. Nevertheless, customer segmentation approaches in marketing show the practical relevance of the concept of consumer types (Wedel and Kamakura 2000). Similarities of consumer types can be determined based on the set of commonly preferred products also for multidimensional differentiated products. Compared to our round-based simulation, real buying processes are non-deterministic and the number and time of product purchases differ strongly across consumers. Furthermore, our algorithm that rates products in the amount of the consumer surplus also abstracts from reality.

However, our study follows a microeconomic research paradigm. We did not investigate existing markets, but instead addressed the fundamental question of whether the recommendation accuracy is affected by the structure of consumers' preferences. We aim at a theoretical analysis of the general effects of different features of the structure of consumers'

**Fig. 5** Efficiency as a function of the number of consumer types for $t_1 = 0.375$ and $t_m = 0.625$




Fig. 6 Efficiency as a function of the distribution of consumers

preferences. For that reason, we intentionally kept our model as simple as possible.

Our investigation of the three scenarios reveals that the structure of consumers' preferences significantly affects recommendation accuracy. Therefore, online retailers need to know the underlying consumer preferences. Our theoretical results motivate empirical studies in future research to judge the suitability of collaborative filtering recommender systems in existing markets.

Conclusion

Previous research has not examined the effect of the structure of consumers' preferences on the accuracy of recommender systems. This paper introduced a simulation model that enables the evaluation of recommender systems for different structures of consumers' preferences. The consumer preferences are modelled based on Hotelling's linear city model. We developed the *efficiency* as a model-specific metric to assess the accuracy of recommender systems. The efficiency defines which percentage of the consumer surplus achieved by a perfect recommendation is attained using the recommender system. We analyze how the structure of consumers' preferences affects the efficiency of a widely used collaborative filtering system by applying our simulation model. In detail, we investigate three features of the structure of consumers' preferences: the number and similarity of consumer types as well as the distribution of consumers to consumer types.

This is the first study to address the fundamental question of whether recommendation accuracy is affected by the structure of consumers' preferences. We answered the research question using a theoretical analysis and a microeconomic research paradigm. Our study found that the structure of consumers' preferences has a considerable effect on the accuracy of recommender systems. The *similarity of consumer types* has the largest effect: starting from distinct consumer types, the efficiency decreases significantly with growing similarity until a medium similarity is reached. Subsequently, the efficiency increases as the similarity grows further, so the efficiency is on a very high level for consumer types with nearly

identical preferences. The increasing *number of consumer types* generally leads to a growing efficiency. The *distribution of consumers to consumer types* has an impact on the recommendation accuracy only for a high concentration of consumers on a particular consumer type. This impact depends on the similarity of consumer types. The efficiency grows for low and medium similarity as the concentration of consumers on a particular consumer type increases. Contrarily, the efficiency decreases with increasing concentration if consumer types are very similar. In general, the dependency of the recommendation accuracy on the structure of consumers' preferences is explained by the impact of the features considered on the number of commonly and non-commonly preferred products and the number of irrelevant product recommendations (products not belonging to a consumer's preference spectrum).

Our study's findings are relevant for practice, especially for online retailers. Our investigation can help discover whether a collaborative filtering recommender system is suitable for a given market. The simulation revealed that even a random product recommendation outperforms a collaborative filtering system for certain structures of consumers' preferences. Consequently, online retailers need to know consumer preferences in specific markets to decide if a collaborative filtering algorithm or an alternative search system (e.g., another type of a recommender system, product configurator, advanced search engine) is preferable. Hence, our theoretical results provide a basis for future empirical studies on the suitability of collaborative filtering systems.

Moreover, our simulation model is relevant for future theoretical research. It can be used to analyze how the structure of consumers' preferences affects the accuracy of other types of recommender systems, e.g., item-based collaborative filtering (Adomavicius and Tuzhilin 2005; Jannach et al. 2011). Furthermore, it can be used to evaluate extensions and new features (e.g. introduction of consideration sets, adding randomly chosen products) of existing recommender systems. In this context, our simulation model allows a comprehensive evaluation varying the structure of consumers' preferences. Finally, since the product price influences the number of products within consumers' preference

spectrums and the similarity of consumers, our simulation model can be applied to investigate the optimal pricing strategy in a market using recommender systems.

References

- 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(6), 734–749. doi:10.1109/TKDE.2005.99 .
- Adomavicius, G., Bockstedt, J., Curley, S., & Zhang, J. (2011). Recommender Systems, Consumer Preferences, and Anchoring Effects. In *Proceedings of the RecSys 2011 Workshop on Human Decision Making in Recommender Systems* (pp. 35–42).
- Anderson, C. (2006). *The long tail: why the future of business is selling less of more* (1st ed.). New York: Hyperion.
- Bennett, J., & Lanning, S. (2007). The Netflix Prize. In *Proceedings of KDD Cup and Workshop* (pp. 3–6).
- Bergemann, D., & Ozmen, D. (2006). Optimal Pricing with Recommender Systems. In *EC '06, Proceedings of the 7th ACM Conference on Electronic Commerce* (pp. 43–51). New York: ACM. doi:10.1145/1134707.1134713 .
- Davidson, J., Liebal, B., Liu, J., Nandy, P., & van Vleet, T. (2010). The YouTube Video Recommendation System. In X. Amatriain, M. Torrens, P. Resnick, & M. Zanker (Eds.), *Proceedings of the Fourth ACM Conference on Recommender Systems* (pp. 293–296). doi:10.1145/1864708.1864770 .
- Desrosiers, C., & Karypis, G. (2011). A Comprehensive Survey of Neighborhood-based Recommendation Methods. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 107–144). Springer. doi:10.1007/978-0-387-85820-3_4 .
- Ekstrand, M. D. (2010). Collaborative Filtering Recommender Systems. *FNT in Human–Computer Interaction (Foundations and Trends in Human–Computer Interaction)*, 4(2), 81–173. doi:10.1561/11000000009 .
- Felfernig, A., Friedrich, G., & Schmidt-Thieme, L. (2007). Recommender Systems. *IEEE Intelligent Systems*, 22, 18–21. doi:10.1109/MIS.2007.52.
- Fleder, D., & Hosanagar, K. (2007). Recommender Systems and Their Impact on Sales Diversity. In *EC '07, Proceedings of the 8th ACM Conference on Electronic Commerce* (pp. 192–199). New York: ACM. doi:10.1145/1250910.1250939 .
- Gemmis, M. D., Iaquinta, L., Lops, P., Musto, C., Narducci, F., & Semeraro, G. (2011). Learning Preference Models in Recommender Systems. In J. Fürnkranz & E. Hüllermeier (Eds.), *Preference Learning* (1st ed., pp. 387–407). s.l.: Springer-Verlag. doi:10.1007/978-3-642-14125-6_18 .
- Herlocker, J. L., Konstan, J. A., Borchers, A., & Riedl, J. (1999). An Algorithmic Framework for Performing Collaborative Filtering. In *SIGIR '99, Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 230–237). New York: ACM. doi:10.1145/312624.312682.
- Herlocker, J. L., Konstan, J. A., Terveen, L. G., & Riedl, J. T. (2004). Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems (TOIS)*, 22(1), 5–53. doi:10.1145/963770.963772 .
- Hervas-Drane, A. (2015). Recommended For You: The Effect of Word of Mouth on Sales Concentration. *International Journal of Research in Marketing*, 32(2), 207–218. doi:10.1016/j.ijresmar.2015.02.005 .
- Hinz, O., & Eckert, J. (2010). The Impact of Search and Recommendation Systems on Sales in Electronic Commerce. *BISE (Business & Information Systems Engineering)*, 2(2), 67–77. doi:10.1007/s12599-010-0092-x .
- Hotelling, H. (1929). Stability in Competition. *The Economic Journal*, 39(153), 41–57.
- Jannach, D., Zanker, M., Felfernig, A., & Friedrich, G. (2011). *Recommender systems: an introduction*. Cambridge: Cambridge University Press.
- Jawaheer, G., Weller, P., & Kostkova, P. (2014). Modeling User Preferences in Recommender Systems: A Classification Framework for Explicit and Implicit User Feedback. *ACM Transactions on Interactive Intelligent Systems*, 4(2), 1–26. doi:10.1145/2512208 .
- Karatzoglou, A., & Weimer, M. (2011). Collaborative Preference Learning. In J. Fürnkranz & E. Hüllermeier (Eds.), *Preference Learning* (1st ed., pp. 409–427). s.l.: Springer-Verlag. doi:10.1007/978-3-642-14125-6_19.
- Levin, A. M., Levin, I. P., & Heath, C. E. (2003). Product Category Dependent Consumer Preferences for Online and Offline Shopping Features and Their Influence on Multi-Channel Retail Alliances. *Journal of Electronic Commerce Research*, 4(3), 85–93.
- Linden, G., Smith, B., & York, J. (2003). Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing*, 7(1), 76–80. doi:10.1109/MIC.2003.1167344 .
- Lops, P., Gemmis, M. d., & Semeraro, G. (2011). Content-based Recommender Systems: State of the Art and Trends. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 73–105). Springer.
- Miller, B. N., Albert, I., Lam, S. K., Konstan, J. A., & Riedl, J. (2003). MovieLens Unplugged: Experiences with an Occasionally Connected Recommender System. In D. Leake, L. Johnson, & E. Andre (Eds.), *Proceedings of the 8th International Conference on Intelligent User Interfaces* (pp. 263–266). doi:10.1145/604045.604094 .
- Pfeiffer, J., & Scholz, M. (2013). A Low-Effort Recommendation System with High Accuracy. *BISE (Business & Information Systems Engineering)*, 5(6), 397–408. doi:10.1007/s12599-013-0295-z .
- Rashid, A. M., Albert, I., Cosley, D., Lam, S. K., McNee, S. M., Konstan, J. A., & Riedl, J. (2002). Getting to Know You: Learning New User Preferences in Recommender Systems. In *IUI '02, Proceedings of the 7th International Conference on Intelligent User Interfaces* (pp. 127–134). New York, NY, USA: ACM. doi:10.1145/502716.502737 .
- Resnick, P., & Varian, H. R. (1997). Recommender Systems. *Communications of the ACM*, 40(3), 56–58. doi:10.1145/245108.245121.
- Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to Recommender Systems Handbook. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 1–35). Springer.
- Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based Collaborative Filtering Recommendation Algorithms. In *WWW '01, Proceedings of the 10th International Conference on World Wide Web* (pp. 285–295). New York: ACM. doi:10.1145/371920.372071.
- Schafer, J. B., Konstan, J., & Riedl, J. (1999). Recommender Systems in E-Commerce. In *EC '99, Proceedings of the 1st ACM Conference on Electronic Commerce* (pp. 158–166). New York: ACM. doi:10.1145/336992.337035 .
- Scholz, M., Dorner, V., Franz, M., & Hinz, O. (2015). Measuring consumers' willingness to pay with utility-based recommendation systems. *Decision Support Systems*, 72, 60–71. doi:10.1016/j.dss.2015.02.006 .
- Su, X., & Khoshgoftaar, T. M. (2009). A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence*, 4. doi:10.1155/2009/421,425 .

- Tirole, J. (1988). *The theory of industrial organization*. Cambridge: MIT Press.
- Wedel, M., & Kamakura, W. A. (2000). *Market Segmentation: Conceptual and Methodological Foundations* (2nd ed.). *International series in quantitative marketing: Vol. 7*. Boston: Kluwer Academic.
- Wu, L.-L., Joung, Y.-J., & Chiang, T.-E. (2011). Recommendation Systems and Sales Concentration: The Moderating Effects of Consumers' Product Awareness and Acceptance to Recommendations. In: *HICSS '11, Proceedings of the 2011 44th Hawaii International Conference on System Sciences* (pp. 1–10). Washington, DC: IEEE Computer Society. doi:[10.1109/HICSS.2011.357](https://doi.org/10.1109/HICSS.2011.357) .
- Zhou, T., Kuscsik, Z., Liu, J.-G., Medo, M., Wakeling, J. R., & Zhang, Y.-C. (2010). Solving the Apparent Diversity-Accuracy Dilemma of Recommender Systems. *Proceedings of the National Academy of Sciences*, *107*(10), 4511–4515. doi:[10.1073/pnas.1000488107](https://doi.org/10.1073/pnas.1000488107) .