




Contents lists available at ScienceDirect

Decision Support Systems

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

Fused latent models for assessing product return propensity in online commerce

Yanjie Fu^a, Guannan Liu^{b,*}, Spiros Papadimitriou^c, Hui Xiong^c, Xiaolin Li^{e,*}, Guoqing Chen^d

^aDepartment of Computer Science, Missouri University of Science and Technology, Rolla, MO 65409, USA

^bDepartment of Information Systems, School of Economics and Management, Beihang University, Beijing 100191, China

^cRutgers Business School, Rutgers University, NJ, USA

^dDepartment of Management Science and Engineering, School of Economics and Management, Tsinghua University, Beijing 100084, China

^eSchool of Management, Nanjing University, Nanjing, Jiangsu 210093, China

ARTICLE INFO

Article history:

Received 6 July 2015

Received in revised form 6 July 2016

Accepted 2 August 2016

Available online xxx

Keywords:

Product return

Online shopping

Propensity assessment

ABSTRACT

In online shopping, product returns are very common. Liberal return policies have been widely used to attract shoppers. However, returns often increase expense and inconvenience for all parties involved: customers, retailers, and manufacturers. Despite the large fraction of purchases that are returned, there are few systematic studies to explain the underlying forces that drive return requests, and to assess the return propensity at the level of individual purchases (i.e., a particular customer purchasing a particular product), rather than in aggregate. To this end, in this paper, we provide a systematic framework for personalized predictions of return propensity. These predictions can help retailers enhance inventory management, improve customer relationships, and reduce return fraud and abuse. Specifically, we treat product returns as a result of inconsistency arising during a commercial transaction. We decompose this inconsistency into two components, one for the buying phase (e.g., product does not match description) and another for the shipping phase (e.g., product damaged during shipping). Along these lines, we introduce a generalized return propensity latent model (RPLM). We further propose a complete framework, called fused return propensity latent model (FRPLM), to jointly model the correlation among user profiles, product features, and return propensity. We present comprehensive experimental results with real-world data to demonstrate the effectiveness of the proposed method for assessing return propensity.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Online shopping is convenient for consumers, with fast product search, quick payment, and timely delivery to their doorstep. As a result, it has become part of most consumers' lifestyle, leading to an explosive growth of e-commerce. According to the report from iResearch center¹, the e-commerce market of China has reached 2.5 trillion US dollars, which is an increase of 21.2% compared to one year prior.

In such a huge market, the high rate of product returns has become a nontrivial issue in online shopping. According to the report from Fits.me [1], the return rates for online sales are much higher compared to "offline" sales, in almost every business sector. The Wall Street Journal reported that up to a third of all Internet purchases

are returned by customers [2]. Unlike customers of brick-and-mortar retailers, online shoppers cannot touch and feel a product, to determine how well it fits their tastes and needs [3]. Instead, they have to rely on textual descriptions, photos, and, in general, representations that may not be sufficiently accurate or media-rich. Therefore, it is more difficult for customers to make accurate decisions [4], with potentially higher risk of dissatisfaction when they receive the product. If this is the case, most consumers may choose to return the product, according to the merchant's return policy.

Although liberal return policies can motivate consumers to purchase from online retailers and can help build customer loyalty, without proper management and control they can lead to substantial losses for consumers, retailers, and manufacturers. For example, retailers need to collect and ship the unwanted products back to the manufacturers, in order to be refurbished for resale, or sold to a third party for residual value, or disposed along with damaged products. Moreover, consumers typically have to undertake the cost of shipping, and incur delay in receiving their desired product. All of these are potential reasons for customer churn. Therefore, assessing

* Corresponding authors.

E-mail addresses: liugn@buaa.edu.cn (G. Liu), lixl@nju.edu.cn (X. Li).

¹ <http://www.iresearchchina.com>

the propensity of product returns with respect to each customer's purchase has potentially very high value to online retailers. Specifically, if retailers can assess potential product returns in advance, they can better manage customer relationships, reduce return fraud and abuse, as well as enhance product and inventory management.

In prior literature, researchers have empirically identified the factors that can impact product returns in online channels [5–7], including web technologies [6] and consumers' reviews [7]. Previous research also studies product return policies in influencing customer purchasing behaviors [8–13]. Additionally, other studies [5,14,15] aim at forecasting product returns at the product category level, in order to optimize inventory control. However, all prior work has rarely addressed the issue of assessing the personalized return propensity for each individual purchase. There is a lack of studies to systematically understand the key driving forces driving consumers' product returns based on historical data, and to predict the potential returns.

To this end, in this paper, we provide a data-driven approach to assess the return propensity for each individual purchase initiated by a consumer. Our specific goal is, for each customer–product pair, to predict the propensity of the customer to return that product, if ordered. It is tempting to approach this as a “recommendation” problem, where “recommendation” refers to the technologies and methods that allow online merchants to recommend a list of products by predicting how a consumer would rate each product. To achieve this goal, recommendation algorithms construct the customer–product rating matrix, where each entry represents a specific user's rating for a product. We could similarly construct a customer–product “return” matrix, where each entry is instead a return frequency of a customer–product pair, and further predict return propensity based on the matrix.

However, although this view leads to an elegant and simple formulation, unfortunately standard factorization-based recommendation solutions fail to provide good results in this setting. First, in addition to a customer's tastes and perception of a product, the likelihood of a return can be affected by the shipping process. Second, the return propensity cannot be measured uniformly across all customers and products, but also depends on both consumer profiles (e.g., gender, salary, and location) and product characteristics (e.g., complexity or vulnerability to shipping risks). Finally, in addition to the sparsity problem in recommendation (i.e., a given customer will have ordered only a tiny fraction of the retailer's inventory), the return propensity distribution is heavily skewed (see Fig. 1): a few products are re-ordered one or more times, some are returned and never re-ordered, but most products are never returned. Besides being skewed, the return propensity distribution is also multimodal. Therefore, the product–customer return propensity matrix has a large number of zeros, which are distinct from missing (i.e., unobserved) values. Next, we outline the intuition behind the technical development of our contributions, to deal with these challenges.

Generally, a customer requests to return a product because of mismatch, or *inconsistency*, between the expected and the received product. This can arise at any point of the online shopping process, in which two main phases have been identified: (1) the buying phase, in which consumers search, order, and pay; (2) the shipping phase, in which retailers and carriers pack, ship, and deliver [8]. As most returns are due to product problems in the buying phase and shipping problems in the shipping phase, we decompose the inconsistency into two types, corresponding to these two phases.

In the buying phase, consumers search and browse products that might meet their needs, forming a cognitive expectation, which we call the “expected product.” This will, hopefully, match the actual product but, due to potential inaccuracy or misunderstanding of online product descriptions, inconsistencies may arise. In this situation, a return is likely to be requested. Furthermore, some customers




Fig. 1. The product return propensity shows a power law distribution, i.e., return propensity of most customer–product pairs is zero while only a small proportion has positive return propensity.

may receive products with quality or design deficiencies, compared to what they expected. Finally, the shipping phase entails various risks, such as delayed delivery, and deficient, damaged, or wrong received products, again leading to inconsistency between actual and received products.

To capture these two types of inconsistency, we propose a generalized return propensity latent model (RPLM), which introduces latent vectors that represent: (1) consumers' cognitive preferences towards products (“expected product”); (2) actual product shipped, over the same latent space; and (3) received product. Interactions between the first two latent vectors capture the first type of inconsistency, whereas interactions between the last two capture the second type. We combine the two consistencies to estimate return propensity through a linear model.

Moreover, return propensity also depends on consumer profiles and product characteristics. To enhance the modeling of return propensity, we fuse both consumer and product features into the generalized RPLM model. Specifically, as consumers and products can be grouped into different clusters in terms of their features, we introduce two latent variables, “customer segment” and “product category”, to represent the clusters of customers and of products, respectively, in the feature space. Customers in the same customer segment not only share similar customer profiles, but also share similar latent vectors of “expected product”; products in the same product category not only share similar product characteristics, but also similar latent vectors of “actual product” and “received product”. In this way, we relate consumer profiles and product characteristics to the two aforementioned types of inconsistency.

To summarize, in this paper, we systematically investigate product return records from online shopping under a data-driven framework, and propose a predictive model to assess the return propensity for each consumer–product pair in online purchases. Specifically, we first introduce the RPLM framework, and subsequently extend it to the complete FRPLM framework. Finally, we conduct a comprehensive evaluation on a real-world dataset obtained from an e-commerce company. The experimental results demonstrate the effectiveness of the proposed method for assessing product return propensity in online shopping.

2. Related work

There are several research areas that are generally relevant to our research focus. These include (1) product returns from a management perspective, (2) recommender systems, and (3) latent factor models in various business applications.

2.1. Product return management and prediction

Researchers in marketing and operations management have studied product returns from the perspective of both online retailers and manufacturers. Pasternack et al. [16] examined the strategy of pricing and return policies and developed a hierarchical model for the pricing decision faced by a producer of a commodity with a short shelf or demand life. The work in [17] suggests that online retailers should either institute a policy of free product returns or examine their customer profiles to determine their customers' responses to free returns. Mukhopadhyay et al. [18] found that a policy of modularization as well as offering generous return terms would increase revenue, but also increase the cost due to higher likelihood of return and higher cost of design. They also developed a profit maximization model to jointly obtain optimal policies for return policy and modularity level, in terms of certain market reaction parameters. The work in [19] develops several theoretical models to examine the impact of online distributor's return policy, product quality, and pricing strategy on the customer's purchase and return decisions, finding that decisions about the return policy are mutual and complementary with product quality and pricing strategies. Davis et al. [20] found that a retailer is more likely to offer a low-hassle return policy when: 1) its products' benefits cannot be consumed in a short period of time; 2) its product line offers opportunities for cross-selling; and 3) it can obtain a high salvage value for returned merchandise. The work in [21] develops a structured model to identify the optimal return policy, so they can trade off between the return cost and the sales demand and finds considerable variation in the value of returns across customers and categories. Petersen et al. [8] demonstrated the role of product returns in the three-step exchange process, and showed the consequences of product returns on future behaviors of consumers and firms. The work in [7] evaluates the impact of online product reviews on product returns, by fusing multi-source information including product characteristics, product reviews, customer characteristics, customer activities, and so on. De et al. [6] empirically showed that web technologies applied in online platforms correlate to product returns. Furthermore, several models have been devised to predict potential product returns. The work in [22] proposes a split adjusted hazard model to predict the return rate, using the information of price and the time after purchase. The work in [5] forecasts the propensity of product return at the product category level.

Our work is different from prior research on designing optimal return policy and predicting product returns. Prior literature of predicting product returns mainly studies the problems of predicting the return probability of a single product type (e.g., NIVEA Smooth Sensation Body Lotion) or a single product category (e.g., body lotion), while we estimate the personalized return propensity of a customer for a product at individual level. Moreover, we model the return propensity as a result of the inconsistency arising in the buyer–retailer exchange process and provide a new perspective in understanding the reasons (problems in shopping or shipping phases) underlying product returns.

2.2. Recommender systems and matrix factorization

Recommender systems are a subclass of information filtering systems that predict the rating or preference that a user would give to an item. Matrix factorization (MF) has recently become a widely used approach in recommender systems. In these models, the rating matrix is decomposed as latent structures over users and items, and the loss between the prediction and real rating is minimized according to a pre-defined distance measure.

Variations of MF adapt the key modeling idea to different scenarios. For example, SVD++ [23] decomposes a user–item rating matrix

into matrices that represents latent user–user features and item–item features, but also combines the ideas of user-based or item-based neighborhood models. Non-negative Matrix Factorization (NMF) [24] adds a non-negative regularization into MF, where the values of user and item latent features are non-negative. Probabilistic Matrix Factorization (PMF) [25] exploits a sparsity regularization, where the values of user and item latent features are drawn from zero mean Gaussian distributions. Bayesian Probabilistic Matrix Factorization (BPMF) [26] extends PMF and employs conjugate prior distributions to regularize the means and variants of the Gaussian distributions on user and item latent features. Bayesian Non-negative Matrix Factorization [27] is a probabilistic version of NMF, where the non-negative constraints are implemented by non-negative probabilistic distributions, such as exponential distributions.

One drawback of the MF-style approach is the so-called ramp-up problem, in which MF performs poorly when the recommender system is not initialized with a sufficient number of ratings. Therefore, more studies [28–31] incorporated rich side information, referred as the attribute information of both products and users, for example, the gender, credit of users, etc., into latent factor models, by extracting observable features of users and items, or even textual data. For example, Ma et al. [32] propose a matrix factorization model with social relations as regularization. Bao et al. [33] fuse the temporal and social information in PMF model to predict users' interests in microblog. Ge et al. [34] incorporate cost and time budget information in the PMF model to recommend cost-effective travel tours for consumers.

2.3. Latent factor models in business applications

Beyond recommender systems, researchers have applied latent factor models (LFM) to various business applications, such as price prediction, risk and financial management [35–37]. The models generally assumed latent structures over observed data, and make predictions based on the latent space. Creal et al. [35] proposed a dynamic latent factor model to predict the credit risks from time series financial data. The study of [37] predicts housing prices with mobility data incorporated in the framework of LFM. Latent factor models have also been used to model travel data, with the latent factors interpreted as cost or distance [36].

Our work is distinguished from the classic techniques of matrix factorization and recommender systems, in which the return propensity of a customer for a product is factorized by a customer vector and a product vector. We adapt the idea of collaborative matrix factorization, introduce the vectors of expected product, shipped product, and received product, and model the inconsistencies in both shopping and shipping stages. More importantly, we incorporate two latent clustering processes (i.e., customer clustering and product clustering) into the modeling, and fuse customer profiles and product characteristics to enhance the collaborative matrix factorization model.

3. Preliminaries

In this section, we first formally state the problem of assessing product return propensity, and then we outline the baseline method of probabilistic matrix factorization [38].

3.1. Problem statement

Our objective is to assess the return propensity for each individual online purchase using historical sales and return data. We formulate the problem as predicting the possibility of product return, given a customer and a purchased product. More formally, given a set of customers \mathcal{I} , a set of products \mathcal{J} in an online shopping website, and historical product order and return records, we aim to assess the

return probability $y_{ij} \in Y$ of any customer $i \in I$ buying any product $j \in J$. Some important mathematical notations are listed in Table 1.

3.2. Probabilistic matrix factorization

We aim to assessing return propensity of any customer for any product. This is similar to the task of predicting a user's personalized rating for an item in recommender systems. Therefore, we outline probabilistic matrix factorization (PMF), which is a widely used latent factor method and serves as the baseline predictive model. Consider the simplest case, in which tuples (i, j, y_{ij}) are given, where $y_{ij} \in Y$ is the return propensity of customer $i \in \mathcal{I}$ for product $j \in \mathcal{J}$. The matrix of return propensities Y can be factorized into u_i and v_j , which represent the latent vector of the customer i and the latent vector of the product j , respectively. Furthermore, the inner product of vectors u_i and v_j encodes the interaction between customer and product, and can be interpreted as the inconsistency perceived by customer i between the expected product and the received product. The missing elements of Y can be completed using the inferred latent vectors.

4. Generalized return propensity latent model framework

In this section, we propose our generalized return propensity latent model (RPLM), which characterizes the inconsistencies that may arise during the customer–seller exchange process. Our complete FPRLM framework is then presented in the next section.

4.1. Model intuition

Different from “offline” shopping, inconsistencies are more likely to arise between consumers' expected product and the received product, due to the longer and more indirect chain of product exchange. We generally assume that online retailers are genuine and product descriptions on their website are not intentionally inaccurate or misleading. Therefore, a customer would request a product return only when there is inconsistency between expected and received product. According to the study in [8], online shopping has two key phases: (1) the buying phase, and (2) the shipping phase. We therefore decompose the overall inconsistency into two components: (1) the inconsistency in the buying phase and (2) the inconsistency in the shipping phase. Specifically, we model the product return propensity based on the following intuitions.

Intuition 1: In the buying phase, there exists inconsistency between consumers' expected product and the actual product shipped. Consumers may form cognitive bias while browsing the product only through the information displayed on the shopping

sites and issues that may arise exclusively from that (e.g., constraints such as color differences on the computer screen, misjudged size or dimensions, and so on). Thus, the product shipped by the retailers may diverge from what the consumers originally expect.

Intuition 2: In the shipping phase, there exists inconsistency between the shipped product and the received product. The products bought online are mostly shipped by third-party delivery companies and there exist potential risks for damage, delay, or other uncontrolled factors during shipping. Thus, the consumers may receive a product with deficiencies, which differs from the originally shipped product.

Intuition 3: The level of inconsistency between the expected and the received product can determine the product return propensity.

4.2. Model description

Based on the above intuitions, we decompose the inconsistency between expected and received products into two types: (1) Type-I inconsistency (α): the inconsistency between the product as expected by customers and the shipped product, and (2) Type-II inconsistency (β): the inconsistency between the shipped product and received product.

In the buying phase, customer i may browse the retailer's description of product j , and form an expectation for this product, which is represented by a latent vector e_i . The elements of e_i are drawn from the probability distribution $P(e_i|\Psi_e)$. According to *Intuition 1*, the actual product shipped may be inconsistent with consumers' expectations. We thus introduce another latent vector s_j , which can be interpreted as the shipped product j , mapped into the same latent space. The elements of s_j are drawn from the probability distribution $P(s_j|\Psi_s)$. We regard the dot product of e_i and s_j as the inconsistency between the expected and shipped product, i.e.,

$$\alpha_{ij} = e_i^T s_j \tag{1}$$

According to *Intuition 2*, inconsistency can also arise between the shipped and received product. We introduce a latent vector r_j to denote the latent representation of the received product j . The elements of this vector are drawn from the probability distribution $P(r_j|\Psi_r)$. We regard the dot product of s_j and r_j as the inconsistency between the shipped and received product, and thus,

$$\beta_j = s_j^T r_j \tag{2}$$

Finally, the return propensity y_{ij} of customer i when purchasing product j is determined by the overall level of inconsistency, following *Intuition 3*. We combine the two types of inconsistency as $\hat{y}_{ij} \propto \alpha_{ij} + \beta_j$, and the propensity is drawn from the probability distribution $P(Y|\hat{Y})$. Here, the range α and β can be arbitrary, and depends on the probability functions used. Summing up, the decomposition of the product return propensity matrix is shown in Fig. 2, in which three lower-rank latent factors are introduced to represent different states/facets of products during the shopping phases.

5. Fused return propensity latent model

In the previous RPLM, the inconsistency is generally captured by the interactions between the latent representation of expected product, shipped product, and received product. However, the inconsistency between the expected and received product cannot be measured in a uniformly standard way, across all customers and products. Different consumers may perceive things differently,

Table 1
Mathematical notations.

Symbol	Description
i	Customer $i \in \{1, \dots, I\}$
j	Product $j \in \{1, \dots, J\}$
y_{ij}	Return propensity of customer i for product j
\hat{y}_{ij}	Predicted return propensity of customer i for product j
p_i	Customer profile information, can be one or multiple elements
f_j	Product features, can be one or multiple elements
g	Customer segment, $g \in 1, \dots, M$
c	Product category, $c \in 1, \dots, N$
e_i	Customer expected product vector, $e_i \in \mathbb{R}^K$
s_j	Shipped product vector, $s_j \in \mathbb{R}^K$
r_j	Received product vector, $r_j \in \mathbb{R}^K$
K	Dimension of the latent space

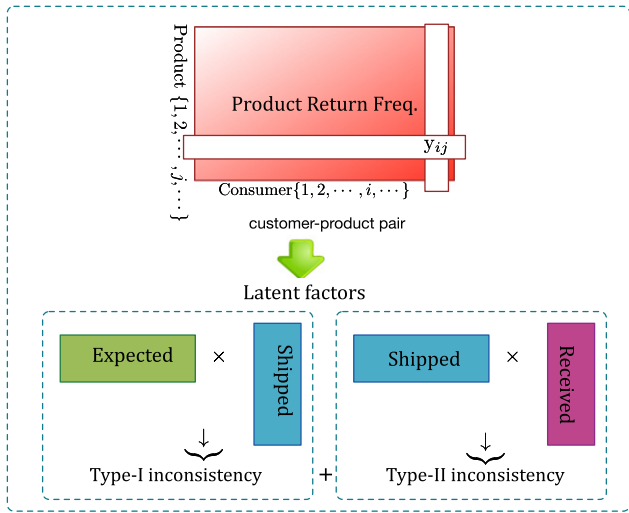


Fig. 2. The framework for decomposing the product returns.

leading to variations in the inconsistency they may feel. Different products may also vary in terms of consumers' cognitive bias and of shipping risks.

Fig. 3 shows that return ratios and counts differ substantially with respect to consumer profiles and to product features. For example, females are more prone to return than males; customers with low credit scores return more frequently than those with high scores; and products without warranty or invoice are more likely to be returned. Therefore, we propose to incorporate the information of consumer profiles and product features, to jointly capture their influence on the return propensity, in order to improve predictive accuracy. This leads to enhanced assessment of propensity. We therefore propose the Fused Return Propensity Latent Model (FRPLM), which is an enhanced specification of the generalized RPLM.

5.1. Model specification

Each product return event corresponds to a user-product pair, and thereby this event is associated with particular consumer profiles and product features, which are correlated with the return propensity. First, we introduce a latent categorical variable called *customer segment* g to represent the clustering of customers in feature space. Each *customer segment* generates the profile features, including gender, location, and historical credit scores, as well as the expected product latent vector, which can be regarded as consumers' perception of products. Similarly, we also introduce a latent categorical variable *product category* c to represent the clustering of products in feature space. This variable generates the product features, including price, discount status, and warranty status, as well as the shipped product and received product latent vectors. By introducing these two additional latent variables, *customer segment* and *product category*, all the side information that can influence a consumer's return propensity for a purchased product is fused into the latent factor model.

More specifically, we denote each customer-product pair with related features by a tuple $\{i, j, \mathbf{p}_i, \mathbf{e}_i, \mathbf{f}_j, \mathbf{s}_j, \mathbf{r}_j, y_{ij}\}$. We first draw the latent customer segment g from the multinomial distribution η , for each customer i . Based on the segment g we draw: (1) the extracted customer profiles \mathbf{p}_i following the multinomial distribution ϵ_g ; (2) the consumer expected product latent vector following a multivariate normal distribution $\mathcal{N}(\mathbf{e}_i | \boldsymbol{\mu}_g, \boldsymbol{\sigma}_g)$.

On the other hand, for each product j , we sample the latent product category c from the multinomial distribution ϵ . Similarly,

according the specific product category c , we draw: (1) product features \mathbf{f}_j from the multinomial distribution ρ_c (2) shipped product latent vector \mathbf{s}_j following the multivariate normal distribution $\mathcal{N}(\mathbf{s}_j | \boldsymbol{\chi}_c, \boldsymbol{\pi}_c)$; and (3) the received product latent vector \mathbf{r}_j following the multivariate normal distribution $\mathcal{N}(\mathbf{r}_j | \boldsymbol{\chi}_c, \boldsymbol{\omega}_c)$.

Finally, the return propensity y_{ij} of a customer i for a product j is $y_{ij} \propto \mathbf{e}_i^T \mathbf{s}_j + \mathbf{s}_j^T \mathbf{r}_j$. We treat the return events of the customer-product pair (i, j) as observations sampled with a certain probability, i.e., return propensity, y_{ij} .

5.2. Model inference

For the sake of simplicity, we assume that the covariances of multivariate normal distributions (i.e., $\boldsymbol{\sigma}, \boldsymbol{\pi}, \boldsymbol{\omega}$) are diagonal, thus the elements of the consumer expected product vector, the shipped product vector, and the received product vector are independent of each other. Also, we only show the model inference for one customer profile element $p_{i,m}$, and one product feature element $f_{j,m}$, since the inference process for the other elements is the same.

Let us denote all the parameters by $\Psi = \{\eta, \epsilon, \boldsymbol{\mu}, \boldsymbol{\sigma}, \mathbf{e}, \theta, \rho, \boldsymbol{\lambda}, \boldsymbol{\pi}, \boldsymbol{\chi}, \boldsymbol{\omega}, \mathbf{s}, \mathbf{r}, \}$ where I, J, M, N are the numbers of consumers, products, latent customer segments, and latent product categories, respectively, $\boldsymbol{\mu} = \{\mu\}_1^M$, $\boldsymbol{\sigma} = \{\sigma\}_1^M$, $\mathbf{e} = \{e_i\}_1^I$, $\rho = \{\rho\}_1^N$, $\boldsymbol{\lambda} = \{\lambda\}_1^N$, $\boldsymbol{\pi} = \{\pi\}_1^N$, $\boldsymbol{\chi} = \{\chi\}_1^N$, $\boldsymbol{\omega} = \{\omega\}_1^N$, $\mathbf{s} = \{s_j\}_1^J$, $\mathbf{r} = \{r_j\}_1^J$, the latent assignments (i.e., missing data) $\Upsilon = \{G, C\}$ where $G = \{g_i\}_1^I$, $C = \{c_j\}_1^J$, and the observed data collection $\mathcal{D} = \{Y, P, F\}$ where $Y = \{y_{ij}\}_{1,1}^I, P = \{p_i\}_1^I$, $F = \{f_j\}_1^J$.

The joint distribution of the model is

$$P(\mathcal{D}, \Upsilon, \Psi) = \prod_{j=1}^J P(\mathbf{s}_j | \boldsymbol{\lambda}_c, \boldsymbol{\pi}_c) P(\mathbf{r}_j | \boldsymbol{\chi}_c, \boldsymbol{\omega}_c) P(\mathbf{f}_j | \rho_c) P(c_j | \theta) \times \prod_{i=1}^I P(\mathbf{e}_i | \boldsymbol{\mu}_g, \boldsymbol{\sigma}_g) P(p_i | \epsilon_g) P(g_i | \eta) \times \prod_{i=1}^I \prod_{j=1}^J P(y_{ij} | \hat{y}_{ij}, \kappa) \quad (3)$$

where κ is the variance of the Gaussian distribution of y_{ij} . With the formulated complete likelihood, the objective is to find the parameters Ψ that maximize the likelihood. We use Expectation Maximization (EM) for that purpose.

E-Step: In the E-step, we first compute the posteriors of the latent variables. Specifically, we have two latent variables, the customer segment g and product category c , which are mutually independent. For each customer i , the latent customer segment is drawn from $g_i \sim P(g_i | \mathcal{D}, \Upsilon^n, \Psi^n)$, where Υ^n and Ψ^n respectively denote the latent assignments and the estimated parameters in the last iteration. By Bayesian inference, we have:

$$P(g_i | \mathcal{D}, \Upsilon^n, \Psi^n) \propto P(\mathbf{e}_i | \boldsymbol{\mu}_g, \boldsymbol{\sigma}_g) P(p_i | \epsilon_g) P(g_i | \eta) \prod_{j=1}^J P(y_{ij} | \hat{y}_{ij}, \kappa) \times \prod_{j=1}^J P(\mathbf{s}_j | \boldsymbol{\lambda}_c, \boldsymbol{\pi}_c) P(\mathbf{r}_j | \boldsymbol{\chi}_c, \boldsymbol{\omega}_c) P(\mathbf{f}_j | \rho_c) P(c_j | \theta) \quad (4)$$

Similarly, for each product j , the associated product category c_j is drawn from $c \sim P(c | \mathcal{D}, \Upsilon^n, \Psi^n)$. By Bayesian inference, we have:

$$P(c | \mathcal{D}, \Upsilon^n, \Psi^n) \propto P(\mathbf{s}_j | \boldsymbol{\lambda}_c, \boldsymbol{\pi}_c) P(\mathbf{r}_j | \boldsymbol{\chi}_c, \boldsymbol{\omega}_c) P(\mathbf{f}_j | \rho_c) P(c_j | \theta) \times \prod_{i=1}^I P(y_{ij} | \hat{y}_{ij}, \kappa) \prod_{i=1}^I P(\mathbf{e}_i | \boldsymbol{\mu}_g, \boldsymbol{\sigma}_g) P(p_i | \epsilon_g) P(g_i | \eta) \quad (5)$$

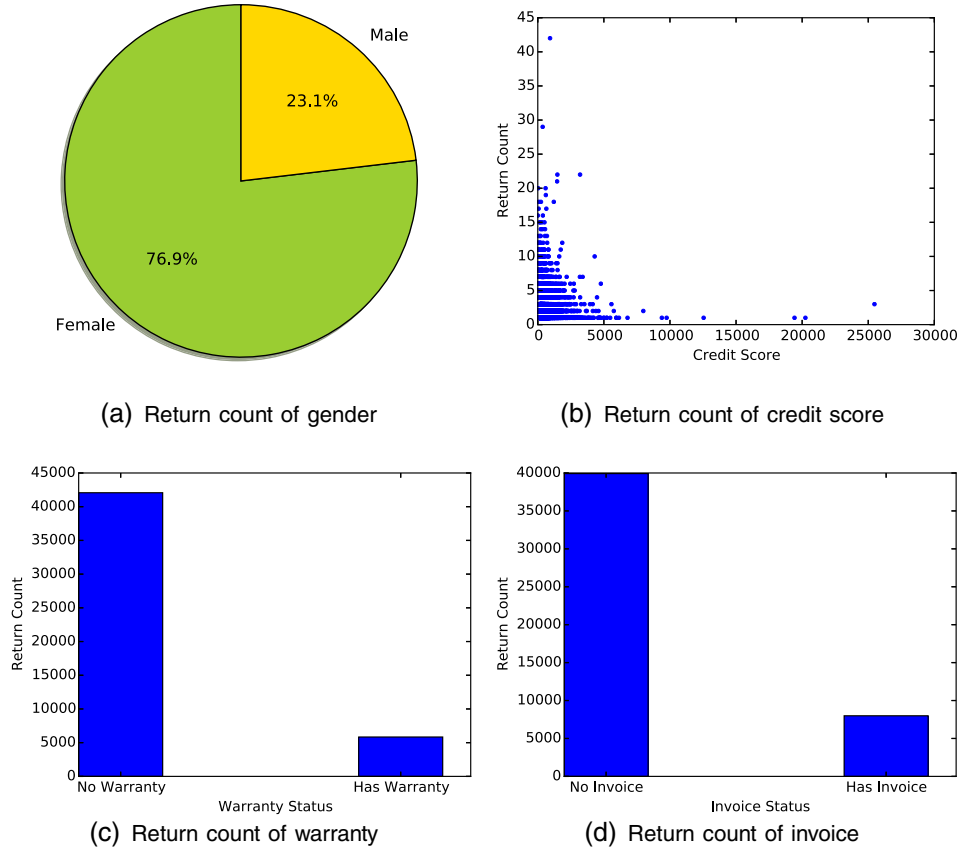


Fig. 3. Correlation analysis between return frequency and customer/product features. (a) Women return more; (b) customers with low credit scores return more; (c, d) products without warranty or invoice are more likely to be returned.

With the posteriors of the latent variables, we can obtain the conditional expectation as

$$Q(\Psi|\Psi^n) = E(P(D, \Upsilon|\Psi)|D, \Psi^n) = \prod_{j=1}^J \prod_{i=1}^I P(y_{ij}|\hat{y}_{ij}, \kappa) \times \prod_{j=1}^J \sum_{c_j=1}^N P(\mathbf{s}_j|\lambda_{c_j=1}^N \boldsymbol{\pi}_{c_j}) P(\mathbf{r}_j|\mathcal{X}_{c_j}, \boldsymbol{\omega}_{c_j}) P(f_j|\rho_{c_j}) P(c_j|\mathcal{D}, \Upsilon^n, \Psi^n) \times \prod_{i=1}^I \sum_{g_i=1}^M P(\mathbf{e}_i|\boldsymbol{\mu}_{g_i=1}^M, \boldsymbol{\sigma}_{g_i}) P(p_i|\epsilon_{g_i}) P(g_i|\mathcal{D}, \Upsilon^n, \Psi^n) \quad (6)$$

M-Step: In the E-step, we have derived the conditional expectation Q . By applying Jensen's inequality, the logarithm of Q can be approximated by a lower bound \mathbb{Q} , which is

$$\ln Q(\Psi|\Psi^n) \geq \mathbb{Q} = \sum_{j=1}^J \sum_{i=1}^I \ln P(y_{ij}|\hat{y}_{ij}, \kappa) + \sum_{j=1}^J \sum_{c_j=1}^N P(c_j|\mathcal{D}, \Upsilon^n, \Psi^n) \times [\ln P(\mathbf{s}_j|\lambda_{c_j}, \boldsymbol{\pi}_{c_j}) + \ln P(\mathbf{r}_j|\mathcal{X}_{c_j}, \boldsymbol{\omega}_{c_j}) + \ln P(f_j|\rho_{c_j})] + \sum_{i=1}^I \sum_{g_i=1}^M P(g_i|\mathcal{D}, \Upsilon^n, \Psi^n) [\ln P(\mathbf{e}_i|\boldsymbol{\mu}_{g_i}, \boldsymbol{\sigma}_{g_i}) + \ln P(p_i|\epsilon_{g_i})] \quad (7)$$

Assuming that the dimensions of all latent vectors are K , we can rewrite \mathbb{Q} by expanding the corresponding probability density functions (e.g., Gaussian or multinomial distribution).

$$\mathbb{Q} = \sum_{j=1}^J \sum_{c_j=1}^N \left(-\frac{1}{2} \ln \kappa^2 - \frac{[y_{ij} - \hat{y}_{ij}]^2}{2\kappa^2} \right) + \sum_{j=1}^J \sum_{c_j=1}^N \sum_{k=1}^K \left(-\frac{1}{2} \ln \pi_{c_j,k}^2 - \frac{(\mathbf{s}_{jk} - \lambda_{c_j,k})^2}{2\pi_{c_j,k}^2} \right) P(c_j|\mathcal{D}, \Upsilon^n, \Psi^n) + \sum_{j=1}^J \sum_{c_j=1}^N \sum_{k=1}^K \left(-\frac{1}{2} \ln \omega_{c_j,k}^2 - \frac{(\mathbf{r}_{jk} - \chi_{c_j,k})^2}{2\omega_{c_j,k}^2} \right) P(c_j|\mathcal{D}, \Upsilon^n, \Psi^n) + \sum_{j=1}^J \sum_{c_j=1}^N \ln \rho_{c_j} P(c_j|\mathcal{D}, \Upsilon^n, \Psi^n) + \sum_{i=1}^I \sum_{g_i=1}^M \sum_{k=1}^K \left(-\frac{1}{2} \ln \sigma_{g_i,k}^2 - \frac{(\mathbf{e}_{ik} - \boldsymbol{\mu}_{g_i,k})^2}{2\sigma_{g_i,k}^2} \right) P(g_i|\mathcal{D}, \Upsilon^n, \Psi^n) + \sum_{i=1}^I \sum_{g_i=1}^M \ln \epsilon_{g_i} P(g_i|\mathcal{D}, \Upsilon^n, \Psi^n) \quad (8)$$

Since \mathbb{Q} is smooth and differentiable, we apply a gradient descent method to maximize the Q function, and then update the parameters Ψ ,

$$\Psi^{(n+1)} = \Psi^{(n)} + \epsilon \frac{\partial(\mathbb{Q})}{\partial\Psi}, \tag{9}$$

where ϵ is the learning rate. By iterating between E-step and M-step until the log likelihood converges, we can obtain a good estimate of the parameters.

5.3. Return assessment

After obtaining the parameters Φ , we can construct the predictive function for assessing the propensity of a customer i to return a purchased product j , i.e., $\mathbb{E}(y_{ij}|\Phi) = \mathbf{e}_i \mathbf{s}_j^T + \mathbf{s}_j \mathbf{r}_i^T$. Given a customer-product pair (i, j) , we may predict its return propensity accordingly. The larger the $\mathbb{E}(y_{ij}|\Phi)$ is, the more likely user i is to return product j .

6. Experiment results

In this section, we demonstrate the effectiveness of our proposed FRPLM framework on real data from an online retailer.

6.1. Experiment setup

We obtained data from an online retailer in Taobao, the largest B2C platform in China, owned by Alibaba Corporation. The main products of the retailer are in cosmetics and daily necessities. The data includes four parts: (1) consumer profile data, (2) product feature data, (3) order data, and (4) return and refund data. The consumer profile data includes customer ID, gender, VIP level, credit score, location (zip code, city, state, district, etc.), birthday, and job. The product feature data includes product ID, name, function, warranty, invoice, discount, whether recommended, and website URL. The order data includes order ID, customer ID, product ID, unit price, number of products, discount, final payment, pay time, refund status, shipping and handling fee. The refund data includes order ID, order status, product status, carrier, return time, reason category, detailed reason, and refund amount. Table 2 lists main data statistics.

In this study, we fuse the consumer profile and product features to enhance product return assessment. Continuous features such as credit scores, are converted into categorical values by equal-depth discretization.

For the return propensity of each customer-product pair (Y) in the training set, the historical return records are preprocessed to derive continuous values in the range $[0, 1]$. We counted the frequency that customer i buys product j , as well as the frequency that i returns j . Thus, the return propensity of customer-product pair (i, j) is computed as the ratio between the two frequencies, $y_{ij} = \frac{\# \text{ of return for } i \text{ and } j}{\# \text{ of purchase for } i \text{ and } j}$. Fig. 1 shows the distribution of return propensity.

Finally, the return propensity of customer-product pairs y_{ij} , together with the consumer profiles and product features were fed into the proposed model. In more detail, we randomly divided the data into 70% for training and 30% for testing, and then trained the

model with the stopping criteria that the maximum number of iterations is more than 100 or the relative tolerance of likelihood $\frac{lik^t - lik^{t-1}}{lik^{t-1}}$ is less than 10^{-5} . We initially set the number of customer segments $M = 3$, the number of product categories $N = 3$, the number of latent features $K = 3$, and randomly initialized the parameters of model. In particular, we randomly initialized the mean parameters of Gaussian distributions with expectation equal to 0; we randomly initialized the variance parameters of Gaussian distributions with expectation equal to 100. By setting larger variance values, the proposed model can search and identify optimal parameters in a wider range of real values.

6.2. Evaluation metrics

We evaluate our proposed FRPLM method in terms of overall accuracy, confusion matrix, precision, and recall.

6.2.1. Overall accuracy

To evaluate the overall accuracy, we adopt two metrics: (1) the Mean Absolute Error, $MAE = \sum_{i,j} |y_{ij} - \hat{y}_{ij}|/N$; (2) the Root Mean Square Error, $RMSE = \sqrt{\sum_{i,j} (y_{ij} - \hat{y}_{ij})^2}/N$, where y_{ij} and \hat{y}_{ij} denote the observed return propensity and the predicted return propensity, respectively, and N denotes the test dataset size. The smaller MAE or RMSE are, the more precise a prediction is.

6.2.2. Precision and recall

Retailers are interested in predicting customer-product pairs with high return propensity. Since only a small fraction (4.8%) of customer-product pairs have return propensity greater than 0.1 (see Fig. 1), we choose 0.1 as a threshold, treating values ≥ 0.1 as “high return propensity” and values < 0.1 as “lower return propensity”. Given a top- N list of customer-product pairs E_N sorted in a descending order according to the predicted return propensity, precision and recall can be defined as $\text{Precision}@N = \frac{|E_N \cap E_{\geq 0.1}|}{N}$ and $\text{Recall}@N = \frac{|E_N \cap E_{\geq 0.1}|}{|E_{\geq 0.1}|}$, where $E_{\geq 0.1}$ is the set of customer-product pairs with return propensity values greater than or equal to 0.1.

6.2.3. Receiver operating characteristic (ROC) and area under curve (AUC).

Fig. 1 shows that the distribution of return propensity is extremely unbalanced: the return propensity for most customer-product pairs is zero, while only a small portion have high return propensity. In such imbalanced case, when every customer-product pair is predicted to be 0, the RMSE and MAE values will be small. Nevertheless, such prediction is meaningless for retailers, as it cannot help them identify customer-product pairs that are risky in terms of return propensity. In practice, managers care about whether a customer would return a particular product. Thus, we treat this task as a binary classification problem, and we classify a customer-product pair as either low return propensity or high return propensity according to a predefined threshold. We present ROC curves and AUC values to better illustrate the overall effectiveness of return predictions on imbalanced data. A receiver operating characteristic (ROC) curve is a plot that illustrates the performance of a binary classifier as its discrimination threshold is varied. The curve is created by plotting the true positive rate (TPR) against the false positive rate (FPR) for various threshold settings. Here, $TPR = TP/(TP + FN)$ and $FPR = FP/(FP + TN)$ where TP , FN , FP , and TN represent the number of true positives, false negatives, false positives, and true negatives, respectively.

6.3. Baseline algorithms

As noted before, we can view the problem of assessing the return propensity for each individual purchase as similar to the task of recommender systems, which predict personalized ratings for each

Table 2
Statistics of the experiment data.

Data sources	Properties	Statistics
Customer	Number of customers	34,360
Product	Number of products	6406
Order	Number of orders	315,881
	Average orders per user	9.19
Return and refund	Number of return requests	50,250
	Average return products per user	1.64

user–item pair. Matrix factorization is a widely used collaborative filtering method in recommender systems, and thus we compared our model against several methods based on matrix factorization based: (1) *Probabilistic Matrix Factorization (PMF)* [38]: The widely used SVD finds the matrix $R \approx U^T V$ of given rank which minimizes the sum-squared distance to the target matrix R . PMF can be treated as a probabilistic extension of the SVD model, which (2) *Nonnegative Matrix Factorization (NMF)* [39]: Given a non-negative matrix R , NMF aims to find non-negative matrix factors U and V such that $R \approx U^T V$. NMF can be treated as a non-negative extension of the SVD model.

Aside from matrix factorization methods, we also compared our method with two memory-based methods: (3) *User Frequency based Method*: Given a customer–product pair, we compute the return frequency and order frequency of this customer, and use the ratio of return to order frequency as the predicted propensity. (4) *Item Frequency based Method*: Similarly, given a customer–product pair, we compute the return frequency and order frequency of this product, and use the ratio of return to order frequency as predicted propensity.

Moreover, we further compared with the method (5) *RPLM* proposed in Section 4, which does not fuse any user profiles or product features, to validate the benefits of fusing side information about customers and products.

We perform experiments on a x64 machine with i7 3.40GHz Intel CPU (with 4 cores) and 24GB RAM. The operating system is Microsoft Windows 7.

6.4. Parameter setting

Next, we investigate the sensitivity with respect to different parameter settings, in terms of RMSE and MAE.

6.4.1. Number of iterations

First, we report the RMSE with versus iteration number during model training. Fig. 4 shows that the RMSE values gradually decrease and finally converge.

6.4.2. Dimensions of latent factors

Second, we report the RMSE and MAE with respect to different dimensions K of latent factors, including the latent vectors of the expected, shipped, and received product, which is displayed in Fig. 5a and b. If the number of latent features K is larger than 15, both RMSE and MAE will significantly increase, since the model may suffer from overfitting when the dimension is too large. Therefore, we recommend to choose a value of K less than or equal to 15, for example, 3 or 5.

6.4.3. Numbers of latent customer segment and product category

Third, we report the RMSE and MAE with respect to different numbers of latent customer segments ($M = 2, 3, 5, 10, 15, 20$). Fig. 6a shows that RMSE varies over different number of latent customer segments. We achieve lower RMSE values at the settings of $M = 2$ or 3. Fig. 6b presents the MAE values with respect to different numbers of latent segments. We can see that the lowest MAE values are achieved when the settings of M are neither too large nor too small, e.g., $M = 5$. Similarly, we report the RMSE and MAE with respect to different number of latent product category ($N = 2, 3, 5, 10, 15, 20$) in Fig. 7a and b. We can see that both the RMSE values and MAE values fluctuate within a small range.

6.5. Study of model effectiveness

We report the overall predictive effectiveness of our method compared to all the baseline algorithms, in terms of all the previously described metrics. We repeat the model training five times, and compute the average performance to obtain the evaluation results. We set the number of customer segments to $M = 3$, the number of product categories to $N = 3$, and the dimensions of latent factors to $K = 3$ and 5.

Fig. 8 compares performance in terms of precision and recall when the dimensions are set to $K = 3$. Overall, FRPLM outperforms the other baseline methods. To better capture the predictive ability of the proposed model, we also visualize the ROC curves and AUC values. Fig. 9 shows that FRPLM achieves larger AUC values, compared to baseline algorithms.

The following implications follow from the above comparison. First, matrix factorization methods (i.e. PMF and NMF) fail to identify customer–product pairs of high return propensity in the top- N ranking, and therefore generate biased and unbalanced predictions, which typically predict the return propensity of most customer–product pairs to be zero (PMF) or positive (NMF). Thus, the results do not convey information of practical managerial value. Second, the memory-based methods FreqUser and FreqItem provide a much more balanced and robust prediction. However, the performance cannot compete, in terms of precision and recall, with our proposed methods. Third, RPLM (without fusing any side information) is more balanced compared to matrix factorization methods, but it still cannot compete with our proposed FRPLM method. Thus, fused consumer and product features substantially contribute to better prediction. To sum up, our proposed FRPLM factorizes the return propensity as the interactions between several latent factors, which can be interpreted as different, yet personalized, “manifestations” of a product throughout the entire online shopping process. Moreover, we intelligently fuse correlated side information by introducing latent categorical variables. Therefore, FRPLM outperforms other baseline algorithms in terms of top- N ranking measures, while




Fig. 4. RMSE on the training data versus iteration. The error of the proposed model rapidly decreases as we iteratively update the parameters.




Fig. 5. The RMSE and MAE of the proposed model with respect to different dimensions of latent factors.




Fig. 6. The RMSE and MAE of the proposed model with respect to different numbers of customer segments.

simultaneously providing better practical guidance, as shown by the confusion matrix.

6.6. Study of return reasons

Furthermore, we present a case study to show the capability of identifying the return reasons based on our proposed method. Specifically, we select several customer–product pairs in which the customer buys multiple products and returns a portion of the products due to various reasons. These reasons are generally recorded in the system, either reported by the customer or processed by the

retailers. We list the return reasons for two customer–product pairs in Table 3. We see that both of the customers have made several purchases for the particular product (13 times respectively); however the return frequency and the reasons are different. The customer ‘DF0001’ returned multiple times, and most of the reasons are due to shipping problems; while the customer ‘MX0025’ returned only once, due to the reason that the product caused allergies. Then, we computed the defined inconsistency for the two customer–product pairs from the learned parameters of the model. According to the proposed model, the inconsistency can arise in both the buying phase (α_{ij}) and shipping phase (β_j), which can be estimated by




Fig. 7. The RMSE and MAE of the proposed model with respect to different numbers of product categories.




Fig. 8. The Precision@N and Recall@N of the proposed model and baseline algorithms when K = 3.

Eqs. (1) and (2), respectively. For the overall estimated inconsistency $\hat{y}_{ij} = \alpha_{ij} + \beta_j$, the 'DF0001-17630425695' pair is obviously greater than the 'MX0025-14986823369' pair, which is an example of the predictive power of our model. For the 'DF0001-17630425695' pair, the estimated β inconsistency is greater than α , revealing that the return reasons for this particular pair mostly happen in the shipping phase, which corresponds to the listed return reason in Table 3.

On the other hand, for 'MX0025-14986823369' pair, the situation is reversed, with α greater than β , indicating that the inconsistency mainly arises in the buying phase (or, in other words, the expected product is relatively consistent with the actual product). This is also consistent with the reason that the received product may have some quality problems (e.g., cause allergy problems).

6.7. Managerial implications

As the experiments demonstrate, our proposed model can identify the most risky product return events, and can also partly explain the reasons underlying each product return. A return is due to inconsistency between the expected and the received product, which arises throughout all stages of the online shopping process.

More specifically, our models decompose this inconsistency into two types, and can leverage the inferred parameters to identify which type contributes more to a specific product return incident. Therefore, through this analysis, online merchants can take informed actions to better manage their products and consumers. For example, if “type-I” inconsistency (with large α_{ij} values) often arises for a particular product, managers should pay more attention to the descriptions of the products, and use more media-rich representations (e.g., animations or videos) to mitigate such inconsistency. Particularly, the merchants can profile a consumer’s intention in returning a product and help the merchants better manage their customer relationships. Considering the consumers’ profiles, we can analyze which group of consumer are more likely to return particular products with particular reasons. For example, the merchants may discover that “type-I” inconsistency for some female consumers is higher when buying clothes, possibly because some female consumers may be very particular when choosing clothes. Therefore, their “expectation” for the product could be high, and thus the actual

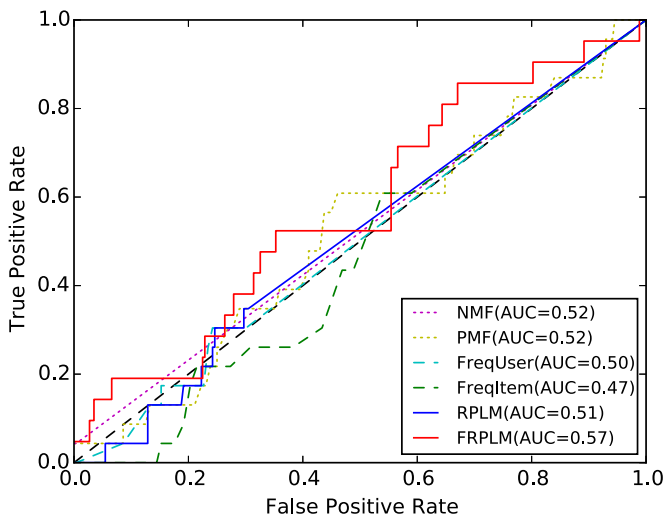


Fig. 9. The ROC curves and AUCs of the proposed methods and the baseline algorithms.

Table 3

Return reasons of two customers. CID represents customer identification number and PID represents product identification number.

CID	PID	Purchase freq.	Order date	Return reason	α	β	Overall inconsistency
DF0001	17630425695	13	2013-04-27	Not received	0.07551	0.10105	0.17656
			2013-04-28	Not received			
			2013-04-28	Not received			
			2013-04-30	Not received			
			2013-04-30	Not received			
MX0025	14986823369	13	2013-05-25	Description on web site was not accurate	0.06547	0.04586	0.11133
			2013-12-06	The product causes allergy.			

“shipped product” may not meet their expectations, resulting in a return.

If ‘type-II’ inconsistency (with large β_j values) often arises (either instead of or in addition to type-I), managers should check out the characteristics of the products. The products may be fragile but not well-packaged during shipping, thus they should take actions to improve the packaging and choose more reliable third-party carriers.

Furthermore, our model can also help identify consumers that abuse the liberal return policies. For example, if a consumer returns frequently, and the return reasons often diverge from those of most return cases, then managers should closely watch the behavior of such consumers and take measures, when necessary.

7. Conclusions

In this paper, we investigated what are the underlying forces that drive product return requests, and how to assess return propensity of specific customers for specific products. Specifically, we assumed that product returns are due to the inconsistency between expected and received product. To model this inconsistency, we devised a generalized return propensity latent model (RPLM) framework to model the inconsistency between expected product and shipped product, and the inconsistency between shipped product and received product, which arises during the buying and shipping phases, respectively. Furthermore, we proposed an enhanced specification, called fused return propensity latent model (FRPLM), to jointly model the correlation among user profiles, product features, and return propensity. Finally, we conducted extensive experiments on real-world data from an online shopping site, which contain orders, returns, and refunds, as well as side information, including as user profiles and product features. As demonstrated by the experimental results, a view of inconsistency generated along with retailer–customer exchange process can better assess return propensity, with interpretable explanations. The performance improvement of our proposed method is substantial, compared to benchmark methods.

References

- Online sales growth puts retailers at risk of over-stating christmas sales in their new year trading reports, <http://fits.me/online-sales-growth-puts-retailers-at-risk-of-over-stating-christmas-sales-in-their-new-year-trading-reports/>.
- Online retailers tackle high rate of customer returns, <http://abcnews.go.com/Business/online-shopping-transactions-returned/story?id=21312312>.
- M. Maity, M. Dass, Consumer decision-making across modern traditional channels: e-commerce, m-commerce, in-store, *Decision Support Systems* 61 (2014) 34–46.
- E. Ofek, Z. Katona, M. Sarvary, bricks and clicks: the impact of product returns on the strategies of multichannel retailers, *Marketing Science* 30 (1) (2011) 42–60.
- B. Toktay, Forecasting product returns INSEAD, 2001.
- P. De, Y. Hu, M.S. Rahman, Product-oriented web technologies product returns: an exploratory study, *Information Systems Research* 24 (4) (2013) 998–1010.
- N. Sahoo, S. Srinivasan, C. Dellarocas, The Impact of Online Product Reviews on Product Returns and Net Sales, 2013 Workshop on Information Systems Economics, Milan, Italy, 2013.
- J.A. Petersen, V. Kumar, Are product returns a necessary evil? Antecedents and consequences, *Journal of Marketing* 73 (3) (2009) 35–51.
- V. Padmanabhan, I.P. Png, Manufacturer’s return policies and retail competition, *Marketing Science* 16 (1) (1997) 81–94.
- N.N. Bechwati, W.S. Siegal, The impact of the prechoice process on product returns, *Journal of Marketing Research* 42 (3) (2005) 358–367.
- W. Chu, E. Gerstner, J.D. Hess, Managing dissatisfaction how to decrease customer opportunism by partial refunds, *Journal of Service Research* 1 (2) (1998) 140–155.
- S.L. Wood, Remote purchase environments: the influence of return policy leniency on two-stage decision processes, *Journal of Marketing Research* 38 (2) (2001) 157–169.
- R.T. Rust, K.N. Lemon, V.A. Zeithaml, Return on marketing: using customer equity to focus marketing strategy, *Journal of Marketing* 68 (1) (2004) 109–127.
- T. Clotey, W. Benton, R. Srivastava, Forecasting product returns for remanufacturing operations, *Decision Sciences* 43 (4) (2012) 589–614.
- M.P. de Brito, R. Dekker, Modelling product returns in inventory control – exploring the validity of general assumptions, *International Journal of Production Economics* 81 (2003) 225–241.
- B.A. Pasternack, Optimal pricing and return policies for perishable commodities, *Marketing science* 27 (1) (2008) 133–140.
- A.B. Bower, J.G. Maxham, III, Return shipping policies of online retailers: normative assumptions and the long-term consequences of free and return, *Journal of Marketing* 76 (5) (2012) 110–124.
- S.K. Mukhopadhyay, R. Setoputro, Optimal return policy and modular design for build-to-order products, *Journal of Operations Management* 23 (5) (2005) 496–506.
- Y. Li, L. Xu, D. Li, Examining relationships between the return policy, product quality, and pricing strategy in online direct selling, *International Journal of Production Economics* 144 (2) (2013) 451–460.
- S. Davis, M. Hagerty, E. Gerstner, Return policies and the optimal level of hassle, *Journal of Economics and Business* 50 (5) (1998) 445–460.
- E.T. Anderson, K. Hansen, D. Simester, The option value of returns: theory and empirical evidence, *Marketing Science* 28 (3) (2009) 405–423.
- J.D. Hess, G.E. Mayhew, Modeling merchandise returns in direct marketing, *Journal of Interactive Marketing* 11 (2) (1997) 20–35.
- Y. Koren, Factorization meets the neighborhood: a multifaceted collaborative filtering model, *KDD 08*, 2008.
- D.D. Lee, H.S. Seung, Algorithms for non-negative matrix factorization, *NIPS 00*, 2000.
- R. Salakhutdinov, A. Mnih, Probabilistic matrix factorization, *NIPS*, 2008.
- R. Salakhutdinov, A. Mnih, Bayesian probabilistic matrix factorization using Markov chain Monte Carlo, *ICML 08*, 2008.
- M.N. Schmidt, O. Winther, L.K. Hansen, Bayesian non-negative matrix factorization, *ICA 09*, 2009.
- T. Chen, Z. Zheng, Q. Lu, W. Zhang, Y. Yu, Feature-based matrix factorization, *arXiv preprint arXiv:1109.2271*, 2011.
- D. Agarwal, B.-C. Chen, Regression-based latent factor models, *KDD 09*, 2009.
- B.-C. Chen, J. Guo, B. Tseng, J. Yang, User reputation in a comment rating environment, *KDD 11*, 2011.
- L. Zhang, D. Agarwal, B.-C. Chen, Generalizing matrix factorization through flexible regression priors, *RecSys 11*, 2011.
- H. Ma, D. Zhou, C. Liu, M.R. Lyu, I. King, Recommender systems with social regularization, *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, 2011, pp. 287–296.
- H. Bao, Q. Li, S.S. Liao, S. Song, H. Gao, A new temporal and social pmf-based method to predict users’ interests in micro-blogging, *Decision Support Systems* 55 (3) (2013) 698–709.
- Y. Ge, H. Xiong, A. Tuzhilin, Q. Liu, Cost-aware collaborative filtering for travel tour recommendations, *ACM Transactions on Information Systems* 32 (1) (2014) 1–31. <http://dl.acm.org/citation.cfm?doid=2576772.2559169>. <http://dx.doi.org/10.1145/2559169>.
- D. Creal, B. Schwaab, S.J. Koopman, A. Lucas, Observation-driven mixed-measurement dynamic factor models with an application to credit risk, *Review of Economics and Statistics* 96 (5) (2014) 898–915.

- [36] M. Guerzhoy, A. Hertzmann, Learning latent factor models of human travel, NIPS Workshop on Social Network and Social Media Analysis: Methods, Models and Applications, 2012.
- [37] Y. Fu, H. Xiong, Y. Ge, Z. Yao, Y. Zheng, Z.-H. Zhou, Exploiting geographic dependencies for real estate appraisal: a mutual perspective of clustering and ranking, KDD14, 2014.
- [38] A. Mnih, R. Salakhutdinov, Probabilistic matrix factorization, *Advances in Neural Information Processing Systems*, 2007. pp. 1257–1264.
- [39] D.D. Lee, H.S. Seung, Algorithms for non-negative matrix factorization, *Advances in Neural Information Processing Systems*, 2001. pp. 556–562.

Yanjie Fu received the B.E. degree from the University of Science and Technology of China (USTC), Hefei, China, 2008, the M.E. degree from the Chinese Academy of Sciences (CAS), Beijing, China, 2011, and the Ph.D. degree from Rutgers University. He is currently an assistant professor at Missouri University of Science and Technology. His research interests include data mining, urban computing, mobile intelligence, and personalization techniques. He has published in refereed journals and conference proceedings, such as TKDE, TKDD, TMC, KDD, ICDM, SDM.

Guannan Liu is currently an Assistant Professor in the Department of Information Systems with Beihang University, Beijing, China. He received the Ph.D. degree from Tsinghua University, China, 2016. His research interests include data mining, social networks, and business intelligence. His work has been published in the journal of Decision Support Systems, Neurocomputing, etc., and also in the conference proceedings such as KDD, ICDM, etc.

Spiros Papadimitriou is an Assistant Professor at the Department of Management Science & Information Systems at Rutgers Business School. Previously, he was a research scientist at Google, and a research staff member at IBM Research. His main interests are large scale data analysis, time series, graphs, and clustering. He has published more than forty papers on these topics and has three invited journal publications in best paper issues, several book chapters and he has filed multiple patents. He has also given a number of invited talks, keynotes, and tutorials. He was a Siebel scholarship recipient in 2005 and received the best paper award in SDM 2008.

Hui Xiong is currently a Full Professor and Vice Chair of the Management Science and Information Systems Department, and the Director of Rutgers Center for Information Assurance at the Rutgers, the State University of New Jersey, where he received a two-year early promotion/tenure (2009), the Rutgers University Board of Trustees Research Fellowship for Scholarly Excellence (2009), and the ICDM-2011 Best Research Paper Award (2011). He received the B.E. degree from the University of Science and Technology of China (USTC), China, the M.S. degree from the National University of Singapore (NUS), Singapore, and the Ph.D. degree from the University of Minnesota (UMN), USA. His general area of research is data and knowledge engineering, with a focus on developing effective and efficient data analysis techniques for emerging data intensive applications. He has published prolifically in refereed journals and conference proceedings (3 books, 40+ journal papers, and 60+ conference papers). He is a co-Editor-in-Chief of Encyclopedia of GIS, an Associate Editor of IEEE Transactions on Data and Knowledge Engineering (TKDE) and the Knowledge and Information Systems (KAIS) journal. He has served regularly on the organization and program committees of numerous conferences, including as a Program Co-Chair of the Industrial and Government Track for the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining and a Program Co-Chair for the 2013 IEEE International Conference on Data Mining (ICDM-2013). He is a senior member of the ACM and IEEE.

Xiaolin Li is currently an Associate Professor in the School of Management at Nanjing University, China. Her main research interests include business intelligence, data mining, and decision-making. She has published a number of papers in refereed journals and conference proceedings, such as Information Sciences (INS), ECML, and PAKDD. She also has served as a reviewer and PC member for numerous journals and conferences, such as IEEE TEC, KAIS, and KDD'15.

Guoqing Chen received the Ph.D. degree in Managerial Informatics from the Catholic University of Leuven, Leuven, Belgium, in 1992. He is currently a Professor of Information Systems with the School of Economics and Management, Tsinghua University, Beijing, China. His current research interests include data mining and business intelligence, e-business, and fuzzy logic. His work has been published in journals such as Decision Support Systems, Information Sciences, IEEE Transactions on Fuzzy Systems, Fuzzy Sets and Systems, IEEE Intelligent Systems, Communications of the ACM.