Beyond Movie Recommendations:
Solving the Continuous Cold Start Problem in E-commerce Recommendations

Julia Kiseleva[1]　　　Alexander Tuzhilin[4]　　　Jaap Kamps[3]
Melanie J.I. Mueller [2]　　　Lucas Bernardi[2]　　　Chad Davis[2]
Ivan Kovacek[2]　　　Mats Stafseng Einarsen[2]　　　Djoerd Hiemstra[5]

[1]Eindhoven University of Technology, Eindhoven, The Netherlands
[2]Booking.com, Amsterdam, The Netherlands
[3]University of Amsterdam, Amsterdam, The Netherlands
[4]Stern School of Business, New York University, New York, USA
[5]University of Twente, Enschede, The Netherlands

[1]j.kiseleva@tue.nl　　[3]kamps@uva.nl　　[4]atuzhili@stern.nyu.edu　　[5]d.hiemstra@utwente.nl
[2]{melanie.mueller, lucas.bernardi, chad.davis, ivan.kovacek, mats.einarsen}@booking.com

ABSTRACT

Many e-commerce websites use recommender systems or personalized rankers to personalize search results based on their previous interactions. However, a large fraction of users has no prior interactions, making it impossible to use collaborative filtering or rely on user history for personalization. Even the most active users may visit only a few times a year and may have volatile needs or different personas, making their personal history a sparse and noisy signal at best. This paper investigates how, when we cannot rely on the user history, the large scale availability of other user interactions still allows us to build meaningful profiles from the contextual data and whether such contextual profiles are useful to customize the ranking, exemplified by data from a major online travel agent Booking.com.

Our main findings are threefold: First, we characterize the *Continuous Cold Start Problem* (*CoCoS*) from the viewpoint of typical e-commerce applications. Second, as explicit situational context is not available in typical real world applications, implicit cues from transaction logs used at scale can capture essential features of situational context. Third, contextual user profiles can be created offline, resulting in a set of smaller models compared to a single huge non-contextual model, making contextual ranking available with negligible CPU and memory footprint. Finally we conclude that, in an online A/B test on live users, our contextual ranker increased user engagement substantially over a non-contextual baseline, with click-through-rate (CTR) increased by 20%. This clearly demonstrates the value of contextual user profiles in a real world application.

1. INTRODUCTION

In addition to the handful of general web search engines, there are millions of online e-commerce websites driving the online economy [17]. Many of these e-commerce websites are built around personalized search and recommendations systems. Amazon.com recommends books, Booking.com recommends accommodations and destinations, Netflix recommends movies, Reddit recommends news stories and so on. Recommender systems predict unknown ratings based on past or/and current information about users and items, such as past user ratings, user profiles, item descriptions. If this information is not available for new users or items, the recommender system runs into the *Standard Cold Start Problem*: it does not know what to recommend until the new, 'cold' user or item gets 'warmed-up', i.e. until enough information has been received to produce recommendations. For example, which hotels should be recommended to someone who visits Booking.com for the first time? If the recommender system is based on the history of users click' in the past, the first recommendations can only be made after the user has clicked on a couple of hotels on the website.

Several approaches have been proposed to deal with the cold-start problem, such as utilizing baselines for cold users [32], combining collaborative filtering with content-based recommenders in hybrid systems [49], eliciting ratings from new users [43], promoting diversity in recommendations [23], or exploiting the social network of users [50]. In particular, content-based approaches have been very successful in dealing with cold-start problems in collaborative filtering [48, 49]. However, these approaches deal explicitly with 'cold' users or items, and provide a 'fix' until enough information has been gathered to apply the core recommender system. Thus, rather than providing unified recommendations for 'cold' and 'warm' users, they temporarily bridge the period during which the user or item is 'cold' until it is 'warm'. This can be very successful in situations in which this warm-up period is short, and when warmed-up users or items stay warm.

However, in many practical e-commerce applications, users or items remain 'cold' for a long time, and can even 'cool down' again, leading to the *Continuous Cold Start Problem* (*CoCoS*). For example in Booking.com, many users visit and book infrequently because they have only one or two vacations per year, leading to a prolonged cold-start and extreme sparsity of collaborative filtering matrices, see Figure 1 (A). In addition, even 'long term warm' users can cool down as they change their needs over time [27], e.g. com-
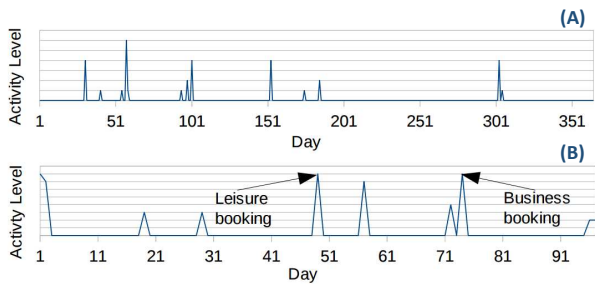
Figure 1: Continuously 'cold' users at Booking.com. Activity levels of two randomly chosen users over time. (A): The top user has only rare activity throughout a year. (B): the bottom user exhibits different personas by making a leisure and a business booking without much activity in between.
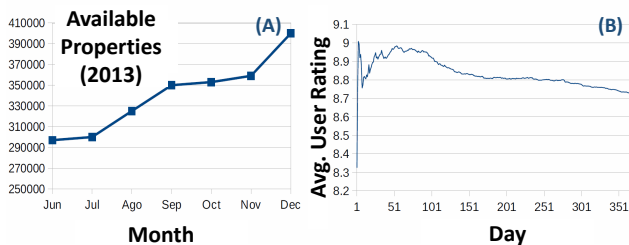

Figure 2: Continuously cold items at Booking.com. (A): Thousands of new accommodations are added every month. (B): The user ratings of a randomly chosen hotel change continuously over the year.

ing from Booking.com of youth hostels for backpacking to booking of resorts for family vacations. Such 'cool-downs' can happen more frequently and rapidly for users who book accommodations for different travel purposes, e.g. for leisure holidays and business trips as shown in Figure 1 (B). Moreover, we have a mirror problem in the items to recommend: new items appear frequently leading to many items without prior interactions as shown in Figure 2 (A) for accommodations at Booking.com, and items can change their characteristics as shown in Figure 2 (B), making historical interactions a noisy signal. The *CoCoS* is ignored in the literature despite its relevance in industrial applications. Classical approaches to the cold-start problem fail in the case of a *CoCoS*, since they assume that users get warmed up in a reasonable time and stay warm after that.

This paper proposes a new approach of using contextual user profiles for personalized search and recommendations in the context of a major online travel agent, Booking.com, in particular using the Destination Finder. Situational context provides powerful cues about user preferences that hold the promise to improve the quality of recommendations over the use of traditional long term interests [e.g., 4, 9, 11]. In this setup, rankings are computed based on the current context of the current visitor and the behavior of other users in similar contexts [e.g., 5, 22, 53]. This type of data is readily available in most e-commerce settings. This approach naturally addresses sparsity by clustering users into contexts. Since context is determined on a per-action basis, user volatility and multiple personas can be addressed robustly.

Working in a real world setting comes with specific challenges for search and recommendation systems [33]. First, in an online

service, context is shallow but available at scale. Context can be almost anything—ranging from explicit user profiles to data about moods and attitudes—but explicit user context is typically not available in online services. There is an abundance of situational context (day, time, device, etc) in server logs which may hold important implicit contextual cues. Hence, although rich contextual information is not available for a large fraction of users, the large scale availability of implicit situational context may still allow us to capture essential context features. Second, if it's not fast it isn't working. Due to the volume of traffic, offline processing—done once for all users—comes at marginal costs, but online processing—done separately for each user—can be excessively expensive. Clearly, response times have to be sub-second, but even doubling the CPU or memory footprint comes at massive costs. Hence we cannot include implicit contextual features directly or build an adaptive model for each unique user, but we can build profiles offline and map incoming users to one of the profiles at negligible online processing costs.

We are trying to answer the following main research question: *Can we automatically detect contextual user profiles and does customized ranking with these profiles improve travel search and recommendation?* We breakdown of our general research problem into four specific research questions:

- **RQ1:** *How to characterize the continuous cold start problem in travel recommendation?*

We introduce and characterize the *Continuous Cold Start Problem* (*CoCoS*) that happens when users or items remain 'cold' for a long time, and can even 'cool down' again after some time.

- **RQ2:** *How to define and discover contextual user profiles from multi-criteria ranking data in an unsupervised setup?*

We combine multi-criteria ranking data with the n-dimensional contextual space in order to discover contextual user profiles.

- **RQ3:** *How to apply contextual user profiles for the ranking of travel destinations in a continuous cold start setting?*

We propose a novel approach exploiting contextual user profiles which are defined as *'closely connected'* regions of an n-dimensional contextual space.

- **RQ4:** *How effective are contextual profiles for real-world users of the destination finder system in terms of user engagement measures?*

We set up a large-scale online A/B testing evaluation with live traffic from Booking.com, and demonstrate how contextual travel ranking leads to a significant increase in user engagement.

The remainder of the paper is organized as follows. In Section 2 we discuss the most relevant prior work, and position our paper with respect to it. The problem setup is introduced in Section 3. As our approach is generally applicable to any multi-criteria ranking data associated with standard contextual information from web logs, Section 4 outlines our approach as a general framework for discovering and using contextual user profiles. Next, in Section 5, we detail the specific application to our online travel agent service. In Section 6, we describe the results of the online evaluation of the approach in an A/B test with live traffic. Finally, Section 7 concludes our work in this paper and highlights future directions.

2. BACKGROUND AND RELATED WORK

In this section, we review related work in the following two areas. First, we summarize previous work on the attempts to solve *CoCoS*. Second, we review approaches to build situational recommendations.

2.1 Cold Start Problem

In classical formulations of Recommender Systems (RS), the recommendation problem relies on *ratings* (R) as a mechanism of capturing user (U) preferences for different items (I). The problem of estimating unknown ratings is formalized as follows: $F : U \times I \rightarrow R$. Due to practical applications, RS have been an expanding research area since the first papers on collaborative filtering in the 1990s [46, 51]. Many different recommendation approaches have been developed since then, in particular content-based and hybrid approaches have supplemented the original collaborative approaches [2]. For instance, RS based on latent factor models have been effectively used to understand user interests and predict future actions [6, 7]. Such models work by projecting users and items into a lower-dimensional space, thereby grouping similar users and items together and subsequently computing similarities between them. This approach can run into data sparsity problems and into *CoCoS* when new items continuously appear. Although, to our knowledge, the *CoCoS* as defined in this work has not been directly addressed in the literature, several approaches are promising.

Tang et al. [57] propose a context-aware recommender system, implemented as a contextual multi-armed bandits problem. Although the authors report extensive offline evaluation (log based and simulation based) with acceptable CTR, no comparison is made from a cold-start problem standpoint.

Sun et al. [54] explicitly attack the user volatility problem. They propose a dynamic extension of matrix factorization where the user latent space is modeled by a state space model fitted by a Kalman filter. Generative data presenting user preference transitions is used for evaluation. Improvements of RMSE when compared to time SVD [34] are reported. Consistent results are reported in [18], after offline evaluation using real data.

Tavakol and Brefeld [58] propose a topic driven recommender system. At the user session level, the user intent is modeled as a topic distribution over all the possible item attributes. As the user interacts with the system, the user intent is predicted and recommendations are computed using the corresponding topic distribution. The topic prediction is solved by factored Markov decision processes. Evaluation on an e-commerce data set shows improvements when compared to collaborative filtering methods in terms of average rank.

This paper builds on our initial discussion of the cold start recommendation problem in ecommerce practice [14], and extends the initial experiments recommendation [31] by looking at ways to exploit the implicit context for increasing the effectiveness of travel recommendations in the real-world setting of Booking.com.

2.2 Context-Aware Recommendations

The radical departure from classical, two-dimensional RS is context-aware recommendation system (CARS) [3], which attract an increasing attention in academic work [20, 21, 53]. Rating prediction in CARS relies primarily on the information of *how* (which rating, e.g. a user giving '3' of '5' stars to an item) and *who* (which user, e.g. gender, mood or nationality) rated *what* (which item, e.g. movie, news article, or hotel). This additional information is called *context*. The general formulation of CARS rating prediction takes into account the context dimension C as follows [3]:

$$F : U \times I \times C \rightarrow R. \qquad (1)$$

Defining context is an important research question in itself. The structured definition of context was introduced in [16]. *Multidimensional context* C is defined as a group of contextual feature-category pairs:

$$C = \{(F_n : \{v_m\}_{m=1}^{M})\}_{n=1}^{N}, \qquad (2)$$

where F_n are contextual features, and v_m are categories for F_n. For example, the contextual feature *location* has the contextual categories 'USA', the 'Netherlands' etc. *Contextual categories* are often predefined by taxonomies [12, 22, 63]. Alternatively, an unsupervised technique is used to discover contextual information [30, 38]. Moreover, context discovery can be formulated as an optimization problem [29] or a feature selection problem [60, 61].

Incorporating contextual information into CARS can be viewed as a separate area of research, and can be classified into three groups: *pre-filtering, post-filtering and contextual modeling* [3]. In the pre-filtering approach, contextual conditions are projected on the items, thereby essentially reducing the problem to a classical RS problem. Adomavicius et al. [4] introduce a multidimensional approach taking various contextual aspects into account in collaborative filtering. They use a reduction based approach mapping a three-dimensional prediction function (of Equation 2) to a two-dimensional one. Baltrunas and Ricci [9, 11] introduce item splitting for dealing with context by generating new items, where context sensitive items are duplicated and the ratings divided over the respective contextual conditions, reducing it to a classical RS problem. This approach is expanded by Baltrunas and Ricci [10] and evaluated on synthetic and real world data sets.

Contextual information is initially ignored for post-filtering approaches, which also can be referred to as contextualization of the recommendation output [42]. The ratings are predicted using any traditional two dimensional RS set-up on the entire data. Then, the resulting set of recommendations is adjusted (contextualized) for each user using the contextual information.

A common context modeling approach is to use contextual information to expand the feature set, thus treating context as a predictive feature. For example, Rendle et al. [45] proposed a novel approach applying Factorization Machines [44] to model contextual information and provide context-aware rating predictions, *using context explicitly* specified by a user to expand *the set of predictive features*.

Tensor Factorization, which is a generalization of Matrix Factorization, allows a flexible and generic integration of contextual information by modeling the data as a User-Item-Context N-dimensional tensor instead of the traditional 2D [28, 52]. In terms of an interactive system, the paper [41] has shown that it was useful to consider the history of user interactions, more specifically changes in these entities. In the paper [20], a co-occurrence analysis is used to mine the top frequent tags for songs from social tagging web sites, and topic modelling is used to determine a set of latent topics for each song. Recently, more techniques for context modeling were developed [13, 21, 56].

In multi-criteria RS [1, 5, 35] (MCRS) the rating function has the following form:

$$F : U \times I \rightarrow r_0 \times r_1 \cdots \times r_n. \qquad (3)$$

The *overall rating* r_0 for an item shows how well the user likes this item, while criteria ratings $r_1, \ldots, r_n$ provide more insight and explain which aspects of the item she likes. MCRS predicts the overall rating for an item based on the past ratings, using both overall and individual criteria ratings, and recommends to users the item with the best overall score. According to [1], there are two basic approaches to compute the final rating prediction in the case when the overall rating is known. First, in *similarity based* approaches, the similarity between users is calculated based on their detailed ratings (e.g. Euclidean distance, Chebyshev distance, or

Pearson correlation). Second, in *aggregation function based* approaches, we exploit the assumption of a relationship between the overall and the criteria ratings, $r_0 = f(r_1, \ldots, r_k)$ (e.g. multiple linear regression techniques can be used). These two approaches have been significantly improved in [26] by using Support Vector regression and combining user- and item-based regression models with a weighted approach. Liu et al. [37] assumed that the overall rating highly correlates with criteria ratings that are particularly significant for individuals.

RS methods are not easy to apply for large scale industrial applications. A large scale application of an unsupervised RS is presented in [24], where the authors apply topic modeling techniques to discover user preferences for items in an online store. They apply Locality Sensitive Hashing techniques to overcome performance issues when computing recommendations.

To summarize, the key distinction of our work compared to previous efforts is twofold: First, we introduce new Continuous Cold Start (*CoCoS*) settings that is common in e-commerce. Second, we propose the discovery of contextual user profiles (CUPs) within a *CoCoS* setting. CUPs are used both to build customized context-aware rankers (which can be done offline), and to map incoming users to the closest contextual user profile to provide contextual recommendations.

3. PROBLEM SETUP

In this section we will study our **RQ1:** *How to characterize the continuous cold start problem in travel recommendation?* First, we characterize the *Continues Cold Start Problem* (*CoCoS*) in Section 3.1. Second, we introduce a Booking.com service Destination Finder that 'suffers' from *CoCoS* in Section 3.2. It will be our platform for experimentation in the remainder of the paper.

3.1 Characterizing Continuous Cold Start

CoCoS can in principle arise on both the user side and the items side. We characterize it using the following four features: **S**: data *sparsity*, related to the original cold-start problem; **V**: *volatility*, or the degree of variation in the object of interest; **I**: object *identity*, due to different technical [39] or law regulation related problems complicating correct identification; **P**: '*personas*', or the different types of behavior expressed by one user in different situations.

The *User Continuous Cold Start Problem* (*UCoCoS*) can be characterized by:

- **S**: new or rare users;
- **V**: users' interests change over time;
- **I**: a failure to match data from the same user;
- **P**: users have different interests at different, possibly close-by points in time.

New users arrive frequently as shown in Figure 1(A), or may appear new when they do not log in or use a different device so we would fail to match their identity. Some websites are prone to extreme sparsity in user activity when items are purchased only rarely, such as travel, cars etc. Most users change their interests over time (volatility), e.g. movie preferences evolve, or travel needs change. On even shorter timescales, users have different personas. Depending on their mood or their social context, they might be interested in watching different movies. Depending on the weather or their travel purpose, they may want to book different types of trips as presented in Figure 1 (B).

Similarly we characterize *Item Continues Cold Start Problem* (*ICoCoS*):

- **S**: new or rare items;
- **V**: item properties or value change over time;
- **I**: a failure to match data from the same item;
- **P**: an item appeals to different types of users.

New items appear frequently in e-commerce catalogues, as shown in Figure 2 (A) for accommodations at Booking.com. Some items are interesting only to niche audiences, or sold only rarely, for example books or movies on specialized topics. Items can be volatile if their properties change over time, such as a phone that becomes outdated once a newer model is released, or a hotel that undergoes a renovation. Figure 2 (B) shows fluctuations of the review score of a hotel at Booking.com. Some items have different 'personas' in that they target several user groups, such as a hotel that caters to business as well as leisure travellers. When several sellers can add items to an e-commerce catalogue, or when several catalogues are combined, correctly matching items can be problematic so we run into an item identity problem.

3.2 Optimizing Destination List within CoCoS

To motivate our problem set-up, we introduce a Booking.com service which allows to find travel destinations based on users' preferred activities: the Destination Finder. Consider a user who knows what activities she wants to do during her holidays, and is looking for travel destinations matching these activities. This process is a complex exploratory recommendation task in which users start by entering activities in the search box as shown in Figure 3. The service returns a ranked list of recommended destinations [8].

The underlying data is based on 'endorsements' of users that have booked a hotel at some destination via the online travel agent in the past. After the users visited the destination, they are asked to endorse the place using a set of endorsements. Initially, the set of endorsements was extracted from users' free-text reviews using a topic-modeling technique such as LDA [15, 40]. Nowadays, the set of endorsements consists of 256 activities such as 'Beach,' 'Nightlife,' 'Shopping,' etc. These endorsements imply that a user liked a destination for particular characteristics. Two examples of the collected sets of endorsements for two destinations 'Bangkok' and 'London' are shown in Figure 4. As an example of the multi-criteria endorsement data, consider three endorsements: $e_1 = $ 'Beach', $e_2 = $ 'Shopping', and $e_3 = $ 'Family Friendly' and assume that a user u_j after visiting a destination d_k (e.g.'London') provides the review $r_i(u_j, d_k)$ as:

$$r_i(u_j, d_k) = (0, 1, 0). \qquad (4)$$

This means our user ranks London for the 'Shopping' activity only. However, we cannot conclude that London is not 'Family Friendly', i.e. negative user opinions are hidden. In contrast to the ratings data of the traditional recommender systems setup, we are dealing with multi-criteria ranking data. Destination Finder is a good example of the service which is working under the *CoCoS* settings from both sides: users and items.

UCoCoS at Destination Finder It is used to plan holidays, so many users visit it infrequently because they have only one or two vacations per year, leading to the sparsity problem. Since users interact with service rarely—many changes can happen and they might shift their preferences from backpacking activities to family friendly places. Users can use different devices to search over Destination Finder without login to the system, so user matching is an actual problem. Users can express different types of preference while planning trips, e.g. they might go to a family friendly resort while traveling with children and look for 'Shark Diving'
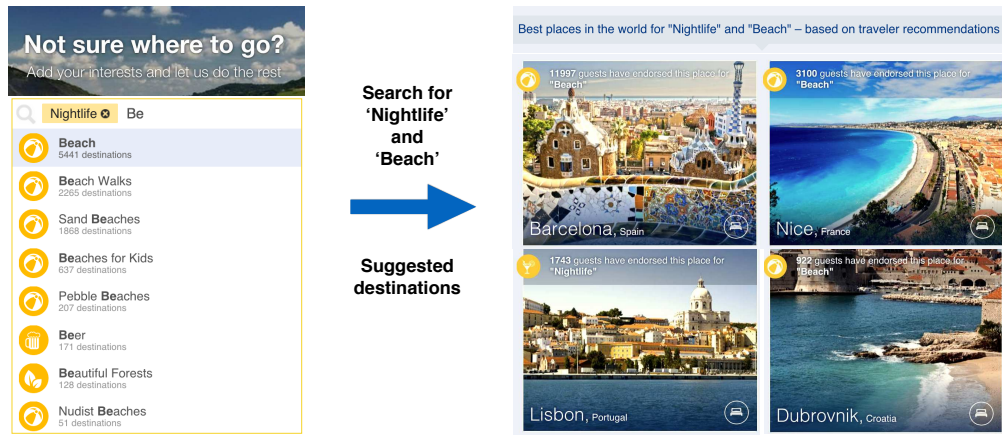
Figure 3: Example of Destination Finder use: a user searching for 'Nightlife' and 'Beach' obtains a ranked list of recommended destinations (top 4 are shown).
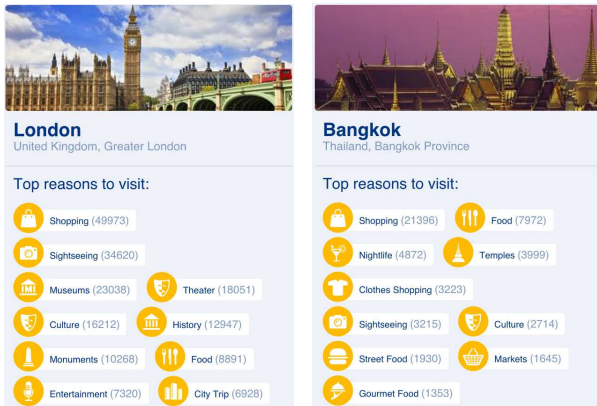

Figure 4: The Destination Finder endorsement pages of London and Bangkok.

while planning holidays alone, so we need to deal with different user 'personas'.

ICoCoS at Destination Finder The list of destinations is growing continuously over time because users share their experience about new places, so we run into the item sparsity problem. User reviews for destination depends on contextual information. For example, the resort 'The Hague at North Sea' is widely endorsed for the activity 'Beach' during summer, but not during winter, so we run into the item volatility. Moreover, destination might change over time, e.g. a new aquarium is build and users start to endorse a place for it. Some destinations have different 'personas' in which they target several user groups, such as a destination which can be family friendly but at the same time has rich night live. Therefore, we have places that are expressing different 'personas'.

These aspects of *CoCoS* at Destination Finder can be addressed partially by taking context into account. We propose that the described multi-criteria endorsements can be enhanced by contextual information. We build a contextual ranker for recommending destinations, whereas the current live systems uses an advanced non-contextualized ranker.

To summarize, we introduced the continuous cold start problem,

and characterized the user and item sides of the *CoCoS*. We also introduced the Destination Finder setup that we used in this paper: (**1**) we have a set of geographical destinations such as 'Paris', 'London', 'Amsterdam' etc.; (**2**) each destination is ranked by users who visited the place using a set of endorsements under some situation (which can be described by a set of contexts). In the setting of *CoCoS*, our main goal is to find ways to map any incoming user, without assuming prior history or explicit profiles, to some cluster of like-minded previous users using only contextual data. In the next section, we will discuss how to discover such contextual user profiles.

4. MULTIDIMENSIONAL CONTEXTUAL USER PROFILES

In this section we will study our **RQ2:** *How to define and discover contextual user profiles from multi-criteria ranking data in an unsupervised setup?* We present an overview of our framework for discovering multidimensional contextual user profiles (CUPs), as outlined in Figure 5. It has two main stages: *offline* (**A**), and *online* (**B**). The discovery of multidimensional CUPs (**A.1**) happens during the offline stage and is described in Section 4.1. The process of using discovered CUPs is as follows: (**A.2**) during the *offline stage*, we apply the set of discovered CUPs to learn a customized ranker; and (**B**) during the *online stage*, we assign incoming users to one of the CUPs. The process of using CUPs is presented in Section 4.3. Section 4 defines CUPs in a generic way. In Section 5 we show how the framework can be applied to the Destination Finder.

4.1 Defining Contextual User Profiles

Apart from the reviews, as defined in Equation 4, there is additional contextual information about the *situation* in which users made their choice (to consider or not to consider the suggested destination), e.g. the geographical location, the time (when a user is using Destination Finder), the users' device type, or the referral (where is a user coming from). We adopt the definition of the context as described in Equation 2.

In many real world RS it is not feasible to track user identity information u_j for several reasons: (**1**) privacy issues: only a limited part of the user interaction history can be stored; (**2**) the cold-start problem: when a new user comes without prior history of interaction with the system; (**3**) a user does not have to be logged in: so we cannot make use of his interaction history. However, we would like
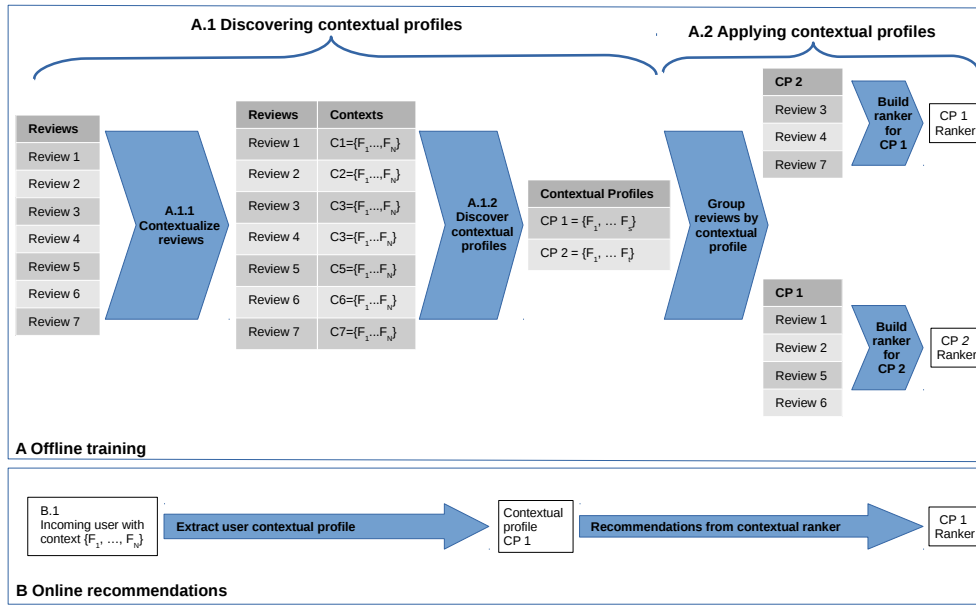
Figure 5: An overall framework for discovering multidimensional contextual user profiles.
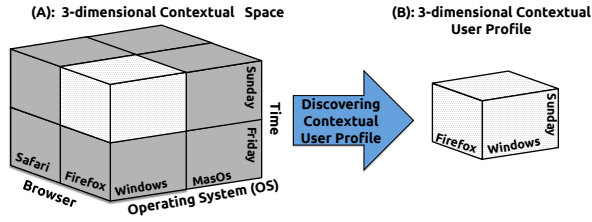

Figure 6: An example for discovering a contextual user profile from 3-dimensional contextual space. The $3D$ contextual space can be visualized as a cube (A), of which the contextual user profile is a cube region (B).

to predict user preferences in order to supply him with suitable recommendations. Therefore, we want to detect a list of typical user situations using contextual information. This type of situational information we call CUPs.

Contextual information can be represented as a n-dimensional space where the dimensions are the set of contextual features, $\{F_n\}_{n=1}^N$, and the coordinates for each dimension are the contextual categories, $\{v_m\}_{m=1}^M$. For example, the contextual feature F_1, 'User Device', is represented by the following contextual categories:

$$F_1 = \{v_1 = \text{'Mobile'}, v_2 = \text{'Tablet'}, v_3 = \text{'PC'}\}. \quad (5)$$

To simplify the notation we rewrite Equation 5 as:

$$F_1 = \{F_{11} = \text{'Mobile'}, F_{12} = \text{'Tablet'}, F_{13} = \text{'PC'}\}. \quad (6)$$

The 3-dimensional example (cube) of contextual space is presented in Figure 6 (A) where we have three dimensions: $\{F_1 = \text{'OS'}, F_2 = \text{'Browser'}, F_3 = \text{'Time'}\}$.

A *contextual user profile* is a *region* in the n-dimensional contextual space that represents 'typical' user behavior. When a user visits our service we can map him to one of the CUPs and use this insight into his preferences to improve the quality of the service, i.e. serving better travel recommendations in the Destination Finder.

4.2 Discovering Contextual User Profiles

We now discuss in more detail the process of discovering CUPs, as outlined in Figure 5 (A.1). The review entities as defined in Equation 4 can be contextualized, i.e., extended by multidimensional contextual information C as depicted in Figure 5 (A.1.1). We use the context definition presented in Equation 6. The contextual review r_i has the following form:

$$r_i(u_j, d_k) = (e_1, \dots, e_X, F_{11}, \dots, F_{NM}), \quad (7)$$

where:

1. u_j is user information that is not stored explicitly, but in our setup we have contextual information regarding how a review is made;

2. d_k is a destination which a user u_j ranks using multi-criteria endorsements;

3. $e_1, \dots, e_X$ are endorsements represented as binary values;

4. $F_{11}, \dots, F_{NM}$ are contextual features represented in a binary way. For example, if a user is using a device with 'Windows' as OS and a 'Firefox' browser on Sunday, then the context vector is $(1, 0, 0, 1, 0, 1)$.

In our setup we combine CARS and MCRS presented in Equation 1 and 3 accordingly. A key difference to standard settings is that we are dealing with sparse multi-criteria ranking data, not with ratings. Therefore, negative user opinions are hidden from us.

Our assumption is that users give similar endorsements in similar situations, and that we can represent it by a subspace of contexts. In order to enrich the contextual space, we use the review entities with endorsements as an additional dimension to the n-dimensional contextual space. Some technique can be applied to discover *'closely connected'* regions in the contextual space. After finding the contextual regions in the extended $(n+1)$-dimensional cube we eliminate the endorsement dimension in order to derive CUPs which consist solely of contexts. This allows us to map new incoming users to CUPs.

The CUP is represented as an agglomeration of a discovered region. For example, if a clustering technique is applied then a cluster center would be an example of CUP, as we will explain in Section 5.2. In the example in Figure 5 (A.1.2), we discover two CUPs: CP_p and CP_q. The choice of the clustering method depends on the type of application. We detail the application to the Destination Finder in Section 5.2. Next, we discuss how the discovered CUPs can be used for ranking suggested destinations.

4.3 Using Contextual User Profiles

The process of using discovered CUPs can be divided into two main parts, see Figure 5: (A.2) *offline* application of CUPs; and (B) *online* mapping of an incoming user to one of the CUPs.

During the offline stage, the set of CUPs can be used for splitting reviews in order to build a set of contextual rankers $\{R_l\}_{l=1}^{L}$ where L is the number of discovered CUPs. Our assumption is that a set of contextual rankers serves 'better' (more suitable) results than a base ranker R_b which is trained based on all reviews.

During the online stage, an incoming user is mapped to one of the CUPs. A user is represented by a vector of contexts as shown in Figure 5 (B.1). In order to map a user to one of the CUPs, CS_1 or CS_2, we can employ any distance metric D. The user would be assigned to the *'closest'* CUP, which is CP_1 in our example in Figure 5 (B). Then the user is supplied with a contextual ranker R_1 which corresponds to CS_1.

To summarize, we presented a general framework for discovering and using contextual user profiles. In principle, any contextual features can be used, including relatively shallow implicit situational context available in any online context. Also any ratings, reviews or other multi-criteria ranking data can be used, including travel endorsements. In the next section, we apply the framework to the Destination Finder application described in Section 3.

5. CONTEXTUAL TRAVEL RECOMMENDATIONS

In this section we will study our **RQ3: How to apply contextual user profiles for the ranking of travel destinations in a continuous cold start setting?** We present an example how our framework for discovering contextual user profiles (CUPs) from Section 4 can be applied to the Destination Finder. First, we describe the data used for our experimental pipeline in Section 5.1. Second, we use a clustering technique to discover contextual user profiles (CUPs) in Section 5.2. Third, we present in Section 5.3: **(1)** how these CUPs can be used within a ranking technique based on Naive Bayes; **(2)** how the customized rankings are deployed for online user traffic. We use standard clustering and ranking methods, such as k-means and Naive Bayes, which scale well to the volume of data available. These methods are sufficient to answer our main question about the value of context-aware recommendations. Further optimization is left for future work.

5.1 Data

In the offline training stage, we use reviews collected within the year 2014. The final set contains in total 5,138,494 reviews. We derive two types of data from web logs as contextual information:

- **user agent data** which is presented by *four dimensions* such as *'Device Type'* with 5 contextual categories (*mobile, tablet* etc.), *'OS'* with 27 contextual categories (*Windows 8.1, Android, Linux, OS X* etc.), *'Browser'* with 114 contextual categories (*Internet Explorer 6, Firefox 30, Firefox 34, Safari 7* etc.), and *'Traffic Type'* with 16 contextual categories (*web, mobile browser, application* etc.);

- **time data** which is *one dimensional*: the day of the week (*Monday, Tuesday* etc.).

This type of contextual information is available in all typical web logs, and can be used to contextualize the reviews as presented in Figure 5 (A.1.1). In total, the contextual space has 5 dimensions with 397 coordinates. In the online testing stage, we run our experiment on live user traffic for 26,868 users.

5.2 Clustering Contextualized Reviews

We use a clustering technique to discover CUPs as shown in Figure 5 (A.1.2). We apply k-means clustering [25] over the set of contextualized reviews as presented in Equation 7. The number of clusters is selected based on Silhouette validation [47], which results in 20 clusters as the optimal number.

After obtaining the final set of clusters, we eliminate the endorsement dimension by projecting on the contextual space. We analyze the set of contexts that is associated with the clusters in order to derive the set of CUPs. Because of the projection on the contextual space, clusters may overlap in some contextual categories.

The cluster centers represent the set of discovered CUPs. We calculate weights for the coordinates of the cluster centres as the ratio of the *(number of times the coordinate F_{nm} appears within cluster C_i)* divided by the *(number of times the coordinate F_{nm} appears within all clusters)*. This weight w_{ij} (where i is a cluster identifier and j is an identifier of a coordinate F_{nm}) shows how strongly the contextual category F_{nm} is associated with cluster i: The closer w_{ij} to 1, the stronger the association.

We employ a pruning technique over the obtained list of CUPs in order to clean up some obvious noise. If w_{ij} is too small for some contextual category F_{nm}, then this category is distributed widely over all CUPs and it does not enhance our definition of CUP. After trails of experiments we empirically determine a threshold:If F_{nm} has $w_{ij} < 0.2$, we do not include it into the CUPs. For example, sometimes contextual categories such as 'Monday', 'Tuesday' are removed because apparently they do not reflect any 'specific' behavior. By applying this pruning technique we ended up with 17 clusters. We present an example of two pruned CUPs in Table 1, which correspond to intuitions about similar users based on context. It may not be a priori clear why such a cluster provides meaningful context, but the clustering informs us that they have distinct interests and preferences.

Next, we will describe how the discovered CUPs can be applied to destination ranking.

5.3 Using Contextual User Profiles for Destination Ranking

As a primary ranking technique we use a Naive Bayes approach. We will describe its application with an example. Let us consider a user running the searching for 'Beach'. We need to return a ranked list of destinations. For instance, the ranking score for the destination 'Miami' is calculated using the following formula:

$$P(\text{Miami}, \text{Beach}) = P(\text{Miami}) \times P(\text{Beach}|\text{Miami}); \quad (8)$$

where $P(\text{Beach}|\text{Miami})$ is the probability that the destination Miami gets the endorsement 'Beach'. $P(\text{Miami})$ is a prior knowledge about Miami. In the simplistic case the prior would be a ratio of the number of endorsements for Miami to the total number of endorsements in our database.

If a user uses a second endorsement (e.g. + 'Food') the ranking score is calculated in the following way:

$$P(\text{Miami}, \text{Beach}, \text{Food}) = P(\text{Miami}) \times P(\text{Beach}|\text{Miami}) \\ \times P(\text{Food}|\text{Miami}); \quad (9)$$

Table 1: An example of two obtained cluster centers from real data. Cluster i can be characterized as 'users coming from mobile devices' and Cluster $i+1$ as 'users coming from windows-based devices on Fridays and Sundays'.

	Cluster		
i-1	i	i+1	i+2
...	iPhone.OS.7.Chrome	Windows.Phone	...
	iPhone.OS.5.Chrome	Windows.Vista	
	iPhone.OS.6.Chrome	Friday	
	Android.2.2	Sunday	
	Android.2.2.Tablet		
	Android.3.1.Tablet		
	Android.4.0.Tablet		
	Android.4.4.Tablet		
	Android.2.1.Tablet		
	Android.3.0.Tablet		
	Android.4.1		
	Android.4.3.Tablet		

If our user provides n endorsements, Equation 9 becomes a standard Naive Bayes formula.

We split our set of reviews according to the obtained clusters. Then we train a set of contextual rankers using the same approach as described in Equation 9 to obtain the customized rankers $R(C_i)_{i=1}^{17}$. This process can be mapped to the general framework presented in Figure 5 (A.2).

During the online stage, which is shown at general framework work-flow in Figure 5 (B), an incoming user to the Destination Finder is mapped to the closest CUP. As we use only situational context that does not change per session, we only have to assign our user to the nearest cluster once, and there is no need to update the assignment during the session. Then we use a ranker $R(C_i)$ which corresponds to CUP.

As a distance metric we use Euclidean distance, which deals well with the different nature of some of the clusters (e.g., some clusters capture aspects of the day of the week, and others capture aspects of the used devices). More advanced mapping of users as mixtures of CUPs is left to future work, as our main goal in this paper is to determine the impact of contextual ranking.

To summarize, we described the use of the framework for discovering contextual user profiles for the Destination Finder. We contextualized reviews with user agent and time data. Our main goal is to determine the impact of contextual ranking, hence we use standard clustering and ranking methods. Specifically, we use k-means for clustering and Naive Bayes for ranking and we map incoming users to the nearest cluster based on euclidean distance. In the next section, we will present our experimental pipeline which involves online A/B testing at a major travel agent.

6. EXPERIMENTS AND RESULTS

In this section, we will study our **RQ4:** *How effective are contextual profiles for real-world users of the destination finder system in terms of user engagement measures?* To test the effectiveness of contextualization, we perform experiments on users of Booking.com where an instance of the Destination Finder is running.

6.1 Research Methodology

We take advantage of a production A/B testing environment at a major online travel agency. A/B testing randomly splits users to see either the baseline or the new variant version of the website, which allows to measure the impact of the new version directly on real users [33, 55]. As baseline we use a non-contextualized ranker corresponding to the live system. This is an optimized system, trained on a massive volume of traffic, and far superior to standard baselines such as popularity [8].

As our primary *evaluation metric* in the A/B test, we use clicks-per-user and click-through-rate (CTR) [36]. As explained in the motivation, we are dealing with an exploratory task and therefore aim to increase customer engagement. More clicks-per-user and higher CTR are signals that users click more on the suggested destinations and interact more with the system.

6.2 Results

Table 2 shows the results of our A/B test. We see that the contextual ranker does not significantly change conversion compared to the baseline non-contextual ranker, i.e. the probability for a user to click at least once remains the same. Thus, our recommendations do not influence the basic user intent of using the Destination Finder. In contrast, the contextual ranker significantly increases further user engagement after the first click: The CTR increases by absolute 3.7%, and both CTR and clicks-per-user increase dramatically by relative 20% and 23%, respectively. Our contextual recommendations invite users to perform more searches and click on more recommendations, both per search and per user. In total, users are significantly more engaged with the Destination Finder when presented with contextual recommendations.

We achieved this substantial increase in clicks with a simple contextualization using straightforward k-means clustering of reviews and a Naive Bayes ranker. Most computations can be done offline, and only simple calculations have to be performed online. Thus, our model could be trained on large data within reasonable time, and did not negatively impact wallclock and CPU time for the Destination Finder web pages in the online A/B test. This is crucial for a webscale production environment [33].

To summarize, we compared our contextual travel recommendations against the same non-contextualized ranker. This allowed us to compare the effect of contextualization independently of the underlying ranking. This is a hard baseline corresponding to the current live system applied to the exact same data. We observe a dramatic increase in user engagement, with click-through rates and clicks by users increasing by 20%. The simplicity of our contextual models enables us to achieve this engagement without significantly increasing online CPU and memory usage. The experiments clearly demonstrate the value of contextual profiles in a real world application.

7. CONCLUSIONS

This paper investigated the common case in e-commerce websites relying on search and recommendation to satisfy their user's needs, yet standard personalization and recommender systems rely on rich user profiles but the majority of users are new or visit highly infrequently—we face a continuous cold start recommendation problem. We specifically studied this problem in the context of one of the largest travel websites, Booking.com, and its Destination Finder service.

Our first research question was **RQ1:** *How to characterize the continuous cold start problem in travel recommendation?* We introduced and characterized the *Continues Cold Start Problem* (*CoCoS*) that happens when users (*UCoCoS*) or/and items (*ICoCoS*) remain 'cold' for a long time, and can even 'cool down' again after some time due to some external signals.

Our second research question was **RQ2:** *How to define and dis-*

Table 2: Results of the Destination Finder A/B testing based on the number of unique users, searches and clicks. The contextual ranker does not significantly change conversion (probability to click at least once), but significantly increases clicks-per-user and click-though-rate (CTR). Significance is assessed as non-overlapping 95% confidence intervals.

Ranker	Users	Searches	Clicks	Conversion	Clicks/user	CTR
Baseline	13,306	34,463	6,373	21.7±0.7%	0.479±0.012	18.5±0.4%
Contextual	13,562	35,505	7,866	21.3±0.7%	0.580±0.013	22.2±0.4%

cover contextual user profiles from multi-criteria ranking data in an unsupervised setup? We presented a general framework for discovering and using contextual user profiles. Since we work in settings of *CoCoS* clients visit infrequently and have volatile interests, we cannot rely on historical user interactions. Mining situational profiles to which we can map an incoming user is an effective way to deal with data sparsity and changing user interests. In principle, any contextual features can be used, including relatively shallow implicit situational context available in any online context. Also any ratings, reviews or other multi-criteria ranking data can be used, including travel endorsements. Similar endorsement data is being used in a venue recommendation benchmark [59].

Our third research question was *RQ3: How to apply contextual user profiles for the ranking of travel destinations in a continuous cold start setting?* We used the general framework for discovering contextual user profiles for the Destination Finder. As explicit situational context is not available in typical real world application, implicit cues from transaction logs used at scale can capture the essential features of situational context. We contextualized reviews with user agent and time data. Our main goal is to determine the impact of contextual ranking, hence we used standard methods, specifically k-means for clustering and Naive Bayes for ranking. We mapped incoming users to the nearest cluster based on Euclidean distance.

Our fourth research question was *RQ4: How effective are contextual profiles for real-world users of the destination finder system in terms of user engagement measures?* We compared our contextual travel recommendations to a non-contextual ranker. This is a hard baseline corresponding to the current live system. Contextual user profiles can be created offline, resulting in a set of smaller models compared to the single, huge, non-contextual model, making contextual ranking available with negligible CPU or memory footprint. We observed an increase in user engagement, with higher click-through rates (20%) and higher clicks per user (21%).

Our general conclusion is that our contextual ranking approach shows a dramatic increase in user engagement over a non-contextual baseline, clearly demonstrating the value of contextualized profiles in a real world application that suffers from *CoCoS*. We focused on an e-commerce setting, applicable to millions of online companies, where the continuous cold start is the rule rather than the exception. But also in settings such as the internet search engines where interactions are frequent and rich profiles are typically available, our approach has large potential value. The problem of fast changing content is well-known [19]. Perhaps the fraction of new users is small, yet they may be important enough to warrant extra effort, think of new users considering a search engine switch [62].

Our future work is to further investigate the following directions. First, we plan to extend the contextual space, for example using the geographical location of the user. However, this is not straightforward since simple splitting using some ontological knowledge, e.g. country, can lead to very skewed distributions of traffic within the contextual features and fails to capture deeper relations in the data. More generally, we plan to look into unsupervised techniques for the context discovery, over a wider range of contextual conditions

including aspects of the session at hand. Second, it is promising to extend our method of mapping incoming users to one of the discovered CUPs to a 'fuzzy' mapping in which a user can be assigned to two or more CUPs. This will allow to serve a personalized ranking based on the resulting mixture weights in the model, while still maintaining online efficiency. Third, we will look into possibilities of more efficient and accurate CUPs discovery techniques, looking also in adaptive models that take into account long term trends such as seasonal differences.

REFERENCES

[1] G. Adomavicius and Y. Kwon. New recommendation techniques for multicriteria rating systems. *EXPERT*, 22(3): 48–55, 2007.

[2] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *TKDE*, 17:734–749, 2005.

[3] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In *Recommender Systems Handbook*, pages 217–253, 2011.

[4] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *TOIS*, 23(1):103–145, 2005.

[5] G. Adomavicius, N. Manouselis, and Y. Kwon. *Multi-Criteria Recommender Systems*, volume 768-803. Recommender Systems Handbook, Springer, 2011.

[6] D. Agarwal and B.-C. Chen. Regression-based latent factor models. In *KDD*, pages 19–28, 2009.

[7] D. Agarwal and B.-C. Chen. flda: matrix factorization through latent dirichlet allocation. In *WSDM*, pages 91–100, 2010.

[8] Anonymous. Citation withhold to preserve anonymity. In *SIGIR*, 2015.

[9] L. Baltrunas and F. Ricci. Context-based splitting of item ratings in collaborative filtering. In *RecSys*, pages 245–248, 2009.

[10] L. Baltrunas and F. Ricci. Experimental evaluation of context-dependent collaborative filtering using item splitting. *User Modeling and User-Adapted Interaction*, 24:7–34, 2014.

[11] L. Baltrunas and F. Ricci. Context-dependent items generation in collaborative filtering. In *CARS*, 2009.

[12] T. Bao, H. Cao, E. Chen, J. Tian, and H. Xiong. An unsupervised approach to modeling personalized contexts of mobile users. *Knowl. Inf. Syst.*, 31(2):345–370, 2012.

[13] F. Belém, R. L. T. Santos, J. M. Almeida, and M. A. Gonçalves. Topic diversity in tag recommendation. In *RecSys*, pages 141–148, 2013.

[14] L. Bernardi, J. Kamps, J. Kiseleva, and M. J. I. Müller. The continuous cold-start problem in e-commerce recommender systems. In *CBRecSys*, pages 30–33, 2015.

[15] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3:993–1022, 2003.

[16] H. Cao, T. Bao, Q. Yang, E. Chen, and J. Tian. An effective approach for mining mobile user habits. In *CIKM*, pages 1677–1680, 2010.

[17] Census Bureau. E-stats: Measuring the electronic economy. https://www.census.gov/econ/estats/, 2015.

[18] F. C. T. Chua, R. J. Oentaryo, and E.-P. Lim. Modeling temporal adoptions using dynamic matrix factorization. In *ICDM*, pages 91–100, 2013.

[19] A. Dong, Y. Chang, Z. Zheng, G. Mishne, J. Bai, R. Zhang, K. Buchner, C. Liao, and F. Diaz. Towards recency ranking in web search. In *WSDM*, pages 11–20, 2010.

[20] N. Hariri, B. Mobasher, and R. D. Burke. Context-aware music recommendation based on latent topic sequential patterns. In *RecSys*, pages 131–138, 2012.

[21] N. Hariri, B. Mobasher, and R. D. Burke. Query-driven context aware recommendation. In *RecSys*, pages 9–16, 2013.

[22] A. Hawalah and M. Fasli. Utilizing contextual ontological user profiles for personalized recommendations. *Expert Syst. Appl. (ESWA)*, 41(10):4777–4797, 2014.

[23] Y.-C. Ho, Y.-T. Chiang, and J. Hsu Yung-Jen. Who likes it more?: mining worth-recommending items from long tails by modeling relative preference. In *WSDM*, pages 253–262, 2014.

[24] D. J. Hu, R. Hall, and J. Attenberg. Style in the long tail: Discovering unique interests with latent variable models in large scale social e-commerce. In *KDD*, pages 1640–1649, 2014.

[25] R. Jancey. Multidimensional group analysis. *Australian Journal of Botany*, 14(1):127–130, 1966.

[26] D. Jannach, Z. Karakaya, and F. Gedikli. Accuracy improvements for multi-criteria recommender system. In *EC*, pages 674–689, 2012.

[27] K. Kapoor, K. Subbian, J. Srivastava, and P. Schrater. Just in time recommendations: Modeling the dynamics of boredom in activity streams. In *WSDM*, pages 233–242, 2015.

[28] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *RecSys*, pages 79–86, 2010.

[29] J. Kiseleva, H. T. Lam, M. Pechenizkiy, and T. Calders. Predicting current user intent with contextual markov models. In *ICDM Workshops*, 2013.

[30] J. Kiseleva, H. T. Lam, M. Pechenizkiy, and T. Calders. Discovering temporal hidden contexts in web sessions for user trail prediction. In *TempWeb@WWW'2013*, pages 1067–1074, 2013.

[31] J. Kiseleva, M. J. I. Müller, L. Bernardi, C. Davis, I. Kovacek, M. S. Einarsen, J. Kamps, A. Tuzhilin, and D. Hiemstra. Where to go on your next trip?: Optimizing travel destinations based on user preferences. In *SIGIR (Industry Track)*, pages 1097–1100, 2015.

[32] D. Kluver and J. A. Konstan. Evaluating recommender behavior for new users. In *RecSys*, pages 121–128, 2014.

[33] R. Kohavi, A. Deng, R. Longbotham, and Y. Xu. Seven rules of thumb for web site experimenters. In *KDD*, pages 1857–1866, 2014.

[34] Y. Koren. Collaborative filtering with temporal dynamics. *Communications of the ACM*, 53(4):89–97, 2010.

[35] K. Lakiotaki, N. F. Matsatsinis, and A. Tsoukias. Multicriteria user modeling in recommender systems. *IEEE Intelligent System*, 26(2):64–76, 2011.

[36] M. Lalmas, H. O'Brien, and E. Yom-Tov. Measuring user engagement. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 6(4):1–132, 2014.

[37] L. Liu, N. Mehandjiev, and D.-L. Xu. Multi-criteria service recommendation based on user criteria preferences. In *RecSys*, pages 77–84, 2011.

[38] H. Ma, H. Cao, Q. Yang, E. Chen, and J. Tian. A habit mining approach for discovering similar mobile users. In *WWW*, pages 231–240, 2012.

[39] G. D. Montanez, R. W. White, and X. Huang. Cross-device search. In *CIKM*, pages 1669–1678, 2014.

[40] A. Noulas and M. S. Einarsen. User engagement through topic modelling in travel. In *Second Workshop on User Engagement Optimization*, 2014.

[41] C. Palmisano, A. Tuzhilin, and M. Gorgoglione. Using context to improve predictive modeling of customers in personalization applications. *TKDE*, 2008, November.

[42] U. Panniello, A. Tuzhilin, M. Gorgoglione, C. Palmisano, and A. Pedone. Experimental comparison of pre- vs. post-filtering approaches in context-aware recommender systems. In *RecSys*, pages 265–268, 2009.

[43] A. M. Rashid, I. Albert, D. Cosley, S. K. Lam, S. M. McNee, J. A. Konstan, and J. Riedl. Getting to know you: Learning new user preferences in recommender systems. In *IUI*, pages 127–134, 2002.

[44] S. Rendle. Factorization machines. In *ICDM*, pages 995–1000, 2009.

[45] S. Rendle, Z. Gantner, C. Freudenthaler, and L. Schmidt-Thieme. Fast context-aware recommendations with factorization machines. In *SIGIR*, pages 635–644, 2011.

[46] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. Grouplens: An open architecture for collaborative filtering of netnews. In *CSCW*, pages 175–186, 1994.

[47] P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *JCAM*, 20: 53–65, 1987.

[48] M. Saveski and A. Mantrach. Item cold-start recommendations: Learning local collective embeddings. In *RecSys*, pages 89–96, 2014.

[49] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock. Methods and metrics for cold-start recommendations. In *SIGIR*, pages 253–260, 2002.

[50] S. Sedhain, S. Sanner, D. Braziunas, L. Xie, and J. Christensen. Social collaborative filtering for cold-start recommendations. In *RecSys*, pages 345–348, 2014.

[51] U. Shardanand and P. Maes. Social information filtering: Algorithms for automating 'word of mouth'. In *CHI*, pages 210–217, 1995.

[52] Y. Shi, A. Karatzoglou, L. Baltrunas, M. Larson, A. Hanjalic, and N. Oliver. Tfmap: optimizing map for top-n context-aware recommendation. In *SIGIR*, pages 155–164, 2012.

[53] Y. Shi, A. Karatzoglou, L. Baltrunas, M. Larson, and A. Hanjalic. Cars2: Learning context-aware representations for context-aware recommendations. In *CIKM*, pages 291–300, 2014.

[54] J. Z. Sun, K. R. Varshney, and K. Subbian. Dynamic matrix factorization: A state space approach. In *ICASSP*, pages 1897–1900, 2012.

[55] D. Tang, A. Agarwal, D. O'Brien, and M. Meyer. Overlapping experiment infrastructure: More, better, faster experimentation. In *KDD*, pages 17–26, 2010.

[56] J. Tang, H. Gao, X. Hu, and H. Liu. Context-aware review helpfulness rating prediction. In *RecSys*, pages 1–8, 2013.

[57] L. Tang, Y. Jiang, L. Li, and T. Li. Ensemble contextual bandits for personalized recommendation. In *RecSys*, pages 73–80, 2014.

[58] M. Tavakol and U. Brefeld. Factored mdps for detecting topics of user sessions. In *RecSys*, pages 33–40, 2014.

[59] TREC. Text retrieval conference: Contextual suggestion track. https://sites.google.com/site/treccontext/, 2015.

[60] P. Turney. The identification of context-sensitive features: A formal definition of context for concept learning. *CoRR*, 2002.

[61] B. Vargas-Govea, J. G. González-Serna, and R. Ponce-Medellín. Effects of relevant contextual features in

the performance of a restaurant recommender system. In *CARS*, 2011.

[62] R. W. White and S. T. Dumais. Characterizing and predicting search engine switching behavior. In *CIKM*, pages 87–96, 2009.

[63] H. Zhu, H. Cao, E. Chen, H. Xiong, and J. Tian. Exploiting enriched contextual information for mobile app classification. In *CIKM*, pages 1617–1621, 2012.