Knowledge-Based Systems 104 (2016) 165–178

Contents lists available at ScienceDirect

Knowledge-Based Systems

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

Towards latent context-aware recommendation systems

Moshe Unger*, Ariel Bar, Bracha Shapira, Lior Rokach

Ben-Gurion University of the Negev and Telekom Innovation Laboratories at BGU, P.O box 84105, Beer Sheva, Israel

A R T I C L E I N F O

Article history:
Received 22 October 2015
Revised 25 March 2016
Accepted 20 April 2016
Available online 27 April 2016

Keywords:
Recommendation
Recommender systems
Context-aware recommender systems
Context
Matrix factorization
Deep learning

A B S T R A C T

The emergence and penetration of smart mobile devices has given rise to the development of context-aware systems that utilize sensors to collect available data about users in order to improve various user services. Recently, the use of context-aware recommender systems (CARS) aimed at recommending items to users has expanded, particularly those that consider user context. Adding context to recommendation systems is challenging, because the addition of various environmental contexts to the recommendation process results in the expansion of its dimensionality, and thus increases sparsity. Therefore, existing CARS tend to incorporate a small set of pre-defined explicit contexts which do not necessary represent user context or reflect the optimal set of features for the recommendation process. We suggest a novel approach centered on representing environmental features as low dimensional unsupervised latent contexts. We extract data from a rich set of mobile sensors in order to infer unexplored user contexts in an unsupervised manner. The latent contexts are hidden context patterns modeled as numeric vectors which are efficiently extracted from raw sensor data. The latent contexts are automatically learned for each user utilizing unsupervised deep learning techniques and PCA on the data collected from the user's mobile phone. Integrating the data extracted from high dimensional sensors into a new latent context-aware recommendation algorithm results in up to a 20% increase in recommendation accuracy.

1. Introduction

Context-aware computing aims at tailoring services to the user's circumstances and surroundings. Context-aware computing was first mentioned by Schilit et al. [33] as "software that adapts according to its location of use, the collection of nearby people and objects, as well as changes to those objects over time." A system is considered context-aware if it can extract, interpret, and use contextual information to adapt its functionality to users' immediate contexts [6]. As "context" refers to aspects of users' immediate locations, activities or circumstances, users in pervasive environments need services and recommendations that suit their contexts and preferences.

Context-aware recommender systems (CARS) aim at recommending items while considering the user's context. Adomavicius et al. [1] suggested three main approaches to incorporating context into recommender systems: pre-filtering, post-filtering, and contextual modeling. While the pre- and post- approaches filter the recommended items before or after the recommendation list is computed and do not incorporate context into the recommendation model, the contextual modeling approach is preferred, since it applies the contextual information directly to the modeling technique as part of the rating estimation. The contextual modeling approach is challenging, because the addition of various types of environmental contexts to the recommendation process results in the expansion of its dimensionality (i.e., users, items, and contexts). This entails that each rating of an item must relate to each category of context, thus sparsity is increased as training is required for every triple of <user, item, context>. Labeled data regarding users' contexts and their preferences in each context must then be collected in order to train the system for supervised learning. This is, of course, almost infeasible, and therefore presents a sparsity challenge.

In order to resolve this problem, the dimensionality of the context representation must be reduced. Baltrunas [3,5] and others [9,12,32,36,40] modeled situations and circumstances as explicit specific contexts in order to limit the dimensionality space. The specific contexts describe the circumstances of the information collection, e.g., weather conditions ("sunny," "cloudy," "raining," etc.) or precise location conditions ("at home," "at work," etc.). While in this setting the set of contexts is both small enough to handle and sufficient to prevent sparsity, it may not take into consideration other important environmental features and does not necessarily represent an optimal set of features for the recommendation process.

We propose a novel approach that uses a rich set of mobile sensors for the recommendation process (e.g., Wi-Fi networks,

* Corresponding author.
E-mail address: mosheun@post.bgu.ac.il (M. Unger).

http://dx.doi.org/10.1016/j.knosys.2016.04.020

accelerometers, light, microphones, etc.). We incorporate all of these environmental features into the recommendation engine in order to consider various types of unsupervised contexts. In order to address the sparsity challenge, we automatically selected a small set of the best features to handle. Our contribution is twofold:

First, we suggest several methods for reducing the dimensionality space by extracting latent context from the data collected from mobile device sensors. In our approach, latent contexts are comprised of unsupervised hidden context patterns which are modeled as numeric vectors and extracted efficiently from the raw sensor data. The latent contexts are learned automatically by applying unsupervised deep learning techniques and principal component analysis (PCA) to the collected data.

Explicit context can be better explained by human experts and users than latent context, since it describes known user situations. However, the motivation for using latent context stems from privacy, data availability, and usability considerations. Using explicit context for a recommendation service may raise privacy issues, since the exact context of the user in known to the service, which is not the case for latent context. While latent context can be obtained automatically by applying unsupervised learning techniques on available raw data (e.g., mobile sensors), obtaining explicit context is a resource-demanding task and may interfere with the user' activities.

Secondly, we show how to utilize these latent contexts and describe a novel recommendation technique that uses latent contexts and improves the accuracy of state-of-the-art CARS. In addition, since latent features represent the essence of the high dimensional sensors, and usually improve models' accuracy [24], we suggest a hybrid model which adds latent contextual features to the set of explicit context features.

Extensive experiments to evaluate the suggested models were conducted over a period of four weeks with a specially developed Android application which displays POI (points of interest) recommendations to users in order to evaluate the new algorithms. The participants provided positive and negative feedback about the recommendations, while we recorded a rich set of sensor data from their mobile devices. During the experiments, we also examined the effect of explicit and latent contexts on the accuracy of context-aware recommender systems. The experiments show that all of our suggested models outperform traditional CARS in terms of RMSE (root mean square error) on the test set.

The rest of this paper is structured as follows: Section 2 describes related work, and Section 3 describes an updated recommendation rating model, which includes latent contexts. Furthermore, we describe methods for extracting these contexts, utilizing deep learning and PCA. Section 4 presents our new recommendation technique that is based on explicit and latent contexts. Section 5 presents the setup of the field experiments and discusses the evaluation of the collected data. Finally, in Section 6 we discuss the results and outline plans for future research.

2. Related work

2.1. CARS (Context-aware recommender systems)

The area of context-aware recommender systems (CARS) deals with modeling and predicting user tastes and preferences by incorporating available contextual information into the recommendation process as explicit additional categories of data[1]. In CARS, these long-term preferences and user interests are usually expressed as ratings and are modeled as the function of not only items and users, but also of context. In other words, ratings are defined with the rating function: R: *User* $\times$ *Item* $\times$ *Context* $\rightarrow$ *Rating*. In this approach, there is a pre-defined finite set of contextual types in a given application and each of these types has a well-defined structure.

There are several definitions for contextual information and the definitions' form of representation for CARS. One option is to define the contextual information as a set of contextual dimensions K in which each contextual dimension is defined by a set of n attributes $K = (K^1, \ldots, K^n)$ having a hierarchical structure and capturing a particular type of context—such as time or location. An additional definition for the contextual information refers to its dimensions $D^1, D^2, \ldots, D^n$, as a multi-dimensional model. These dimensions are handled as a Cartesian product $S = D_1 \times D_2 \times \ldots D_n$. Moreover, let *Rating* be a rating domain representing the ordered set of all possible rating values. Then, the rating function is defined over the space $D_1 \times D_2 \times \ldots D_n$ as: $R : D_1 \times D_2 \times \ldots D_n \;\rightarrow Rating$.

In Adomavicius and Tuzhilin [1], three main paradigms for incorporating contextual information into the recommendation process are presented: (1) contextual pre-filtering, where context is used for selecting the relevant set of rating data before computing predictions; (2) contextual post-filtering, where the resulting set of recommendations is adjusted for each user using the predictions generated by a traditional model; and (3) contextual modeling, in which contextual information is directly incorporated into the prediction model as part of the rating estimation.

Previous works [4,8] used a pre-filtering approach in their context-aware recommendation process. In their work, they used only explicit, pre-defined contexts, which could be less accurate in expressing users' behavior, since a more accurate context prediction could be achieved by utilizing different contexts [20]. Other works [40,25] used a post-filtering approach in order to re-rank the recommendations. Xu et al. [40] proposed a travel location recommendation method based on topic distribution of travel histories in a given context (season and weather). They utilized user travel histories in order to build a user-user similarity model and used post-filtering in order to filter out recommendations that do not meet the contextual constraints. However, considering contextual conditions in the prediction model yields better results than typical RS methods that do not consider context [1]. Moreover, using only season and weather contexts does not necessarily represent an optimal set of features for the recommendation process.

In an experimental comparison of traditional pre-filtering versus post-filtering approaches, presented by Panniello et al. [29], post-filtering approaches regularly outperform pre-filtering methodologies. In order to decide which approach should be used in a particular recommendation application, various post-filtering methods should be compared with the pre-filtering approach, resulting a laborious and time consuming task. While such contextual approaches may work in practice, they have the shortcoming that all steps in the process need supervision and fine-tuning. The main limitation of the pre- and post-filtering techniques comes from the difficulty to obtain ratings in all of the possible explicit contextual situations in order to build a robust and contextualized rating prediction model. Our approach of latent contexts automatically incorporates various types of environmental features, such as weather, location, and time into the recommendation model. Since extracting those contexts is done with unsupervised learning techniques, it also reveals the connections between the different features.

An alternative approach to the pro and pre- filtering is to incorporate the context into the prediction model, a method referred to as the contextual approach [1]. Baltrunas et al. [4,5] suggested a context based splitting approach in which ratings of certain items are split according to the value of an item-dependent contextual condition. Several metrics were suggested for the splitting criteria, and the info-gain proved to be the best criteria for this task. This pre-filtering technique is limited to a single binary context and thus cannot model relations among several contexts.