A Holistic Approach to Proximity Marketing

Dimitris Poulopoulos and Athina Kalampogia
QIVOS, Athens, Greece

Keywords: Proximity Marketing, Recommendation Systems.

Abstract: Proximity marketing, as a promotional technique, can benefit shopping centres and malls in terms of revenue, as well as customer loyalty, by analysing the customers' data and using their profiles to better address their needs and target advertising and promotional campaigns. To this end, retailers exploit cellular technology to send marketing messages to users' mobile devices, that are near a specific area of one of its stores. With more than six billion mobile phones in the hands of consumers today, every consumer with a smartphone is potentially susceptible to a proximity marketing campaign.

1 INTRODUCTION

IKEA is the world's largest furniture retailer, with a presence in over 45 countries. Every day, thousands of customers roam the stores' exhibitions considering new purchases for their homes or workplaces. In this work, we present the description and design of the proximity marketing solution implemented for the Greek branch of the organisation.

IKEA has several ways of gathering and storing data derived from its customers. Those data are later analysed to extract knowledge, that will assist the brand address the customers' expectations. QIVOS has developed a loyalty scheme for IKEA, where many of its customers are registered. Through this program, consumers are equipped with a loyalty member plastic card and collect points upon their purchases. By the time a customer collects a specific amount of points, IKEA can offer specific discounts to the customer's receipt. However, through this loyalty scheme, IKEA is not only offering discounts to its customers but also gathering data about the customers' consumer behaviour, that is further analysed and used for marketing reasons.

Currently, the loyalty program is reactive. In CloudDBAppliance, the loyalty scheme is transforming into a proactive, mobile-first loyalty program. Moreover, QIVOS, in collaboration with IKEA, installed beacon devices in different areas in one of the brand's stores to collect information in real-time about the customers' specific location.

To undertake this transformation, offer innovative services to the customers, facilitate their purchasing decisions and exploit the new sources of data in real-time, QIVOS has invested in implementing a novel incremental recommendation engine.

Chris Anderson, in his 2004 article entitled "The Long Tail", said that we are leaving the age of information and entering the age of recommendation (Anderson, 2006). Unless we have a way to filter the information overload that we absorb every day and retain only what is important to us, data reduce to noise.

Moreover, we seem to diverge from the notion of search and embrace that of discovery. The difference is, that when searching, a user is actively looking for something. Discovery suggests that something the user did not know existed or even did not know how to ask for it, finds him (O'Brien, 2006). A recommender or recommendation system aims to predict the behavioural patterns of the user, his preferences or dislikes, and utterly provide personalised recommendations on items the user would be likely to interact.

When explicit feedback from the users is available, the system tries to solve a surrogate problem, where we view ratings as a proxy to preference. The problem then is usually solved as a regression problem; thus, the algorithm is trying to predict a real-valued rating to complete the missing values in the user-item matrix.

However, most of the times - like in our case - we do not have access to explicit user feedback. In such scenarios, the system works passively in the background trying to collect significant features. Those features can be used to implicitly discover user

Poulopoulos, D. and Kalampogia, A.
A Holistic Approach to Proximity Marketing.
DOI: 10.5220/0007861106290634
In *Proceedings of the 9th International Conference on Cloud Computing and Services Science (CLOSER 2019)*, pages 629-634
ISBN: 978-989-758-365-0

habits, preferences and behaviour, with some level of confidence. In the case of binary implicit feedback, we simplify the user-item matrix to a Boolean matrix, where "true" values correspond to positive user-item interaction, and "false" values indicate no interaction. In any case, implicit feedback has consequences to both algorithmic design and evaluation measures.

Our approach covers the following challenging requirements:

- *Real-time:* Recommendations should arrive in real-time, in less than 30 seconds, while the ideal goal is under 10 seconds. This requirement derives from the fact that customers that pass a specific location rarely return to pick up a late recommended product. Thus, the time requirement drives us to either scale up complex solutions or turn to more straightforward techniques.
- *Incremental Learning:* For a real-time scenario, an ideal recommendation engine updates its parameters incrementally and adapts to the ever-changing needs and behaviours of customers.
- *Proximity:* IKEA exhibitions define a precise path that every customer follows. For example, the customer passes through the living room sector, then enters the bathroom region, and finally, ends up exploring kitchen products. Thus, customers should be able to quickly match what they already have in their basket with products from the same category.
- *Up-selling:* One of the requirements is to detect related products to what a customer has already in the basket, and recommend those that, although a bit pricier, present an opportunity due to some in-store special offer, or a stock policy.

This work is structured as follows: We present a brief description of the use-case in Section 2. Section 3 presents a bird's eye view of the system. Section 4 provides the requirements, and Section 5 concludes this paper.

2 USE CASE DESCRIPTION

Shopping centres and malls can benefit (i.e. in terms of revenue) from analysing its customers' data, and using the concept of proximity marketing to send marketing messages to users' that are near a specific area of interest. With more than six billion mobile phones in the hands of consumers today[1] every consumer with a smartphone is potentially susceptible to a proximity marketing campaign.

2.1 Overview

IKEA is a major retailer, that participates in CloudDBAppliance project as a partner, with the interest of trying a cloud-based proximity marketing solution, developed by QIVOS, a company specialised in cloud customer solutions for large enterprises, especially in the retail sector.

To this end, the proposed scenario takes place in IKEA stores with the aim of analysing customers' data deriving from multiple sources (e.g. Bluetooth enabled devices), with an aim to offer personalised content, encourage specific behaviours, enhance the shopping experience, facilitate the purchase decision and predict the needs of its customers. The goal is to have a platform able to get insights from all IKEA stores, segment the clients, and produce personalised offers based on the information extracted globally from all stores, in real-time.

To achieve this, a recommendation engine will leverage the vast amounts of data, that will surge from real-time sensors and historical customer profiles. Moreover, to adapt to such a dynamic system, it is essential to refer to the notion of incremental learning and develop the engine accordingly.

2.2 Problem Statement

In our scenario, user preferences change frequently, and new data continuously arrive in a real-time manner. A recommender system should ideally adapt to these changes as they happen, modifying its model to always speak for the current status, while requiring a single pass through the data. This is the idea of incremental learning.

While most recommender systems utilise some variation of collaborative filtering, they suffer, most of the time, from scalability and efficiency problems, as the computations needed to grow polynomially with the number of users and items in a database.

To address this problem, researchers have proposed several approximation methods; Breese et al. (Breese, 1998) and Ungar et al. (Ungar, 1998) employ Bayesian network and clustering approaches.

[1] Proximity Marketing| What Is Proximity Marketing?" *Marketing Schools*. N.p., n.d. Web. 20 Mar. 2019 http://www.marketing-schools.org/types-of-marketing/proximity-marketing.html.

In (Sarwar, 2000, Deerwester, 1990) Sarwar et al. and Deerwester et al. perform a dimensionality reduction technique for the user-item interaction matrix, by applying folding in Singular Value Decomposition (SVD). Other researchers also turned their attention to data reduction by removing irrelevant and redundant elements (Zeng, 2004, Yu, 2002), and content boosted collaborative filtering methods, where they score each item relevancy by partitioning the item space according to categories \cite{popescul2001probabilistic}. Finally, greedy algorithms that randomly sample users, or discard popular and unpopular items have also been proposed.

Although such approximation methods improve run-time performance, they do it at the expense of accuracy. This is, for example, the case of clustering based methods, and although different optimizations have been proposed using several fine-grained segments (Jung, 2001), the cost of computation approaches this of classic collaborative filtering approaches. Moreover, there are other disadvantages that one might consider. For example, Bayesian networks work fine in environments where user behaviour changes slowly with respect to the time required to build the model, but they are not practical in environments where changes happen rapidly. Considering all this, there seems to be a trade of between recommendation quality and performance efficiency. This is the problem that incremental learning tries to alleviate, by composing highly scalable algorithms, that have much faster run-times with no accuracy degradation.

2.3 Data Acquisition

IKEA has several ways of gathering and storing data from its customers. To begin with, in that specific scenario, IKEA customers will be equipped through their smartphones with a mobile application, in which customers will take note of the purchases that they are willing to make. Thus, this “wish-list” will be digitised, allowing IKEA to know exactly what its customers will probably buy, and what their itinerary will be while in store.

What is more, most of IKEA customers are registered to the IKEA loyalty program, through which they are equipped with members’ plastic cards and collect points upon their purchases. By the time a customer collects a specific amount of points, these points can be transformed into discount vouchers to the customer’s receipt. However, through this loyalty program, IKEA is not only offering discounts to its customers but also gathering data about the customers’ consumer behaviour, that is further analysed and used for marketing reasons. Currently, the loyalty program is reactive. In CloudDBAppliance, the loyalty program will be transformed into a proactive, mobile loyalty program, and to enrich customers’ data, will leverage beacons that will be installed in different areas in the IKEA stores, to collect information in real-time about the customers’ specific location.

2.4 Methodology

We have developed several ways of collecting and managing customers’ data, to assist specially designed mechanisms to process and analyse this information in real-time, so as to predict customers’ needs and suggest additional purchases, offers and coupons.

More specifically, by the time a customer with a balanced consumer behaviour will add to her/his shopping cart an item, real-time analytics will be performed on:

- Data concerning the customer’s consumer behaviour according to his/her previous purchases
- Other customers’ consumer behaviour according to what kind of similar items they have purchased in combination with that product

Thus, similar consumer patterns will be identified and forecasted between customers with similar behaviours, and suggestions will be provided to the customers’ device, through predictive real-time analytical mechanisms, about products that other customers purchased along with the selected product.

Furthermore, through the installed beacon devices, IKEA will track, in real-time, the “geographic location” of all its customers. As a result, real-time analytical algorithms will be executed including data deriving from:

- A customer’s current location into the store
- A customer’s current shopping cart
- Past consumer behaviour of customers’ that purchased same or related products

In that scenario, real-time predictions will be made utilizing the aforementioned vast amounts of data, in such a way that customers will be suggested through their devices about items that match/correspond to the items that they have already added to their shopping cart, and are located just a few steps away from them in the store.

3 SYSTEM OVERVIEW

Figure 1 depicts a high-level overview of the IKEA proximity marketing system.

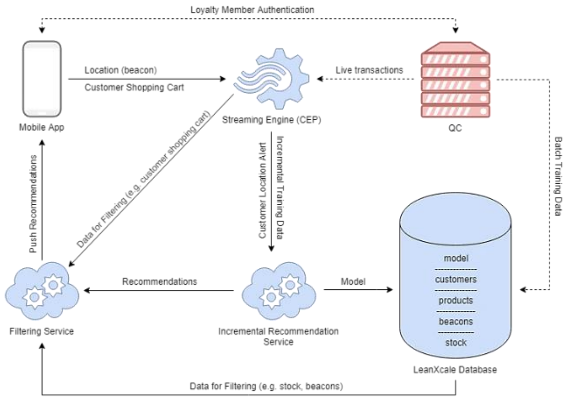

Figure 1: IKEA proximity marketing conceptual diagram. Data are retrieved by the operational database and fed to the incremental recommendation service. In real-time the customer's location and shopping cart are fed as extra input to the model and post-filtering service, to create an incremental context-aware recommendation engine.

Section 3.1 presents the conceptual architecture depicted in Figure 1 and the initial architecture of the incremental recommendation service.

3.1 Conceptual Architecture

The process starts with the IKEA mobile app, which sends back the beacon identifier as well as the shopping cart for each update (insert, update, delete item). This data, along with live transactions coming from QIVOS Cloud (QC), feeds the Streaming Engine.

The Streaming Engine, in turn, promotes live transactions to the recommendation engine, which uses them for incremental training. It also promotes the user's location and shopping cart in the filtering service to clean up irrelevant suggestions.

The incremental recommendation service, in addition to live data, initially receives a batch of historical training data, for bootstrapping, from the operational database.

Finally, the request for recommendations comes from the streaming engine, in the form of an alert, whenever a user changes position in the store. Thus, by deploying a continuous query on the beacon streaming data, we can identify the position of the customer at any point in time. Whenever we have something of significance, for example, the customer changes a showroom in the exhibition (e.g. moves from kitchens to bedrooms), an alert is triggered, and new recommendations are produced, concerning the new context.

3.2 Recommendation Engine Architecture

Initially, we feed the customer's historical data, stored in the database, in the pre-processing unit. The unit must shape them in a form that the recommendation algorithm accepts. Moreover, it logically transforms the data so that the algorithm can discover intricate patterns behind customers' purchases and behaviour.

The training of the algorithm is done incrementally, in an online manner. In real-time, we also feed the customer's location and shopping cart into the model. By the time a customer adds something to the shopping cart or moves to a new in-store location, the system produces candidate items.

The primary job of the algorithm is to create a list of items, that might be of interest to the user and are located nearby. This list is passed to a ranking algorithm, that sorts these candidate items to create personalised recommendations.

The algorithm is inspired by the work done in language models (Mikolov, 2013, Collobert, 2008, Bengio, 2003), as well as research done on incremental collaborative filtering approaches (Papagelis, 2005, Vinagre, 2014, Miranda, 2008) and is summarised in Figure 2. Its job is to find the correlations between the items that customers choose together, as well as the reasoning behind these purchases (e.g., same colour, brand, style, etc.). The research that has explicitly done on the recommendation algorithm is a subject of a future publication.

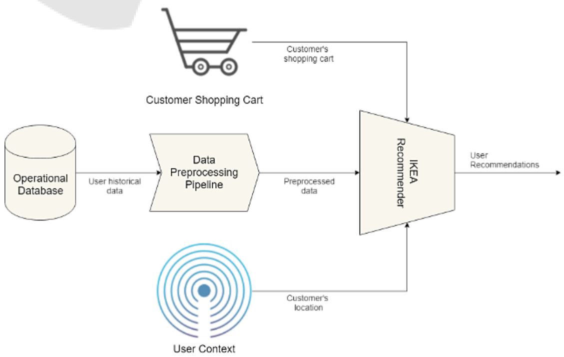

Figure 2: Recommendation engine conceptual diagram. Data are retrieved by the operational database and fed to the data pre-processing unit. In real-time the customer's location and shopping cart are fed as extra input to the model, to create an incremental context-aware recommendation engine.

During training, we make use of offline metrics like precision, recall, and ranking loss, but the real value of a recommender is not only to predict the held-out data in a test set but also to discover additional items that might be of interest to the customers. This is to balance the exploration-exploitation trade-off and offer a diverse set of recommendations, suggestion products that customers are unaware of their existence. Consequently, we can only draw a safe conclusion using specifically designed A/B tests in production.

4 REQUIREMENTS

Most of the requirements of the system have to do with the real-time nature of the business logic.

Ideally, recommendations should arrive in under 10 seconds, considering the huge traffic that could occur in a store during rush hour. This requirement derives from the fact that customers that pass a specific location rarely return to pick up a late recommended product. Thus, the time requirement drives us to either scale up complex solutions or turn to more straightforward techniques.

Our approach is to build an incremental recommendation service, that after a bootstrap training session, adapts dynamically to the ever-changing customer needs and behaviours. The training of the system happens in an online manner, where each new customer-item interaction is integrated directly into the model.

Moreover, IKEA exhibitions define a precise path that every customer follows. For example, the customer passes through the living room sector, then enters the bathroom region, and finally, ends up exploring kitchen products. Thus, customers should be able to quickly match what they already have in their basket with products from the same category.

This has also some implications in the post-filtering service. This means that we need to consider items that the customers have already bought, items that are in the customers' current shopping carts and items that are out of stock.

Finally, there is up-selling. Thus, one of the requirements is to detect related products to what a customer has already in the basket, and recommend those that, although a bit pricier, present an opportunity due to some in-store special offer, or a stock policy. This could drive the customer's average basket up and have a huge impact on the IKEA's revenue.

5 CONCLUSIONS

IKEA is the world's largest furniture retailer, with a presence in over 45 countries. Every day, thousands of customers roam the stores' exhibitions, consider new purchases for their homes or workplaces.

In this work, we presented the description and design of the proximity marketing solution implemented for the Greek branch of the organisation, which is one of the three use-cases of the CloudDBAppliance project. The primary objective is the realisation of a cloud-based proximity marketing solution, developed by QIVOS, tailored to large enterprises, especially in the retail sector. The implementation depends on heterogenous data from various sources, aggregated so as to be fed into an incremental recommendation engine, that adapts in real-time to the changing consumer behaviour. Future work is aiming on parallelizing the algorithm, making it more scalable, to address the needs of ever growing datasets.

ACKNOWLEDGEMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number 732051, CloudDBAppliance project.

REFERENCES

Alexandrin Popescul, David M Pennock, and Steve Lawrence, 2001. Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. In Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence. Morgan Kaufmann Publishers Inc., pp. 437–444.

Badrul Sarwar et al., 2000. Application of dimensionality reduction in recommender system-a case study. Tech. rep. Minnesota Univ Minneapolis Dept. of Computer Science, 2000.

Catarina Miranda and Alípio M Jorge, 2008. Incremental collaborative filtering for binary ratings. In Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology-Volume 01, pp. 389–392.

Chris Anderson, 2006, The long tail: Why the future of business is selling more for less. Hyperion.

Chun Zeng et al., 2004. Similarity measure and instance selection for collaborative filtering". In International Journal of Electronic Commerce 8.4, pp. 115–129.

Jeffrey M O'Brien, 2006. The race to create a 'smart'google. In Fortune Magazine.

João Vinagre, Alípio Mário Jorge, and João Gama, 2014. Fast incremental matrix factorization for recommendation with positive-only feedback. In International Conference on User Modeling, Adaptation, and Personalization, pp. 459–470.
John S Breese, David Heckerman, and Carl Kadie, 1998. Empirical analysis of predictive algorithms for collaborative filtering. In Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence. Morgan Kaufmann Publishers Inc., pp. 43–52.
Kai Yu et al., 2002. Instance selection techniques for memory-based collaborative filtering. In Proceedings of the 2002 SIAM International Conference on Data Mining. SIAM, pp. 59–74.
Lyle H Ungar and Dean P Foster, 1998. Clustering methods for collaborative filtering. In AAAI workshop on recommendation systems. Vol. 1., pp. 114–129.
Manos Papagelis et al., 2005. Incremental collaborative filtering for highly-scalable recommendation algorithms. In International Symposium on Methodologies for Intelligent Systems, pp. 553–561.
Ronan Collobert and Jason Weston, 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning". In Proceedings of the 25th international conference on Machine learning, pp. 160–167.
Scott Deerwester et al., 1990. Indexing by latent semantic analysis. In Journal of the American society for information science, pp. 391–407.
Sung-young Jung and Taek-Soo Kim, 2001. An incremental similarity computation method in agglomerative hierarchical clustering. In Journal of Korean Institute of Intelligent Systems 11.7, pp. 579–583.
Tomas Mikolov et al., 2013. Distributed representations of words and phrases and their compositionality". In Advances in neural information processing systems, pp. 3111–3119.
Yoshua Bengio et al., 2003. A neural probabilistic language model. In Journal of machine learning research, pp. 1137–1155.