

The Continuous Cold Start Problem in e-Commerce Recommender Systems

Lucas Bernardi¹, Jaap Kamps², Julia Kiseleva³, Melanie J.I. Mueller¹

¹Booking.com, Amsterdam, Netherlands. Email: {lucas.bernardi, melanie.mueller}@booking.com

²University of Amsterdam, Amsterdam, Netherlands. Email: kamps@uva.nl

³Eindhoven University of Technology, Eindhoven, Netherlands. Email: j.kiseleva@tue.nl

ABSTRACT

Many e-commerce websites use recommender systems to recommend items to users. When a user or item is new, the system may fail because not enough information is available on this user or item. Various solutions to this ‘cold-start problem’ have been proposed in the literature. However, many real-life e-commerce applications suffer from an aggravated, recurring version of cold-start even for known users or items, since many users visit the website rarely, change their interests over time, or exhibit different personas. This paper exposes the *Continuous Cold Start (CoCoS)* problem and its consequences for content- and context-based recommendation from the viewpoint of typical e-commerce applications, illustrated with examples from a major travel recommendation website, Booking.com.

General Terms

CoCoS: continuous cold start

Keywords

Recommender systems, continuous cold-start problem, industrial applications

1. INTRODUCTION

Many e-commerce websites are built around serving personalized recommendations to users. Amazon.com recommends books, Booking.com recommends accommodations, Netflix recommends movies, Reddit recommends news stories, etc. Two examples of recommendations of accommodations and destinations at Booking.com are shown in Figure 1. This widescale adoption of recommender systems online, and the challenges faced by industrial applications, have been a driving force in the development of recommender systems. The research area has been expanding since the first papers on collaborative filtering in the 1990s [12, 16]. Many different recommendation approaches have been developed since then, in particular content-based and hybrid

approaches have supplemented the original collaborative filtering techniques [1].

In the most basic formulation, the task of a recommender system is to predict ratings for items that have not been seen by the user. Using these predicted ratings, the system decides which new items to recommend to the user. Recommender systems base the prediction of unknown ratings on past or current information about the users and items, such as past user ratings, user profiles, item descriptions etc. If this information is not available for new users or items, the recommender system runs into the so-called *cold-start problem*: It does not know what to recommend until the new, ‘cold’, user or item has ‘warmed-up’, i.e. until enough information has been generated to produce recommendations. For example, which accommodations should be recommended to someone who visits Booking.com for the first time? If the recommender system is based on which accommodations users have clicked on in the past, the first recommendations can only be made after the user has clicked on a couple of accommodations on the website.

Several approaches have been proposed and successfully applied to deal with the cold-start problem, such as utilizing baselines for cold users [8], combining collaborative filtering with content-based recommenders in hybrid systems [14], eliciting ratings from new users [11], or, more recently, exploiting the social network of users [6, 15]. In particular, content-based approaches have been very successful in dealing with cold-start problems in collaborative filtering [3, 4, 13, 14].

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 there are no warm users [3], or in situations when the warm-up period is short and 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 a *continuous* cold-start (CoCoS). In the example of Booking.com, many users visit and book infrequently since they go on holiday only once or twice a year, leading to a prolonged cold-start and extreme sparsity of collaborative filtering matrices, see Fig. 2 (top). In addition, even warm long-term users can cool down as they change their needs over time, e.g. going from booking youth hostels for road trips to resorts for family vacations. Such cool-




Figure 1: Examples of recommender systems on Booking.com. User-to-user collaborative filtering (left): recommend accommodations viewed by similar users to a user who just looked at ‘Hotel Sacher Wien’. Item-to-item content-based recommendations (right): recommend destinations similar to a particular destination, Vienna.

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 Fig. 2 (bottom). These continuous cold-start problems are rarely addressed in the literature despite their relevance in industrial applications. Classical approaches to the cold-start problem fail in the case of CoCoS, since they assume that users warm up in a reasonable time and stay warm after that.

In the remainder of the paper, we will elaborate on how CoCoS appears in e-commerce websites (Sec. 2), outline some approaches to the CoCoS problem (Sec. 3), and end with a discussion about possible future directions (Sec. 4).

2. CONTINUOUS COLD-START

Cold-start problems can in principle arise on both the user side and the items side.

2.1 User Continuous Cold-Start

We first focus on the user side of CoCoS, which can arise in the following cases:

Classical cold-start / sparsity: new or rare users

Volatility: user interest changes over time

Personas: user has different interests at different, possibly close-by points in time

Identity: failure to match data from the same user

All cases arise commonly in e-commerce websites. New users arrive frequently (classical cold-start), or may appear new when they don’t log in or use a different device (failed identity match). Some websites are prone to very low levels of user activity when items are purchased only rarely, such as travel, cars etc., leading to sparsity problems for recommender systems. 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, see Figure 2 for examples from Booking.com.

These issues arise for collaborative filtering as well as content-based or hybrid approaches, since both user ratings or activities as well as user profiles might be missing, become outdated over time, or not be relevant to the current user persona.

2.2 Item Continuous Cold-Start

In a symmetric way, these CoCoS problems also arise for items:

Classical cold-start / sparsity: new or rare items

Volatility: item properties or value changes over time

Personas: item appeals to different types of users

Identity: failure to match data from the same item

New items appear frequently in e-commerce catalogues, as shown in Figure 3 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. In the context of news or conversions, item volatility is also known as topic drift [9]. Figure 3 on the right 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 (identity problem).

3. ADDRESSING COLD-START

Many approaches have been proposed to deal with the classical cold-start problem of new or rare users or items [11]. However, they mostly fail to address the more difficult CoCoS.

The most popular strategy to address the classic cold-start problem is the hybrid approach where collaborative filtering and content-based models are combined, see [14] as an example. If one of the two methods fails due to a new user or item, the other method is used to ‘fill-in’. The most basic assumption is that similar users will like similar items. Similarity of users is measured by their purchase history when warm, and by their user profile when cold. Conversely, similarities between items is computed by the set of users that purchased them when warm, and by their content when cold. In CoCoS, users change their interests, so both collaborative filtering and user-profile-based approaches can fail, since looking at the past and similarities can be misleading. Items also suffer from volatility, although to a lesser degree, which makes the standard hybrid approach also problematic for




Figure 2: Continuously cold users at Booking.com. Activity levels of two randomly chosen users of Booking.com over time. The top user exhibits only rare activity throughout a year, and the bottom user has two different personas, making a leisure and a business booking, without much activity inbetween.




Figure 3: Continuously cold items at Booking.com. Thousands of new accommodations are added to Booking.com every month (left). The user ratings of a hotel can change continuously (right).

items. Hybrid approaches also ignore the issue of multiple personas.

Although, to our knowledge, the continuous cold-start problem as defined in this work has not been directly addressed in the literature, several approaches are promising.

Tang et al. [19] 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. [18] 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 timeSVD [10] are reported. Consistent results are reported in [5], after offline

evaluation using real data.

Tavakol and Brefeld [20] 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.

4. DISCUSSION

In this manuscript, we have described how CoCoS, the continuous cold-start problem, is a common issue for e-commerce applications. Industrial recommender systems do not only have to deal with ‘cold’ (new or rare) users and items, but also with known users or items that repeatedly ‘cool

down'. Reasons for the recurring cool-downs include the volatility in user interests or item values, different personas depending on user context or item target audience, or identification problems due to logged-out users or items from different catalogues. Despite the practical relevance of CoCoS, common literature approaches do not deal well with this issue.

We consider several directions as particularly promising to deal with CoCoS. Traditional approaches to solve cold-start problems try to employ collaborative filtering based on pseudo or inferred clicks. Recommendations based on social networks are an interesting new development that can supplement missing information based on the social graph. For example, recommendations based on Facebook likes are proposed in [15]. Beyond the difficulty to get access to social data, the application to user volatility or multiple personas remains challenging. Online user intent prediction can be used to estimate a user's current profile on the fly. When a user visits the website, his browsing behavior is used to estimate his intent after a few clicks, which are then used to compute recommendations accordingly. However, this still delays recommendations until enough clicks have occurred, which can be problematic if quick recommendations are needed. For example, in last-minute bookings, users may be pressed to book an accommodation quickly, leading to very short sessions.

More promising approaches employ content based or contextual recommendation. Content based recommendations can be very effective based on very little signal: just an initial query or single interaction can be exploited to find an initial item or set of items and exploit relations between items to make effective recommendations. In particular context aware recommendations are one of the most promising strategies when it comes to solving CoCoS. In this setup, recommendations are computed based on the current context of the current visitor and the behaviour of other users in similar contexts [see 2, 7, 17] for examples. Context is defined as a set of features such as location, time, weather, device, etc. Often this data is readily available in most commercial implementations of recommender systems. This approach naturally addresses sparsity by clustering users into contexts. Since context is determined in a per-action basis, user volatility and multiple personas can be addressed robustly. On the other hand, context aware recommenders cannot address the item side of the problem and they might also suffer from cold-start problems in the case of a cold context that has never seen before by the system.

References

- [1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17:734–749, 2005.
- [2] G. Adomavicius and A. Tuzhilin. Context-aware recommender systems. In *Recommender Systems Handbook*, pages 217–253, 2011.
- [3] M. Aharon, N. Aizenberg, E. Bortnikov, R. Lempel, R. Adadi, T. Benyamini, L. Levin, R. Roth, and O. Serfaty. OFF-set: One-pass factorization of feature sets for online recommendation in persistent cold start settings. In *Proceedings of the 7th ACM Conference on Recommender Systems*, pages 375–378, 2013.
- [4] S. Bykau, F. Koutrika, and Y. Velegrakis. Coping with the persistent cold-start problem. In *Personalized Access, Profile Management, and Context Awareness in Databases*, 2013.
- [5] F. C. T. Chua, R. J. Oentaryo, and E.-P. Lim. Modeling temporal adoptions using dynamic matrix factorization. In *IEEE 13th International Conference on Data Mining (ICDM)*, pages 91–100, 2013.
- [6] I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, and S. Ofek-Koifman. Personalized recommendation of social software items based on social relations. In *Proceedings of the Third ACM Conference on Recommender Systems*, pages 53–60, 2009.
- [7] A. Hawalah and M. Fasli. Utilizing contextual ontological user profiles for personalized recommendations. *Expert Syst. Appl. (ESWA)*, 41:4777–4797, 2014.
- [8] D. Kluver and J. A. Konstan. Evaluating recommender behavior for new users. In *Proceedings of the 8th ACM Conference on Recommender Systems*, pages 121–128, 2014.
- [9] D. Knights, M. Mozer, and N. Nicolov. Detecting topic drift with compound topic models. In *Proceedings of the Third International ICWSM Conference*, pages 242–245, 2009.
- [10] Y. Koren. Collaborative filtering with temporal dynamics. *Communications of the ACM*, 53:89–97, 2010.
- [11] 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 *Proceedings of the 7th International Conference on Intelligent User Interfaces*, pages 127–134, 2002.
- [12] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. Grouplens: An open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work*, pages 175–186, 1994.
- [13] M. Saveski and A. Mantrach. Item cold-start recommendations: Learning local collective embeddings. In *Proceedings of the 8th ACM Conference on Recommender Systems*, pages 89–96, 2014.
- [14] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 253–260, 2002.
- [15] S. Sedhain, S. Sanner, D. Braziunas, L. Xie, and J. Christensen. Social collaborative filtering for cold-start recommendations. In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 345–348, 2014.
- [16] U. Shardanand and P. Maes. Social information filtering: Algorithms for automating 'word of mouth'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 210–217, 1995.
- [17] Y. Shi, A. Karatzoglou, L. Baltrunas, M. Larson, and A. Hanjalic. Cars2: Learning context-aware representations for context-aware recommendations. In *Proceeding of CIKM*, pages 291–300, 2014.
- [18] J. Z. Sun, K. R. Varshney, and K. Subbian. Dynamic matrix factorization: A state space approach. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1897–1900, 2012.
- [19] L. Tang, Y. Jiang, L. Li, and T. Li. Ensemble contextual bandits for personalized recommendation. In *Proceedings of the 8th ACM Conference on Recommender Systems*, pages 73–80, 2014.
- [20] M. Tavakol and U. Brefeld. Factored mdps for detecting topics of user sessions. In *Proceedings of the 8th ACM Conference on Recommender Systems*, pages 33–40, 2014.