



The Synergies of Context and Data Aging in Recommendations

Anna Dalla Vecchia[✉], Niccolò Marastoni[✉], Barbara Oliboni[✉],
and Elisa Quintarelli[✉]

University of Verona, Strada Le Grazie, 15, Verona, Italy

`anna.dallavecchia@studenti.univr.it`,

`{niccolo.marastoni,barbara.oliboni,elisa.quintarelli}@univr.it`

Abstract. In this paper, we investigate the synergies of data aging and contextual information in data mining techniques used to infer frequent, up-to-date, and contextual user behaviours that enable making recommendations on actions to take or avoid in order to fulfill a specific positive goal. We conduct experiments in two different domains: wearable devices and smart TVs.

Keywords: Context-awareness · Data Aging · Data mining · Recommendations

1 Introduction

Mobile technology has become an essential element of modern life. The use of sensors and mobile devices provides a constant and steady stream of new data. Such information can range from wearable devices data, to data coming from IoT home/office devices, and to data related to media consumption habits.

Sensor data is intrinsically temporal in nature, since events happen in succession, and thus temporal data can be analyzed to extract sequences of frequent events (patterns). Looking at a user's historical actions can give insight into their future actions, and events that occurred more recently can be more valuable than older ones for determining new events. In this case, data aging techniques can improve the accuracy of predictions for the future by focusing the attention on recent patterns.

Another aspect that can greatly impact a user's choices is the context in which they find themselves when making those choices. This means that considering additional contextual information could improve the quality of the recommendations. The context, as we refer to it in this paper, encompasses all information that identifies the situation in which a recommendation is made [3]. This contextual data, while providing information that is external to the event, is inherently personalized for each individual user, as users can be influenced by different contextual factors. It therefore cannot be imposed at design time nor will it be equivalent for all users. Contextual information belongs to different categories, among them we can cite the temporal perspective and the one

related to the user [9]. One example of a temporal perspective is whether the day is a weekday or on the weekend; one example of a user-related perspective is whether the user is with their family or by themselves.

In this paper we extend our preliminary proposal [8] and we show that aging and contextual data can work together and improve the performance of recommender systems, both in terms of accuracy and computational complexity. In particular, we combine data aging with different contexts, in order to gauge their respective impact on the overall performance of the algorithm.

As general use cases, we consider two different datasets for evaluating our proposal. The first dataset contains logs of data collected by Fitbit devices on physical activity levels and sleeps scores. The second use case is a dataset including logs of TV programs watched by users. In order to enrich the considered data with contextual information, we encode different contextual perspectives according to the considered dataset, including the temporal one which we refine to distinguish between vacation and normal days when possible. As for the TV programs, besides analysing data w.r.t. temporal perspectives, we also use details about the people watching TV, distinguishing between lone child, lone adult, child with family, adult with family.

The paper is organized as follows. Section 2 introduces ALBA, the extension of the LookBack Apriori (LBA) algorithm with the aging mechanism. Section 3 describes how we consider contextual data, and Sect. 4 presents the validation of the proposed approach. Section 5 summarizes the contribution of the paper.

2 ALBA: Adding Aging to LookBack Apriori

In this section, we introduce the Aged LookBack Apriori algorithm (ALBA), an extension of Apriori [4] that increases the importance of recent information through the introduction of an aging mechanism.

ALBA starts from the LookBack Apriori algorithm (LBA) [8] that generates frequent and totally ordered sequence rules from temporal data while considering i) a possibly predefined temporal window τ_w , concatenating each transaction with the previous ones in such a way to be able to look back at $\tau_w - 1$ time units, and ii) the relative time of each item w.r.t. the current time instant. The mined rules have the form $r : I_{-(\tau_w-1)} \wedge \dots \wedge I_{-2} \wedge I_{-1} \rightarrow I_0 [s_r, c_r]$.

It is important to highlight that we have developed our approach without starting from well know algorithms, like GSP [6] because we consider relevant the sequential relationship between each itemset.

In ALBA, differently from LBA, in order to compute the support of the items at each iteration of the Apriori algorithm, the system builds a matrix of dimensions $n \times m$, where n is the number of transactions and m is the number of different items in the considered temporal dataset. The value of the element at row i and column l will be 1 if item l appears in day i , 0 otherwise. With this approach, the support of item l can be easily computed by summing all the elements of its l -th column and dividing the result by n . To penalize the older items, each $i - th$, with $i < \tau_w$, is multiplied by an aging factor $\alpha_i = \frac{i}{n - \tau_w + 1}$,

where τ_w is the temporal window and n is the number of transactions in the dataset. This guarantees that the items in the temporal window will still be represented with 1, while older items slowly decay but never quite reach 0.

3 Context Modeling

The importance of considering contextual features in recommender systems to improve the relevance of provided suggestions is widely recognized [2, 3, 9]. Indeed, users' behaviours, preferences, and decisions are often affected by external factors used to characterize the environment or situation they are acting in, e.g. spatio-temporal information, weather, and social conditions, presence of other people. Although there is no standard definition of context, it is commonly accepted that context includes any information that can be used to characterize the situation of an entity, where an entity is anything considered relevant for the interaction between a user and an application [1]. This is mainly composed of highly dynamic features that change their values over time, thus requiring an efficient way to handle them in order to be able to provide fresh and relevant recommendations on the base of the user current context. Since contextual information is not always available, or complete, hierarchical contextual models have been proposed in the literature [5, 7], as they allow to represent the context at different levels of granularity.

According to the systematic framework introduced in [9], in this work we consider the contextual modeling paradigm, i.e. the contextual features are integrated into our recommendation model and are considered during the preference computation process. This means that each transaction is composed of itemsets I_{t_i} that contain, besides other items, the item C_{t_i} related to the context at time unit t_i . Context information has to be acquired before being integrated into the preference computation; sometimes it may be gathered from the same smart devices or sensors providing data related to the objects to monitor (e.g. the temporal information can be acquired from a wearable device), but other times it has to be obtained from external sources, as in the case of weather or social information.

4 Evaluation

In this section we present some experiments to validate our approach. We will consider the following aspects: (i) the benefits of data aging to LBA, (ii) the benefits of adding contextual information and (iii) the advantages of modeling contexts with different granularity levels.

The first set of experiments focuses on wearable device data: we analyze logs of physical activity levels and sleep scores extracted from Fitbit devices, enhancing the process with an aging mechanism and contextual information, in order to provide contextual, personalized, and timely recommendations to users. To further evaluate the synergies between the aging mechanism and the relevance of context during the recommendation process we apply our proposal

to a Smart TV scenario, where the recommender system suggests the next genre to watch based on historical and contextualized data.

Fitbit. This scenario involves physical activity and sleep score logs obtained from Fitbit devices and it features two datasets: PMdata and Custom. PMdata consists of logs from 16 users and the Custom dataset was collected from 4 users specifically for this study. From both of these datasets we make use of the logs pertaining to “light”, “medium” and “heavy” activity, along with the sleep score for each day. These features are recorded by Fitbit as minutes spent in each activity type, thus we discretize them in order to obtain categorical data as described in [8]. During this discretization process the activity levels are further split into 3 sub-levels according to set thresholds, e.g. a heavy activity (*HA*) can be encoded into three possible labels: *HA* : 1, *HA* : 2 and *HA* : 3.

Auditel. In this scenario, we study the TV watching habits of 7 users, aged between 5 and 77 years old. The data has been collected by the Auditel company and it consists of logs of TV programs watched, where each data point contains, among other things, the genre of the program itself, a timestamp of the exact start time and end time, and the specific user. An interesting feature of this dataset is the “family” field which indicates the age range of the user (“child” or “adult”) and if the user was watching the program alone or with members of the family. The 7 users have been selected from the thousands included in the dataset by virtue of the richness of their logs and familiar context. For these experiments we only consider the data points containing more than 10 min of screen time, filtering out channel surfing and disliked programs. Each data point contains one program viewing and it is encoded by a two-letter label representing one of the five possible genres: documentary, serious, fun, reality, sport.

4.1 Contexts

The first temporal contextual dimension considered in this study is the **Part of the week** which characterizes each day as either a weekday or part of the weekend. This is used in all the datasets, as the smallest time unit in the Fitbit datasets is the day, so each data point can be easily enriched by the value *WD* if it represents a day that falls into the weekend. Moreover, Auditel data allows for smaller time granularities, so we add **Portion** as a second temporal contextual dimension, called *TS* in this section, which tracks the time slot in which the individual data point (a program viewing) takes place, be it morning, afternoon or night time. As mentioned earlier, the Auditel dataset comes with the “family” feature, so we consider the first individual contextual dimension related to the **Type** of a group, called here *FA*, and use it to track four different values: lone child, lone adult, child with family, adult with family. Lastly, we add the **Situation** dimension, reported as *VA*, to distinguish the vacation days from normal days. Unfortunately, this is only available on the Custom dataset, as it requires the collection of private data from each user.

4.2 Methodology

We reserve 80% of the logs of each user as a training set to generate the rules and then use the remaining 20% as a test set. We then run the experiments on the training set with multiple configurations of LookBackApriori plus aging and context, and generate a set of rules for each user in the dataset with no specified goal. For the Fitbit dataset we only extract rules that have the sleep score (our target) in the consequent, while on the Auditel dataset the consequent will contain the genre of the last TV program in the frequent itemset.

The test dataset is then divided in multiple queries that are then checked for accuracy against the rules learned in the training phase. For both scenarios we first order the extracted rules by confidence and support, but in the Fitbit experiments we prioritize rules that are “complete”, meaning that they contain itemsets for each timestamp. This is done to preserve the feasibility of the original proposal, since we need to be able to suggest what types of activity will lead to better sleep, having a knowledge of the complete history of user’s past activities in the Fitbit log.

4.3 Fitbit Validation

To evaluate the impact of aging and context on the PMdata dataset we first ran the experiments with: 1) the original LBA, 2) LBA plus the aging mechanism, 3) LBA with the *WD* context (the only available context for this dataset) and then 4) LBA with both aging and the context. The temporal window has been set to 3, as that is the value that gives the best results while keeping the run time acceptable.




Fig. 1. Accuracies for all PMdata users.

The results can be seen in Fig. 1. From these we can observe that both the aging mechanism and the *WD* context substantially improve the average

accuracy of the LBA algorithm when used separately. The combination of *WD* and aging further improves the average accuracy due to a few outliers, but the performance for the individual users is not always better than the previous experiment. The experiments on the Custom dataset were ran with the same configurations as PMdata but with the additional context *VA*, which is only available in this dataset. This adds four configurations: *VA*, the combination of *VA* and *WD*, and both previous configurations with aging.

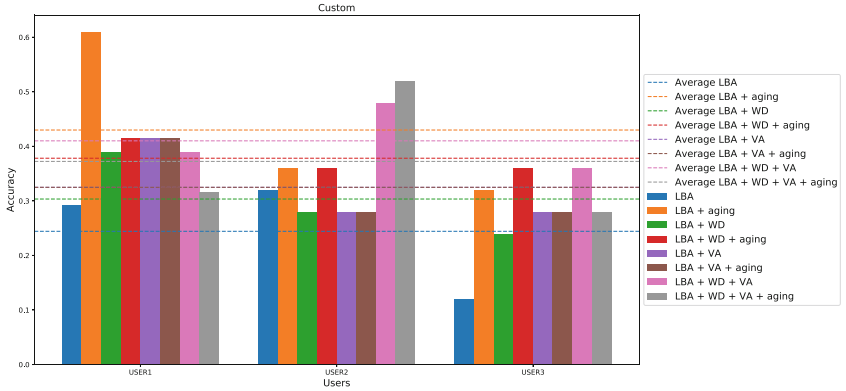


Fig. 2. Accuracies for all Custom users.

The results, shown in Fig. 2, show that the aging mechanism improves LBA average accuracy slightly better than the combination of the *VA* and *WD* contexts. It is also interesting to note that the average accuracy does not tell the whole story, as the experiments on different users in the Custom dataset react differently to data aging and the addition of context. For example, the performance on User 1 is strongly, and positively, affected by the aging mechanism. This is easy to explain, as User 1 started logging data before all the others, resulting in a much bigger dataset that is better suited to aging. We posit that the aging mechanism does not benefit as much those temporal datasets that span shorter periods, and the experiments in the next section further support this proposition. Another example is User 2, which is one of the clearest examples of context synergy in these experiments: the accuracy drops slightly at the addition of the contexts *VA* and *WD* separately, but their combination (with aging) leads to a 63% increase

4.4 Auditel Validation

We run the experiments on the selected users of the Auditel dataset, starting with LBA and then adding the available contexts in all possible configurations. Interestingly, the aging mechanism, while being very effective in the Fitbit experiments, does not improve the results for the Auditel scenario. We suspect that

the reasons for this are to be found mainly in the smaller size of the Auditel dataset: since only 3 months of data are available, the aging mechanism is penalizing data that might still be relevant, as it is not particularly old. Another possible reason is the higher granularity of the data: Auditel items always represent events happening in the same day, thus the aging mechanism wrongly penalizes (although not excessively) events that are temporally very close. For this reason the aging mechanism is excluded from the presented experiments, as it would only lead to redundancy.




Fig. 3. Accuracies for some Auditel users.

The temporal window for these experiments has been set to 6, as we found that further increasing the value, while leading to better accuracy, incurs in a big running time penalty. The best performing context in this scenario is *TS*, representing the time slot in which the program is televised. As shown in Fig. 3, the addition of *TS* increases the accuracy of the predictions for every user in the dataset. While the boost to the average accuracy is around 38%, the individual rise can be up to 104% (see user 1750 in Fig. 3). This is not surprising, as the nature of TV programs tends to change within the different periods of the day (i.e. serious programs are often at night, while kids program are more common in the afternoon). An interesting finding in this scenario is that the weekday context (*WD*) when used alone, actually decreases the accuracy of the algorithm in every user, while it improves the accuracy if used in conjunction to the *TS* context. The improvement cannot be entirely attributed to the already successful *TS* context, as the results of *TS + WD* are superior than those with only *TS*. This is another example of synergy between different contexts. The familiar context (*FA*) is the only context that does not enhance the performance of the recommender system; in fact it often does the opposite, even when paired with more successful contexts.

5 Conclusions and Future Work

In this paper, we combine contextual information and an aging mechanism in data mining techniques to evaluate their impact in the context of recommendation systems. The results of our experiments show that adding context to the recommender system leads to an increase of accuracy in most cases. The aging mechanism proves to be a very positive addition in the Fitbit experiments, improving all results but one, often by quite a large margin. Conversely, aging never improves the results on Auditel data. Overall, it seems like the decision on whether to add contextual data or aging (or neither) to the recommender system needs to be analyzed on a case-to-case basis, as different problems, different scenarios, and different data will require specific solutions.

References

1. Abowd, G.D., Dey, A.K., Brown, P.J., Davies, N., Smith, M., Steggles, P.: Towards a better understanding of context and context-awareness. In: Gellersen, H.-W. (ed.) HUC 1999. LNCS, vol. 1707, pp. 304–307. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48157-5_29
2. Adomavicius, G., Mobasher, B., Ricci, F., Tuzhilin, A.: Context-aware recommender systems. *AI Mag.* **32**(3), 67–80 (2011)
3. Adomavicius, G., Sankaranarayanan, R., Sen, S., Tuzhilin, A.: Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.* **23**(1), 103–145 (2005)
4. Agrawal, R., Srikant, R.: Fast algorithms for mining association rules in large databases. In: Bocca, J.B., Jarke, M., Zaniolo, C. (eds.) *Proceedings of VLDB 1994*, pp. 487–499 (1994)
5. Bolchini, C., Quintarelli, E., Tanca, L.: CARVE: context-aware automatic view definition over relational databases. *Inf. Syst.* **38**(1), 45–67 (2013)
6. Harms, S.K., Deogun, J.S.: Sequential association rule mining with time lags. *J. Intell. Inf. Syst.* **22**(1), 7–22 (2004)
7. Hosseinzadeh Aghdam, M.: Context-aware recommender systems using hierarchical hidden Markov model. *Phys. A Stat. Mech. Appl.* **518**, 89–98 (2019)
8. Marastoni, N., Oliboni, B., Quintarelli, E.: Explainable recommendations for wearable sensor data. In: *International Conference on Big Data Analytics and Knowledge Discovery*, pp. 241–246 (2022)
9. Villegas, N.M., Sánchez, C., Díaz-Cely, J., Tamura, G.: Characterizing context-aware recommender systems: a systematic literature review. *Knowl. Based Syst.* **140**, 173–200 (2018)