Normalization of Timeseries for Improving Recommendations

Alexandru Ciobanu[1], Asmaa Haja[2], and Andreas Lommatzsch[2]

[1]Technische Universität Berlin, Straße des 17. Juni 135, D-10625 Berlin, Germany
alexandru.ciobanu@campus.tu-berlin.de
[2]DAI-Labor, Technische Universität Berlin, Ernst-Reuter-Platz 7, D-10587 Berlin, Germany
{firstname.surname}@dai-labor.de

Abstract. Recommender systems typically analyze the user behavior to predict the items most likely matching the user preferences. While analyzing the user preferences the observed data must be normalized considering characteristic usage pattern as well as the lifecycles of the items. The normalization is especially relevant in domains characterized by frequently updated item sets and fast changes in the set of users, typical in news recommender scenarios.
In this paper, we study approaches for normalizing the collected data in a news recommender system. We analyze live usage data from several German news portals available in the NewsREEL challenge. We discuss the influence of seasonality, day of week and hour of day patterns. We show that applying time series decomposition and normalization methods enable us to model the relevant aspects resulting in the observed item popularity. The derived models are the basis for precisely predicting the interest in concrete items in the future.

1 Introduction

Publishers and users generate an enormous amount of content every minute. As a result, users struggle to find relevant, interesting content. Recommender systems have emerged as tools to supports users to find relevant information. Their success has enabled online shops and entertainment systems to develop innovative business models [4].

Recommender systems consider existing user-item interactions. They derive models to predict for individual users the most likely relevant items of interest. Therein, recommender systems uncover group interests and trends. Recommending news articles or social media posts entails multiple challenges: First, the items' lifecycles expire quickly. In other words, users' perception of items tend to be less and less relevant over time. Second, users exhibit characteristic routines when consuming news. Consequently, highly relevant items can attract lesser attention at night compared to rather obscure items in rush hours. As a result, recommender systems have to pay attention to these routines to arrive at meaningful predictions.

In this paper, we study methods for analyzing time series data describing the user-news item interactions. Separating different types of biases furthers our understanding of news items' lifecycles as well as users' preference. Therefore, our approach is to decompose time series data into their characteristic components, which improves computing news recommendations by extrapolating observed user preferences. Our analysis

shows that this approach provides more accurate recommendation models and leads to a better understanding of items' lifecycles.

The remaining paper is structured as follows. In Sec. 2, we present a concrete application scenario and discuss the dataset used for the evaluation. Sec. 3 analyzes existing models and similar systems. Our approach is explained in Sec. 4. Subsequently, we discuss the evaluation results obtained by applying our method in the NEWSREEL scenario. Finally, Sec. 6 concludes and gives an outlook to future work.

2 Problem Description

Several scenarios demand developing recommendation algorithms for streams of news items and social media posts. In this paper, we focus on recommending news on selected major German news portals. This section explains the analyzed scenario in detail, presents the dataset, and discusses the optimization objectives.

2.1 News Recommendation

News is ubiquitously available in many different forms. Digital media have reduced both publication and consumption costs over the most recent decade. As a result, news portals have gained tremendous popularity. The vast amount of news and the broad spectrum of topics challenge readers to find interesting items. Recommender systems assist users in finding interesting news articles that match their individual preferences. Advanced news recommender systems consider trends and users' habits to ensure the delivery of relevant recommendations. Consequently, news recommender systems must take into account the items' lifecycles and characteristic patterns describing news consumption. Powerful, context-aware news recommender systems emerge from understanding these patterns.

2.2 The NewsREEL Challenge

We take advantage of the NEWSREEL challenge to study interactions with news in a real-world scenario. The NEWSREEL challenge allows researchers to evaluate innovative recommender algorithms both online and offline on real-life data. The competition is hosted by PLISTA[1], a global service provider for data-driven native advertising and recommendations. The Open Recommendation Platform (ORP) facilitates data exchange in the scope of NEWSREEL. ORP supplies participants with information concerning published news articles and user-item interactions. Besides, participants receive recommendation requests which they must respond to in real time. The system forwards the recommendations and presents them to the reader. Clicking on recommended items signals a positive response. The team registering the highest Click-Through-Rate (CTR) in a predefined timeframe wins the challenge.

[1] https://www.plista.com

2.3 The NewsREEL Data Resources

NEWSREEL supplies data in the form of a continuous stream of published news items and user-item interactions. Registered participants may either work interactively with the live system or replay previously recorded streams.

The Message Types The Open Recommendation Platform provides messages as JSON objects. Participants encounter three types of messages: `recommendation-request`, `item-update`, and `event-notification`. Messages of type `item-update` inform participants about changes to the item collections. These changes include adding new items as well as adjusting existing items. Messages of type `event-notification` inform participants about user-item interactions. These interactions include reading news and clicking on recommendations. Participants do not need to reply to `item-updates` or `event-notifications`. Messages of type `recommendation-requests` demand a response consisting of a list of recommended items. Participants may collect data and conduct experiments to gain insight into news consumption and domain-specific characteristics.

Dataset Characteristics Previously, news recommender systems have been studied in detail for static datasets [3, 6]. Methods including collaborative filtering and popularity-based approaches rank highly among well-known methods. They rely on the analysis of user-item interactions. Studying the approaches in detail, we observed, that users show strongly varying preferences. Moreover, we encounter varying performances depending on the daytime. This is the motivation for us to study the influence of time/date in this work. We create time series from web server logs and user behavior contained in them. In particular, we study the frequency of impressions and published items depending on daytime and weekday. Subsequently, we decompose and normalize the time series by removing seasonal and cyclic effects. As a result, we can compare articles despite varying in activity.

2.4 Metrics

We consider the proportion of accepted recommendations as accuracy measure. The Click-Through-Rate (CTR) approximates this objective well. CTR, in our scenario, corresponds to the quotient of clicked recommendations by the total number of recommendations presented to users.

In addition, we analyze how well the models can predict activities in the near future. In other words, we study how accurately we can extrapolate the time series. The predictions form the basis for the recommender model.

2.5 Discussion

In this section, we presented the analyzed scenario, the dataset, and the relevant metrics. Before presenting our approach, we review related research and discuss how existing methods should be adapted to suit our scenario.

3 Related Work

This section reviews existing work. We analyze potentially relevant approaches and discuss the relevance of our scenario.

3.1 Analyzing Temporal Patterns Specific for Web Portals

The analysis of the traffic on web portals usually shows specific patterns. Authors have their habits in publishing new items; users read items depending on their circadian rhythm. The number of visitors to web portals usually shows a high variance. The IBM Research Division published a statistical overview of views by daytime for websites in 1999 [7]. It includes large popular websites such as *CNN*, *Schwab* and *Olympic Games (Nagano, Japan)* reaching a wide audience. The authors have analyzed logs containing more than 56.8 million requests per day. Web server performance optimization, which was their goal, required a detailed look at access patterns resulting in several time series.

Strategies incorporating temporal information improve the recommender performance. Analysis conducted based on the NEWSREEL challenge showed that users seem to be interested in the top news during the day; at nights users are more interested in related, but older news items [8]. Most popular recommenders [11] seems to perform best for short user sessions; when users spend more time on the news portals (e.g. at night), sequence-based approaches reach a high CTR. The approaches evaluated in NEWSREEL handle temporal information as a context parameter but do not analyzes the lifecycle of items [10].

3.2 Time Series

Time series formally model sequences of observations such as website visits over time. Many real-world applications assume time series to exhibit steadiness; the analysis of past observations enables prediction on how the time series will evolve in the future. In contrast to traditional approaches (handling the observed world as a static setting), time series-based approaches explicitly consider the dynamics and the changes [2]. Time series tend to come in the form of lists of measurements with nominal, ordinal, or cardinal values. Fig. 1 shows yearly temperature changes as time series. Most time series entail a sequence of discreetly taken measurements. Still, visualizations tend to show time series as continuous approximations assuming steadiness. These graphs help us to understand the data and its underlying structure. Mathematical models describe time series including their deterministic and stochastic elements. Researchers, businesses, governments, and many more use time series for forecasting in various domains. Several techniques exist for time series analysis, e.g the industry uses time series decomposition widely. Techniques to deal with time series have been developed and enhanced since the start of the last century [5].

3.3 Time Series Decomposition

Time series models tend to comprise three components: *Trend*, *Seasonal Variation*, and *Noise*. The *Trend* part includes a systematic change of the mean level. Fig. 1 visualizes

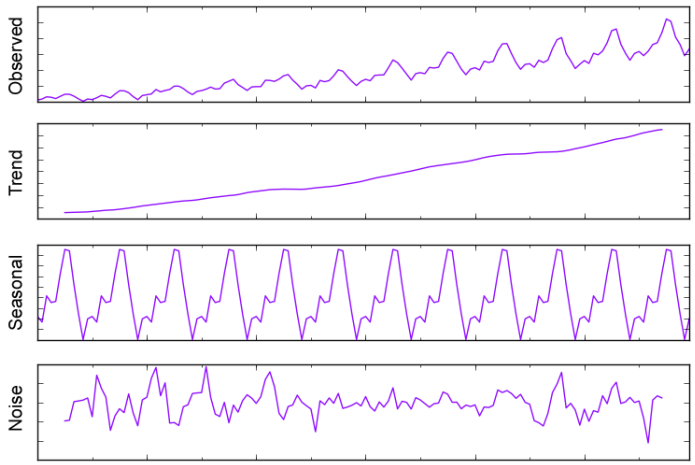

Fig. 1. Decomposition of an observed time series into seasonal, trend, and noise component [1].

how the observed time series decomposes into the three building blocks. The analysis shows that the observed quantity, the temperature in this case, increases in the long term, despite fluctuations in the short intervals. Calculations of general trends include moving averages or similar smoothing algorithms. The *seasonal variation* of one point in time can be described as the average of all data assigned to time points with the same label. We refer to the remaining part as *noise*. ARIMA [12] and ARCH [5] represent more sophisticated time series models. In our analysis we focus on simple, robust approaches.

3.4 Recommendations based on Time Series

The analysis of time series can be used for computing recommendations. Verbitskiy et al. [13] have successfully implemented a recommender analyzing the popularity of recently published news items. The recommender relies on historical data. Still, it neither detects trending topics nor predicts future interests. Thus, a bias in the observations may reduce the precision. Lommatzsch et al. [9] have suggested extrapolating time series in NEWSREEL. Still, they have neglected modeling the different components of time series which impedes the prediction horizon.

3.5 Discussion

Time series has been established as the standard model for observations over time. Decomposing time series promises reliable models with high predictive power. Most-popular recommenders have delivered high precision in previous editions of NEWSREEL. Consequently, we conduct experiments to verify that normalizing and decomposing observations yield reliability and predictive power.

4 Approach

This section presents our approach in detail. First, we discuss the decomposition of the recorded time series. Subsequently, we explain how our model adapts to meet challenges specific to the news recommendation scenario. We discuss the expected precision, the computational complexity, and details of our implementation.

4.1 Comparing the Statistics of the NewsREEL Platform

NEWSREEL's data stream carries information about items and user-item interactions. In our analysis, we use the baseline recommender provided by NEWSREEL's organizers. We aggregate the data by binning each hour. We obtain four time series describing views (impressions), displayed recommendations, clicked recommendations, and item updates. In addition, we experimentally also create streams using bins on a daily basis for the long-term analysis.

We considered six months, varying the granularity of binning, so we can analyze different levels of seasonality. For instance, we can detect shifts between daytime and night or working day and weekend or holiday. Fig. 2 shows the number of views over the course of three days for three publishers. We observe repeated patterns in all three cases. Although, the smallest publisher, with around 1 000 impressions per hour, exhibits a higher level of variance.

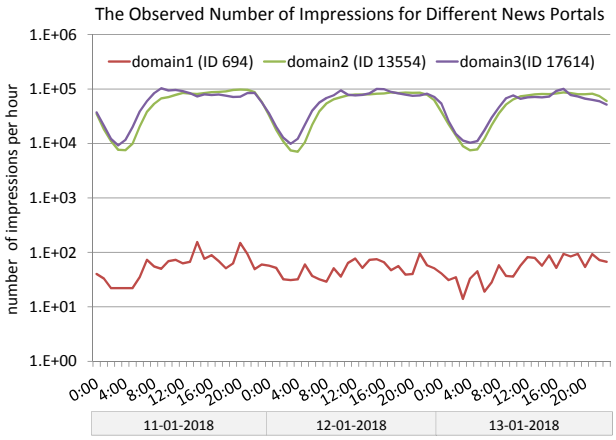

Fig. 2. Impressions per hour on different news portals

4.2 Normalization of Observed Number of Events

High variance in the number of impression and click events complicates detecting specific trends on an item basis. Thus, we apply different normalization approaches. Normalization removes specific patterns from the dataset and allows analyzing the trends

hidden by dominant patterns. In the news recommendation scenario, highly interesting articles published at night may appear uninteresting due to limited user activity at that time. Likewise, boring articles published during rush hours may seem relevant. Therefore, news recommender systems must pay attention to this phenomenon. We do this by concentrating on the seasonality and the identification of random noise to improve item relevance prediction. In order to calculate seasonality patterns and normalization functions, the arithmetic mean of the observed access numbers are computed. As a result, for each hour of the day or own label, a scalar factor is given on that specific domain. This can be used to normalize item numbers. The resulting time series describes the article figures relative to domain performance.

Discussion We have adopted a promising time series normalization approach in order to meet the specific demands of NEWSREEL. Pre-calculated weights for every seasonality point enable a computationally inexpensive adjustment of figures related to news articles, i.e. number of item views and number of clicked recommendations. The developed solution works efficiently and can be integrated into existing recommenders. We focus on a subset of all articles due to the power law distribution of popularity in given portals. Articles with few views cause high variance in the normalized weights which necessitates the selection. The presented algorithm, applied to articles with many views, reflects a more realistic life cycle representation.

5 Evaluation

We evaluate our approach by analyzing popular news items' impressions. Figure 3 shows the number of impressions for the top news item in the selected period. In addition, it shows the total number of views for the selected publisher. The observed number of impressions indicates a steep decrease in popularity about five hours into the timeslot. In contrast, the normalized average shows a lesser decrease in popularity within the first 24 hours. The normalized average reveals that decreasing overall activity affects the perceived popularity. Removing the effect induced by activity shifts provides a clearer picture of articles' lifecycles. A linear decrease of interest can be observed after the peak. Normalizing observations facilitates predicting items' future popularity more precisely.

In addition to the items' lifecycle, we have analyzed the twenty most popular articles from one publisher on a monthly basis. We observe pattern consistent to the previous findings on an hourly level. Some publishers release interesting items at the same time of the day. Publishers tend to push articles on finance in the morning as opposed to sports. Sports competitions, especially football matches, tend to take place in the evening. Figure 4 exemplifies this phenomenon. It shows both the number of observed and normalized views relative to the hour of publication. We observe a decline at the beginning of the night which concurs with the hypothesis. Publishers, who cover a general news spectrum, show less of this tendency.

The normalized time series are the basis for a better understanding of the item lifecycles and the optimization of recommender models. Participants in the NEWSREEL challenge could integrate our findings into their ideas. Considering normalized popularity yields the potential to achieve higher CTR. Besides, publishers gain further insights.

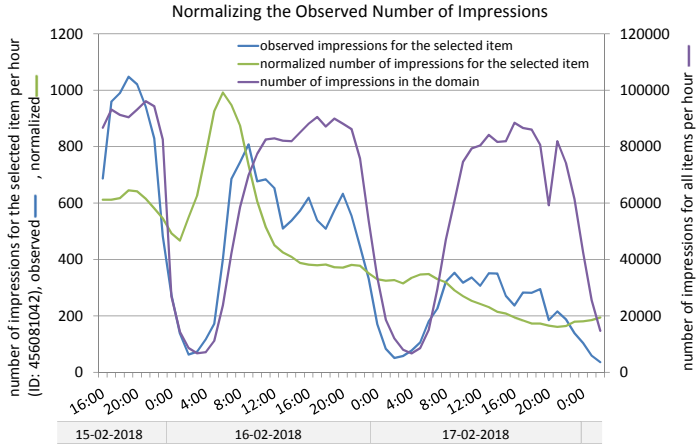

Fig. 3. Item life cycle normalized by hourly domain impressions

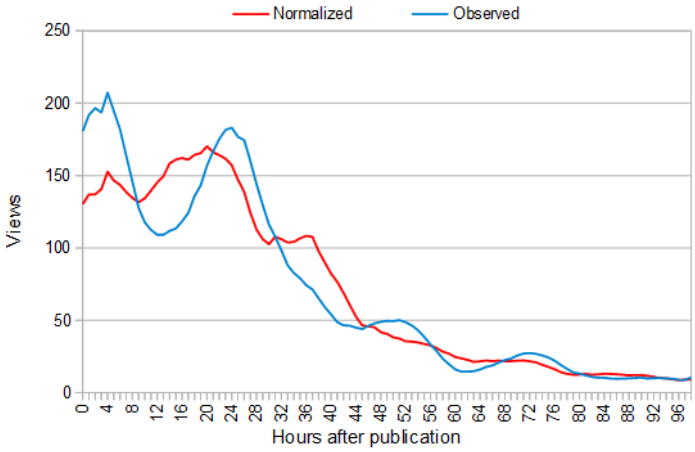

Fig. 4. Average item views after hours of publication

The observations justify removing expired articles and placing promising new articles more visible.

6 Conclusion

We have created a system that uses a pipeline-based architecture to generate different time series in the NEWSREEL scenario. We have normalized the observed number of item impressions by the overall domain activity. The result is an adjusted item lifecycle model. It describes the common patterns of specific article types. This knowledge can be used for computing the expected number of impressions for new articles depending

on the lifetime of an item. Moreover, we have introduced a quantification of real article interest, in addition to the often misleading raw access numbers.

Future work includes the integration of normalization into the existing CLEF NEWSREEL recommender engine. The integration yields the potential to achieve higher CTR for trending items. Publishers can compare articles' success to enhance their publishing guidelines. The gained insights could help to optimize publishing schedules. For instance, publishers could better anticipate popularity peaks. In addition, they can adopt the proposed approach to determine when to remove or modify articles to provide users with a better experience.

References

1. J. Brownlee. *How to Decompose Time Series Data into Trend and Seasonality*. Machinelearningmastery.com, Jan 2017.
2. M. Dettling. Applied time series analysis. *Applied time series analysis*, 2013.
3. F. Garcin, B. Faltings, O. Donatsch, A. Alazzawi, C. Bruttin, and A. Huber. Offline and online evaluation of news recommender systems at swissinfo.ch. In *Proceedings of the 8th ACM Conference on Recommender Systems*, RecSys '14, pages 169–176, New York, NY, USA, 2014. ACM.
4. C. A. Gomez-Uribe and N. Hunt. The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4):13, 2016.
5. J. D. Hamilton. *Time series analysis*, volume 2. Princeton university press Princeton, 1994.
6. F. Hopfgartner, T. Brodt, J. Seiler, B. Kille, A. Lommatzsch, M. Larson, R. Turrin, and A. Serény. Benchmarking news recommendations: The clef newsreel use case. *SIGIR Forum*, 49(2):129–136, Jan. 2016.
7. A. K. Iyengar, M. S. Squillante, and L. Zhang. Analysis and characterization of large-scale web server access patterns and performance. *World Wide Web*, 2(1):85–100, 1999.
8. A. Lommatzsch. Real-time news recommendation using context-aware ensembles. In *Proc. of the 36th European Conference on Information Retrieval*, volume 8416 of *Lecture Notes in Computer Science*, pages 51–62. Springer, 2014.
9. A. Lommatzsch, B. Kille, and S. Albayrak. Incorporating context and trends in news recommender systems. In *Proceedings of the International Conference on Web Intelligence*, WI '17, pages 1062–1068, New York, NY, USA, 2017. ACM.
10. A. Lommatzsch, B. Kille, F. Hopfgartner, M. Larson, T. B. Özlem Özgöbek, and J. Seiler. Clef 2017 newsreel overview: A stream-based recommender task for evaluation and education. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction; Proceedings of the 8th International Conference of the CLEF Association, CLEF 2017*, LNCS, vol. 10456, pages 239–254. Springer International Publishing, 2017.
11. C. A. Ludmann. Recommending news articles in the CLEF news recommendation evaluation lab with the data stream management system odysseus. In *Working Notes of CLEF 2017 - Conference and Labs of the Evaluation Forum, Dublin, Ireland, September 11-14, 2017.*, volume 1866 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2017.
12. R. McCleary, R. A. Hay, E. E. Meidinger, and D. McDowall. *Applied time series analysis for the social sciences*. Sage Publications Beverly Hills, CA, 1980.
13. I. Verbitskiy, P. Probst, and A. Lommatzsch. Development and evaluation of a highly scalable news recommender system. In *Working Notes of the 6th International Conference of the CLEF Initiative*. CEUR Workshop Proceedings, 2015. Vol-1391.