Defining Contextual Factors for News Consumption

Benjamin Kille
Technische Universität Berlin
benjamin.kille@tu-berlin.de

Özlem Özgöbek
Norwegian University of Science and Technology
ozlem.ozgobek@ntnu.no

Andreas Lommatzsch
Technische Universität Berlin
andreas.lommatzsch@dai-labor.de

Abstract

News shape public perception of events. The amount and frequency with which publishers release articles continues to increase. News recommender systems are tools designated to support readers in finding the most relevant articles. These systems struggle with challenging conditions. Readers refuse to express their interests explicitly, and publishers cannot reliably track them to infer their preferences. Publishers have to update their models to accommodate recent trends continuously. In this work, we aim to explore the use of contextual information to improve news recommender systems. We mine characteristic patterns and discuss how these findings can help to develop innovative recommendation strategy and better evaluation protocols.

1 Introduction

Publishers have created digital outlets to accommodate customers' demands for quicker access to recent news. The digital revolution has shifted the publication sector toward an "attention economy" [BO12]. Customers, who relinquish subscribing to a publisher, pay instead with their attention to advertisements. These customers appear to represent the majority as publishers struggle to compete and report declining revenues, especially in print. Consequently, publishers look for innovative ways to keep visitors interested.

Readers, on the other hand, have limited time available to engage with digital media. The amount of published content exceeds their attention capacity by far. This incites intense competition among publishers in which publishers seek to direct the most relevant content to readers. They reason that readers will become loyal to services providing their desired contents. Advertisers, on the other hand, demand evidence for users' attention in order to measure the spread of their information. Consequently, publishers need to monitor their platforms to provide the requested evidence.

Publishers have started to employ news recommender systems to gain competitive advantages. Recommender systems assist readers in navigating more conveniently by suggesting a subset of articles. Ideally, the subset matches readers' preferences. Thereby, readers can avoid searching for exciting content themselves and save time.

The news domain meets existing recommendation algorithms with challenging conditions. Many readers visit publishers' websites without registering a user profile. Consequently, publishers struggle to link readers' sessions. They fail to establish expressive user profiles. Further, articles tend to remain relevant for merely a short period. News stories tend to relate to events which publishers cannot anticipate in advance.

These conditions disallow to directly apply standard algorithms such as collaborative filtering or content-based filtering. Standard algorithms assume somewhat static collections of users and items with sufficiently many interactions between both groups. As a result, content-based filtering and collaborative filtering struggle with the so-called cold-start problem. Therein, the recommender systems lack information about users' preferences toward items. This situation occurs for new users or items. Instead, publishers tend to focus on session-based recommendations, breaking news, and popularity-based recommender approaches. Session-based recommendations fail to capture readers' long-term interests. Breaking news and

popularity-based methods relinquish personalization in favor of simplicity.

In this paper, we analyze readers' engagement with news to determine the importance of contextual information. Contextual information, such as the device type, time, and day, gives us clues about users' contextual news reading behavior. Based on the results from these analyses we argue that publishers should pay more attention to contextual aspects.

Our findings support publishers to enhance their news recommendation systems. In particular, publishers can improve their evaluation practices to avoid suboptimal configurations. Publishers monitor the activity on their websites by recording event-based logfiles. They tend to conduct experiments in the form of A/B tests [KL17]. In these tests, the system randomly assigns users to groups and presents each group a fixed variation of the website. The websites may differ in style, layout, or content. We argue that publishers have to pay attention to contextual information to draw meaningful conclusions from their data. Specifically, publishers risk suboptimal configurations if user feedback correlates with contextual settings rather than the website.

This paper's remainder presents the following parts. Section 2 reviews existing approaches for news recommender systems. Section 3 dives deep into the contextual aspects of news engagement. Section 4 discusses the findings of our data analysis concerning the different news publishers. Section 5 concludes and indicates further directions for research on context-aware news recommender systems.

2 Related Work

Users' perception of items changes over time [SJNP12]. They may enjoy a particular item once but find it less valuable later on. Alternatively, the may like some items in specific circumstances. Context features describe circumstances in which users interact with items. Context-aware recommender systems use contextual data to provide more relevant recommendations. In [AT11], the authors discuss the context dimensions within a recommender system environment and how they change over time. Users' ongoing interactions with the system generate context information. Recommender systems use this information to enhance the instantaneous recommendations or user profiles. They include contextual information either via pre-filtering, post-filtering, or modeling. The approaches differ concerning the stage at which contextual information affects the recommendations. Pre-filtering introduces a set of recommendation models each of which serves a particular context. Post-filtering takes the recommendations and excludes items according to the context. Context modeling integrates contextual information into the recommendation models. Context-aware recommendations have entered a variety of domains including movies [KABO10, LA13, LWW15, GRST10], music [BKL+11, WRW12], locations [YSC+13], tourism [ZBM12], and apps [SKB+12].

The news domain confronts context-aware recommender systems with particular challenges. Accordingly, types and usage of contextual features have to be customized [ÖGE14, KJJ18]. The news domain experiences more dynamics than other domains. For instance, publishers keep extending the set of available items by adding new articles. In addition, users' interests rapidly change, and news recommender systems struggle to anticipate these changes. In [JNT10], the authors explain that users' current context affects their choice of news articles to read. Based on these findings, the authors of [LZYL14] propose a model which extends the user profile with contextual features. Recommender systems strive to consider both short-term and long-term reading interests. As a result, they build user profiles for both settings in an attempt to reflect shorter and more extended periods. Similarly, in [WZL+15], the authors present a hybrid context-aware news recommender system which uses explicit and implicit indicators (e.g. location, ratings, and reading time) to calculate the users' long-term and short-term interests. Herein, capturing long-term preferences represents a significant challenge as systems struggle to recognize reoccurring visitors. In [CC09], the authors propose a semantic runtime context complementing the recommendation algorithm. Thereby, they ease matching user and item profiles. In addition, they point out that widely used contextual features including recently engage items, computing platforms, network conditions, the social and physical environment, as well as the location affect performance. Keeping the context information updated presents another challenge to news recommender systems. This is due to the domain's highly dynamic nature with publishers continuously releasing fresh content. In [GDF13], the authors propose a news recommender system based on context trees in order to establish a dynamically changing model. Context trees model sequences of news, topics, and topic distributions. They evolve following users' behavior and news trends. Context trees' structure facilitates responding quickly to incoming requests as paths usually involve only a few vertices. In [Lom14], the author introduces a context-aware ensemble for news recommendation. The ensemble blends different news recommendation algorithms. Simultaneously, the system records context-dependent performances. Subsequently, the ensemble can determine the best candidate algorithm for a given context.

The paper documents what publishers can gain by considering contextual aspects compared to a popularity-based baseline. In [SSZ18], the authors consider a session as the vehicle of a contextual setting. They argue that readers' interest in news depends on their current situation and propose a session-based news recommender system. The system integrates both collaborative and content-based filtering to create short-term interest models. In [SKP13], the authors consider users' location as important contextual factor. They argue that users reading news on their mobile devices are interested more in the news concerning their local surroundings. Consequently, news recommender systems ought to emphasize articles related to close-by events to match the context. In [LSC+10], the authors have access to a broader set of features describing users and articles. The considered features including users' location and articles' categories. Having reduced the feature space's dimensionality, a contextual bandit algorithm consumes the data. The contextual bandit selects recommendation models in a context-dependent fashion.

The consideration of contextual factors is paramount when evaluating recommender algorithms [BBL+16]. Changing contextual features renders experimental results irreproducible. Hence, evaluation protocols for news recommender systems must consider context to avoid falling for suboptimal configurations.

Overall, the existing research shows that contextual features have a high influence on the performance of recommender systems. There is still a need for a better understanding of contextual factors in news recommender systems. In particular, research has yet to determine which contextual features matter in which domain.

3 Data Analysis

Understanding users' needs and habits is the key to provide suitable recommendations. We expect needs and habits to depend on the context such as the device type—desktop, mobile, or tablet—or the time and day. Users with limited time or browsing on mobile devices—with limited screen size to display articles—might prefer short breaking news over long, detailed stories. Some topics may frequently appear at certain times. Groups of readers with similar taste may appear more concentrated at times, for example, readers may like to read sports news at the weekend more than weekdays.

Other contextual features can be taken into account for a more complex user behavior analysis. However, we think that these three contextual features are the most prevalent ones to start such an analysis. They are also the most common contextual information available in existing datasets that we can use without violating users' privacy. We will analyze each of these aspects by looking at log files in order to verify this assumption.

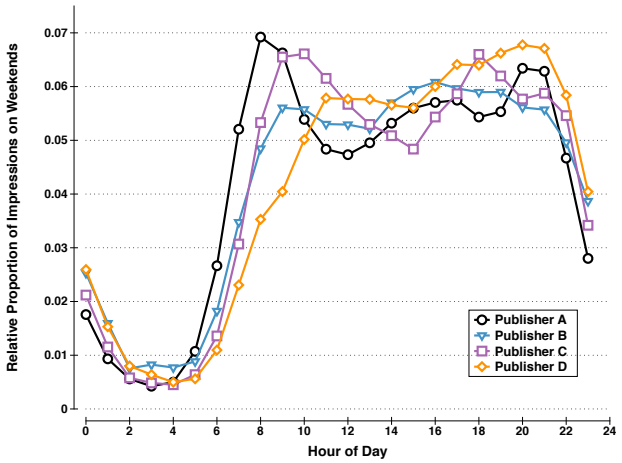

Figure 1: Comparison of the relative proportions of impressions by publishers over the hour of day at weekends. Each curve refers to the relative proportion of impressions for a publisher, only considering impressions occurring on Saturdays and Sundays.

3.1 Approach

Publishers record user engagement on their platforms based on the interaction amid browser and web server. The resulting log files show events along with the context in which they occurred. We analyze four publishers' log files. Three publishers are located in Germany; one is situated in Norway, see Table 1. Three publishers blend regional and general news, whereas one focuses on car-related news. Our analysis considers two contextual dimensions. First, we look at the time at which events took place. Therein, we distinguish between the daytime—measured as the hour of the day—and the weekday. Second, we investigate the kind of device used to read news articles. We consider desktops, mobile devices, and tablets. In addition, we explore how categories affect user engagement.

3.2 Data Description

We introduce the data underpinning our analysis. The log files come from the websites of four publishers referred to as A, B, C, and D. Publisher A's data have been recorded in the time from January 1 to March 31, 2017. The remaining publishers' data have been recorded in March 2017. Whenever a reader has loaded an article from any of the publishers' websites, the system has appended a line to the log file. As a result, we obtain a list of impressions for each of the publish-

ers. Each impression entails a timestamp from which we derive the hour of the day as well as the weekday. Besides, each impression expresses which kind of device had been used to access the news article. Table 2 shows the number as well as the proportion of impressions for each device. Table 3 shows the number and proportion of impressions by weekday.

3.3 User Engagement

Users engage with websites in different ways. Publishers may look at what pages users decide to visit. We refer to this interaction as *impression*. Besides, users click on recommendations, dwell on pages, or write comments. We focus exclusively on impressions. We gauge users' engagement with publishers' services. Impressions occur in various contextual circumstances. We consider the choice of the device, the hour of the day, and weekday as dimensions.

Figure 5 encodes the user activity for various contexts in the form of a heat map. Each row consists of four heat maps linked to a particular kind of device. Each column refers to a particular publisher. The heat maps stretch the time dimensions hour of day and weekday. The colors encode the relative degree of activity in that particular configuration as explained by the legends on each heat map's right-hand side. In other words, the hotter the color, the more this particular context contributes to the overall activity.

All heat maps share one commonality. The nighttime throughout the week shows the lowest level of activity independent of the context or the publisher. The desktop usage tends to concentrate on usual working days. In contrast, at the weekends a smaller number of impressions originated from desktop devices. Tablets' activity appears primarily focused on the evenings with additional activity on the daytime on weekends. Mobile devices, such as smartphones, share this tendency. On the other hand, they spread additional activity throughout the daytime even on working days.

Figure 1 illustrates how publishers' weekend activity changes over the time of day. Again, the nighttime shows comparatively little activity. Still, in the afternoon the levels of activity show distinct patterns.

3.4 Publisher Activity

Publishers release news articles in a context-dependent fashion. Some stories allow journalists to prepare texts in advance. For instance, sport-related organizations, governmental bodies, and publicly traded companies announce schedules for competitions, votes, and shareholders' meetings. Other stories break unexpectedly such as natural catastrophes, traffic accidents, and celebrity deaths.

Figure 6 shows the times when publishers have added new articles. Two clocks refer to the hours of the day for three publishers. The number of new articles is color-coded according to the legend below. All publishers appear to add a large number of articles at night. This phenomenon is particularly striking for publisher D which appears to use the night exclusively to publish content. Publishers of car-related news generally deal with few breaking stories. The remaining publishers add new articles over the course of the day as well.

Figure 3 looks at the number of impressions for articles assigned to different categories. News concerning finance and health show a noticeable peak in the time between 1:00 p.m. and 2:00 p.m. Impressions appear to center around the noon with lesser peaks in the morning and afternoon. Figure 4 illustrates the number of articles published for different categories. Compared to the previous findings, the publications spread more evenly over the working hours. We observe two significant peaks between 10:00 a.m. and 11:00 a.m. as well as between 1:00 p.m. and 2:00 p.m. Figure 2 contrasts the previous illustrations. The ordinate depicts the number of impressions. The abscissa shows the number of articles published. The various categories are color-coded according to the legend. Aside from each point, the figure presents the average number of impressions for each published article. We observe that articles related to finance and health accumulate many more impressions as the remaining categories. Still, the relatively few articles published about literature attract relatively more impressions. Articles related to cars attract the fewest readers.

4 Discussion

Our analysis has shown that user engagement varies considerably with time, device, and publisher. Users hardly engage with news services at night. Working hours show users heavily using desktops in favor of tablets and mobile devices. Evenings and weekends tell a different story. Then, users engage more via mobile devices and tablets. We may expect different device usage over the course of the week and day due to readers' habits and lifestyles. Still, we need to pay attention to these differences to optimize our news recommender systems. We have observed variations among the readers located in different countries. Norwegian readers start reading news earlier in the morning compared to readers in Germany. Conversely, German readers stop later in the evening their exploration of the news landscape. These differences could be the result of culture and lifestyle. Furthermore, our analysis indicates that users' engagement differs between publishers based on the subject and

Table 1: Publisher Description. The domain refers to the publishers' spectrum of topics. The location refers to the publishers' headquarters. The source refers to the data set's origin. The duration refers to the length of the period during which the data have been collected.

Publisher	Domain	Location	Source	Duration
A	general & local	Trondheim, Norway	Adressa [GZL+17]	three months
B	general & local	Berlin, Germany	NewsREEL [KHBH13]	one month
C	general & local	Cologne, Germany	NewsREEL [KHBH13]	one month
D	cars	Berlin, Germany	NewsREEL [KHBH13]	one month

Table 2: Number and proportion of impressions by publisher and device. Publishers are listed as columns, whereas rows refer to devices. The bottom row shows the total number of impressions per publisher.

	Publisher A		Publisher B		Publisher C		Publisher D	
Device	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%
desktop	38 403 480	33.8	5 230 237	34.6	4 770 134	76.3	9 349 859	41.4
mobile	53 906 527	47.5	8 293 081	54.9	139 878	2.2	10 603 232	46.9
tablet	21 269 688	18.7	1 590 700	10.5	1 339 766	21.4	2 640 384	11.7
Σ	113 579 695		15 114 018		6 249 778		22 593 475	

Table 3: Number and proportion of impressions by publisher and weekday. Publishers are listed as columns, whereas rows refer to weekdays. The bottom row shows the total number of impressions per publisher.

	Publisher A		Publisher B		Publisher C		Publisher D	
Weekday	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%	$\nu_{\text{impressions}}$	%
Monday	17 326 061	15.3	1 968 745	13.0	954 698	15.3	3 141 987	13.9
Tuesday	17 574 750	15.5	2 322 038	15.4	958 494	15.3	3 041 217	13.5
Wednesday	17 664 762	15.6	2 511 516	16.6	1 146 296	18.3	3 684 470	16.3
Thursday	17 076 853	15.0	2 576 185	17.0	1 121 812	17.9	3 525 732	15.6
Friday	16 128 126	14.2	2 367 162	15.7	975 581	15.6	3 421 870	15.1
Saturday	12 522 288	11.0	1 627 104	10.8	521 625	8.3	2 762 025	12.2
Sunday	15 286 855	13.5	1 741 268	11.5	571 271	9.1	3 016 174	13.5
Σ	113 579 695		15 114 018		6 249 778		22 593 475	

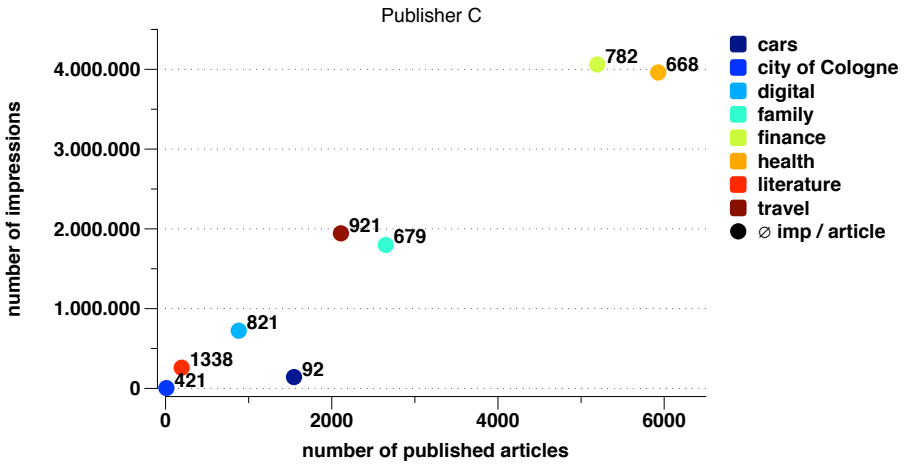

Figure 2: Relation of number of articles published and number of impressions for the categories of publisher C

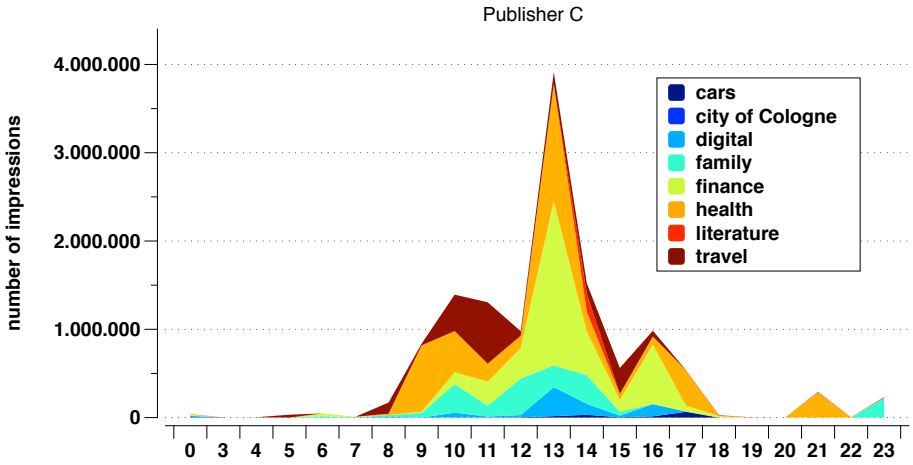

Figure 3: Number of impressions by category for publisher C.

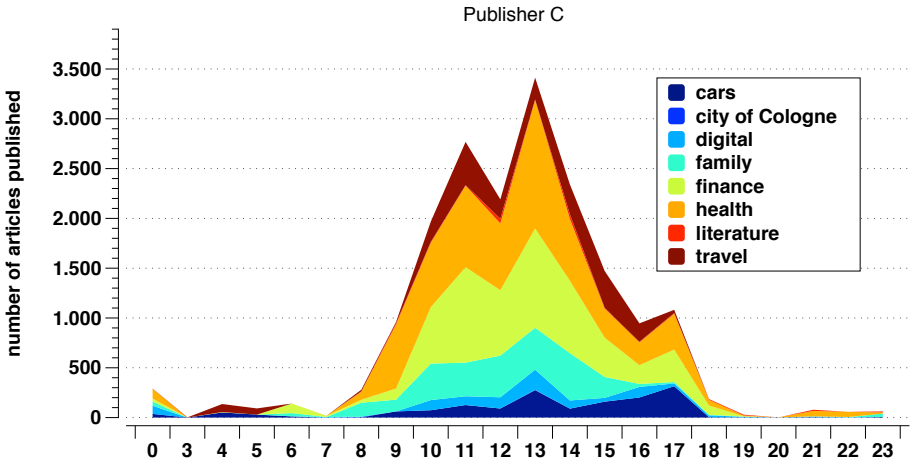

Figure 4: Number of articles published by category for publisher C.

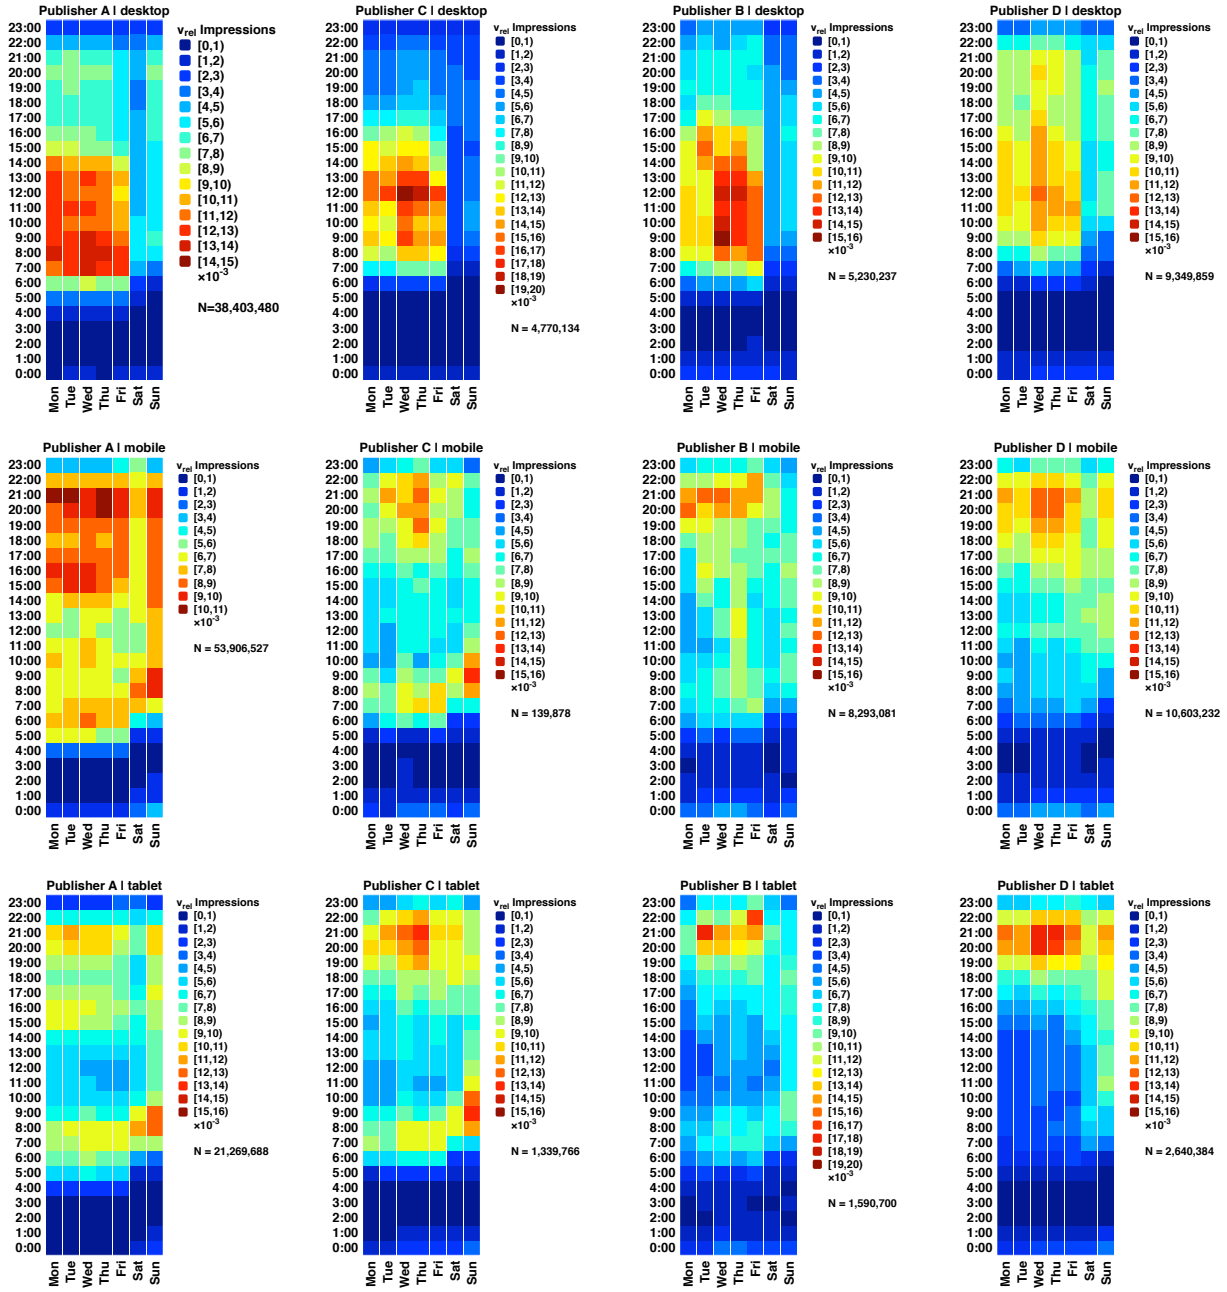

Figure 5: Heat maps showing user activity by device, daytime, and weekday. Each row contains four heat maps related to a specific device. Each column contains three heat maps related to a particular publisher. The number of impressions is color-coded according to the legends on each heat maps right-hand side. The number of impressions in total for the context, or the pair of publisher and device, is shown below the legend.

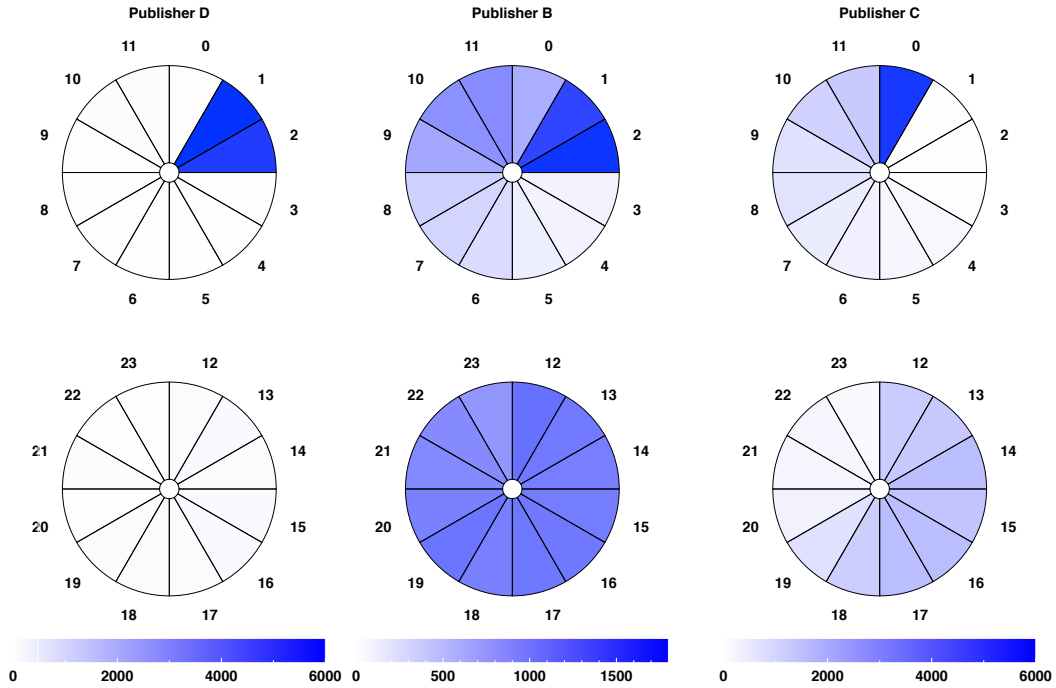

Figure 6: Illustration of the daytime when publishers release articles. For three publishers, each column presents two clocks. The clocks color-code the proportion of articles released over the course of half a day.

type. Topic-specific publishers, such as publisher D, exhibit markedly different engagement patterns compared to general news outlets. As a result, publishers need to pay close attention to their reader base and adapt to their particular needs. Especially, publishers must guarantee that new stories become available before the majority of readers engages the service. Transferring knowledge from different domains represents a challenge. We cannot reliably track users engaging with different publishers. Hence, we struggle to establish links between articles which suitably complement one another.

In addition, our analysis has shown that publishers release articles with varying schedules. Publishers can prepare stories about anticipated events or opinions. News recommender systems have to adapt to situations with radically changing item collections. On the one hand, a push of many prepared stories at night adds uncertainty to the story selection. On the other hand, breaking news stories attract a majority of attention. Publishers may rely heavily on A/B testing protocols to optimize their recommendation services. These evaluation tools monitor the behavior of disjoint groups of users each of which experiences a different system configuration. The longer the data gathering progresses, the more the contextual patterns fade. Hence, publishers optimize their systems for the most wide-spread contexts and abandon potential gains attainable by more fine-grained contextual analyses. Instead, our analysis suggests monitoring user behavior and determining groups with similar preferences. This opens up the opportunity to introduce a low-level personalization which also respects users' anonymity and privacy. Moreover, publishers seek to maximize long-term business goals. Some publishers have managed to establish a large enough group of subscribed users to cover their costs. For them, adding subscribers takes precedence over short-term user engagement. Publishers need to align their recommendation strategies to their business goals. Consequently, recommendation algorithms face different challenges depending on the context. Tuning the recommendation algorithm to users' needs promises to help in achieving long-term goals. Presenting shorter articles on mobile devices and topics popular in particular contexts, represent two directions publishers could turn. Additionally, publishers ought to consider contextual differences as they internationalize their services. Our analysis indicates that culture and lifestyle affect readers' perception of recommendations. Consequently, publishers need to take into account their recommendations' destination.

5 Conclusion and Future Work

Publishers operate in dynamic, digital environments. They compete for the attention of users to monetize their content. News recommender systems facilitate readers' access to information. As a result, publishers continue to optimize their recommendations. Our

analysis has shown that readers' engagement varies considerably in between contextual settings. We have considered the time and device as contextual dimensions. This kind of engagement reflects readers' daily routines. We have observed negligible activity at night, mostly desktop usage on working hours, and mobile device usage in the evenings and on weekends. Publishers release schedules vary considerably. Prepared articles emerge in the night while stories related to breaking events enter irregularly. The example of publisher D shows that news categories differ in popularity. Besides, the categories' popularity changes over the course of the day.

We see multiple directions to extend this line of thought. Publishers ought to devise methods to introduce context-awareness to their systems. Conducting user studies yields clearer insight into readers' requirements. Alternatively, publishers could extend A/B testing to encompass longer periods and simultaneously monitor contextual features. Further analysis of contextual reader engagement could include additional publishers to verify the discussed results. Establishing publisher-specific recommendations demands a more in-depth analysis of contextual information and readers' behavior.

References

[AT11] Gediminas Adomavicius and Alexander Tuzhilin. Context-aware recommender systems. In Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor, editors, *Recommender Systems Handbook*, pages 217–253. Springer, 2011.

[BBL+16] Jöran Beel, Corinna Breitinger, Stefan Langer, Andreas Lommatzsch, and Bela Gipp. Towards reproducibility in recommender-systems research. *User Model. User-Adapt. Interact.*, 26(1):69–101, 2016.

[BKL+11] Linas Baltrunas, Marius Kaminskas, Bernd Ludwig, Omar Moling, Francesco Ricci, Aykan Aydin, Karl-Heinz Lüke, and Roland Schwaiger. Incarmusic: Context-aware music recommendations in a car. In *International Conference on Electronic Commerce and Web Technologies*, pages 89–100. Springer, 2011.

[BO12] Erik Brynjolfsson and JooHee Oh. The attention economy: Measuring the value of free digital services on the internet. In *ICIS*. Association for Information Systems, 2012.

[CC09] Iván Cantador and Pablo Castells. Semantic contextualisation in a news recommender system. *Workshop on Context-Aware Recommender Systems*, 2009.

[GDF13] Florent Garcin, Christos Dimitrakakis, and Boi Faltings. Personalized news recommendation with context trees. In *RecSys*, pages 105–112, New York, 2013. ACM Press.

[GRST10] Zeno Gantner, Steffen Rendle, and Lars Schmidt-Thieme. Factorization models for context-/time-aware movie recommendations. In *Proceedings of the Workshop on Context-Aware Movie Recommendation*, pages 14–19. ACM, 2010.

[GZL+17] Jon Atle Gulla, Lemei Zhang, Peng Liu, Özlem Özgöbek, and Xiaomeng Su. The adressa dataset for news recommendation. In *Proceedings of the International Conference on Web Intelligence*, pages 1042–1048. ACM, 2017.

[JNT10] Jeremy Jancsary, Friedrich Neubarth, and Harald Trost. Towards context-aware personalization and a broad perspective on the semantics of news articles. In *Proceedings of the Fourth ACM Conference on Recommender Systems*, RecSys '10, pages 289–292, New York, 2010. ACM.

[KABO10] Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 79–86. ACM, 2010.

[KHBH13] Benjamin Kille, Frank Hopfgartner, Torben Brodt, and Tobias Heintz. The plista dataset. In *Proceedings of the 2013 International News Recommender Systems Workshop and Challenge*, pages 16–23. ACM, 2013.

[KJJ18] Mozhgan Karimi, Dietmar Jannach, and Michael Jugovac. News recommender systems - Survey and roads ahead. *Information Processing and Management*, 2018.

[KL17] Ron Kohavi and Roger Longbotham. Online controlled experiments and a/b testing. In *Encyclopedia of Machine Learning and Data Mining*, pages 922–929. Springer, 2017.

[LA13] Xin Liu and Karl Aberer. Soco: a social network aided context-aware recommender system. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 781–802. ACM, 2013.

[Lom14] Andreas Lommatzsch. Real-time news recommendation using context-aware ensembles. In *European Conference on Information Retrieval*, pages 51–62. Springer, 2014.

[LSC+10] Lihong Li, Robert E Schapire, Wei Chu, John Langford, and John Langford. A contextual-bandit approach to personalized news article recommendation. In *the 19th International Conference on the World Wide Web*, pages 661–670, New York, 2010. ACM Press.

[LWW15] Qiang Liu, Shu Wu, and Liang Wang. COT: Contextual operating tensor for context-aware recommender systems. In *AAAI*, pages 203–209, 2015.

[LZYL14] Lei Li, Li Zheng, Fan Yang, and Tao Li. Modeling and broadening temporal user interest in personalized news recommendation. *Expert Systems with Applications*, 41(7):3168–3177, 2014.

[ÖGE14] Özlem Özgöbek, Jon Atle Gulla, and Riza Cenk Erdur. A survey on challenges and methods in news recommendation. In *Proceedings of the 10th International Conference on Web Information Systems and Technologies (WEBIST)*, volume 2, pages 278–285, 2014.

[SJNP12] Alan Said, Brijnesh J Jain, Sascha Narr, and Till Plumbaum. Users and noise: The magic barrier of recommender systems. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 237–248. Springer, 2012.

[SKB+12] Yue Shi, Alexandros Karatzoglou, Linas Baltrunas, Martha Larson, Alan Hanjalic, and Nuria Oliver. Tfmap: optimizing map for top-n context-aware recommendation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 155–164. ACM, 2012.

[SKP13] Jeong-woo Son, A-yeong Kim, and Seongbae Park. A Location-Based News Article Recommendation with Explicit Localized Semantic Analysis. *SIGIR*, pages 293–302, 2013.

[SSZ18] Gabriele Sottocornola, Panagiotis Symeonidis, and Markus Zanker. Session-based News Recommendations. In *the 2018 World Wide Web Conference*, pages 1395–1399, New York, 2018. ACM Press.

[WRW12] Xinxi Wang, David Rosenblum, and Ye Wang. Context-aware mobile music recommendation for daily activities. In *Proceedings of the 20th ACM international conference on Multimedia*, pages 99–108. ACM, 2012.

[WZL+15] S. Wang, B. Zou, C. Li, K. Zhao, Q. Liu, and H. Chen. Crown: A context-aware recommender for web news. In *2015 IEEE 31st International Conference on Data Engineering*, pages 1420–1423, Seoul, 2015. IEEE.

[YSC+13] Hongzhi Yin, Yizhou Sun, Bin Cui, Zhiting Hu, and Ling Chen. Lcars: a location-content-aware recommender system. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 221–229. ACM, 2013.

[ZBM12] Yong Zheng, Robin Burke, and Bamshad Mobasher. Differential context relaxation for context-aware travel recommendation. In *International Conference on Electronic Commerce and Web Technologies*, pages 88–99. Springer, 2012.