Predicting the success of news

Using an ML-based language model in predicting the performance of news articles before publishing

Atte Jääskeläinen†
School of Business and Management
LUT University
Finland
atte.jaaskelainen@lut.fi

Elli Taimela††
Fourkind Oy
Finland
elli.taimela@fourkind.com

Tomas Heiskanen††
Fourkind Oy
Finland
tomas.heiskanen@fourkind.com

ABSTRACT

Traditional recommendation systems have limited possibilities to optimise business value in editorial decision making in news production, as they select the recommendations only from the content whose production has been decided editorially in the daily news process or content from existing content inventories. This paper explores an approach to use predictive analytics to make it possible to optimise story assignment and editing in daily editorial work based on selected business objectives already before publishing. In this case study exploration, we use the 'constructive approach' as a method to provide solutions to concrete business problems with a scientific approach. We contribute by experimenting a novel method combining elements from several scientific domains like strategic management and system dynamics. We conclude that with language analysis using recurrent neural networks, we were able to predict the success of a news story published on a digital channel in a way that fulfils the 'weak market test' criteria of the constructive approach. A company with whom the model was developed considered it valuable enough to decide to move it from exploration to be further developed and used in real news production.

†Atte Jääskeläinen is Professor of Practice at LUT University, Finland. †† Elli Taimela is a consultant and Tomas Heiskanen an artificial intelligence consultant at Fourkind Oy, Finland.

AcademicMindtrek '20, January 29–30, 2020, Tampere, Finland

ACM ISBN 978-1-4503-7774-4/20/01.
https://doi.org/10.1145/3377290.3377299

CCS CONCEPTS

•Applied computing~Operations research•General and reference•Applied computing~Arts and humanities

KEYWORDS

Editorial predictive analytics, Recurrent neural networks, Constructive approach, Digital News Media

ACM Reference format:
Atte Jääskeläinen, Elli Taimela and Tomas Heiskanen. 2020. Predicting the success of News. Using an ML-based language model in predicting the performance of news articles before publishing. In Proceedings of Mindtrek Academic Conference, 29.-30. January 2020, Tampere, Finland.

ACM ISBN 978-1-4503-7774-4/20/01.
https://doi.org/10.1145/3377290.3377299

1 Introduction

Traditional news media have long been struggling with the declining business. Since the introduction of internet-based technologies in news distribution, competitive forces have significantly strengthened because technological barriers for competition have lowered and information needs of consumers can be fulfilled with freely available substitutes, like real-time information from social media. [35,65]

For professional news media, this development has forced them to reconsider the value fundamentally their content creates for the customers, including in private news media news end-consumers, advertisers, or both, or even sponsors of content, depending on the business model. Fundamental changes in the experienced customer value have also forced the news media to rebuild the value creation and value capture components of their business models. [12,13,25,30,31,66,67]

An essential part of the value creation logic of a business model is to optimise resource usage. In news media, this has been influenced by the rise of data-informed and data-driven culture in newsrooms, with the help of real-time analytics tools and metrics indicating news usage by the news consumers [11]. When this information of news usage is available, it is possible to react for it and find more valuable ways to create value with the resources available. For example, a business objective may be increasing advertising revenue, which is translated to editorial targets like increasing the number of visits or unique visitors. [6,11,60,61]

For advertisers, who represent a significant income source for traditional news media in both print and broadcasting, new technologies have offered more cost-effective means to reach their target audiences. In marketing, analytics and automation have developed earlier and further than in newsrooms, because the incentives to use all possible means to reach business goals have fostered that and cultural barriers to using these methods have not existed in marketing in the same extent as in the newsrooms. [64]In automated sales and delivery of digital advertising, called 'programmatic' advertising, the whole process from selling to publishing advertisements on a website is controlled by software [83,89]. The majority of digital advertising is now spent on search and social media platforms [26,62].In 2017, advertising connected with web search represented 20% of the world's advertising spending, and the share of Google of this was 85%. In 2018 the share of Google and Facebook of the world's digital advertising was 56% and the share of all global advertising revenue around 25%.[5,7,8,56]

After the problems faced in the digital advertising market caused by intensified competition, business models of news media have moved towards paywalls and other models based on user payments [33,82].

There are lots of examples of recommendation systems employed in news services, including individual targeting [2,19,64]. The algorithms used in these systems are based either on content, user, social connections, or their combinations, called 'hybrid models' [69]. Some organisations are also exploring the usage of optimisation algorithms with target functions that reflect public interest objectives of the newsrooms [28,29,74], and also address the problems related to dynamic changes typical for news content and news user preferences, even predicting user preferences [90].

There are already systems in news media that predict customer behaviour and target both content and marketing actions to users that are predicted to be potentials for cancelling their subscriptions or potentials for subscribing. However, in most approaches focusing in recommending content, automatic targeting uses only the available content, either produced by the newsrooms at the moment or acquired from other sources (like news agencies or content produced in the same group company but by other brands) or archives.

Predictive analytics is a practice to forecast future behaviour and events based on patterns found and analysed in existing data [1:3–17,24]. Predictive analytics methods such as machine learning can be used to support business objectives such as growth and profitability through foreseeing changes in the market and automating the analysis, decision-making, and feedback loops.[73]. In the existing applications within the news media industry, new and historical data is typically analysed to foresee customer behaviour, and apply the insights in, e.g. automated targeting and recommendations to increase conversion. The existing applications typically bring value for media businesses during and after the publication of the news content, but applications in news content creation and other steps before publishing the new content are rare. [2,18,63,84]

Information from advanced data-informed newsrooms has revealed significant variances of how news content is consumed. In many cases, analytics tells that a significant part of the news content published on the website does not find any commercially or otherwise relevant audience, and only a few well-performing articles generate most of the traffic [12–14,25,44,61,70,78]. There can,, therefore, be seen opportunities in predicting and optimising the potential value of the content before its publication. Applying, for example, modern Natural Language Processing methods and machine-learning combined with predictive analytics methods might enable predicting the success of news, to maximise the value creation of newsrooms in the increasingly strengthening competition.

However, this kind of approach has a significant risk of being confronted with traditional journalistic culture and identity, which underlines the independence of the newsroom, public interest and journalistic decision-making based on professional judgment and intuition. There are strong identity-based fears of losing the gate-keeping role of newsrooms they historically had. [20,21,49,64] The data-informed culture has challenged these traditional values of journalism and revealed a conflict between audience preferences and the preferences of the newsroom [89]. Cultural and technological changes are also confronted with other types of tensions, like fear of technology or divided between journalists and journalistic managers [9,10].

Despite these tensions, we argue that if the content that is produced is not consumed, that constructs we argue that there is room to improve business value by augmenting journalistic decision-making with more insights from data and analytics and predicting audience preferences for content. Combining an accurately predicted value according to an agreed objective enables more data-informed decision-making in the newsroom and has the potential to increase business value. [14,25,44,61,70,78] To what extent these predictions are used directly in automation or how much they affect journalistic decision-making is an editorial decision.

Based on these findings and observations, we consider it to be of great interest to study whether it is possible to accurately and automatically predict the value of news content before it is published, to increase the capability to manage the newsroom resources in a way that they are better utilised in the value-creation process. If proven to be possible, this kind of analytics would allow the news managers to assign stories whose probable contribution to the success of the news brand is better than with human assignment, and for journalists that would allow considering means to make changes to the story before it is published, thereby increasing its value.

2 Related work

Most of the research relevant for this study have utilised quantitative methods in analysing the success of news and supporting the editorial decision-making in newsrooms. Predicting the success of online news has interested scholars for more than ten years. The success of a published or non-published content has been tried to be predicted through various metrics such as the number of page views, comments, recommendations or shares in social media, and with different time spans and advanced analytics methods. Many studies have been one-off trials and some have focused on comparing different predictive analytics methods, but some studies have also targeted to build a tool to support editorial decision-making in a real-life newsroom.

Some studies have focused on modelling the behaviour of the audience in the discussion forum linked to the content. Kaltenbrunner et al. used the technical website Slashdot's material and information derived from its discussion community and on temporal patterns. They found repeating oscillatory cycles in audience's commenting behaviour in online news. This work's contribution was mainly in the field of understanding how the audience growth, in general, cumulates over time, finding regular daily and weekly deviations from general approximations. [36,37]

Predictions of the audience's behaviour in discussion forums or commenting sections have also been utilised in building predictive models to support editorial decision-making. Tsakias et al. examined the possibility to predict the number of comments, indicating its importance or impact, of both features of the story and commenting volume shortly after publication, trying to find means to identify stories with potential to grow audience fast and have an effect on front page positioning [85,86]. Tatar et al. used user comments in predicting the popularity and ranked articles based on their predicted popularity, proposing a predictive model to automatic news ranking on websites [81]. Lee et al. created a model to predict the likelihood that content on a discussion forum will be more popular than a set threshold in metrics like lifetime of threads and number of comments using only publicly observable metrics like views, links and comments to a story, but from an observed period counted in days after the publication. [45,46]

Many studies have attempted to predict the success of a news story or other content by analysing the reactions of the audience after publishing the content. Szabo and Huberman modelled views and votes in Youtube and Digg to predict long-term success (30 days) from early hours after publication, reaching greater accuracy in stories whose lifespan was short than with evergreen content [77]. Kim et al. tried to figure out the characteristics of a successful article from early observations and found that the dynamic system, surrounding content in online space greatly impacts the long-term success of the content and limits the performance of using only early data from the interactions. [41,42] Lerman and Hogg tried to sort out elements of web site design from the prediction model using stochastic models and improving the results thereby [48].

The success of news has also been predicted by following the behaviour of the audience in social media channels. Bandari et al. took the challenge of trying to support editorial decision-making prior publication. They used features such as the news source, the category of news, the subjectivity of the language and the named entities to create a multi-dimensional model to predict the success of the story on Twitter before publication, using the number of times the URL is posted and shared. They could reach 84% accuracy with this method [3]. Lerman with Ghosh analysed the role of the network in the spread of the news and influencing the network [47].

Predicting the attractiveness of content has also been studied by modelling the audience's interests. This was addressed by De Francisci Morales et al. using real-time information from the user's Twitter account (social connections, posted content and topic popularity) and modelling the relevance of Yahoo News stories based on that, reaching good accuracy [19].McCreadie et al. used blogosphere as an indication of stories' importance. [54,55]

Many studies have focused on comparing different machine learning or other advanced analytics methods in predicting the success of news. Fernandes et al. compared different data mining and machine learning approaches to predict the popularity of online news using Mashable stories as their material. They found Random Forest and Neural Network as best and reaching an accuracy of 65% with optimal parameters. [27]Namous, Foad and Javed compared different data mining algorithms in predicting the popularity of online news, using pre-analysed features of and found Random Forest and Neural Networks being the best ones with an accuracy of 65% with optimal parameters. [59]

Some studies have also proved that the content itself can be utilised in predicting its success. Ren and Quan also worked with Mashable data, finding Random Forest having the best performance with 70% accuracy, underlining the importance of feature selection in improving the accuracy. They already pinpoint the possibility of using all words of the content as features and using machine learning algorithms in improving the accuracy of the predictions. [69] Hardt and Rambow worked with data from Jyllands-Posten from Denmark and found that text features are predictive of user's behaviour and using the whole text as feature set was more predictive than the headline, or the headline with the teaser. [32]

A few studies have also attempted to build tools to support editorial decision-making related to for example resource allocation in real-life newsrooms. Keneshloo et al. faced the challenge of helping the editors to allocate the resources to the most interesting articles and to support a better reading experience. Their features were harvested in 30 minutes after publication (metadata, contextual or content-based, temporal and social) and the popularity prediction system was deployed in practice in the Washington Post. [40] One of the best-known examples of computational tools assisting newsroom in social media monitoring and also predicting the newsworthiness of the information is Reuters Tracer, whose development has been openly documented in academic papers. [50–52]

Topic and event detection are another relevant research areas of predictive analytics in the field of news media. With predictive analytics looking the topic or event as a feature, it might be helpful for the newsroom both to monitor the environment and find newsworthy events and monitor the behaviour of the audience and predict their interests, perhaps even matching these two approaches. Cucchiarelli et al. have studied extensively the possibility to detect events and topics that fulfil more complex evaluation criteria than the number of clicks. For evaluation, they have developed the concept of serendipity, combining novelty (unexpectedness, surprise) and salience (usefulness, relevance). In their methodology component metrics are combined, and journalists are proposed with topics in breaking news situations in which there is an information need (detected from online news, Twitter and Wikipedia), but not relevant supply for aspects not covered by the news media. [15–18]

Considering the methodology, we find the work by Diakopolos et al. interesting in the sense that they used a design process to produce a use case of a useful prediction tool to spot eyewitness information from social media, arranging interviews with reporters to understand their use of social media and journalistic information needs before developing the tool [22]. Similar features of the approach are found in the work of Zubiaga et al. [91]

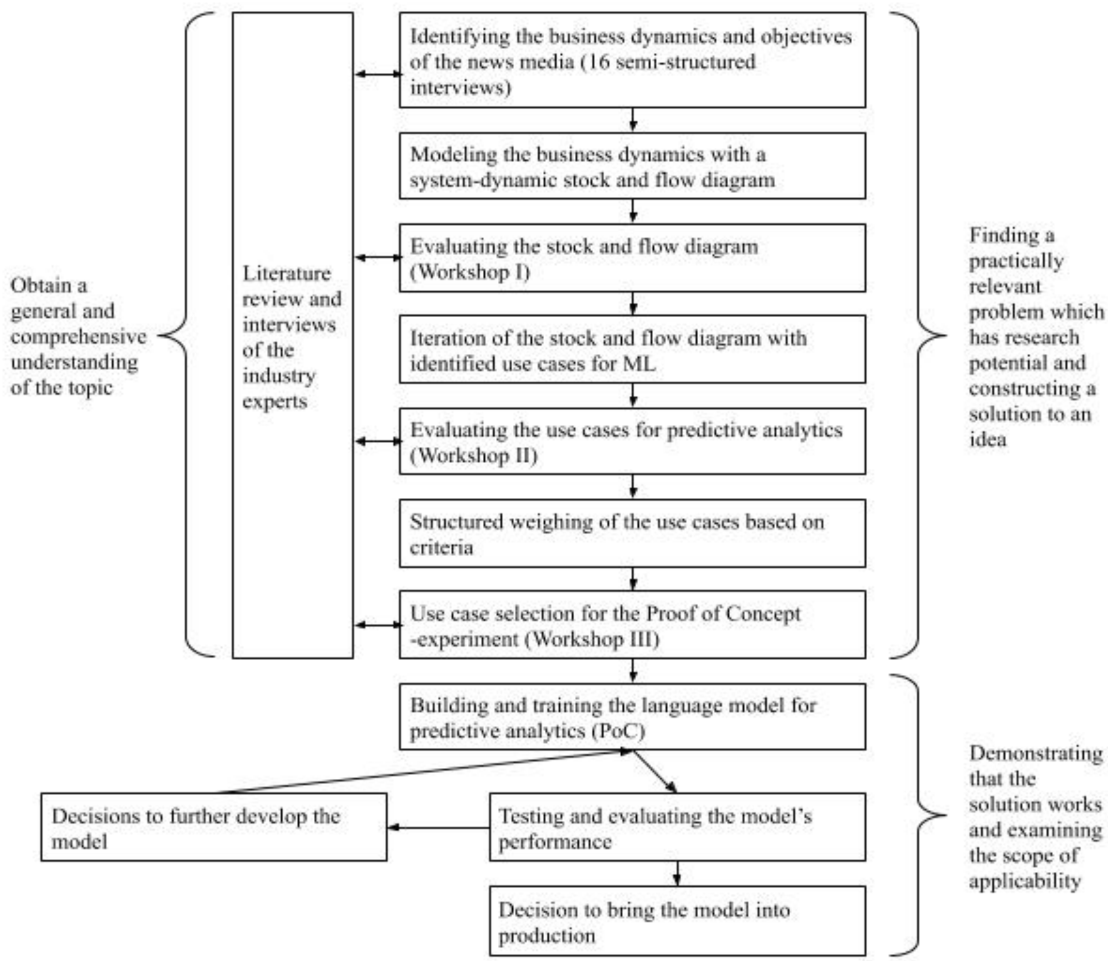

Figure 1: The constructive approach used in this study.

3 Research question and methodology

In this paper, we present the findings of an experimental, qualitative case study that was conducted with a news media company publishing regional newspapers and a tabloid newspaper, with different business models for these two types of papers. The regional papers' critical strategic objective was to increase digital subscriptions, while the tabloid newspaper business model was based on advertising. On top of that, both papers had editorial objectives that were not directly derived from business objectives but were either related to desired brand perceptions and public interest aims of the newsrooms and their management.

3.1 Research question

RQ: How and to what extent can a machine-learning-based language model that predicts the success of a news story before it is published create business value?

We approach the research question with a case study based on the so-called constructive approach, which can be characterised by dividing the research process into phases from finding a practically relevant problem, then proceeding through understanding it towards creating a solution idea, then to demonstrating that the solution works, and then showing the theoretical connections and applicability of the solution [39]

We evaluate the usefulness of the explored and tested machine-learning solution through its applicability in the practical test, which in this case is the decision whether the tested model is accepted and put into production in real-life. That is the so-called 'weak market test'. [39:253]

Figure 1 illustrates the process towards the decision about the use-case to be tested in Proof of Concept, and after that, the weak market test, modifying the constructive approach first introduced by Kasanen et al. [39:246–247]

In this paper, we have modified this approach and sourced from other fields of methodology and created a structured approach first to map the impact of decision-making in the newsroom to the business objectives of the organisation and to find a consensus of the cognitive understanding about those. After that, with these results, we structurally supported the decisions to prioritise the topic of the experiment we were supposed to conduct to test the value-creating ability of the chosen technology. We explain in detail the phases of this approach in the following chapters.

3.2 Finding a practically relevant problem and constructing a solution to an idea: Utilising system-dynamic tools for analysis

The cognitive assumptions of the relations between business dynamics (meaning interconnections between resources and decisions and actions taken in the newsroom) and business objectives of the news media company were identified through semi-structured interviews with key stakeholders and managers in the case company. The 16 interviews were recorded, transcribed and analysed.

The results of the analysis were presented through visualising the causalities identified in the business environment through a system-dynamic stock and flow diagram [75:135–229,76], The diagram visualises the newsroom's publication process, the assumed drivers for revenue, and the assumed causes-and-effects of editorial decision-making in the system. The diagram was presented for the news media company, and feedback was collected for iteration.

This phase was designed and conducted to create a common consensus about the cognitive state in the news media about the causalities between these different factors and to foster a deeper understanding about the possibilities to create business value with predictive analytics. The method was inspired by the research tradition of business cognition and causality maps. [4,38,53,71,87]

Based on the causality analysis and the research group's professional understanding and experience of machine-learning systems, 15 different use case opportunities for predictive analytics were identified. These included opportunities to utilise machine learning in the augmentation of decision-making of the news media before publishing content, during the publication, and after publishing the content. The use case opportunities were visualised and mapped on top of the stock and flow diagram and evaluated with the news media company. Based on the evaluation, the business value, the risks and the investment costs of the use case opportunities were evaluated. Based on this analysis, the news media company's representatives prioritised one use case for the experiment as a Proof of Concept.

The scope of the use case selected for further exploration and to build a Proof of Concept was to estimate the contribution of a news article to a consumer making a digital subscription or staying as a subscriber. The news brand that was selected for exploration was a regional newspaper with high growth targets for their digital subscriptions.

3.3 Demonstrating that the solution works: Building and testing a machine-learning-based language model

In the Proof of Concept, a machine-learning-based model was built to predict the success of a news story before its publication. The success of an article was defined as the following metrics that were chosen to be predicted:9ge model was trained preliminarily first for general features of the Finnish language with the full content of Finnish Wikipedia, and after that about the features of the specific news language with archived stories of all the newspapers of the publisher in question in the experiment. For fine-tuning the model, we also used Universal Language Model Fine-tuning (ULMFiT), a transfer learning method for Natural Language Processing, which also allows the language model to be effectively used in other contexts. [35]

Three predictive linear regression models, one for each success metric chosen for predictions, were built on top of the fine-tuned neural language model by changing the last activation function of the language model. The regression models were based on historical usage data on news articles of the newspaper in question, and they were developed to predict 1) the page views of the articles once digitally published among digital subscribers, 2) the page views of the articles once digitally published among non-subscribers but regular readers of digital content, and 3) the number of subscriptions made immediately after reading the article.

We evaluated the performance of the regression models by MSE between the actual page views and the predicted page views. After testing, the model was decided to be developed further by enriching the model with content metadata and changing the reporting logic for interpreting the results. The reporting logic was changed such that instead of page view count we predicted page view classes that represent ranges for page views. The change was made because the magnitude of the page views was considered more important than the exact page view count. When the exact page view count was predicted, the journalists were concentrating on the difference between the actual and the predicted page views independent of the scale. Since the data was skewed and predicting a high number of page views was a hard task for regression, the model was changed into a classification model. The classes were created using quantiles of data, ensuring that all the classes were equally balanced. The testing of the new versions of the exploration was conducted by classifying the news content based on the estimated and the actual page views and analysing the accuracy of the predictions with a normalised confusion matrix.

The decision to bring the model into production was conducted by the news media company's editorial management, and the decision was based upon the perception of the value that the predictive model can bring to the daily operations of the newsroom.

As this paper is based on a real business case, we agreed on the confidentiality of business sensitive information and news consumption data used in the experiment. Therefore all the details of the tests can not be revealed in this paper. The news media company in question has given consent to the level published in this paper

4 Findings

The key finding of the study is that a machine-learning-based language model that predicts the success of a news story was seen valuable for a news media company as it was seen helpful in encouraging actions to develop further the content before its publication. The actual performance of the model was dependent on the quality and the amount of training data for the model, the selected metrics for the success of the news story, and assumably the number of external factors that relate to the predictability of success. However, the business value of the model perceived by the management of the newsroom was also dependent on what they considered relevant - it did not matter if the model could not predict everything if it predicted the most relevant outcomes from the business perspective.

In general, the business model and revenue streams of the news media company were relevant factors affecting in which use cases predictive analytics was seen valuable by the news media company. During the phase to find a practically relevant business problem for utilization of predictive analytics, machine learning and especially Natural Language Processing were in general seen as valuable methods for predictive analytics for a newsroom. When prioritizing different use cases, the news media case company saw the highest potential in use cases that contributed to decision making before publishing the content, e.g. content creation. The use case selected for the Proof of Concept was seen especially valuable due to the regional newspaper's high growth targets in increasing the number of their subscribers. Because the quality of the content was seen as an essential cause for subscriptions, predictive analytics in content creation was perceived as valuable in bringing new subscribers and keeping existing subscribers.

Another finding was that the metrics to evaluate a language model's performance affected th e business value that the newsroom perceived. At first, the average mean differences between the predicted and actual page views were analyzed, but when the same results were presented through the classification model, they

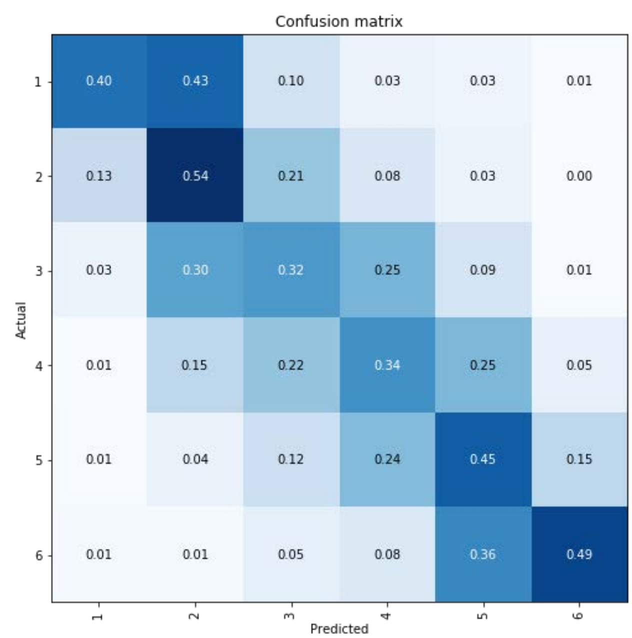

Figure 2: An example of a normalized confusion matrix used in the evaluation of the model. In this case, the predicted page views among registered users were compared with the realized page views.

became more understandable and then the same language model was considered giving more practical business value and more precise signals for action in the newsroom.

In analyzing the performance of the model, the number of page views was still quite small in a regional digital newspaper, and the absolute numbers in the test batch were small. The number of subscriptions made after single articles was so small that a model predicting the number of subscriptions after reading each article was not practically feasible due to a high error rate.

When evaluating the performance of the first version of the model, the MSE between the predicted and actual page views among subscribers was 412 page views, and 75% of predictions fell within that range. In 50% of the tested cases, the model could predict the number of page views within a margin of 182 page views. Most of the stories got a small number of page views, and the model was able to predict those stories with greater accuracy. Among non-subscribers, but regular readers, the average difference was 384 page views, and in 70% of predictions the difference between predicted and actual page views were less than 390 page views.

When a confusion matrix was created from the results of the page views among subscribers, we could see that a vast majority of the actualized number of page views were close to the predictions and very few of the predictions were entirely out of the close range of the actual numbers. This can be seen in the graph above that

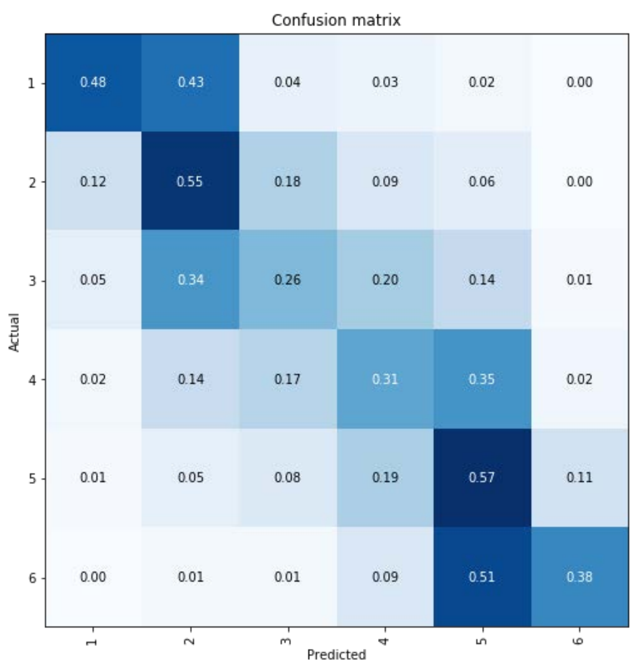

Figure 3: An example of a normalized confusion matrix used in the evaluation of the model. In this case, the predicted page views among regular non-registered users were compared with the realized page views.

represents one of the normalized confusion matrices created during the experiment.

On average with this 6-point scale, the share of precisely right predictions (prediction was in the same category than the actualized page view number) was 41%, but only a tiny share of the predictions was further away than in the next category.

When analyzing the performance of the model among non-registered but regular users, on average with this 6-point scale, the share of precisely right predictions (prediction was in the same category than the actualized page view number) was 42%. Again, the number of completely wrong predictions was very small. The graph in Figure 3 presents the accuracy of the prediction among regular but not registered users.

Another important observation was that with the categories used in these tests, the model was useful in predicting both very small numbers and very high numbers, which both were considered having significant business value.

Since the regression models were built on top of a neural language model that can be considered a black box, feature importances were naturally not able to be analysed. However, we could see that the beginning of the article correlates with the page views more than its ending. When journalists used the regression models in practice, they noticed that changing the title or the first few sentences of an article affected the outcomes of the predictions more than changing the last few sentences of the article.

This small-scale experiment succeeded to create a model that, according to the representatives of the case company, "exceeded expectations". A language model could predict both very high and very low page views on a level that fulfilled the business requirements and therefore, were shown to create business value. During the process, it became evident that the model should be developed further to sort better out other factors, both internal and external, that contribute to the success of the story. This was decided to be done in further development of the model after the finalisation of this paper. During early experiments, we already could see differences between different news media and their audiences' different consumption patterns.

As the experiment is still a work in progress, we can only report conclusions drawn while this paper is being submitted. After choosing, building and testing a model for the initial scope of the Proof-of-Concept, the news media case company has made decisions to

1) further develop the model by adding more metadata (such as publication time, weekday of publication, time of the year, position on the front page, among others) into it, to better differentiate the reasons for the success of an article, and to better differentiate the elements in a story that contribute to the desired outcome from other, often factors external to the story itself,

2) broaden the experiment to one more newspaper, with a different business model and therefore different business objectives, and

3) still further develop the model and bring the model into a production environment, but with a three-class model that suits better to the business needs of the newsroom.

5. Discussion

Although when evaluated purely statistically the model still requires further development, there were already elements in the model that the news organisation considered having practical business value. It was assumed that the model is able to augment human decision in critical decision points in a highly competitive market. It became clear that the way results are presented contributes strongly to the perceived business value for the organisation. When, for example, the same results were presented through the classification model, they became more understandable. There was also significant discussion of how the classes should be defined so that they carry a clear message that is consistent with the editorial management's strategic goals. For example, there is a difference between expressing a prediction that a story belongs to "the best third" of all stories, or expressing that a story has reached a level that is set as a target by the news management.

In evaluating this kind of model, it is obvious and even desirable that the model is not fully accurate. The pure nature of the news is that they should be new and surprising pieces of information. If the success of a news story can fully be predicted from historical patterns, that might be rather a signal of suboptimal operations in the newsroom instead of success in predictions. How using this kind of models affects journalistic judgment in practice is an interesting topic for further research.

Implementing predictive analytics in daily news work in the newsroom may face the same kinds of cultural challenges as introducing real-time analytics, because traditional professional ethics of journalism carry a strong sense of public interest and identity of a gatekeeper, and even the most commercial newsrooms have non-commercial editorial preferences [11,64,79,80]. We will investigate also these issues in our further work.

This paper is based on work that is in progress. Therefore some of the results are preliminary, and the experiment continues. As this is a real business case, some of the results are agreed to be kept confidential. The constructive approach used in this paper might be of interest to the research community, as it is rather rare to collaborate with media companies in real business experiments especially on data that is business sensitive. The approach allowed the usage of the sensitive data and the achievements of the experiments to be reported for both the academic and industry community with the aim of increasing the impact of the results achieved.

ACKNOWLEDGEMENTS

This study received funding from the Media Industry Research Foundation of Finland and the partner media company in question. The authors wish to thank the three anonymous reviewers for their valuable contributions and remarks and Professor Kirsimarja Blomqvist, LUT University, for comments and help with the methodology and Professor Ari Jantunen for comments and help with issues related to organizational cognition. Pauliina Karell, Machine Learning Consultant, Fourkind Oy joined in participating in the further development of the model and contributed in finalising this paper.

REFERENCES

[1] Dean Abbott. 2014. Applied Predictive Analytics: Principles and Techniques for the Professional Data Analyst. John Wiley & Sons.

[2] Dyaa Albakour, David Corney, Julio Gonzalo, Miguel Martinez, Barbara Poblete, and Andreas Vlachos. 2018. Report on the 2nd International Workshop on Recent Trends in News Information Retrieval (NewsIR'18). ACM SIGIR Forum 52, 1 (August 2018), 140–146. DOI:https://doi.org/10.1145/3274784.3274799

[3] Roja Bandari, Sitaram Asur, and Bernardo A. Huberman. 2012. The pulse of news in social media: Forecasting popularity. In Sixth International AAAI Conference on Weblogs and Social Media. Retrieved from https://www.aaai.org/ocs/index.php/ICWSM/ICWSM12/paper/viewPaper/4646

[4] Jukka-Pekka Bergman. 2016. Strategic interpretation on sustainability issues – eliciting cognitive maps of boards of directors. Corporate Governance: The International Journal of Business in Society 16, 1 (January 2016), 162–186. DOI:https://doi.org/10.1108/CG-04-2015-0051

[5] Zoë Bernard. 2018. Facebook accounts for 20% of the global online ad market. Business Insider. Retrieved December 7, 2019 from https://www.businessinsider.com.au/facebook-advertising-market-share-chart-2018-3

[6] Nicole Blanchett Neheli. 2018. News by Numbers. Digital Journalism 6, 8 (September 2018), 1041–1051. DOI:https://doi.org/10.1080/21670811.2018.1504626

[7] Alex Brownsell. 2017. Trend Snapshot: Google and Facebook's digital "duopoly" | WARC. WARC. Retrieved December 7, 2019 from https://www.warc.com/content/paywall/article/warc-exclusive/trend-snapshot-google-and-facebooks-digital-duopoly/113333

[8] B&T Magazine. 2018. Study: Search Now Accounts For 20% Of All Global Ad Spend (With Google Snaring The Lion's Share). B&T. Retrieved December 7, 2019 from https://www.bandt.com.au/study-search-now-accounts-20-global-ad-spend-google-snaring-lions-share/

[9] Mel Bunce. 2019. Management and resistance in the digital newsroom. Journalism 20, 7 (July 2019), 890–905. DOI:https://doi.org/10.1177/1464884916688963

[10] Kalyani Chadha and Rob Wells. 2016. Journalistic Responses to Technological Innovation in Newsrooms. Digital Journalism 4, 8 (November 2016), 1020–1035. DOI:https://doi.org/10.1080/21670811.2015.1123100

[11] Federica Cherubini and Rasmus Kleis Nielsen. 2016. Editorial Analytics: How News Media are Developing and Using Audience Data and Metrics. Reuters Institute for the Study of Journalism. Retrieved January 29, 2019 from https://reutersinstitute.politics.ox.ac.uk/sites/default/files/research/files/Editorial%2520analytics%2520-%2520how%2520news%2520media%2520are%2520developing%2520and%2520using%2520audience%2520data%2520and%2520metrics.pdf

[12] Mădălina Ciobanu. 2015. Why the Financial Times is betting on "audience-first." Journalism.co.uk. Retrieved December 5, 2019 from https://www.journalism.co.uk/news/why-the-financial-times-is-betting-on-an-audience-first-newsroom/s2/a584707/

[13] Mădălina Ciobanu. 2015. Data-informed: How the Guardian and the FT use newsroom analytics. Journalism.co.uk. Retrieved December 5, 2019 from https://www.journalism.co.uk/news/-data-informed-how-the-guardian-and-the-ft-approach-newsroom-analytics/s2/a591604/

[14] Ingrid Cobben. 2019. Why journalists shouldn't be afraid of metrics. Medium. Retrieved November 11, 2019 from https://medium.com/we-are-the-european-journalism-centre/why-journalists-shouldnt-be-afraid-of-metrics-1402a24e780b

[15] Alessandro Cucchiarelli, Christian Morbidoni, Giovanni Stilo, and Paola Velardi. 2017. What to Write? A topic recommender for journalists. In Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism, 19–24. Retrieved from https://www.aclweb.org/anthology/papers/W/W17/W17-4204/

[16] Alessandro Cucchiarelli, Christian Morbidoni, Giovanni Stilo, and Paola Velardi. 2018. What to write and why. Proceedings of the 33rd Annual ACM Symposium on Applied Computing - SAC '18. DOI:https://doi.org/10.1145/3167132.3167274

[17] Alessandro Cucchiarelli, Christian Morbidoni, Giovanni Stilo, and Paola Velardi. 2019. A topic recommender for journalists. Information Retrieval Journal 22, 4–31. DOI:https://doi.org/10.1007/s10791-018-9333-2

[18] Abhinandan S. Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google News Personalization: Scalable Online Collaborative Filtering. In Proceedings of the 16th International Conference on World Wide Web (WWW '07), 271–280. DOI:https://doi.org/10.1145/1242572.1242610

[19] Gianmarco De Francisci Morales, Aristides Gionis, and Claudio Lucchese. 2012. From Chatter to Headlines: Harnessing the Real-time Web for Personalized News Recommendation. In Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12), 153–162. DOI:https://doi.org/10.1145/2124295.2124315

[20] Mark Deuze. 2005. What is journalism?: Professional identity and ideology of journalists reconsidered. Journalism 6, 4 (November 2005), 442–464. DOI:https://doi.org/10.1177/1464884905056815

[21] Mark Deuze and Tamara Witschge. 2018. Beyond journalism: Theorizing the transformation of journalism. Journalism 19, 2 (February 2018), 165–181. DOI:https://doi.org/10.1177/1464884916688550

[22] Nicholas Diakopoulos, Munmun De Choudhury, and Mor Naaman. 2012. Finding and Assessing Social Media Information Sources in the Context of Journalism. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12), 2451–2460. DOI:https://doi.org/10.1145/2207676.2208409

[23] Norman R. Draper and Harry Smith. 1998. Applied regression analysis. John Wiley & Sons.

[24] Wayne W. Eckerson. 2007. Predictive analytics. Extending the Value of Your Data Warehousing Investment. TDWI Best Practices Report 1, (2007), 1–36. Retrieved from https://www.researchgate.net/profile/Waranpong_Boonsiritomachai/publication/307571066_Exploring_business_intelligence_and_its_depth_of_maturity_in_Thai_SMEs/links/585017a708ae4bc8993b6696/Exploring-business-intelligence-and-its-depth-of-maturity-in-Thai-SMEs.pdf

[25] Abigail Edge. 2014. Ophan: Key metrics informing editorial at the Guardian. Journalism.co.uk. Retrieved September 29, 2018 from https://www.journalism.co.uk/news/how-ophan-offers-bespoke-data-to-inform-content-at-the-guardian/s2/a563349/

[26] Tom Evens. 2018. Media economics and transformation in a digital Europe. In Comparative Media Policy, Regulation and Governance in Europe. Unpacking the Policy Cycle, Leen Haenens, Helena Sousa and Josef Trappel (eds.). Intellect, 41–54. Retrieved from https://biblio.ugent.be/publication/8560297/file/8560300

[27] Kelwin Fernandes, Pedro Vinagre, and Paulo Cortez. 2015. A Proactive Intelligent Decision Support System for Predicting the Popularity of Online News. In Progress in Artificial Intelligence, 535–546. DOI:https://doi.org/10.1007/978-3-319-23485-4_53

[28] Ben Fields, Rhianne Jones, and Tim Cowlishaw. 2018. The Case for Public Service Recommender Algorithms. Retrieved from https://piret.gitlab.io/fatrec2018/program/fatrec2018-fields.pdf

[29] Max Goldbart. 2019. BBC Sounds forging "public service algorithm." Broadcast. Retrieved August 15, 2019 from https://www.broadcastnow.co.uk/bbc/bbc-sounds-forging-public-service-algorithm/5139295.article

[30] Manuel Goyanes. 2014. An Empirical Study of Factors that Influence the Willingness to Pay for Online News. Journalism Practice 8, 6 (November 2014), 742–757. DOI:https://doi.org/10.1080/17512786.2014.882056

[31] Manuel Goyanes. 2015. The Value of Proximity: Examining the Willingness to Pay for Online Local News. International Journal of Communication 2015, 9 (2015), 1505–1522. Retrieved from https://ijoc.org/index.php/ijoc/article/view/3388/1383

[32] Daniel Hardt and Owen Rambow. 2017. Predicting user views in online news. In Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism, 7–12. Retrieved from https://www.aclweb.org/anthology/W17-4202.pdf

[33] Teemu Henriksson. 2015. Newspaper Revenues Shift To New Sources. WAN-IFRA World Press Trends. Retrieved December 5, 2019 from https://www.wan-ifra.org/press-releases/2015/06/01/world-press-trends-newspaper-revenues-shift-to-new-sources

[34] Jeremy Howard and Sebastian Ruder. 2018. Universal Language Model Fine-tuning for Text Classification. arXiv [cs.CL]. Retrieved from http://arxiv.org/abs/1801.06146

[35] Atte Jääskeläinen and Servet Yanatma. 2019. The future of national news agencies in Europe - Case study 4: Business model innovation in media-owned national news agencies. London School of Economics and Political Science, Department of Media and Communications, London, UK. DOI:https://doi.org/10.21953/lse.1oelxlquslqm

[36] Andreas Kaltenbrunner, Vicenç Gómez, and Vicente López. 2007. Description and Prediction of Slashdot Activity. DOI:https://doi.org/10.1109/laweb.2007.4383159

[37] Andreas Kaltenbrunner, Vicenç Gómez, Ayman Moghnieh, Rodrigo Meza, Josep Blat, and Vicente López. 2007. Homogeneous temporal activity patterns in a large online communication space. arXiv [cs.NI]. Retrieved from http://arxiv.org/abs/0708.1579

[38] Sarah Kaplan. 2011. Research in Cognition and Strategy: Reflections on Two Decades of Progress and a Look to the Future. Journal of Management Studies 48, 3 (May 2011), 665–695. DOI:https://doi.org/10.1111/j.1467-6486.2010.00983.x

[39] Eero Kasanen, Kari Lukka, and Arto Siitonen. 1993. The constructive approach in management accounting research. Journal of management accounting research 5, 1 (1993), 243–264. Retrieved from https://mycourses.aalto.fi/pluginfile.php/183797/mod_resource/content/1/Kasanen%20et%20al%201993.pdf

[40] Yaser Keneshloo, Shuguang Wang, Eui-Hong (sam) Han, and Naren Ramakrishnan. 2016. Predicting the Popularity of News Articles. In Proceedings of the 2016 SIAM International Conference on Data Mining. Society for Industrial and Applied Mathematics, 441–449. DOI:https://doi.org/10.1137/1.9781611974348.50

[41] S. Kim, S. Kim, and H. Cho. 2011. Predicting the Virtual Temperature of Web-Blog Articles as a Measurement Tool for Online Popularity. In 2011 IEEE 11th International Conference on Computer and Information Technology, 449–454. DOI:https://doi.org/10.1109/CIT.2011.104

[42] Su-Do Kim, Seonyeong Kim, and Hwan-Gue Cho. 2012. A model for popularity dynamics to predict hot articles in discussion blog. Proceedings of the 6th International Conference on Ubiquitous Information Management and Communication - ICUIMC '12. DOI:https://doi.org/10.1145/2184751.2184764

[43] David G. Kleinbaum, Lawrence L. Kupper, Keith E. Muller, and Azhar Nizam. 1988. Applied regression analysis and other multivariable methods. Duxbury Press. Retrieved from https://pdfs.semanticscholar.org/ff7b/2b68a164e355633207dd31c05e00a455d2dd.pdf

[44] Eun-Ju Lee and Edson C. Tandoc. 2017. When News Meets the Audience: How Audience Feedback Online Affects News Production and Consumption. Hum. Commun. Res. 43, 4 (October 2017), 436–449. DOI:https://doi.org/10.1111/hcre.12123

[45] J. G. Lee, S. Moon, and K. Salamatian. 2010. An Approach to Model and Predict the Popularity of Online Contents with Explanatory Factors. In 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, 623–630. DOI:https://doi.org/10.1109/WI-IAT.2010.209

[46] Jong Gun Lee, Sue Moon, and Kavé Salamatian. 2012. Modeling and predicting the popularity of online contents with Cox proportional hazard regression model. Neurocomputing 76, 134–145. DOI:https://doi.org/10.1016/j.neucom.2011.04.040

[47] Kristina Lerman and Rumi Ghosh. 2010. Information Contagion: an Empirical Study of the Spread of News on Digg and Twitter Social Networks. arXiv [cs.CY]. Retrieved from http://arxiv.org/abs/1003.2664

[48] Kristina Lerman and Tad Hogg. 2010. Using a Model of Social Dynamics to Predict Popularity of News. In Proceedings of the 19th International Conference on World Wide Web (WWW '10), 621–630. DOI:https://doi.org/10.1145/1772690.1772754

[49] Seth C. Lewis. 2012. THE TENSION BETWEEN PROFESSIONAL CONTROL AND OPEN PARTICIPATION. Inf. Commun. Soc. 15, 6 (August 2012), 836–866. DOI:https://doi.org/10.1080/1369118X.2012.674150

[50] Quanzhi Li, Sameena Shah, Xiaomo Liu, and Armineh Nourbakhsh. 2017. Data sets: Word embeddings learned from tweets and general data. In Eleventh International AAAI Conference on Web and Social Media. Retrieved from https://www.aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/viewPaper/15672

[51] Quanzhi Li, Sameena Shah, Xiaomo Liu, Armineh Nourbakhsh, and Rui Fang. 2016. TweetSift: Tweet Topic Classification Based on Entity Knowledge Base and Topic Enhanced Word Embedding. In Proceedings of the 25th ACM International on Conference on Information and Knowledge Management (CIKM '16), 2429–2432. DOI:https://doi.org/10.1145/2983323.2983325

[52] X. Liu, A. Nourbakhsh, Q. Li, S. Shah, R. Martin, and J. Duprey. 2017. Reuters tracer: Toward automated news production using large scale social media data.

In 2017 IEEE International Conference on Big Data (Big Data), 1483–1493. DOI:https://doi.org/10.1109/BigData.2017.8258082
[53] Lívia Markíczy and Jeff Goldberg. 1995. A Method for Eliciting and Comparing Causal Maps. Journal of Management 21, 305–333. DOI:https://doi.org/10.1177/014920639502100207
[54] Richard McCreadie, Craig Macdonald, and Iadh Ounis. 2011. A Learned Approach for Ranking News in Real-Time Using the Blogosphere. String Processing and Information Retrieval, 104–116. DOI:https://doi.org/10.1007/978-3-642-24583-1_11
[55] Richard M. C. McCreadie, Craig Macdonald, and Iadh Ounis. 2010. News Article Ranking: Leveraging the Wisdom of Bloggers. In Adaptivity, Personalization and Fusion of Heterogeneous Information (RIAO '10), 40–48. Retrieved from http://dl.acm.org/citation.cfm?id=1937055.1937064
[56] James McDonald. 2019. Regulation, user migration, Amazon… is the duopoly's future assured? | WARC. Retrieved December 7, 2019 from http://www.warc.com/newsandopinion/opinion/regulation_user_migration_amazon_is_the_duopolys_future_assured/3029?utm_source=daily-email-free-link&utm_medium=email&utm_campaign=daily-email-apac-prospects-20190326
[57] Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2017. Regularizing and Optimizing LSTM Language Models. arXiv [cs.CL]. Retrieved from http://arxiv.org/abs/1708.02182
[58] Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. An Analysis of Neural Language Modeling at Multiple Scales. arXiv [cs.CL]. Retrieved from http://arxiv.org/abs/1803.08240
[59] Feras Namous, Ali Rodan, and Yasir Javed. 2018. Online News Popularity Prediction. 2018 Fifth HCT Information Technology Trends (ITT). DOI:https://doi.org/10.1109/ctit.2018.8649529
[60] Jacob L. Nelson. 2018. The Elusive Engagement Metric. Digital Journalism 6, 4 (April 2018), 528–544. DOI:https://doi.org/10.1080/21670811.2018.1445000
[61] Jacob L. Nelson and Edson C. Tandoc Jr. 2019. Doing "well" or doing "good": What audience analytics reveal about journalism's competing goals. Journalism Studies 20, 13 (2019), 1960–1976. DOI:https://doi.org/10.1080/1461670X.2018.1547122
[62] OECD. 2010. News in the Internet Age. New Trends in News Publishing. OECD Publishing. Retrieved from https://market.android.com/details?id=book-BRxPMeSyG08C
[63] Michael J. Pazzani. 1999. A Framework for Collaborative, Content-Based and Demographic Filtering. Artificial Intelligence Review 13, 5 (December 1999), 393–408. DOI:https://doi.org/10.1023/A:1006544522159
[64] Caitlin Petre. 2018. Engineering Consent. Digital Journalism 6, 4 (April 2018), 509–527. DOI:https://doi.org/10.1080/21670811.2018.1444998
[65] Robert G. Picard (Ed.). 2004. Strategic responses to media market changes. Jönköping Business School. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.452.6474&rep=rep1&type=pdf
[66] Robert G. Picard. 2010. A business perspective on challenges facing journalism. The changing business of journalism and its implications for democracy (2010), 17–24. Retrieved from https://reutersinstitute.politics.ox.ac.uk/sites/default/files/research/files/The%2520Changing%2520Business%2520of%2520Journalism%2520and%2520its%2520Implications%2520for%2520Democracy.pdf#page=23
[67] Robert G. Picard. 2011. Mapping Digital Media: Digitization and Media Business Models. Open Society Foundations. Retrieved from http://www.robertpicard.net/files/OSF-Media-Report-Handbook_Digitization_and_Media_Business_Models-final-07-18-2011-WEB.pdf
[68] He Ren and Yang Quan. 2015. Predicting and Evaluating the Popularity of Online News. Stanford University Machine Learning Report (2015). Retrieved from https://pdfs.semanticscholar.org/9e91/6a3469e9e2fc5f0c8f927d7d1d05f5575729.pdf?_ga=2.97936799.1626222008.1576067591-2008714051.1570988539
[69] Maya Sappelli, Dung Manh Chu, Bahadir Cambel, David Graus, and Philippe Bressers. 2018. SMART Journalism: Personalizing, Summarizing, and Recommending Financial Economic News. In The Algorithmic Personalization and News (APEN18) Workshop at ICWSM. Retrieved from https://graus.nu/blog/wp-content/papercite-data/pdf/sappelli2018smart.pdf
[70] Paul Sawers. 2015. Ophan: Inside the Guardian's data-driven newsroom. VentureBeat. Retrieved November 2, 2019 from https://venturebeat.com/2015/07/03/ophan-inside-the-guardians-data-driven-newsroom/
[71] Annibal José Scavarda, Tatiana Bouzdine-Chameeva, Susan Meyer Goldstein, Julie M. Hays, and Arthur V. Hill. 2006. A methodology for constructing collective causal maps. Decision Sciences 37, 2 (2006), 263–283. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5915.2006.00124.x
[72] George A. F. Seber and Alan J. Lee. 2012. Linear regression analysis. John Wiley & Sons. Retrieved from https://www.google.com/books?hl=en&lr=&id=X2Y6OkXl8ysC&oi=fnd&pg=PR5&dq=regression+analysis&ots=sdoPE3kOpw&sig=txxZCyOtMmOTupHadbjzjNvyA3c
[73] Eric Siegel. 2013. Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, Or Die. John Wiley & Sons.
[74] Jannick Kirk Sørensen. 2019. Public Service Media, Diversity and Algorithmic Recommendation: Tensions between Editorial Principles and Algorithms in European PSM Organizations. In ResearchGate. Retrieved October 20, 2019 from https://www.researchgate.net/publication/335276877_Public_Service_Media_Diversity_and_Algorithmic_Recommendation_Tensions_between_Editorial_Principles_and_Algorithms_in_European_PSM_Organizations
[75] John D. Sterman. 2000. Business Dynamics: Systems Thinking and Modeling for a Complex World with CD-ROM. McGraw-Hill Education.
[76] John D. Sterman. 2001. System dynamics modeling: tools for learning in a complex world. Calif. Manage. Rev. 43, 4 (2001), 8–25. Retrieved from https://journals.sagepub.com/doi/pdf/10.2307/41166098
[77] Gabor Szabo and Bernardo A. Huberman. 2008. Predicting the Popularity of Online Content. DOI:https://doi.org/10.2139/ssrn.1295610
[78] Edson C. Tandoc. 2014. Journalism is twerking? How web analytics is changing the process of gatekeeping. New Media & Society 16, 4 (June 2014), 559–575. DOI:https://doi.org/10.1177/1461444814530541
[79] Edson C. Tandoc Jr. 2015. Why web analytics click: Factors affecting the ways journalists use audience metrics. Journalism Studies 16, 6 (2015), 782–799. DOI:https://doi.org/10.1080/1461670X.2014.946309
[80] Edson C. Tandoc and Ryan J. Thomas. 2015. The Ethics of Web Analytics. Digital Journalism 3, 2 (March 2015), 243–258. DOI:https://doi.org/10.1080/21670811.2014.909122
[81] Alexandru Tatar, Panayotis Antoniadis, Marcelo Dias de Amorim, and Serge Fdida. 2014. From popularity prediction to ranking online news. Social Network Analysis and Mining 4, 1 (February 2014), 174. DOI:https://doi.org/10.1007/s13278-014-0174-8
[82] Arvid Tchivzhel. 2015. Predictive analytics in the publishing industry. Analytics Magazine. Retrieved December 5, 2019 from http://analytics-magazine.org/predictive-analytics-in-the-publishing-industry/
[83] Julian Thomas. 2018. Programming, filtering, adblocking: advertising and media automation. Media International Australia 166, 1 (February 2018), 34–43. DOI:https://doi.org/10.1177/1329878X17738787
[84] Towards Data Science. 2018. Top AI and Machine Learning Trends in Media and Entertainment. Medium. Retrieved November 9, 2019 from https://towardsdatascience.com/top-ai-and-machine-learning-trends-in-media-and-entertainment-823f7efea928
[85] Manos Tsagkias, Wouter Weerkamp, and Maarten de Rijke. 2009. Predicting the volume of comments on online news stories. Proceeding of the 18th ACM conference on Information and knowledge management - CIKM '09. DOI:https://doi.org/10.1145/1645953.1646225
[86] Manos Tsagkias, Wouter Weerkamp, and Maarten de Rijke. 2010. News Comments:Exploring, Modeling, and Online Prediction. Lecture Notes in Computer Science, 191–203. DOI:https://doi.org/10.1007/978-3-642-12275-0_19
[87] James P. Walsh. 1995. Managerial and Organizational Cognition: Notes from a Trip Down Memory Lane. Organization Science 6, 280–321. DOI:https://doi.org/10.1287/orsc.6.3.280
[88] Kasper Welbers, Wouter van Atteveldt, Jan Kleinnijenhuis, Nel Ruigrok, and Joep Schaper. 2016. News selection criteria in the digital age: Professional norms versus online audience metrics. Journalism 17, 8 (November 2016), 1037–1053. DOI:https://doi.org/10.1177/1464884915595474
[89] Western iMedia. 2014. Future of journalism, media revenues depend on programmatic buying (no pressure). International News Media Association (INMA). Retrieved December 7, 2019 from https://www.inma.org/blogs/world-congress/post.cfm/future-of-journalism-media-revenues-depend-on-programmatic-buying-no-pressure
[90] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A Deep Reinforcement Learning Framework for News Recommendation. In Proceedings of the 2018 World Wide Web Conference (WWW '18), 167–176. DOI:https://doi.org/10.1145/3178876.3185994
[91] Arkaitz Zubiaga, Heng Ji, and Kevin Knight. 2013. Curating and Contextualizing Twitter Stories to Assist with Social Newsgathering. In Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13), 213–224. DOI:https://doi.org/10.1145/2449396.2449424