

News personalization symposium report

05/05/2018 - Institute for Information Law, University of Amsterdam

N. Helberger (University of Amsterdam)
B. Bodo (University of Amsterdam)
J.K. Sørensen (Aalborg University)
M.Z. van Drunen (University of Amsterdam)



Instituut voor Informatierecht
Institute for Information Law

Contents

Plenary session: Introduction to the Personews project & discussion	3
Breakout session #1: Filter bubbles and other threats	4
Breakout session #2: Values	6
Breakout session #3: Editorial understanding of news personalization	8
Breakout session #4: The use of news personalisation	9
List of participants	14

Plenary session: Introduction to the Personews project & discussion

The symposium built on a series of interviews with the news media which were conducted as part of the PersoNews project. These revealed various shared questions about news personalization and a lack of established wisdom about the best way forward. The symposium created a platform through which editors, journalists, technologists and academics could build a shared understanding of these issues and connect for future collaborations. The discussions during the workshop were so insightful that we decided to share some of the key insights with a broader audience in the form of this report. We would like to thank all the participants of the Amsterdam News Symposium (see list attached) for their active contribution and the good discussion.

In the first half of the workshop, the PersoNews project introduced itself, and the studies it has conducted on the different aspects of news personalization. Prof. Dr. Natali Helberger introduced the research; Dr. Balazs Bodo presented the findings of his interviews with newsrooms; Dr. Judith Mueller presented the users' side of the story; Sarah Eskens spoke about the role of data protection law in news personalization; Max van Drunen discussed the challenges of making news personalization transparent to users; and finally Natali Helberger presented a conceptual framework for assessing diversity in recommendations.

The presentations were followed by a plenary discussion. In the first part, the comments were addressing the current state of news personalization. Some participants reported, for example, that the pool of news materials that news organizations produce is less diverse than assumed, which severely limits the opportunities for personalization, as there is simply not enough material to produce personalized recommendations from. There was an agreement that good personalization requires high quality metadata (especially for audiovisual content), which is often missing and expensive to produce, further limiting the often archived pool of content that is available for personalized recommendations. As for the goal of personalization, some suggested that personalization is often just a simple rearrangement of news articles into a different order than what would be the case with human curated suggestions. In other cases, personalization is about how to better sell stories, but there are professional limits to such efforts, as serious news organizations rarely want to build addictive services for their audiences, and they usually don't want people to be on their network all day. Finally, a number of questions emerged around the impact of personalization, and the news in general, and about the ways to measure that impact, and its changes over time. How to measure the (social) impact of journalism? How to best measure the impact of (personalized) news on social cohesion? - were the most emblematic question which closed this discussion.

In the second half of the plenary session an extensive discussion emerged over the transparency of news personalization. Workshop participants were very interested in the legal framework, and the interpretation of the legal rules, probably because news organizations are

looking for some kind of guidance for daily decision making. There are multiple reasons in consumer law why news organizations should inform news consumers about news personalization. For example, in the Consumer Rights Directive all the information requirements apply to news organizations as information society services. Rules about advertising could be applied to personalization as well. These rules may lead to a number of currently unanswered questions, such as: “Should we tell people why a certain article was recommended? How we should differentiate the content and recommendations? Is it possible to explain for the user why they got recommended a particular article? Is it possible to harmonize the data scientist and the legal perspectives and expectations on transparency? How does the very concept of transparency change depending on who the audience (regulators, users, supervisory board, etc.) changes?”

As an interim solution the participants discussed the possibility to differentiate the GDPR’s “right to explanation”, its scope and depth, depending on the context of the algorithmic decision. A decision on a jail term is a different decision than a decision on a news recommendation, and consequently different depth of explanation is necessary. The expectation is that main use of the right to explanation will mainly be for grave situations, not for digital content. In the same manner, procedural justice has also limited relevance for very trivial decisions. In these cases, the right to explanation will probably be something like due process. Maybe news organizations could at least say to the consumer: this article was recommended to you on the basis of an algorithm.

After the plenary discussion the workshop participants formed four breakout sessions based on their shared interests. We summarized the four sessions below.

Breakout session #1: filter bubbles and other threats

Moderator: Balazs Bodo

Filter bubbles in context

In the breakout session about filter bubbles and other threats of news personalization, there was much discussion about the extent to which news media, especially serious news media should be blamed for creating filter bubbles. There was an agreement among the participants that social media, digital (advertising) platforms, and other commercial entities may have serious filter bubble issues, as they operate large, sophisticated automatic recommender systems with limited human oversight, which are often being abused by malicious third parties en masse. The failures of platforms also put serious news organizations under intense scrutiny for filter bubbles. Yet, the latter are much smaller, operate recommender systems with more human oversight, they operate under strict and well tested professional codes, and they are less likely to be abused by third parties. In short, news organizations are less likely to create filter bubbles, if at all. Many argued, that there needs to be a better distinction of the context in which policy and academia, and the public in general discusses the threat of filter bubbles. At the moment there is little evidence of filter bubbles created by news organizations, and the impression that news has a filter bubble problem is created by the bias of the platforms that present news, rather than the issues at the news organizations themselves.

Need for better evidence

There was also some agreement that this distinction between platforms and news does not need to mean that news media is by definition immune from filter bubbles. But before we can reach any conclusions about the filter bubbles in news, we need a better understanding of this threat. We need better and more conclusive evidence about the nature of news bubbles on news sites, the affected audiences (teens, vs. others, for example), etc. Without such in-depth knowledge it is almost impossible to build sound strategies and counter threats. There was an agreement that a continuous effort is needed to monitor the emergence of potential threats that stem from the use of novel technologies, and it is necessary to trace the impact of technological innovation in and around the news domain on the public sphere, on the quality of the news diet, on the audiences, etc.

The algorithm of someone else

Another major issue that was discussed in this session was the (in)visibility of values that come embedded with recommendation technologies. The academic discourse is well aware of the fact that technologies are hardly neutral, and usually come with a number of values embedded in them, visibly, or invisibly. In the case of news personalization, such embedded values can cause problems, as the technologies news organizations buy, and use to implement their own services are usually developed for other purposes, in other domains. There was much discussion whether a recommender technology which was developed and optimized for an e-

commerce application, or for the delivering of advertising can seamlessly be used to personalize news content. On the same issue, Public Service Media organizations expressed their concerns about the usability of recommender technologies developed for commercial broadcasters, and for-profit media, which have their own set of values, such as optimization for engagement vs optimization for breadth embedded in them. Some argued, that the goals embedded in the technology do not necessarily match the goals of the organization, so those embedded values need to be removed, and new values need to be embedded, or a completely new technology needs to be built which reflects the goals of the organization that implements the personalization technology.

Working with platforms

The last threat the group discussed was the threat of platforms to news delivery. Many issues were identified, such as the loss of source identification and branding during social media sharing (content shared on Facebook gets stripped from the PSM brand); the difficulties of competence building in platform dominated distribution environments; the frequent, unpredictable changes in the algorithm of platforms; the continuous investment necessary to manage audience relations under such circumstances; the difficulty of building long term, loyal audiences for news media on such platforms; or the limited flexibility of platforms to take into account news organizations' interests. Again the divergent strategies and priorities of PSM and commercial media emerged: platforms are hyper-commercial environments which only poorly reflect or take into account the priorities of PSM. The question arose: "In that context what is the right PMS strategy? Not personalizing at all?"

Ways to deal with the issues

In the last part of the session the group discussed strategies to address the aforementioned challenges. The diversity of the underlying organization, the employees and, interestingly enough, the diversity of audiences emerged as key resources that are necessary to successfully face the challenges. Apart from that, there was a clear need for better, more detailed, more systematic, and more reliable empirical evidence in the forms of academic studies, long term overviews, industry surveys, better market and business intelligence. Participants identified the need of a broader discourse about the challenges the industry faces in the media. They also called for more professional venues and opportunities to discuss the challenges. There was a strong need for different fora where media professionals can safely share ideas, and formulate criticism. Opportunities, such as a forum bringing technologists together with editorial staff, so technologists can learn what editorial values need to be embedded in the technology were discussed. There was also some discussion on the need to better educate audiences, to promote media literacy among most vulnerable groups, and to provide technology education about how algorithms work.

Breakout session #2: Values

Moderator: Natali Helberger

Need to revisit central values

In order to decide how to use data and algorithms, for example in the context of news personalisation, the need was felt to revisit the question of the editorial identity of the outlet. This can involve taking a step back, and allocating time and space within a news organisation to define how algorithms and AI can contribute to the editorial mission, and what exactly that mission is. This can be different for different organisations, as the strategy behind the use of personalisation strategies, but also ways in using quantified approaches to measuring the audience can differ between different news outlets. A question that has been raised, for example, was whether public service media and quality media will have a similar approach to personalisation or the use of audience metrics for editorial/strategic decision making than more popular outlets.

This led to the next question: how to identify values. Previously media organizations had a very strong awareness of their values. To make these values work in the machine- and data driven environment, it is useful to remember and, where necessary, revisit these values. Generally, though, there was a feeling that the values that the news media subscribe to so far are pretty general, making them easily adaptable also to the use of recommendations. The more pressing question is maybe: are the existing values and editorial mission statements complete, or do we need to consider new values such as explainability or respect for readers' privacy? An interesting question was whether users could or should be usefully involved in this process? And could AI help? For example, would it be an option to analyze historical content to determine the values, and then use these computational values as a means of informing recommendation strategies?

Giving users voice

One particular value that we discussed in more depth was giving users' a voice. So far, lots of decisions around data and algorithms were felt to be driven by strategic considerations, not so much by the wish to give users more voice. Technological developments, however, have also created new opportunities for materialising voice. Giving users a word could be done, for example, in the form of a dashboard. Right now, we often seem to show an overreliance on inferences from observed data. Should algorithmic recommender design move to more conversational, dynamical profiles, with dashboards as a way of establishing a sustainable dialogue, rather than imposing choices upfront? Another way of giving users more voice could be the use of user-driven personalisation modes. The question is then on which users to focus, and for whom to design the technology: for the few users that are likely to engage actively? Or should the media rather try to engage the broader audience. Clearly, different strategies are possible, and the ultimate answer will depend on the concrete media outlet and its editorial strategy/mission.

Measuring the performance of recommendations

Once values are identified, the next question is how to measure whether recommendations live up to these values. The identity of the news publisher should also be reflected in the recommendations. This may require defining new benchmarks to be able to assess and, where necessary, adjust recommendations. Equally important is it to make sure that there are no discrepancies in the way in which the algorithm is evaluated and the way in which the outcome of the recommendations is evaluated.

There was also general agreement that the right measures or benchmarks are still missing. Having said so, the difficulty of identifying the right benchmarks to measure performance and impact is by no means a new challenge. Insofar, valuable lessons could be learned from existing approaches, including, for example, linear television.

Automatisation and editorial independence

Another value we discussed in more depth was editorial independence. And here an important observation was brought forward, namely that using third party software can also mean using that third parties' interpretation of certain values, such as diversity, quality or impartiality. Therefore, a very important question is to what extent a news medium is able to negotiate conditions with that third party, e.g. that the news media are able to access and review training data, gain detailed insights on recommendation logics, point out conflicting results and demand correction. Put differently, values, such as independence, become also a question of negotiation power. Trade secrecy protection can further render the negotiation position of the news media more difficult. Therefore, some media concluded that to be truly independent they need to build their own recommender. Where does that leave smaller, local news outlets that do not have the resources and knowledge to build own recommendation solutions? More forward looking, an important topic for the near future is finding ways to increase negotiation power vis-a-vis external producers of algorithms.

One way to increase negotiation power of the media vis-a-vis technology producers, and this was an important key take-away, is to share knowledge, insights, best practices, technology, algorithms among the media. So far, algorithms are often considered as proprietary black boxes, but what is really proprietary and a black box (not at least for reasons of privacy protection) are the data, not the algorithms. Often, among certain media companies there already exists an openness to sharing insights and experiences about the algorithm. What would be helpful was the existence of more structural initiatives or platforms for doing so.

Breakout session #3: Editorial understanding of news personalization

Moderator: Max van Drunen

Recommendations as a means of furthering editorial goals

Participants from both the public service and commercial media agreed that that personalization could be a great tool for serving their editorial goals. It was seen as being particularly useful for reaching new (younger) audiences, and for serving niche and underserved segments of the current audience. At the same time, they identified shared technical and organizational obstacles to implementing their editorial understanding of news personalization.

The need to know your audience & your content

Two shared technical obstacles were identified: a lack of (complete) information about the audience, and a lack of information about the content so that it can be appropriately matched to individuals. Participants discussed the benefits of one concrete solution, namely to outsource the creation of metadata about the content to a country with lower labor costs.

Successful implementation of personalization within the organisation

The discussion then turned to the organizational steps which need to be taken to ensure that editorial values are implemented in news personalization. The first is to address the unwillingness to engage with news personalization which sometimes exists on the editorial side. As this was seen as being fueled by fears that algorithms will replace journalists and fail to meet editorial standards, it was suggested that personalization could first be implemented in a smaller/supporting role, or as a new feature instead of a replacement.

Secondly, editorial influence on news personalization would benefit from clearer communication with the other parties involved. Concretely, this might involve clearly outlining the editorial goals which news personalization can achieve, and communicating the way in which news personalization can contribute to them to management and to the audience. It also involves being open to technologists about these goals and the need to be flexible in which goals are prioritized, while at the same time being aware of the limits and potential difficulties involved in changing the technology.

Breakout session #4: the use of news personalisation

Moderator: Jannick Sørensen

Reasons for engaging in personalisation: serving the audience better

The group featured both representatives from private press as well from public service media (PSM). The discussions showed that the objectives for personalisation are not located far from each other in the two cases; rather it is the ways personalisation / recommendation is conducted that normatively are different, as well as the over-arching narrative for the algorithmic recommendation. The myth of 'optimizing' / expanding consumption was countered: For both types of media, it is important to serve the audiences in the best way, and also to establish a longer relationship. For the private media, this could e.g. mean converting readers into subscribers. For PSM it is to reach as far as possible, to serve the societal tasks and contribute to societal cohesion. The discussion of the particular PSM objectives led to a remark whether the personal and the societal which often are constructed to be each other's opposites in terms of relevancy of a news item, in reality are opposites? Is 'value for society' and 'value for the user' a real conflict?

A practical reason for news recommender systems is the small screen size of the mobile phone; the selection of news items must be more precise. The mobile phone, with its location sensors, opens also the possibility of using the user's location as a criteria in the recommendation.

Need to differentiate between “personalisation” and “individualisation”

It was noted that personalisation and individualisation is not the same thing: Different users may via filtering systems get exposed to the same news item, just differently ranked. The relation / difference between 'individualisation' and 'personalisation' is to be examined further. The social function of news - namely "to know what the others know" must be acknowledged, noting that "the others" now is an algorithmically created entity.

Finding the right benchmarks to measure the performance of the recommendation algorithm

The discussions proceeded with the question: How to measure, and what to measure? Relevancy (a main concept in recommender systems) may in reality be difficult to measure. The relevancy could depend on many criteria, e.g. the context. Also, measurements should be read critically; e.g. is the number for how many times an item has been shared not necessarily an expression of its relevancy. An open question emerged: Is there a way to explain / measure the depth of interests? Could these deep interests, expressed by explicit choices, be supported with personalisation?

Another question related to measurements: Are publishers too producer-centric in their measurements - focusing only on figures that are important internally in the news house? How well do measurements reflect the point of view of users? There might be correlations between what users and editors find important, but if measurements are made from an editorial point of view - fulfilling the information needs of the editors - the users' reasons for consuming or not consuming a news item may be overlooked. The question was raised whether we need algorithms that are managed by users, are tools for the users, not tools for the publishers? These would echo an old idea of software as user agents. It can be discussed whether Google search engine fulfils that purpose of being the user's agent.

While some publishers do have many data sources (data signals) available for the calculation of the personalised recommendation, the praxis shows that much of this data is not really relevant to take into the computation. Fewer data signals may be as good - and in terms of processing simpler to handle. One commercial publisher / recommender service explained that the user preferences that users manually indicating when initiating the service as a new user ('on-boarding') are skipped by the system as criteria after the user has been seeing a number of articles: The actual user behaviour gives a better data signal than the initial user preferences. As some publishers - via advertising networks - have access to a lot of personal data about users' consumption and commuting habits, this data could also be included in the calculation of the recommendations. However, again a more simple model for recommending content is more efficient. The abundance of user data does not necessarily produce more interesting recommendations.

Recommendations and PSM - a special case?

As PSM have other values and other success criteria, one could ask whether PSM organisations also need another type of algorithms than those used by commercial media. As the legacy algorithmic recommender systems techniques and methods have shown stability and value, it is not immediately relevant to imagine new algorithms. Instead the configuration of the algorithms may be different. It should reflect the different PSM brands and their identities, e.g. through the KPIs that are applied. However, also private media seek to reflect their brand values in the configuration of the recommender system.

In the case of PSM, the mechanisms of the recommender systems should be more transparent to users, e.g. via visualisations in dashboards. User involvement from the start of the introduction of a recommender system may not make sense, but when the system is running, a dash board can provide users with a systematic overview of the recommendations.

PSM has, maybe even to a bigger extent than private media, an interest in exposing the diversity of the content and to help users find the hidden 'gems'. An idea could however also be to focus users' attention more, by recommending very few items. When it comes to the utilisation of user-data, PSMs are limited, as the trustworthiness sets the limits (as well as legal obligations).

The potential of recommenders to inform the newsroom

The potential in recommender systems are not limited to find items to users. The behavioural data from the recommender system can be used to inform journalists about users' interests. The user behaviour can thus inspire journalists' production of new stories. The performance in the recommender system can also be used to benchmark human editors' performance and predict their choices.

Demonstrating 'value for money' to the management

While recommender systems are a hot topic among publishers these days, it can be discussed whether this interest will remain. Are recommender systems just a fashion? One factor that could determine the future popularity of recommender systems is whether management see an advantage in buying and maintaining the systems. The members of the break out group agreed that it is difficult to show-case the 'value for money' of recommender systems to management. However, the systems may have come to stay. The systems now have become so widely accepted and part of user expectations that it will be logical to continue using them.

One way of showing value to management is to show if 'cross-selling' (exposure diversity) is growing. Loyalty, and in the case of private media the sale of subscriptions may be another sales argument. It is however important to keep users' perception of recommender systems in mind. We need to remember that users hate to be monitored, partly is due to bad experiences with recommendations that were inappropriate caused by bad recommender systems.

In conclusion...

Throughout the different break-out sessions, some common threads of discussion and issues emerged. A very central consideration is the question of how algorithmic recommendations fit in the broader editorial strategy of a news outlet, and what exactly that mission is in an environment driven by user data and AI. Clear is that recommendations can and are being seen as an extension of the editorial activity of at least the quality newsmedia, and as such are driven by the values and objectives a particular news outlet puts central. Clear is also that strategic considerations and the values that one finds relevant and optimises for distinguish the way (quality) news media and online platforms think about, and use algorithms.

It was rather striking to notice the relative absence of economic or marketing arguments in the considerations of most of the participants of the factors they are striving to optimise for, indicating that the way recommendations can add to editorial decision making is and will remain an editorial question, and how to serve the audience and the editorial mission of the platform best, not: the demands from the marketing division. This does not take away that recommendations should bring added value for both, users and editors, and can also be a means to attract and hold subscribers.

When considering the values or metrics that recommendations should be optimised for, diversity and overcoming filterbubbles is an important consideration, both for PSM as well as (commercial) quality news media. Taking the filterbubble discussion one step further and identifying the target groups at risk or the conditions under which filterbubbles could create a potential problem would also help news organisation to identify new innovative strategies of how to counter them. Identifying, measuring and optimising for diverse recommendations is a challenge for many organisations, and one where academics could play a useful role in informing the process.

How to protect one's editorial independence in the presence of third party algorithm producers, powerful platforms and the ever-present insights into what users like and want is another important consideration. Editorial independence has been identified as another important value that is potentially at stake. When using third party software, the values of the designers of that software (e.g. the question of what diversity is, or how to explain rankings) do not necessarily match with the interpretations or values of a particular news outlet. Values can then become a matter of negotiation power, and be a reason for some news outlets to build their own recommenders (providing they have the expertise and means). Insofar, building up in-house expertise can be an important way of also safeguarding editorial independence, at least for those news organisations that can afford doing so.

One major take-away was the added value of sharing insights and knowledge on the best ways of how to use algorithmic recommendations, but also: on how to measure their performance, but also bringing different parties together: media professionals from different organisations, technologists with editorial staff, academics with media professionals, etc.

List of Participants (in no particular order)

Anne-Jel Hoelen (ACM)
Balazs Bodo (UvA, Institute for Information Law)
Bronwyn Jones (BBC)
Cornelius Puschmann (Hans-Bredow-Institut)
Daan Odijk (RTL/Blendle)
Damian Trilling (UvA, Amsterdam School of Communication Research)
David Graus (FD)
Erik Borra (UvA)
Frank Mekkelholt (Persgroep)
Hanna Kimmel (NPO)
J.K. Sorensen (Aalborg University)
Judith Möller (UvA, Amsterdam School of Communication Research)
Karen Yeung (University of Birmingham)
Lara Ankersmit (NOS)
Laura Ellis (BBC)
Lisa Merten (Hans-Bredow-Institut)
Luuk Willekens (NRC)
Mariella Bastian (UvA, Institute for Information Law)
Matthijs van Hunnik (NPO)
Max van Drunen (UvA, Institute for Information Law)
Natali Helberger (UvA, Institute for Information Law)
Neil Thurman (Ludwig-Maximilians University, Munich)
Philippe Bressers (FD)
Rhia Jones (BBC)
Robert Amlung (ZDF)
Sarah Eskens (UvA, Institute for Information Law)
Mathieu Triay (BBC)
Wiebke Loosen (Hans-Bredow-Institut)