



UvA-DARE (Digital Academic Repository)

Democratizing algorithmic news recommenders: how to materialize voice in a technologically saturated media ecosystem

Harambam, J.; Helberger, N.; van Hoboken, J.

DOI

[10.1098/rsta.2018.0088](https://doi.org/10.1098/rsta.2018.0088)

Publication date

2018

Document Version

Final published version

Published in

Philosophical Transactions of the Royal Society A - Mathematical, Physical and Engineering Sciences

[Link to publication](#)

Citation for published version (APA):

Harambam, J., Helberger, N., & van Hoboken, J. (2018). Democratizing algorithmic news recommenders: how to materialize voice in a technologically saturated media ecosystem. *Philosophical Transactions of the Royal Society A - Mathematical, Physical and Engineering Sciences*, 376(2133), [20180088]. <https://doi.org/10.1098/rsta.2018.0088>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Research



Cite this article: Harambam J, Helberger N, van Hoboken J. 2018 Democratizing algorithmic news recommenders: how to materialize voice in a technologically saturated media ecosystem. *Phil. Trans. R. Soc. A* **376**: 20180088. <http://dx.doi.org/10.1098/rsta.2018.0088>

Accepted: 20 July 2018

One contribution of 9 to a theme issue 'Governing artificial intelligence: ethical, legal, and technical opportunities and challenges'.

Subject Areas:

artificial intelligence, human-computer interaction, algorithmic information theory

Keywords:

algorithms, news media, personalization, transparency, voice, persona

Author for correspondence:

Jaron Harambam
e-mail: jaron.harambam@gmail.com

Democratizing algorithmic news recommenders: how to materialize voice in a technologically saturated media ecosystem

Jaron Harambam¹, Natali Helberger¹ and Joris van Hoboken^{1,2}

¹Institute for Information Law (IViR), University of Amsterdam, Amsterdam, The Netherlands

²Law Science Technology and Society (LSTS), Vrije Universiteit Brussel, Brussels, Belgium

 JH, 0000-0002-8286-7147

The deployment of various forms of AI, most notably of machine learning algorithms, radically transforms many domains of social life. In this paper we focus on the news industry, where different algorithms are used to customize news offerings to increasingly specific audience preferences. While this personalization of news enables media organizations to be more receptive to their audience, it can be questioned whether current deployments of *algorithmic news recommenders* (ANR) live up to their emancipatory promise. Like in various other domains, people have little knowledge of what personal data is used and how such algorithmic curation comes about, let alone that they have any concrete ways to influence these data-driven processes. Instead of going down the intricate avenue of trying to make ANR more transparent, we explore in this article ways to give people more influence over the information news recommendation algorithms provide by thinking about and enabling possibilities to express *voice*. After differentiating four ideal typical modalities of expressing voice (alternation, awareness, adjustment and obfuscation) which are illustrated with currently existing empirical examples, we present and argue for *algorithmic recommender personae* as a way for people to take more control over the algorithms that curate people's news provision.

1. Introduction

Few developments stir up societal debate as the rise of artificial intelligence (AI). Its increasing deployment in virtually all domains of life engenders both apocalyptic thoughts of humankind digging its own grave and mythical narratives of a coming new dawn rife with unparalleled possibilities to exceed our human condition. Although we are far from either, it is hard to ignore the profound disruptions AI is already causing in so many societal domains. AI enables organizations to rapidly extract valuable insights from large datasets, discover new solutions through simulations, predict future behaviour of individuals and groups, and facilitate a diverse range of human-machine-environment communications [1]. From agriculture to finance, and from healthcare to public safety, AI is transforming business operations, social processes and local infrastructures. These disruptive changes raise important ethical, legal, social and technical questions about how we can ensure that AI's deployment is beneficial and not harmful to us. Technology is, after all, never an unstoppable or uncontrollable force of nature, but always the product of our making, including the course it may take. Even with AI.

In this paper we focus on transformations in the media ecosystem where the deployment of AI (mostly in the form of machine learning) is becoming increasingly pervasive, and this development is, according to insiders, only likely to continue in the coming years (see, for example, [2]). Next to the challenges of adapting media operations to an online environment (e.g. [3–5]), and the rise of computational or robot journalism (cf. [6–8]), one of the biggest transformations is the data-driven turn towards the audience. The increased possibilities to capture what and how people actually consume news has led media organizations to be more interested in, and responsive to the behaviour (or wishes, interests and needs) of their audiences. This is not just curiosity, but sheer necessity: news organizations need to compete with other information suppliers (social media, blogs, etc.) for the attention of people, so knowing what keeps readers (or *users*, when speaking of online tools and services) engaged is of prime importance (cf. [9,10]). Many newsrooms today therefore use such audience metrics to monitor the reception of their products (and adapt them if necessary), but these also inform various kinds of editorial decision-making [11–15]. Audience metrics are at the same time used by media outlets as input for algorithmic news recommenders (ANR) which target their readers with the news they are (believed to be) interested in (cf. [16,17]). Such algorithmic systems thus enable media outlets to customize their news to highly specific audience segments, possibly leading to fully personalized newspapers, the realization of the so-called 'Daily Me' (cf. [18–20]). Although the rise of algorithms thus impacts both the production *and* the consumption side of news, in this paper we focus on the latter: how ANR influence the selection of information news readers are presented with.

While ANR are presented as tailoring to the interests and preferences of the audience, there is reason to worry that prevalent data-driven inferences might actually silence their individual and collective voices. First, it can be questioned whether ANR and the data-driven inferences of user preferences actually capture what news readers are interested in. Since the algorithms that provide such customized recommendations mostly work on the basis of previous news consumption and other past online behaviour, be it their own and/or of others, there is little opportunity for news consumers to indicate occasional, momentary or future news preferences and needs. But aggregate reading history alone is a poor proxy for reader interest. Even when readers may indicate specific preferences, these are then aggregated in a *general* news taste insensitive to fleeting interests and contingent wishes. Second, since the inferences made by these algorithms are executed in non-transparent ways, news readers have no way to understand why certain recommendations are made to them, what these algorithms are actually optimized

for (including the commercial goals of the newspaper), let alone find out if and how they can intervene in these data-driven processes. Third, as ANR predominantly deliver information that aligns with people's current interests and preferences, they can drive homogeneity and could lower people's chances to encounter different and not yet discovered content, opinions and viewpoints.¹ This would result in increasing societal polarization, the so-called 'echo chambers' [25] and the 'filter bubbles' [26], and consequently, a demise of the public sphere as the interaction of different ideas and people is considered to be crucial to functioning democracies. It should be noted that such discussions have the signs of a moral panic, and current empirical research actually suggests a toning down of such sweeping alarmist claims (cf. [27–32]), especially since people actively select news from many different sources (e.g. [33]). It can furthermore be questioned whether personalized recommendations are merely pernicious for democratic ideals,² or that exposure to a diversity of content and voice is automatically good for pluriform societies.³

The media play a crucial yet complex role in these debates. Although they are generally commercial companies that should serve the interests of the reader (and its stakeholders), the media aspire to live up to their professional ideals and esteem their editorial autonomy as well [41]. But while they are clearly entitled to press freedom, this also entails an important societal responsibility to provide the public with the information it has a right to receive.⁴ After all, the media form an arena for public debate and create the collective realities democracies rest upon (cf. [42–44]). Giving voice and reaching out to the different ideas and opinions in society is therefore at the core of the democratic role of the media (cf. [45–49]), a necessity which is codified as a fundamental human right [50]. But when should editorial judgement of 'what is worth knowing' prevail, and how far should media go in pushing their ideas of what is important news through ANR? And how much of a right to have a say or even complain should readers have? These are old problems, but today's technologically saturated media ecosystem, and in particular the data-driven turn towards the audience, complicates these matters and highlights the complex tensions between editorial autonomy, audience preferences, technological possibilities, corporate interests and democratic responsibilities. The democratic problem we highlight here is not related to the existence of ANR *per se*, but to how they are designed, deployed, and governed. In contrast to common debates where algorithms feature as ontologically stable and transcendent objects of power (cf. [51–53]), we should (and do) have a hand in how such technologies come to play a role in our lives. The critical question is, of course, how to realize this. At the moment, ordinary citizens and news readers have little knowledge about how algorithms define our information landscape (cf. [54–56]), nor do they have many possibilities to deploy them in ways that befit their personal interests and societal needs. This makes them particularly vulnerable to various kinds of manipulation.⁵ Much of this lack of knowledge and influence has to do with the opacity

¹What this homogeneity is about all depends on the context in which ANR are used. If they are used by a specific newspaper then it is mostly about missing certain types of articles (content) one wouldn't otherwise receive. But if the context is a news aggregator, then it is also about missing certain ideological perspectives (cf. [21–24]).

²ANR may, for example, help people deal with today's information-saturated world by pre-selecting what is relevant for each individual (cf. [22,34,35]), help people specialize in the topics they are interested in, hereby fostering the cultivation of 'expert citizens' (cf. [36,37]), or can help foster people's ideological (group) identities, preparing them for agonistic democratic politics [38].

³Some scholars have pointed to the 'backfire effects' of diversity exposure, actually leading to greater polarization: e.g. [39,40].

⁴European Court of Human Rights, Case of the Sunday Times vs The United Kingdom, Application no. 6538/74, Strasbourg, 26 April 1979.

⁵Manipulation takes many forms: from the relatively little dangerous perpetuation of the idea that curating algorithms are neutral technologies which obscures the many contextual factors encoded in them, or the moderately dangerous nudging of people into consuming certain preferred (ideological/moral/political) contents (which is a common business practice in the political economy of web 2.0, see for example: <https://www.theguardian.com/technology/2017/oct/05/smartphone-addiction-silicon-valley-dystopia>), to the very dangerous targeting of highly specific people with perfectly customized messages so as to direct their thoughts and behaviour. The recent Cambridge Analytica/Facebook scandal perfectly attests to this danger: <https://www.nytimes.com/2018/04/12/technology/privacy-researchers-facebook.html>; <https://www.theguardian.com/technology/2017/may/07/the-great-british-brexiteer-hijacked-democracy>. But the broader social media and fake news discussions similarly highlight the manipulative powers of the opaque selections algorithms provide. See also [23,57,58].

of algorithms and the socio-technical systems they are part of (in short: algorithms) (e.g. [59]). As a result, efforts to mitigate the problems associated with ANR therefore focus on more transparency about the precise workings of algorithms (e.g. [60,61]). But algorithms are notoriously difficult to grasp and understand, even for trained people like computer scientists (cf. [62,63]). This is partly because of proprietary reasons, most operators keeping their algorithms secret to maintain competitive advantage, to prevent adversarial learning and gaming (cf. [64,65]). But making the workings of algorithms transparent or explainable is further complicated by the fact that they are often complex entities with a difficult genealogy: they are the product of many interventions, by many different people, with different interests, values and goals (cf. [66–68]). Algorithms are always in flux, because they continuously interact with many different actors (other algorithms, users, developers, aggregate data, etc.) and in various different settings (cf. [69–71]). This makes it hard to pinpoint what exactly can be made transparent. Especially as algorithms are increasingly programmed to improve autonomously, they may become even more elusive, moving beyond the human capacity to grasp and understand.⁶ Recent calls and efforts for more transparency as an ethical principle and a practical means to understand and govern algorithms (cf. [64,72]) may therefore be falling short [73].

Instead of trying to make ANR more transparent, we explore in this article ways to give people more influence over the information news recommendation algorithms provide by thinking about and enabling possibilities to express *voice*. We loosely follow the seminal work of Albert Hirschman who specifies *voice* (communicating one's ideas in order to bring about improvement) as a strategy people have in contrast to *exit* (withdrawal) when they are unsatisfied with provided services [74], and draw more generally from democratic theory and the normative principle that people should have a say over the factors that influence their lives (or at least not be denied or prevented from having such a say) (e.g. [75,76]). We thus define *voice* as *the possibility to exert control over the algorithms that curate people's own news provision*.⁷ In a way, voice can be seen as a way to make the ideals of transparency actionable, and to make them work in everyday practice. Our goal is to advance thinking about how to make ANR truly more responsive to the information wishes and needs of news readers, and to develop more concrete ways by which people can deploy ANR in ways that befit *their* and not (just) corporate interests. This will help media organizations as well to better balance the aforementioned different goals and interests in today's technologically saturated media ecosystem.

We do so by first reviewing the literature on media responsiveness and intervenability in data-driven processes which both make a case for giving people more influence over the data-driven processes that structure people's news provision. We then explore currently existing ways to express voice in relation to ANR in order to build a conceptual frame to think about different modalities of giving voice, which leads us to our alternative: the introduction of algorithmic recommender personae. These are pre-configured and anthropomorphized *types of recommendation algorithms* from which people can choose when browsing (news) sites. Not to be mistaken with idealized (or stereotyped) user types (and their alleged reading behaviour) to which people need to conform, algorithmic recommender personae allow people instead to demand from technologies, such as ANR, to behave in ways that align with their own specific (news) interests at each single moment. Our efforts are part of the larger *Fair News* project, a multidisciplinary research project in the Netherlands funded by the National Science Agenda, in which we study the role of algorithms in the news industry and elaborate the concept of fairness in both theoretical and empirical ways.⁸

⁶This holds particularly for the advances in 'deep learning' which takes AI one step further up. Vast neural networks comprising of different interconnecting layers make up such algorithmic systems whose decisions are almost by definition inscrutable as their ways of knowing, seeing, interpreting may just be radically different from human cognition.

⁷Obviously, voice can mean many types of influence over the algorithms curating news items, like on a more general or legislative level, the options are many, but here we focus on the individual reader and her own algorithmically curated information diet.

⁸The Fair News project is led by IViR researcher Joris van Hoboken, ASCoR Professor of Political Communication Claes de Vreese and Assistant Professor Claudia Hauff at TU Delft. The project involves a collaboration with Dutch newspaper de Volkskrant. Fair News is part of the VWData Initial Impulse Programme (Startimpulsprogramma VW Data), with 'VW

2. Media responsiveness and intervenability

The lack of responsiveness and the powerlessness of the audience has long been described as ‘one of the most difficult problems for media regulation’ [77]. How to balance between people’s interests and editorial ideas of what is important news? For a long time, the media have been veering on the side of editorial judgement, and mechanisms for people to participate, express preferences or to voice dissatisfaction were simply not part of the system [78]. Instead, audience interests were ‘institutionally rather than legally protected’ [79], either through government policies (in the case of the audiovisual media) or self-regulatory frameworks (in the case of the press), instructing the media to take viewers’ interest at heart. Such instructions would range from concrete programme obligations and diversity requirements, or advertising restrictions to pledges of journalistic integrity and professionalism. A central idea behind media diversity policies, for example, is to give the different ideas and groups in society ‘voice’ but also to confront the audience with these ‘other voices’ [80]. The closest mechanisms towards establishing real voice for members of the audience were the ability to send programme complaints, to send letters to the editor, and the establishment of programme councils that would have audience representation. More indirect means, such as audience surveys and polls, existed as well [81], but these measures give a rather limited and at times anecdotal idea of the ‘voice of the audience’ [78,82].

The limited ability of taking into account audience feedback and giving them the opportunity to express voice was at least in parts conditioned by the state of technology, and the lack of interactivity in early mass media. One defining feature of the traditional ‘broadcasting’ approach was the lack of a direct return channel, i.e. an opportunity for the audience to directly express preferences or signal dissatisfaction with a particular programme, other than switching the channel. Digitization in broadcasting technology, but also the proliferation of online formats boosted interactivity as a new means and ideal of engaging with the audience. As a consequence, several waves of audience empowerment followed [83].

Unlike traditional broadcasting and print media, data-driven news personalization can be far more responsive to individual user needs and interests, at least in theory. Hindman goes as far as stating an obligation for journalists to use audience analytics for this purpose [84]. This counts especially for user-driven personalization where users are invited to detail their interests and express explicit preferences for topics, categories, or journalists. Nowadays, many mobile news apps offer sophisticated menus to express interests and preferences. The success of such user-driven forms of personalization as a form of giving users voice is limited, however, partly because of the seeming lack of interests of users [85]. The question is, of course, whether this lack of enthusiasm should be blamed on users, or whether there is room for improvement at the level of the technological and operational design and associated policy frameworks.

The case of data-driven personalized recommendations creates some particular complexities with respect to audience needs and wishes, as a result of the applied logic of inferences and optimization. Here, the media *infer* preferences based on an increasing amount of data that they collect from their users and other sources (e.g. using the BBC iPlayer is nowadays conditional upon users actively signing in and permitting users to share their information, similarly the NPO app). This data is subsequently being analysed and subjected to all kinds of statistical analyses which must help to create an optimal data-driven strategy of how the user can be engaged. Arguably, this could be a new, technologically mediated form of facilitating users’ voice, provided that the interests and preferences that the media inferred, and the resulting algorithmic recommendations, do indeed succeed in articulating the voice of users.

But this assumption alone—that inferred and optimized user preference can also be a form of voice—can be met with serious criticism. As we argued before, it is unclear whether the

Data’ standing for responsible value creation (verantwoorde waardecreatie), and the programme aimed at ‘responsible value creation with Big Data’. The programme is financed by the Netherlands Organisation for Scientific Research (NWO) as part of the recently created National Research Agenda in the Netherlands. The National Research Agenda comprises studies intended to lead to answers to the social and economic issues of the day and the near future.

data that ANR use actually mirrors the true preferences and needs of the user, or that they are outweighed by other signals and optimization benchmarks (e.g. preferences of advertisers, the editorial team, values on the side of technology developers) [86,87]. Zamith points to the importance of distinguishing between construed and real audiences as ‘the former might reflect the latter poorly’ [82]. Another question is if the ‘real’ audience actually knows what it needs or wants [88]. Even if an algorithmic recommendation was successful in accurately predicting the information needs of users, there is still the (ethical) question of whether that should overrule what users (at that moment) believe they want, what editors think news readers should know, and, related to that, what other voices should be shown (diversity).

But even if the media would at some point be able to perfectly know how to engage the user, there would still remain new issues with voice. While the media might be able to better predict with the use of algorithms what people want or need than readers themselves, this does not take away that responsiveness involves the possibility for users to actually have a say. Voice, under these conditions, will probably take on a very different meaning, namely giving users the opportunity to exercise agency (voice) in how their virtual representation is constructed, and how it is used to present them with recommendations that allegedly articulate their interests and preferences. This can take the form of being able to provide some form of feedback on recommendations, have some control over the data that is being used, or the recommendation logic that is being applied. In other words, voice requires transparency, but it should also entail possibilities to influence settings and applied profiles.

Construing the voice of the user is thus closely related to, and interlinked with, the evolution of audience measurement techniques and the governance of user data in reader demographics and usage statistics [89]. Reader statistics used to be built through surveys, and later through new forms of audience surveillance and online tracking. The position of users in relation to such inferred constructions of the ‘voice of the audience’ connect us to another form of voice, which appears as the concept of intervenability in the literature on data protection law and privacy by design. One of the core principles of data protection law is that people should have control over the processing of their personal data. To be more precise, they should have a say in what data is being collected about them, how it is used to forge profiles and what automated decisions are being made. This principle is codified, most recently in the General Data Protection Regulation (GDPR), under the rubric of consent (and the possibility to withdraw one’s consent), transparency, as well as the various data subject rights (the rights to access one’s personal data, the right to ask for correction and/or object to their processing, and the right not to be subject to certain forms of automated decision-making).⁹

In the technical literature on privacy by design, this broad principle of control over personal data processing has been conceptualized as the principle of intervenability in (personal) data-driven processes. As an ENISA study on privacy engineering goals summarizes:

Intervenability ensures intervention is possible concerning all ongoing or planned privacy-relevant data processing, in particular by those persons whose data are processed. The objective of intervenability is the application of corrective measures and counter-balances where necessary.¹⁰

The principle of intervenability can be grounded in the right to informational self-determination, which itself is linked to individual liberty and democratic self-governance. Intervenability safeguards the individual’s effective agency when confronted with data-driven processes and their consequences. Hansen clarifies that while intervenability is primarily focused on safeguards for the data subject, intervenability as a goal is also important for the organizations processing the relevant personal data ([90]; see also [91]). In practice, intervenability requires

⁹Regulation 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).

¹⁰ENISA 2014, Privacy and Data Protection by Design – from policy to engineering.

designing for the possibility for users to give feedback, make complaints and exercise their rights, stop specific data processing operations and intervene or review data-driven decisions and allow for effective oversight.

Intervenability in the data protection framework is linked to the processing of so-called personal data. The use of this normative concept in the context of data-driven news recommendations does not require such a strict limitation. First, it is unlikely that the normative principles in the media context should directly translate into legal obligations, considering the special position of media and the preference for self-regulation following from media freedom. Second, we are less interested here in the flow of personal data *per se*, than in the data flows and processes that power ANR more generally. Even in the situation in which the data is non-personal, the principle and value of voice in the context of the media clearly points in the same direction: the possibility of the audience to exert control over the data-driven processes that shape their news provision. Whereas media theory may have focused most of its attention to processes related to the production, dissemination and consumption of content, data protection principles bring another crucial governance layer (data) into perspective. The data-intensive nature of ANR requires this perspective to come into play as well. It probably goes without saying that intervenability in data-driven processes can be challenging and costly to implement, in particular at scale, considering the nature of the required possibility to intervene. One of the benefits of automation that can be achieved through ANR is precisely that it scales well and allows for personalized news processes to function, *without* the need for (human) intervention. Thus, an important question for research and practice is what forms of intervention at the level of data inputs and processing can be achieved in the context of ANR to further the value of voice, while allowing the benefits of news personalization to materialize as well.

3. Exploring current ways to express voice

The lack of knowledge about and the little influence people have on the information they get to see through algorithmic curation has not solely been an academic concern. A variety of civil society organizations have occupied themselves with these issues, and ordinary citizens have expressed their dissatisfaction as well. The introduction of more advanced forms of algorithmic curation by social media platforms of people's newsfeed has caused, for example, much commotion. When Twitter announced that it was going to 'personalize' the display of tweets based on algorithmically defined 'relevance', the hashtag #RIPTwitter emerged through which users expressed their dismay (<https://www.theguardian.com/technology/2016/jun/07/new-algorithm-driven-instagram-feed-rolled-out-to-the-dismay-of-users>). They tweeted how *they* 'want to decide what is relevant' and feared that this is the 'first step toward becoming Facebook', meaning losing control over the posts they are exposed to (<https://www.makeuseof.com/tag/algorithm-change-death-twitter/>). The point is that people care about having influence over the algorithms that structure the information they get to see. Yet there are few possibilities to actually do so.

In the following we discuss existing ways through which people can express voice or influence the algorithmically curated information they get to see before presenting our own alternative. Building on the work of Hirschman, we differentiate four *ideal typical* modalities of expressing voice: alternation, awareness, adjustment and obfuscation. We explain in conceptual terms what these modalities entail, present examples and evaluate their respective merits and limitations in the current context. We shift between perspectives that involve the input, the computational and the output side of algorithmic curation. Because algorithmic curation in news media is still in its nascent stage, we also discuss the voice of users on social media platforms like Facebook and Twitter, which have become dominant forces in the media ecosystem. The list of possibilities and applications is not exhaustive, nor is it meant to be. The purpose is to map and analytically organize different modalities of expressing voice, so as to foster and advance deployments of

this democratic principle in the algorithmically saturated media ecosystem of today and the foreseeable future.¹¹

(a) Alternation

A first way for people to exert control over the curation practices of algorithms in news provision is simply by switching between different news outlets and media forms, and also by using *multiple* or different recommenders.¹² After all, if people let only one type of algorithm define their information provision, they are quite vulnerable to the opaque workings of that specific curation mechanism, whatever these may be. But when people (can) take advantage of the different personalization algorithms out there, they multiply their possibilities. Alternation operates at a practical level since people can easily choose to visit different news outlets in order to receive different selections of news. It can also be integrated more systemically, for instance through meta-search engines which combine the results from different search engines (e.g. <https://www.searchenginepeople.com/blog/10-meta-search-engines-reviewed-and-compared.html>). While there are to date no systematized ways or unified platforms to alternate between different curating news algorithms, there are a number of other initiatives that illustrate what this could look like. The central point is that alternating between different algorithmic regimes is a means of exerting control over the information people get to see.

In some way, news aggregators and news feeds have performed the principle of alternation in its nascent stage, with RSS readers having gone through an interesting bloom and downfall (see [92]). News aggregators are information services that collect new news items from many different sources in a systematized way (e.g. Google News), while RSS feeds and readers create control over what topics or media outlets users wish to follow.¹³ Other forms of alternation exist by what certain (imaginary) user types find interesting,¹⁴ but nowadays also by what human curators believe to be relevant,¹⁵ and increasingly by algorithms following people's complete browsing history.¹⁶ By receiving news items from many different sources, readers are able to overcome the relatively narrow perspective, or bias, one news outlet might give on a certain topic, and alternate between sources in order to get a wider, more comprehensive view. *MashedMedias*, to give an example, brings together news items from a wide variety of both 'mainstream media' and 'alternative media'. Their 'goal is to provide an unbiased source for news [...] providing exposure to lesser known news sources [...] With people's distrust in the media at an all-time high and there being little to no trust in the mainstream media to report the news fully, accurately or fairly, Mashedmedias.com provides a venue for users to read, research and determine the validity of the news for themselves' (<https://www.mashedmedias.com/about>). The use of news aggregators helps to bypass the algorithms operating at each single news outlet. But as most advanced news aggregators today use sophisticated algorithms to curate peoples' aggregated news feed, the circle is round again. Until a news aggregator or alike arrives.¹⁷

¹¹The authors of this article are keen on being updated on such initiatives, so please send other versions to us by mail to the corresponding author.

¹²Obviously, news readers have great agency to search and gather any kind of information themselves, and we acknowledge that, but here we speak about receiving information through ANR and specifically about ways to take more control over their curating practices. Using completely different ways to gather information is indeed a way to bypass the influence of ANR, but it is not really a way to exert influence over their curating practices.

¹³Think of Feedly (<https://feedly.com/i/welcome>), Alltop (<https://alltop.com/>), TheOldReader (<https://theoldreader.com>) and Awasu (<https://awasu.com/>).

¹⁴This is what the algorithmically personalized news aggregator app News360 does; besides giving personalized recommendations based on each user's reading behaviour, users can choose to read the 'top stories today' for 'Jack, Businessman', 'Mary, Fashionista', and 'Michelle, Web Designer', each of them with a graphic image of such personae.

¹⁵In particular, Digg (<http://digg.com>) is relevant here as the human curators are presented on their website, cultivating a certain culture for consumers to identify with, but there are also other news aggregators like Flipboard and the Dutch *Blendle*.

¹⁶Especially NewsPrompt (<https://newsprompt.co>), SmartNews (<https://www.smartnews.com/en/>) and News360 (<https://news360.com>) are big on algorithmic curation and other forms of personalization.

¹⁷Google recently launched a new feature called 'full coverage' that 'provides a complete picture of how that story is reported from a variety of sources' and 'is the same for everyone—it's an unpersonalized view of events from a range of trusted news

One type of news aggregator is worth mentioning here in more detail, the so-called ‘matrix’ news aggregator, which lets readers alternate not just between different news sources, but between different predefined ideological perspectives, or ‘frames’. These aggregators generally present the *same* news event from a limited number of ideologically different origins next to each other, in a matrix, so that readers can learn about the same event from multiple angles by alternating between these different perspectives. Computer and information scientists from Germany have developed *NewsBird*, an ‘extensible news aggregator’, to ‘reduce the effects of media bias’ by showing news articles on international events from different perspectives (now mostly organized by country) in a comprehensive matrix (see their article: [93]). *AllSides*, a US-based non-profit political news aggregator, starts from a similar assumption that ‘unbiased news doesn’t exist’ and ‘provides [readers with] multiple angles on the same story so you can quickly get the full picture, not just one slant’ (<https://www.allsides.com/unbiased-balanced-news>). Their website presents news items on the same topic from different mainstream media outlets labelled by them on a five-item political left-to-right scale (L-L-C-R-R) so that people are aware of the ideological origins of the news they consume, but also to facilitate reading across ideological divides.¹⁸ Although not really a news aggregator, the Swedish *Filterbubblan* is ‘a tool that lets [readers] see how different debates sound in the three ideological bubbles that characterize Swedish domestic politics, from left to right’, and alternate between these ‘filter bubbles’ with one click (<http://filterbubblan.se/>).

While these efforts are great at reducing predefined biases by letting news readers alternate between different perspectives and sources, they do little to work around hidden biases in algorithmic curation. For that to work, readers need to be able to alternate between different curating algorithms so that they can see how these result in different news provision logics and outcomes and consequentially choose what news outlets and personalization tools to consult in order to bypass each specific algorithmic bias.

(b) Awareness

An important way for people to gain more influence over the curating practices of algorithms is by *being aware* of them. Understanding the workings of algorithms may not be easy, as previously discussed, but raising awareness about personalization practices is an important and necessary first step. Often underestimated in current debates is the ability of individuals to make their own choices as to what news outlets to visit, which people to befriend and follow, and what data-collection practices to allow. This can be related to people’s media literacy more generally. For people to make effective decisions they need to be aware of the personalization practices at work, and so we find various efforts, legally enforced or not, to inform users.

First, depending on the reigning legal framework, and on the ethical and political signature of the organization, some news outlets provide information about their personalization practices themselves. Often, as a result of the personal nature of user data that is powering ANR, these are explained in the privacy policy where news outlets detail what user data they collect and why [94]. A prominent media and publishing house in the Netherlands, *De Persgroep*, explains, for example, that they collect data and place cookies not just ‘to advertise purposefully’, but also ‘to better understand our readers and users. Based on information about demographics, interests and behaviour of our users, we try to improve our products and services and better attune them to the wishes of the users’ (<https://www.persgroep.nl/privacy?cookie=1>). *The New York Times* states that they collect user data to ‘improve services we offer you, including customized recommendations, advertising and currency display, to improve marketing, and to track access and use of the NYT Services across the devices that you use to access the NYT Services’ (<https://help.nytimes.com/hc/en-us/articles/115014892108-Privacy-policy#b>). Providing personalized

sources’. See Trystan Upstill, ‘The new Google News: AI meets human intelligence’, Google Blog, 8 May 2018, <https://www.blog.google/products/news/new-google-news-ai-meets-human-intelligence/>.

¹⁸*AllSides* operates in five more domains, e.g. they developed a ‘balanced’ search engine that gives different instead of the most popular results, and (school) programmes fostering civic dialogue and mutual understanding.

news services seems to go hand in hand with targeted advertising and other commercial activities. Yet while news readers are on such pages informed about the more general user data collection and processing practices, very little is disclosed about *how* that specifically influences the information they get to see.¹⁹

The same can be said about dominant social media platforms like Facebook, which have extensive information in their privacy policy about the data they collect and use in their personalization operations, but tend to say little specifically about how such data influences people's *feed*. Other actors, like research institutes, investigative journalists, civil society organizations and private commentators have increasingly stepped into their place by sharing general knowledge about algorithmic curation,²⁰ but these are similarly unable to tell individual users how that affects the information they get to see. Some initiatives, however, have sprung up in recent years to raise more awareness among Internet users about the specific news they consume. Although they focus less on algorithmic curation *per se*, they do give people the ability to reflect on their news diet, especially insofar as they might be caught in an ideological 'filter bubble'. Researchers from University of Michigan's School of Information, for example, developed *Balancer*, a Chrome extension that 'analyses your Web browsing to show you the political slant of your reading history. If you get way out of balance, we may even give you reading suggestions' (<http://www.balancestudy.org/balancer/>). This plugin shows users whether they have read more articles from news sites categorized as liberal or conservative, and hopes to stimulate users to read more from the other side.

In sum, raising awareness about algorithmic curation is widely done, including through user data privacy statements, and goes some way in giving news readers influence over the algorithms structuring their news provision. It would be an interesting empirical question to see how these efforts resonate with people's own cultural, or 'folk', understandings of such technologies and how to game them. We furthermore foresee that the GDPR, which raises the bar on transparency and user control over personal data processing, will have a positive impact on awareness. However, as long as most information is general, not intuitive and not specific about its workings or results, it may do little to help news readers make concrete choices about how to engage with different ANR.

(c) Adjustment

A third way in which people can be given influence over the curation practices of algorithms is through offering the possibility to adjust them according to their interests and wishes. This is, moreover, not just a moral matter, there are commercial interests pushing this imperative as well. After all, if people can positively influence the technologies they use, then it is likely that they will have greater satisfaction using them, and thus use them more, while generating more data for the companies about user interests (see, for example, [95]). To some degree, users are offered possibilities to adjust the algorithms that structure their news feed already, and we will explore these in order to see what they entail.

Most news outlets have developed, to date, no formal ways to influence their curating algorithms [94]. Perhaps this is too much to ask, since most of these outlets have only started to explore personalization for their readers, and currently predominantly offer a 'recommended for you' section on their websites. Based on what types of articles people read, they give suggestions of other articles they might be interested in. Sometimes users can actively indicate the journalists they wish to follow or the topics they are interested in, similar to the way in

¹⁹One exception is perhaps *LabRdr*, an experimental offline news app by *The Guardian Mobile Innovations Lab* that aims to be transparent about what data it uses and how: 'We want to be incredibly transparent about how we are using the data we gather about your commute patterns and the things you choose to read. To that end, for each permission we request, and for each piece of data it allows us to track through the app, we also explain how we'll use it to aim to improve your commute reading. We've also created a section — the Log — where you see what data we've collected on your use of the app'. But again, how this influences the articles you get to see is not entirely clear.

²⁰Such analyses and commentaries literally abound in thousands, simply search for 'how does XX news feed work?' or 'how to influence your XXX feed'.

which news aggregators give some control over the curating algorithm. At best, news outlets give an option to ‘opt in/out’ from certain or all of their personalization services,²¹ but this is a rather limited way of expressing voice. Some, like former *The New York Times* public editor Liz Spayd, argue to more actively include readers and their opinions: ‘think of Pandora radio, where the company [...] lets the listener give the thumbs up or down to specific songs, thus changing the algorithm’ (<https://www.nytimes.com/2017/03/18/public-editor/a-community-of-one-the-times-gets-tailored.html>). Although this would be a great step forward, such methods may mostly serve further optimization, and leave other recommender values, such as diversity or serendipity, untouched.

And what about social media platforms? As the introduction of this section highlighted, many people care about having control over their social media news feed. While Instagram remained insensitive to the discontents of users when they introduced algorithmic curation, and continues to offer no option to adjust it, Twitter responded by letting users decide whether they would like to personalize their timeline or not (https://blog.twitter.com/marketing/en_us/a/2016/an-improved-timeline-for-consumers-and-brands.html). Facebook has more extensive options to adjust the settings that influence people’s news feed. In addition to allowing users to change their news feed from the relevancy-based ‘top stories’ to the chronological ‘most recent’, they let users prioritize who or what they want to see first by (de)following certain friends, pages and groups. The same counts for their ads, or ‘suggested posts’, where users can indicate if they are not interested in such ads. Facebook also enables users to create different lists of ‘friends’, e.g. family, colleagues, etc., which lead to different news feeds between which people can choose. Although these settings give people some influence over the information they get to see, there remains little opportunity to control (or know) the actual weighting of these variables in the algorithmic systems. How such setting changes affect the news feeds of people remains therefore a relative black box. The irony is that while users have little control over their news feeds, these platforms offer powerful filtering tools to advertisers to target the user groups they are interested in. This makes clear again that when people don’t have to pay for a service, they may be better understood as the product being sold (cf. <https://lifehacker.com/5697167/if-youre-not-paying-for-it-youre-the-product>). And it may make little commercial sense to give your products voice, especially if there are two billion of them. When media business models rely on solely advertising, it can be questioned whether there is enough incentive to enable more user control, especially the types of control that would run counter to the business’ commercial optimization benchmarks.

A recently developed (external) tool to control the algorithmic curation of people’s news feeds is Gobo, ‘a social media aggregator with filters you control’, developed by Ethan Zuckerman and colleagues at the MIT Media Lab’s Center for Civic Media (<https://www.media.mit.edu/projects/gobo/overview/>). It works by aggregating posts from people’s contacts on Twitter and Facebook, but lets *users* decide which ones they want to see, instead of these social media platforms. They do so by offering multiple adjustable filters to their aggregator: *politics*, *seriousness*, *rudeness*, *virality*, *gender* and *brands*. The politics slider, for example, lets people get out of their ‘filter bubble’ by moving it from ‘my perspective’ to ‘lots of perspectives’, and the ‘gender’ slider makes it possible for people to hear more from women ‘who often get shouted down in online dialogues’. The easy-to-use sliders make it possible for people to play with these filters to find out what effect they have on their news feed. It is a technology to think and play with, as they say. The project is building more filters for people to further customize their news feeds, but as Gobo is an open experiment in developing civic social media, they have made it possible for others to develop and add filters as well. Finally, Gobo ‘aims to be completely transparent, showing you why each post was included in your feed’, something that is not the case with the aforementioned platforms. Gobo’s underlying mission is ‘to open a conversation about who gets to filter what you see on the Web’ as they believe in ‘escaping echo chambers’ and that ‘seeing a

²¹This is, for example, what the BBC offers: <http://www.bbc.com/usingthebbc/account/about-your-personalisation-settings/>; or, as the aforementioned experimental offline news app *LabRdr* offers, users can opt in to certain personalization features and alerts, but also lets its readers ‘adjust the settings to get a longer or shorter reading package’.

wider picture of the news is a critical piece of democratic society' (<http://www.ethanzuckerman.com/blog/2017/11/16/who-filters-your-news-why-we-built-gobo-social/>).

In summary, the possibilities to adjust the algorithms that structure people's news feeds are still rather limited, but some economic and democratic motives to expand these are definitely there. The challenge (for companies) is to balance different interests appropriately by keeping users satisfied with their offerings by giving them more control, while supporting their business models. Whether we can expect sufficient innovation from commercial companies to support voice remains to be seen, but it should be possible to find technological solutions that may fulfil both goals.

(d) Obfuscation

For personalized news curation to work, algorithmic systems need user data, for without information on what news readers like (to read), they have little notion about *how* to provide individual recommendations. A last, and perhaps counterintuitive, way for people to gain more influence over ANR is by mobilizing against the data-driven processes through obfuscation. This is a strategy conceptualized by Helen Nissenbaum and others, defined as 'the deliberate addition of ambiguous, confusing, or misleading information to interfere with surveillance or data collection' [96]. The idea behind it is that privacy is guaranteed through making a lot of noise around people's own (online) data-producing activities so that these disappear in a cloud of meaningless information. In the specific context of ANR, this could mean a variety of things, starting from the simple handwork of clicking on all news items to sharing one's news profile or device with others, all with the purpose of hindering the construction of a personalized news taste and the profiling of individuals as certain types of readers. We realize that obfuscation may run against some of the goals and benefits of personalization, but it is nevertheless a distinct way to gain influence over the curating practices of algorithms, albeit disruptively. Since no formalized strategies of obfuscation in the context of algorithmic news personalization exist yet, we discuss other more established methods of obfuscation here, both as example, and as inspiration for thinking about potential productive ways to deploy obfuscation as a strategy to gain more control over the news people get recommended.

The first fully developed obfuscation tool is the *TrackMeNot* plugin for Web browsers. As creators Howe and Nissenbaum explain, it 'runs as a low-priority background process that periodically issues randomized search-queries to popular search engines. It hides users' actual search trails in a cloud of "ghost" queries, significantly increasing the difficulty of aggregating such data into accurate or identifying user profiles' [97]. Similarly, obfuscating is the tool they developed together with Zer-Aviv, called *AdNauseam*, which is targeted mostly at the advertising industry: 'the plugin works with existing adblockers to block ads on visited pages, but then quietly clicks each ad in the background, polluting user profiles and creating mistrust between advertisers and the networks they pay for clicks' [98]. Both tools serve not only people's privacy, but also enable them to express their discontent with the pervasive surveillance technologies that keep track of their online behaviour, whether for commercial or political ends. More commercial versions are the many VPN connectors which obfuscate your location and go as such against 'geo-targeting' and local censorship. As one rather popular service explains, 'there isn't a simple fix to the filter bubble phenomenon, but the use of IPVanish VPN will lessen its impact. When you connect to IPVanish, your Web traffic is funnelled through an encrypted tunnel and your IP address is altered. This limits data collection and breaks down location-based tracking' (<https://www.ipvanish.com/why-vpn.php>). In short, these obfuscation strategies engage directly with the interfaces and protocols of the technologies they wish to subvert, turning their logic against them.

Now, what to make out of this strategy for our case here? As explained before, the central aim of these obfuscation strategies is to resist and intervene with the datafication of people's online behaviour and their consequential categorization in certain news reader profiles. But what if this exactly is what news readers need in order to benefit from news

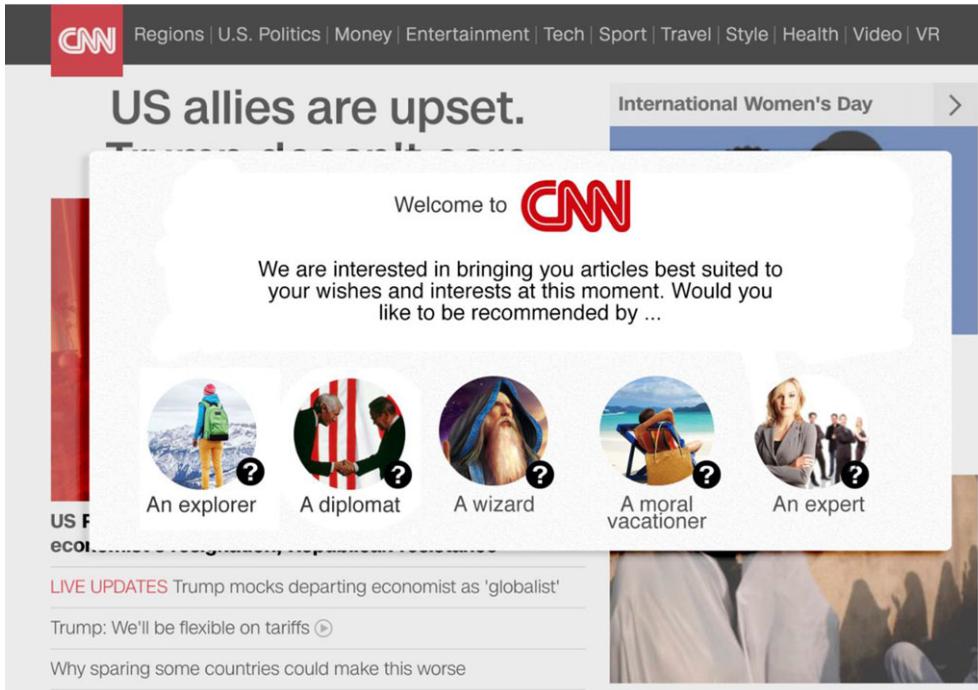
personalization? Obfuscation may therefore be a great form of resistance against large-scale surveillance systems, but it disavows and disables the possibilities for algorithmically defined personal news recommendations. Or so it seems. Next to resistance, the other main objective of obfuscation is to take back control, to reassert agency over the user data that is used, from the technological systems that leave little room for users to engage with and reply to. Perhaps we can think of ways to *selectively* disrupt the surveillance and categorization practices at work, in order to still benefit from the advantages of getting personalized news recommendations, but this time on people's own terms? This brings such obfuscation strategies closer to 'gaming', the strategic interaction with technological systems and their underlying logics in order to bend them to one's own ends or interests. An early example of this in the media algorithms context is the so-called 'Google bomb', in which a linking campaign can return a politically humorous result for a particular query. In a recent article, Gürses *et al.* point to the strategic activities of disadvantaged user of optimization systems, such as Uber drivers, who try to correct, reverse or improve 'unfair' outcomes by feeding misinformation into the system. Gürses *et al.* show how such subversive activities can be formalized as a design principle which they call 'protective optimization technologies' [99]. We think it is safe to expect a lot more in this area of data activism and mobilization against our algorithmic overlords.

4. The case for algorithmic recommender personae

When thinking about concrete ways to materialize voice in the context of algorithmic news recommenders, all four aforementioned modalities play an important informative role. We regard it important that users can easily alternate between different types of recommendation regimes, that they are (made) aware of how these work and what data they use, that they can adjust them to whatever needs they have. While we consider obfuscation a great strategy of resistance, it will remain challenging in the context of news recommenders how to translate this into real productive possibilities. So far, the Gobo tool offers the best concrete way users can exert control over curating algorithms. It is built on the assumption that (digital) technologies should be working in our favour, and that (social) media platforms have a civic and public utility function. These are dedications we endorse, and we share Gobo's concern that it is undesirable and dangerous that algorithms (and the commercial parties behind them) control what we see, while we don't have control over these algorithms ourselves. ANR should be working to help people with their information needs, not just to 'optimize user retention and engagement so that our attention can be sold to advertisers' (<http://www.ethanzuckerman.com/blog/2017/11/16/who-filters-your-news-why-we-built-gobo-social/>). While Gobo is a great example of how ANR can be made more responsive to user demands by offering adjustable filtering sliders, we doubt whether many users would actually experiment and play with all these settings. Yes, there will be some tech-savvy people who might have the patience and curiosity to adjust these sliders each time they go on social media, but experience in commercial and academic settings shows that most users won't spend much time tinkering with settings. Especially as we expect the number of sliders to increase, and choice will thus increasingly abound, people might shy away from using any slider at all. That's the paradox of choice: the more to choose from, the less we actually seem to do. As much as we like Gobo, and as much as we believe it should be developed further and acquire a prominent place in today's (social) media ecosystem, we argue that there should be something else, something easier and quicker to use.

Gobo prompted us to think further. The question that guided us in this article was how to retain the complexity of multiple adjustable filtering settings as a means to express voice, while enabling such choices to be exercised in an easy and more intuitive manner. Following this imperative, we have come up with the idea of algorithmic recommender personae. These are pre-configured and anthropomorphized *types of recommendation algorithms* from which people can choose from when browsing (news) sites. With one click, they get different *sorts* of (news) recommendations, enabling them to easily switch between different 'versions of the world' based on different optimization strategies, depending on people's momentary (news) mood and

purpose. These recommender personae thus give people the ability to influence their information provision with the simplicity of clicking on one specific recommender persona. Informed by theoretical understandings of knowledge engagements and technological feasibilities, we currently distinguish between the Explorer (news from unexplored territory), the Diplomat (news from the other side), the Wizard (surprising news), the Moral Vacationer (guilty pleasures), or the Expert (specialized news based on previous consumption). We will explain these personae in more detail further in this paper, but underneath is an image of what a news outlet's landing page deploying them could look like. While the expressive images give people a quick idea about what types of recommendations are provided by that persona, the question mark hides more specific information about each particular recommending principle.



The use of personae is not novel. We partly took our inspiration from the different avatars gamers can choose from when playing various kinds of role-playing video games [100], each pertaining different characteristics and qualities. But personae are also widely used in various design processes. Starting in the field of software development, personae were introduced as a tool to help design teams do their work better by keeping target users in mind. Following Alan Cooper who developed and popularized this tool in the early 2000s [101], 'personas are not real people, but they are based on the behaviours and motivations of real people we have observed and represent them throughout the design process' [102]. In other words, they are ideal typical constructions of prospective users based on empirical, mostly qualitative, research.²² Because personae are anthropomorphized, the argument goes, developers are better able to understand and empathize with the people who are meant to use their products. In this way, developers can and should think, talk and interact with and about these personae during the design process. As this design tool turns out to be highly productive, it is nowadays used in many different fields: from marketing [104] and product design [105], to healthcare [106].

While in concept similar, our use of personae to anthropomorphize algorithmic recommenders is different from these sorts of applications. First, because our usage is not meant to help develop a future product, but is a product in itself. The users of our personae are therefore not producers,

²²Interestingly in this context, some have recently argued to construct persona's automatically from social media profiles by means of algorithmic machine learning [103].

but news readers. A more important difference, however, is that our users need not identify with the personae in order to have them work, whereas in other situations this generally is the case. This is a crucial difference which becomes clearer when we take an empirical example from the news sector we discussed earlier. The highly personalized news aggregator *News360* uses such user personae. In addition to recommendations personalized for ‘You’ (based on people’s own interests and/or online reading behaviour), users can choose to read the ‘top stories today’ for ‘Jack, Businessman’, ‘Mary, Fashionista’, and ‘Michelle, Web Designer’, each of them with a graphic image of such personae. While this is an easy way to alternate between different news preferences, such *user type* personae force people to identify with these types, while failing to inform people about what *kind* of recommendations they will receive other than through stereotypical assumptions of what these types would read. This is problematic because some news readers might want to read financial news, but would not (like to) identify at all with a businessman, or they might like to read entertainment gossip but would not (want to) identify with a fashionista. The point is that giving users the possibility to read recommendations for a certain (stereotypical) persona brings about unnecessary requirements for identification, and may not be responsive to their individual news interests at all. To avoid this problem, we propose to develop *algorithmic recommender types* instead of *news reader types*. The choice is not what kind of news readers people want to identify with, but from what kind of recommender persona they would like to receive recommendations from. This is a crucial difference, which, we hope, could facilitate more and easier alternations between different recommender algorithms based on people’s information interests at each separate occasion. Now let us explain the recommender personae we currently have in mind.

The *Explorer*, offering people ‘news from unexplored territories’, is inspired by the notion of diversity. As explained before, the central concern around news recommenders is that they lessen the possibilities for people to encounter news and opinions that fall outside their ordinary online histories and reading habits. This is, however, not just an academic or political concern: news consumers emphasize that they find it important to be able to hear and read about topics they haven’t thought of, viewpoints they don’t quite understand, and perspectives that are unknown. Diversity in recommender design needs, however, to go beyond mere variance or serendipity, and would benefit from a more diverse conceptualization of diversity [22]. One can think of diversity in terms of ideological background, professional origin, writing and presentation style, article length, and so on. A growing literature on diversity in relation to the digital news environment can support conceptualizing diversity further and developing the relevant algorithms accordingly [22].

The *Diplomat*, bringing people ‘news from the other side’, is inspired by the notion of intellectual diplomacy. To combat polarization and fragmentation, often associated with algorithmic curation, this recommendation algorithm provides readers with (positive) information from or about ideological or political opponents. Building on the work of sociologists Zygmunt Bauman and Bruno Latour, who separately developed this role of *diplomacy* as the key practice for intellectuals in a multifaceted postmodern world (cf. [107,108]), this recommender persona strives to live up to that role of mediating between different societal groups and their worldviews. The goal is not just to recommend articles about or from opponents, but as far as contents allow, to include more news articles that are based on positive or constructive journalism so as to foster mutual understanding and empathy (cf. [109,110]).

The *Wizard*, giving people ‘surprising news’, is inspired by the notion of serendipity. News personalization is often said to reduce the chances for people to encounter information they had never thought of before. Often used as and confused with pure chance or randomness, serendipity is a concept that comprises of more conceptual richness and variety as well. Although research on this concept is still nascent,²³ we build from such works to think about more diverse ways to recommend unknown articles, other than random variation so that algorithms can learn better (which is the dominant design rationale for including unrelated items). There are multiple

²³Most notably, Yaqub [111], cf. <https://erc.europa.eu/projects-figures/stories/looking-one-thing-finding-another>.

concrete metrics to base such recommendations on, think of popularity among other user types or other geographical areas.

The *Moral Vacationer*, offering people ‘guilty pleasures’, is inspired by the notion of escapism. Consuming news can be emotionally tough. Getting informed about all that goes wrong in the world may be important, but it also asks much of news readers who often get affected by it. Oftentimes, they may want to escape from the heaviness of most ordinary news items and would like to consume lighter, or simply entertaining pieces. Inspired by Don Weenink’s usage of the ‘moral vacation’ to conceptualize the occasional inhibited behaviour of ‘normal’ youngsters to escape the disciplining forces in their everyday lives [112], we imagine a recommender persona that offers pleasurable moments away from the seriousness of ordinary news. Depending on the user, these may include entertainment stories, music reports, funny stuff, positive news, cultural reviews, but one can think of other items that may emotionally satisfy as well.

The *Expert*, at last, recommends ‘specialized news’, and is inspired by the notion of (citizen) expertism. Often suppressed by moral panics around filter bubbles, news personalization can benefit democratic and professional participation as it helps people to specialize themselves in, and deepen their knowledge of certain, perhaps rather specific, subjects [37]. This recommender analyses the topics people most frequently consume and puts forward other items on these subjects which they have not yet found, but which contain more background analyses, opinions and empirical studies. The emphasis is very much on bringing people high quality and specialized content which would improve their understanding of desired topics.

These five algorithmic recommender personae are obviously not exhaustive of the many different possibilities they can offer, but we believe these personae are a good start to explore and develop further.²⁴ Our plan is to further advance the technical and theoretical underpinnings of these personae, so that we can build and beta-test them with news readers in order to empirically test their functionality (e.g. do they give users what they want?) and usefulness (do news readers experience a greater sense of control, more engagement and more satisfaction?).

5. Conclusion

The deployment of various forms of AI, most notably of machine learning algorithms, radically transforms many domains of social life. In this paper we have focused on the news industry, where different algorithms are used to customize news contents to increasingly specific audiences. Based on their online reading behaviour and other (self-assigned) characteristics, such algorithms select news items that are thought to be relevant for (or best optimized to engage) each specific user. While this personalization of news enables media organizations, at least in theory, to be more receptive to the needs and wishes of their readers and can be a great aid for people to specialize themselves as ‘expert citizens’, it can be questioned whether current deployments of ANR live up to their emancipatory promises. Like in various other domains, people have little knowledge of what personal data is used and how such algorithmic curation comes about, let alone that they have any concrete ways to influence these data-driven processes, nor to deploy them for the specific purposes they have in mind. Whose interests are actually being served here?

To give people more influence over the curating practices of such algorithms, we have explored in this article ways to express voice in today’s media ecosystem. After differentiating four ideal typical modalities of expressing voice (alternation, awareness, adjustment and obfuscation) which are illustrated with currently existing empirical examples, we have presented and argued for *algorithmic recommender personae* as a way for people to get more control over the algorithms that curate their news provision. These pre-configured and anthropomorphized types of recommendation algorithms enable people to intuitively express their wishes of different sortings of news articles depending on their specific (news) mood and purpose. The crucial and defining characteristic of these recommender personae is that they do not ask people to identify

²⁴We have deliberately not included any clearly negative personae (show me the news that makes me angry), or personae that aim to further antidemocratic goals (show me only the news that proves that I am right).

with certain stereotypical news reader types (as is sometimes done). Instead, the user is given the opportunity to demand from algorithmic systems to behave in a certain desired way. Obviously, the range and the specific possibilities of these different types of recommender algorithms are limited and defined by those who develop them, but we believe it remains a powerful and empowering move to ask technology to comply with us, the people, instead of us having to comply with the structures that are given.

This specific, and the wider, lack of control over curating algorithms cannot be properly understood without taking the broader political economy of the Internet into account. Whereas originally promised to emancipate and liberate people by offering free and open exchanges of information and technologies, the Internet is in the last decade increasingly dominated by strong commercial market players which have little interest in maintaining decentralized and interoperable technological infrastructures where users have much autonomy and agency. The rise and fall of RSS, a great example of an open, user-oriented technology, within the Google ecosystem can be seen in this light. The big tech players of today may simply have more interest in drawing users into and keeping them in their platform ecosystems, thus generating more data, traffic and revenue, than to allow people to deploy technologies in ways that maximizes voice. However, in the context of the news industry, and especially from quality newspapers and public service media, we may still hope to expect a different take. Such organizations have, after all, a strong public mission to contribute to an open and diverse debate on societal issues. The use of recommender personae is perfectly fit to facilitate user influence, while guaranteeing traffic and revenue through increasing user engagement for such media organizations. We hope to see more of these balancing initiatives, as they would most likely have the biggest chances for survival.

Data accessibility. This article has no additional data.

Competing interests. We declare we have no competing interests.

Funding. This study is part of the Fair News project, which is financed by the Netherlands Organisation for Scientific Research (NWO) as part of the recently created National Research Agenda in the Netherlands.

References

1. One Hundred Year Study on Artificial Intelligence (AI100), Stanford University. See <https://ai100.stanford.edu> (accessed 9 July 2018).
2. The 2018 Reuters Digital News Report, page 29. See <http://www.digitalnewsreport.org/X>
3. Macnamara J. 2010 *The 21st century media (r) evolution: emergent communication practices*. New York, NY: Peter Lang.
4. Meyer P. 2009 *The vanishing newspaper: saving journalism in the information age*. Columbia, MO: University of Missouri Press.
5. Van der Haak B, Parks M, Castells M. 2012 The future of journalism: networked journalism. *Int. J. Commun.* **6**, 16.
6. Anderson CW. 2013 Towards a sociology of computational and algorithmic journalism. *New Media Soc.* **15**, 1005–1021.
7. Coddington M. 2015 Clarifying journalism's quantitative turn: a typology for evaluating data journalism, computational journalism, and computer-assisted reporting. *Dig. Journ.* **3**, 331–348.
8. Dörr KN. 2016 Mapping the field of algorithmic journalism. *Dig. Journ.* **4**, 700–722.
9. Crogan P, Kinsley S. 2012 Paying attention: towards a critique of the attention economy. *Culture Machine* **13**, 1–29.
10. Davenport TH, Beck JC. 2001 *The attention economy: understanding the new currency of business*. Boston, MA: Harvard Business Press.
11. Anderson CW. 2011 Between creative and quantified audiences: web metrics and changing patterns of newswork in local US newsrooms. *Journalism* **12**, 550–566.
12. Christin A. 2018 Counting clicks: quantification and variation in web journalism in the United States and France. *Am. J. Sociol.* **123**, 1382–1415.
13. Lee EJ, Tandoc Jr EC. 2017 When news meets the audience: How audience feedback online affects news production and consumption. *Hum. Commun. Res.* **43**, 436–449.

14. Petre C. 2015 The traffic factories: metrics at chartbeat, gawker media, and the New York Times. Tow Center for Digital Journalism.
15. Welbers K, Van Atteveldt W, Kleinnijenhuis J, Ruigrok N, Schaper J. 2016 News selection criteria in the digital age: professional norms versus online audience metrics. *Journalism* **17**, 1037–1053.
16. Thurman N. 2011 Making ‘The Daily Me’: technology, economics and habit in the mainstream assimilation of personalized news. *Journalism* **12**, 395–415.
17. Gillespie T. 2014 The relevance of algorithms. In *Media technologies: essays on communication, materiality, and society* (eds T Gillespie, PJ Boczkowski, KA Foot), ch. 9, pp. 167–194. Cambridge, MA: MIT Press.
18. Harper C. 1998 *And that's the way it will be: news and information in a digital world*. New York, NY: NYU Press.
19. Negroponte N. 1995 *Being digital*. New York, NY: Alfred Knopf.
20. Sunstein C. 2009 *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
21. Bozdag E, van den Hoven J. 2015 Breaking the filter bubble: democracy and design. *Ethics Inf. Technol.* **17**, 249–265.
22. Helberger N, Karppinen K, D’Acunतो L. 2018 Exposure diversity as a design principle for recommender systems. *Inf. Commun. Soc.* **21**, 191–207.
23. Hindman M. 2008 *The myth of digital democracy*. Princeton, NJ: Princeton University Press.
24. Stroud NJ. 2011 *Niche news: The politics of news choice*. Oxford, UK: Oxford University Press.
25. Sunstein C. 2017 *#Republic: divided democracy in the age of social media*. Princeton, NJ: Princeton University Press.
26. Pariser E. 2011 *The filter bubble: what the Internet is hiding from you*. London, UK: Penguin.
27. Flaxman S, Goel S, Rao JM. 2016 Filter bubbles, echo chambers, and online news consumption. *Public Opin. Q.* **80**, 298–320.
28. Fletcher R, Nielsen RK. 2017 Are news audiences increasingly fragmented? A cross-national comparative analysis of cross-platform news audience fragmentation and duplication. *J. Commun.* **67**, 476–498.
29. Haim M, Graefe A, Brosius HB. 2018 Burst of the filter bubble? Effects of personalization on the diversity of Google News. *Dig. Journ.* **6**, 330–343.
30. Möller J, Trilling D, Helberger N, van Es B. 2018 Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. *Inf. Commun. Soc.* **21**, 1–19.
31. Nguyen TT, Hui PM, Harper FM, Terveen L, Konstan JA. 2014 Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, pp. 677–686. ACM.
32. Zuiderveen BF, Trilling D, Moeller J, Bodó B, de Vreese CH, Helberger N. 2016 Should we worry about filter bubbles? *Internet Policy Review. Journal on Internet Regulation* **5**, 1–16.
33. Thorson K, Wells C. 2015 Curated flows: a framework for mapping media exposure in the digital age. *Commun. Theory* **26**, 309–328.
34. Kovach B. 2010 *Blur: how to know what's true in the age of information overload*. New York, NY: Bloomsbury.
35. Shenk D. 1998 *Data smog: surviving the information glut*. San Francisco, CA: Harper.
36. Bang H. 2005 Among everyday makers and expert citizens. In *Remaking governance: Peoples, politics and the public sphere* (ed. J Newman), pp. 159–178. Bristol, UK: Policy Press.
37. Li Y, Marsh D. 2008 New forms of political participation: searching for expert citizens and everyday makers. *Br. J. Polit. Sci.* **38**, 247–272.
38. Mouffe C. 2011 *On the political*. London, UK: Routledge.
39. Wojcieszak M. 2011 When deliberation divides: processes underlying mobilization to collective action. *Commun. Monogr.* **78**, 324–346.
40. Garrett RK, Gvirsman SD, Johnson BK, Tsfati Y, Neo R, Dal A. 2014 Implications of pro- and counterattitudinal information exposure for affective polarization. *Hum. Commun. Res.* **40**, 309–332.
41. Schudson M. 2003 *The sociology of news, contemporary societies*. New York, NY: Norton.
42. Hampton M. 2010 The Fourth Estate ideal in journalism history. In *The Routledge companion to news and journalism* (ed. S Allan). London, UK: Routledge.
43. Muhlmann G. 2010 *Journalism for democracy*. Cambridge, UK: Polity Press.

44. Starr P. 2005 *The creation of the media: political origins of modern communication*. New York, NY: Basic Books, Inc.
45. Hardy J. 2014 *Critical political economy of the media: an introduction*. London, UK: Routledge.
46. Helberger N. 2011 Diversity by design. *J. Inf. Policy* **1**, 441–469.
47. Kari K. 2013 *Rethinking media pluralism*. New York, NY: Fordham University Press.
48. Napoli PM. 2011 Exposure diversity reconsidered. *J. Inf. Policy* **1**, 246–259.
49. Recommendations and Declarations of the Committee of Ministers of the Council of Europe in the field of media and information society. 2016 See <https://rm.coe.int/1680645b44C>.
50. Eskens S, Helberger N, Moeller J. 2017 Challenged by news personalisation: five perspectives on the right to receive information. *J. Media Law* **9**, 259–284.
51. Barocas S, Hood S, Ziewitz M. 2013 Governing algorithms: a provocation piece. See https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2245322
52. Seaver N. 2017 Algorithms as culture: some tactics for the ethnography of algorithmic systems. *Big Data Soc.* **4**, 2053951717738104.
53. Ziewitz M. 2016 Governing algorithms: myth, mess, and methods. *Sci. Technol. Human Values* **41**, 3–16.
54. Eslami M *et al.* 2015 'I always assumed that I wasn't really that close to [Her]': reasoning about invisible algorithms in news feeds. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pp. 153–162. Seoul, Republic of Korea: ACM.
55. Fletcher R, Nielsen RK. 2018 Generalised scepticism: how people navigate news on social media. *Inf. Commun. Soc.* (doi:10.1080/1369118X.2018.1450887)
56. Bucher T. 2017 The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms. *Inf. Commun. Soc.* **20**, 30–44.
57. Morozov E. 2011 *The net delusion: how not to liberate the world*. London, UK: Penguin.
58. Tufekci Z. 2017 *Twitter and tear gas: the power and fragility of networked protest*. New Haven, CT: Yale University Press.
59. Burrell J. 2016 How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data Soc.* **3**, 1–12.
60. Diakopoulos N, Koliska M. 2016 Algorithmic transparency in the news media. *Dig. Journ.* **5**, 809–828.
61. ter Hoeve M, Heruer M, Odijk D, Schuth A, de Rijke M. 2017 Do news consumers want explanations for personalized news rankings? In *FATREC Workshop on Responsible Recommendation Proceedings, Como, Italy*.
62. Gillespie T. 2014 The relevance of algorithms. In *Media technologies: essays on communication, materiality, and society* (eds T Gillespie, PJ Boczkowski, KA Foot), pp. 167–194. Cambridge, MA: MIT Press.
63. Kitchin R. 2017 Thinking critically about and researching algorithms. *Inf. Commun. Soc.* **20**, 14–29.
64. Pasquale F. 2015 *The black box society: the secret algorithms that control money and information*. Cambridge, MA: Harvard University Press.
65. Sandvig C, Hamilton K, Karahalios K, Langbort C. 2014 Auditing algorithms: research methods for detecting discrimination on internet platforms. In *Data and discrimination: converting critical concerns into productive inquiry, Preconf. 64th Annual Meeting of the International Communication Association 2014, Seattle, WA, 22 May*.
66. Chun WHK. 2011 *Programmed visions: software and memory*. Cambridge, MA: MIT Press.
67. Montfort N *et al.* 2012 *10 PRINT CHR \$(205.5+ RND (1)): GOTO 10*. Cambridge, MA: MIT Press.
68. Parikka J. 2013 *What is media archaeology?* Cambridge, UK: Polity Press.
69. Bucher T. 2012 Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media Soc.* **14**, 1164–1180.
70. Seaver N. 2017 Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data Soc.* (doi:10.1177/2053951717738104)
71. Ziewitz M. 2017 A not quite random walk: Experimenting with the ethnomethods of the algorithm. *Big Data Soc.* (doi:10.1177/2053951717738105)
72. Diakopoulos N. 2016 Accountability in algorithmic decision making. *Commun. ACM* **59**, 56–62.

73. Ananny M, Crawford K. 2016 Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New media Soc.* **20**, 1461444816676645.
74. Hirschman AO. 1970 *Exit, voice, and loyalty: responses to decline in firms, organizations, and states*, vol. 25. Cambridge, MA: Harvard University Press.
75. Robert D. 2000 *On democracy*. New Haven, CT: Yale University Press.
76. David H. 2006 *Models of democracy*, 3rd edn. Stanford, CA: Stanford University Press.
77. Gibbons T. 1998 *Regulating the media*, 2nd edn. London, UK: Sweet & Maxwell.
78. Nguyen A. 2013 Online news audiences: the challenges of web metrics. In *Journalism: new challenges* (eds S Allan, K Fowler-Watt), pp. 146–161. Poole, UK: CJCR: Centre for Journalism & Communication Research, Bournemouth University.
79. Barendt E. 1993 *Broadcasting law: a comparative study*. Oxford, UK: Clarendon Press.
80. Manin B. 1987 On legitimacy and political deliberation. *Polit. Theory* **15**, 338–368.
81. Holznagel B, Jungfleisch C. 2007 The protection of viewer rights in Europe. In *Broadcasters and citizens in Europe: trends in media accountability and viewer participation* (eds P Baldi, U Hasebrink), pp. 53–74. Brighton, UK: Intellect Books.
82. Zamith R. 2018 Quantified audiences in news production. *Dig. Journ.* **6**, 418–435.
83. Helberger N. 2008 From eyeball to media literate viewer—toying with audience empowerment in the audio-visual media service directive. *Entertain. Law Rev.* **6**, 128.
84. Hindman M. 2017 Journalism ethics and digital audience data. In *Remaking the news. Essays on the future of journalism scholarship in the digital age* (eds PJ Boczkowski, CW Anderson), pp. 177–193. Cambridge, MA: MIT Press.
85. Soerensen JK. 2013 Public service broadcasting goes personal: the failure of personalised PSB web pages. *MedieKultur* **55**, 43–71.
86. Belair-Gagnon V, Holton AE. 2018 Boundary work, interloper media, and analytics in newsrooms. *Dig. Journ.* **6**, 492–508. (doi:10.1080/21670811.2018.1445001)
87. Ferrer-Conill R, Tandoc EC. 2018 The audience-oriented editor. *Dig. Journ.* **6**, 436–453. (doi:10.1080/21670811.2018.1440972)
88. Tandoc E, Thomas R. 2015 The ethics of web analytics. Implications of using audience metrics in news construction. *Dig. Journ.* **3**, 243–258.
89. Napoli PM. 2011 *Audience evolution: new technologies and the transformation of media audiences*. New York: NY: Columbia University Press.
90. Hansen. 2015 The Art of Intervenability for Privacy Engineering, Workshop ‘Data Protection, Privacy, and Transparency’ (DPPT’15) Hamburg, 26 May, 2015. See https://www.datenschutzzentrum.de/uploads/vortraege/20150226_ArtOfIntervenability_Hansen_final.pdf.
91. Meis R, Heisel M. 2016 Understanding the Privacy Goal Intervenability. See <https://pdfs.semanticscholar.org/5a86/321bf451a86fb1ea11ac73eee21bd7c7f07e.pdf>
92. Louis T. 2013 ‘Google Kills Reader, Helps RSS’, Forbes 29 June 2013. See <https://www.forbes.com/sites/tristanlouis/2013/06/29/google-kills-reader-helps-rss>.
93. Hamborg F, Meuschke N, Gipp B. 2017 Matrix-based news aggregation: exploring different news perspectives. In *Digital Libraries (JCDL), 2017 ACM/IEEE Joint Conference on*, pp. 1–10. IEEE, or their Github for its open code. See <https://github.com/fhamborg/NewsBirdServer>
94. Makhortykh M, Harambam J. 2018 Decoding the cryptic: How news media communicate personalization practices, in preparation.
95. Gürses S, Van Hoboken J. 2017 Privacy after the agile turn. In *Cambridge handbook of consumer privacy* (eds J Polonetsky, O Tene, E Selinger). Cambridge, UK: Cambridge University Press.
96. Brunton F, Nissenbaum H. 2015 *Obfuscation: a user’s guide for privacy and protest*. Cambridge, MA: MIT Press.
97. Howe DC. 2015 Surveillance countermeasures: expressive privacy via obfuscation. *APRJA, A Peer-Reviewed Journal About Datafied Research* **4**(1). See <https://cs.nyu.edu/trackmenot/>
98. Howe DC. 2015 Surveillance countermeasures: Expressive privacy via obfuscation. *APRJA* **4**(1). See <https://adnauseam.io/>
99. Gürses S, Overdorf R, Balsa E. 2018 POTs: The revolution will not be optimized? arXiv:1806.02711v2 [cs.CY].
100. Harambam J, Aupers S, Houtman D. 2011 Game over? Negotiating modern capitalism in virtual game worlds. *Eur. J. Cult. Stud.* **14**, 299–319.
101. Cooper A. 2004 *The inmates are running the asylum: why high-tech products drive us crazy and how to restore the sanity*, revised edn. Indianapolis, IN: Sams Publishing.

