Selected Papers of AoIR 2016:
The 17th Annual Conference of the
Association of Internet Researchers
Berlin, Germany / 5-8 October 2016

PANEL PROPOSAL: CAUGHT IN A FEEDBACK LOOP? ALGORITHMIC PERSONALIZATION AND DIGITAL TRACES

Wiebke Loosen
Christian Sandvig
Marco T Bastos
Cornelius Puschmann
Uwe Hasebrink
Sascha Hölig
Lisa Merten
Jan--Hinrik Schmidt
Katharina Kinder--Kurlanda
Katrin Weller

Summary

Algorithmically calculated decisions about relevance and news play an increasingly important role in how we perceive the world. This panel introduces new theoretical and methodological approaches that explore algorithmic public spheres and the digital trace data that enables them, focusing on shifts in editorial decision making, user choices in algorithmically driven environments, the epistemological implications involved in researching these topics, and suggesting potential remedies for the normative problems of a contemporary computer-generated view of the world.

The terms “algorithm”, “big data”, “digital traces” are increasingly used as convenient blanket labels to address a range of developments which reshape our understanding of fundamental concepts such as “public”, “relevance”, and “news”. Algorithms working on large amounts of user- and system-generated data construct spheres of public communication, for example, by identifying and connecting users with compatible attributes, interests, and activity patterns (e.g. Beam, 2014), or by filtering content based on algorithmically constructed indicators of relevance (e.g. Eslami et al., 2015). Thus, it is increasingly important to investigate the practices, mechanisms, power structures, and dynamics of such algorithmic public spheres. Approaches to combine the study of the construction and inscribed mechanisms of algorithms with a perspective on their societal consequences already include the idea of “algorithmic accountability“ (Diakopoulos, 2014), “algorithmic ideology“ (Mager, 2012), or “algorithmic harms” (Tufekci, 2015).

Loosen, W., Sandvig, C. Bastos, M.T., Puschmann, C., Hasebrink, U., Hölig, S., Merten, L., Schmidt, J-H., Kinder--Kurlanda, K., Weller, K. (2016, October 5-8). Caught in a feedback loop? Algorithmic personalization and digital traces. Panel at AoIR 2016: The 17th Annual Conference of the Association of Internet Researchers. Berlin, Germany: AoIR. Retrieved from http://spir.aoir.org.

If, as Ananny (2015) argues, algorithms possess the ability to convene people by inferring associations and the power to suggest probable actions, this also makes necessary a reformulation of questions that are at the heart of research on journalism and editorial decision making: How do algorithms define relevance? What are the criteria behind their selection mechanisms and how “objective” are these criteria? There are also novel methodological challenges: How are these mechanisms to be studied if we do not have direct access to the algorithms involved, but can only infer their working from digital traces that are made visible or accessible? And, more generally, how should digital traces be interpreted at scale?

This panel explores the whole spectrum of algorithmic construction principles, editorial decision making, and user choices in digital environments while acknowledging their interrelatedness: Author 1 asks what design intervention is necessary to ensure that algorithmic systems allow later scrutiny by a third party and sees the answer in a “glass-boxing” approach. Author 2 explores the tension between editorial decision making and automatized processes of news filtering, applying the method of the ‘algorithm audit’ to characterize the consequences of recommender systems on news consumption. Author 3 introduces the concept of the “Twitter friend repertoires”, using digital traces to compare the information management of different groups of users in an algorithmically driven environment. Finally, Author 4 discusses the epistemological dimension of researching algorithms which are often based on Big Data, examining challenges posed by the quality, provenance and explanatory power of these new types of data. The panel thus follows algorithms from their origin in design to different application scenarios and finally deconstructs the data sources that underlie them.

PAPER 1: GLASS-BOXING THE ALGORITHMIC PUBLIC SPHERE

Christian Sandvig, University of Michigan

In recent years, media systems of all kinds have been transformed to include automatic selection as a basic part of their operation. News, advertising, the features of a computer's user interface, prices, content recommendations, and more are chosen and even produced dynamically by computer programs (Ziewitz, 2016). At the same time, most of the non-media infrastructure underpinning contemporary US society has become increasingly mediated by the Internet: infrastructure has been upgraded to a "smart" iteration that features embedded computing power, telecommunications links, interactive interfaces, and real-time control (Graham & Marvin, 2001). At the center of these developments sit algorithms (meaning a process or set of rules followed by a computer) that provide functions like relevance judgments, social sorting, credit scoring, matchmaking, market segmentation, personalization, and the management of traffic flows from ads to cars. Making media and infrastructure newly computational has made these systems more powerful, but also much more opaque to public scrutiny and understanding (Pasquale, 2015).

What little has been discovered about the operation of these systems provides cause for alarm. Indeed, the shift to computation may be disastrous for fairness, democracy, and even safety, wholly transforming much of mediated communication. Scholars have argued that along with any benefits, algorithmic systems are also fostering widespread fraud; illegal discrimination by race, gender, disability, and age; false advertising; payola; anticompetitive practices; and political cynicism; while also making society more unequal and degrading the quality of news and information in systematic ways.

Nonetheless, creating different algorithmic systems that involve either changing the motives of multinational corporations or funding, designing, and launching a viable public or private alternative seem like Herculean tasks. But before an alternative future can be proposed, we first need some understanding of what algorithmic systems are actually doing now. Yet even this more modest redress also seems unworkable: algorithmic systems are often proprietary, dynamic, and complex, therefore an "algorithmic literacy" or "transparency" agenda could founder on the basis of impracticality alone (Pasquale, 2011). What is to be done?

Algorithm Audits

Some hope may be found in a relatively obscure social scientific research method: the audit study (Mincy, 1993). Audit studies – also called correspondence studies – are field experiments in which (in one common design) a fictitious correspondence is created purporting to be from a job applicant seeking employment, targeted at a real employer. In these studies, two or more identical resumes are prepared. "The race…of the fictitious applicant is then signaled through one or more cues" such as the fictitious applicant's name, which might be manipulated between the two conditions of "Emily" and "Lakisha" to signal "Caucasian" vs. "African-American" (Pager, 2007). This is then used as a measure of discrimination to determine if identically qualified applicants receive differential treatment based on race.

Instead of a “direct” audit of employers or landlords—the classic targets for this research design—researchers have proposed that the normative concerns that have been raised involving algorithms can be studied via audits of online platforms (Author 1 et al., 2015). They have recommended a program of research to audit important Internet-based intermediaries (e.g., YouTube, Google, Facebook, Zillow, Netflix, LinkedIn, and so on) to identify harmful or worrying algorithm. Crucially, the algorithm audit does not require the cooperation of the platform. Just as a landlord audit study cannot determine why a landlord is racist, an algorithm audit cannot actually reverse-engineer an algorithm. However, it can indicate harmful behavior.

Glass-Boxing the Algorithm

Algorithm Audits have now discovered (for example) that Gmail scrutinizes e-mail text for signs of depression (Lecuyer et al. 2014), women are less likely to be shown ads for high-paying jobs (Datta et al. 2015), and the locations of national borders are personalized on online maps (Soeller et al. 2016). Despite these advances in auditing, attempts to audit algorithms have also encountered significant challenges: anti-hacking laws, expense, complexity, attempts by platforms to avoid auditing, and unclear professional incentives and standards for researchers and auditors.

In this paper, I review the normatively worrying harms discovered by prior work on algorithmic curation, then ask: What design intervention is necessary in each case to ensure that the problem is easily discoverable by a third party? In engineering an alternative to “black-boxing” is “glass-boxing.” To “glass-box” a component, the designer makes decisions at the product’s creation that allow its operation to be better understood later. This glass-boxing approach would acknowledge that the literal reverse engineering of algorithmic systems is likely impossible, and embrace the idea that not everything must be known in order to discover a harmful consequence. While glass-boxing is useful for auditors, it can also have significant benefits on individual users of algorithmic systems (Eslami et al 2016).

I argue that certain inputs and particular internal processes have a normatively privileged position and must be designed at the outset to permit later scrutiny. Glass-boxing presumes that there is some public-interest infrastructure (such as researchers, governments, nonprofit organizations, or user collectives) that can implement later examination of glass-boxed platforms, and repercussions for any trouble discovered (competitive pressure, negative public option, or prosecution). Implementing a glass-boxing regime may be onerous, but I argue that it is the least difficult and most practical solution to the negative features of “calculated publics” (Gillespie, 2012).

PAPER 2: DECENTRALIZED BUT COORDINATED: NEWS READERSHIPS ACROSS ONLINE SOCIAL NETWORKS AND ISSUES OF PERSONALIZATION

Marco T. Bastos, City University London
Cornelius Puschmann, Hans Bredow Institute for Media Research

Editorial decision making is undergoing considerable changes due to the automation of journalism and the continuous feedback of social media audiences (Author, 2012). On the one hand, newsrooms struggle to strike a balance between news that editors understand to be important and news that appeal to their digital readership. On the other hand, aggregated social media data allow newsrooms to incorporate algorithms that can both generate and target news stories to specific audiences, a tension that challenges established notions of journalism as a service to the public (van Dalen, 2012; Bright & Nicholls, 2013). This embattled architecture of digital journalism requires news outlets to develop strategies to catch audience attention and attract readership beyond the limits of their websites (dos Reis et al., 2015), often by incorporating algorithms to filter, republish, and select news items.

The ensuing debate has led journalism scholars to argue that algorithmic personalization of media content limits users' choices by trapping them in filter bubbles, and that the application of algorithmic filtering methods to news has adverse effects on the quality of public debate (Couldry & Turow, 2014, Pariser, 2012). Some scholars suggest that personalization has the potential to be 'corrupt' (Author, 2014), for example when the relation between user-generated content and advertisements is deliberately blurred by a platform provider, a problem that arguably could also impact news. Interestingly, this criticism is noted and discussed not only in media and communication research, but also in fields that actively develop recommender systems, such as computer science and information systems research. For example, Zhou et al note the risk that "more and more users will be exposed to a narrowing band of popular objects, while niche items that might be very relevant will be overlooked" (2010, p. 4511), and suggest measures for countering this effect, and Zeng et al contend that "though they are helpful in filtering information, recommendation algorithms may impose reinforcing influence on the system, by guidance to one's choices which influences subsequent recommendations and hence choices of others" (2012, p. 18005). However, such academic considerations may carry little overall weight in relation to applications driven by economic arguments.

Personalization is not restricted to national media outlets like *The New York Times* (US), *The Guardian* (UK), or *Spiegel Online* (Germany). In fact, tabloid journalism has been subjected to a much faster transition from a context of strong editorial identity to content curation that is both user-generated and created by paid staff members. While quality press can still ring-fence news editors' choices in determining what is in the public interest and should be published, viral news sites have blossomed out of social networking technologies and algorithmic news filtering aimed at a fragmented audience with varying degrees of media sophistication. In contrast to quality press, in which news editors are only beginning to contend with the potential balkanization of readership according to interests of like-minded groups, viral websites explore crowd mechanism and develop controversial stories by meticulously testing potential content on the

website and upholding or removing material according to the number of clicks retrieved from early testers. As Upworthy races for the social media attention, the team breaks down the content into “seeds” and “nuggets”—content to feature on the site and a list of 25 potential headlines that are streamlined for “click testing” (Rohani, 2014). The distribution and potential success of any viral news site depends directly on peer-to-peer engagement on social networking sites, which then feeds into the variables that calibrate the algorithmic decisions that decide which content should appear on top, which items should receive only mild exposure, and which content should not appear at all.

This talk consists of two components that investigate editorial and algorithmic decisions regarding the relevance of journalistic content. In the first part, we examine a set of news articles (n=16,829) retrieved in October 2012 by querying the public APIs of *The New York Times* and *The Guardian* and by tracking the diffusion of each article on the following social media platforms: Twitter, Facebook, Google+, Delicious, Pinterest, and StumbleUpon. The findings show significant differences in topical preferences between news editors and social media users, with social media users favoring opinion pieces, along with national, local, and world news. We discuss these differences in terms of headline composition and sentiment polarity of news headlines that are likely to have boosted the online outreach of news articles. Secondly, we address the feedback effects introduced into news personalization by means of longitudinal platform data from *Spiegel Online*, a major German news website. Drawing on the method of an algorithm audit (Author et al., 2014), we aim to characterize the influence of recommender systems on news consumption on the platform, focusing particularly on how to interpret ambiguous signals, such as user clicks, on individual news items.

The data collected from *The New York Times*, *The Guardian*, and *Spiegel Online* shed light on the changing architecture of journalism and its struggle to reinvent itself in the digital age. Although the strategies employed by news outlets to boost the reach of news items show varying degrees of algorithm adoption, they can be broadly classified in an ordinal scale from human gatekeeping functions towards decentralized, distributed, and eventually automatized processes of news filtering.

PAPER 3: TWITTER FRIEND REPERTOIRES: INFERRING SOURCES OF INFORMATION MANAGEMENT FROM DIGITAL TRACES

Jan--Hinrik Schmidt, Hans Bredow Institute for Media Research
Wiebke Loosen, Hans Bredow Institute for Media Research
Uwe Hasebrink, Hans Bredow Institute for Media Research
Sascha Hölig, Hans Bredow Institute for Media Research
Lisa Merten, Hans Bredow Institute for Media Research

Twitter is a prominent example of a media service in which professional, participatory, and algorithmic modes of filtering and distributing content converge. Most related research has followed a "broadcast communication" paradigm, i.e. looked at practices, networks and consequences of Twitter as a tool to disseminate information. By contrast, few studies analyze Twitter as a tool for information management even though seeking and filtering information is usually acknowledged as an important motivation for Twitter use, especially among journalists (e.g. Broersma & Graham, 2013). In this paper we aim to contribute to an audience-related perspective by adapting the concept of "media repertoires" from audience research to the study of Twitter. We propose the concept of "Twitter friend repertoires", which refers to distinct patterns in the combination of accounts people follow on Twitter.

The concept of "Twitter friend repertoires" shares basic assumptions with the broader idea of "media repertoires" (Author 3 & Domeyer, 2012), namely a user-centered perspective originating in audience research, an interest in the relationality of repertoire components, and an attempt to capture the entirety of the media elements a person chooses to use. Whereas media-centered audience research considers the usage of single media, the starting point for a repertoire-oriented approach is the specific combination of different media as arranged by a single user. As such, Twitter is seen as a platform which contains countless individual "channels" (i.e. the Twitter accounts and their individual timelines) and in itself represents a repertoire notion as Twitter users actively decide which accounts to follow from a vast choice of possibilities. This selection forms the Twitter newsfeed from which s(he) receives her/his information on Twitter. In this way users compose their individual "friend repertoire", and thus, their sources.

Given the diversified structure of Twitter users, a description of friend repertoires that exclusively relies on the level of single accounts would not allow for a clear overview of the underlying structure of this repertoire. Instead these repertoires should be described along broader categories of friends. In order to identify and describe Twitter friend repertoires we have started to develop a database that provides categorical information on certain accounts, e.g. if they belong to a media company, if they are an individual journalist, or if they are one of many different kinds of organizations, resulting in a (continuously updated) collection of the Twitter accounts belonging to publicly relevant speakers in Germany (i.e. media outlets, members of parliament and political parties, major companies, celebrities etc.).

Our methodology involves extracting the Twitter friends of a chosen group of users and checking these against the list (at the time of writing) of about 7,200 previously identified

accounts. This combination of qualitative (database) and data-driven approaches (Twitter data scraping) creates the necessary context to understand patterns of information repertoires. In this way, enriched data makes it possible to calculate individual metrics for each user in the sample, e.g. “share of publicly relevant accounts among all friends” or “number of top politicians vs. celebrities followed by the user”. On an aggregate level, characteristics of the whole sample (e.g. “average share of publicly relevant accounts”) as well as groups based on shared patterns of Twitter repertoires (e.g. contrasting users who follow a disproportionate number of media accounts vs. users who follow mainly personal accounts) can be identified.

This approach allows us to move beyond a focus on Twitter’s “structural layers of communication” (Bruns & Moe, 2014) that characterizes the majority of Twitter research by focusing instead on hashtag-driven issue publics or interpersonal communication via replies. To achieve a more holistic approach, we enrich Twitter data with data that include relevant categories to better understand what kind of sources users and certain user groups follow and combine within their information repertoires. Moreover, the identification and comparison of certain user groups (e.g. journalists vs. politicians) and their Twitter friend repertoires could, to some extent, serve as a “substitute variable” for the independent variables that are often absent but are needed to describe differences in Twitter use and shared patterns and practices of information management more closely.

In order to demonstrate the insights produced by our approach we will present a comparison of the Twitter friend repertoires of five different groups of users, representing different contexts as well as different degrees of professionalism of use:

1. candidates for a regional parliament,
2. audience members of a public service TV newscast,
3. audience members of a national newspaper,
4. members of the German 2014 World Cup Football Team, and
5. political journalists.

The first four of these groups have been analyzed in a previous study with data collected at the end of 2014 (Author, 2015; Author, in review) which demonstrated distinct patterns in users’ Twitter friend repertoires. For example, political candidates and football players both exhibit a very pronounced in-group-orientation, i.e. their Twitter friend repertoires show a marked preference for accounts from other candidates and top politicians or, in the case of football players, for other professional players. The repertoires of the two groups of media audience, in contrast, exhibit preferences for other media-related accounts, with the TV news audience on Twitter tending to follow other broadcast-related accounts, while the newspaper audience tends toward other print-related accounts. Further data will be collected in summer 2016 to provide a more up to date and comprehensive dataset of Twitter friends for the five groups of users for the resulting article and its presentation at the conference. We regard these analyses as only the first steps towards theory-driven repertoire-oriented Twitter research: An important step towards a better understanding of these repertoires would be to identify the Twitter friend repertoires for a random sample of all (German) Twitter accounts which could serve as a benchmark to compare other samples against. Future research could also provide a more thorough analysis of individual patterns which could lead to

typologies of Twitter users according to their Twitter friend repertoires; with the changes of Twitter friend repertoires over time; and with the relation between Twitter friend repertoires and actual tweeting behavior. Yet another line of research might complement such quantitative approaches with more qualitative instruments, for example, by interviewing Twitter users about their practices regarding Twitter as a tool for information management.

PAPER 4: ANALYSING DIGITAL TRACES: THE EPISTEMOLOGICAL DIMENSION OF ALGORITHMS AND BIG DATA

Katharina Kinder--Kurlanda, GESIS Leibniz Institute for the Social Sciences
Katrin Weller, GESIS Leibniz Institute for the Social Sciences

Algorithms that sort and rank information are starting to change public spheres, e.g. by filtering news or by individualizing content. Often these algorithms' work is based on various types of 'big data'. Online big data can be user-generated content as well as activity, trace or location data collected from various internet platforms. Algorithms thus observe, count and sort highly heterogeneous and transient data. Big data promises to allow new insights into social and political processes, making it possible to analyze events in situ, at an unprecedented level of detail, and very close to actual everyday situations. Big data, however, also poses various challenges in analysis due to its size and heterogeneity. We are therefore also witnessing a renewed discussion of methods and epistemology, particularly within the social sciences (e.g. Kitchin, 2014; Frické, 2014). Topics being discussed are, for example, whether opinions voiced online are representative, or how the explanatory power of datasets can be assessed if data access is restricted by platform providers (boyd & Crawford, 2012). While researchers are developing new analytic approaches, big data's quality, provenance and explanatory power are decidedly uncertain. Researchers are only starting to begin to understand the complexities of big data and how they may be taken into account in analyses in meaningful ways. There is therefore a high degree of epistemological uncertainty when working with big data. The algorithms that shape public spheres are therefore operating on a tenuous and often uncertain basis that researchers are only beginning to understand and that can be opaque to users. In addition, algorithms rely on dynamic content that may change exactly due to internet users (as content generators) being presented with the results of the algorithmic calculations. Based on a qualitative study of social media researchers this paper addresses challenges posed by the uncertainty and heterogeneousness of big data for the study of algorithms that are shaping public spheres. Dependencies of algorithms and big data in creating epistemological challenges are explored.

We conducted observations and 40 semi-structured interviews with researchers at several international conferences. Interviewed researchers were working with data gathered directly from social media platforms such as Facebook, Twitter, Reddit or YouTube. We explored researchers' reflections of everyday research practices when engaged in data collection, analysis, and publication. We aimed to capture social media researchers' a) motivations for working with social media data, b) their methodological approaches, and c) the perceived challenges, pitfalls and drawbacks. We explored various practical aspects around 'dealing with data' in social media research which enabled us to capture researchers' everyday problem-solving.

Our study saw researchers finding individual ways to address epistemological challenges in order to realize social media big data's potential for researching social processes on and off the web. (Author, 2015). In particular, interdisciplinary collaborations were seen to allow a pooling of skills and resources but then again posed new methodological and epistemological challenges (Author, 2014). Here we are

focusing on three areas of epistemological concern in big data research: 1.) accessibility and quality of these data, 2.) new methodologies, and 3.) research ethics.

First, we found that social media researchers are currently to a considerable degree concerned with challenges of data access. Challenges are posed by technical restrictions applied to proprietary social media data. These restrictions force researchers to iteratively adjust research questions on the basis of data availability, often leading to inductive approaches to be applied regardless of the initial research question or epistemological viewpoint. Researchers also need to possess (or be able to make use of others') technical skills. An example is the ability to write Python scripts for gathering and cleaning data in order to interact with APIs through which platform providers such as Twitter make data available to researchers. So although researchers may be generally aware of the need for strategies to improve research quality and validity, they often do not yet find themselves in a position to implement such strategies as they are consumed with the struggle for data access. Our study thus revealed a considerable impact of external constraints on researchers' attempts to achieve validity and better research quality.

Second, big data methodologies often defied traditional ways of achieving research validity, for example, through reproducibility. To allow the analysis of large amounts of data, approaches are being found that allow to find patterns in large datasets, often by automatically analyzing content (e.g. filtering for frequencies of predefined keywords) or by quantifying activity patterns (e.g. counting frequencies of interactions between users or between users and platform features such as 'likes' on Facebook). These approaches require data to be cleaned and processed in order to allow for analysis to take place. Spam may need to be filtered out, mistakenly included posts or information may need to be deleted, and other mistakes 'cleaned' in order to allow processing. There are currently no agreed upon rules or standards for how to make transparent cleaning decisions in order to ensure that research counts as sound and valid.

Third, from a research ethics perspective, researchers were seen to be challenged by the lack of users' informed consent to studies based on online content (Author, 2014) and by the fact that they are dealing with data that is often difficult if not impossible to anonymize (Zimmer, 2010). While formally consent may have been given by users agreeing to a platform's terms of service, users may not be in fact be aware of or even want to be part of research projects (Hutton & Henderson, 2015). As more and more 'embedded researchers' are starting to cross the lines between academic and non-academic data science, public negotiations are taking place of what researchers and platform providers should and should not be allowed to do with user data and how far they are allowed to infringe upon spaces such as Facebook which users may perceive as a personal and private space.

The exploration of these three areas not only enables a description of the shifting epistemological foundations of big data, but also allows critically questioning the algorithms that rely on these data in order to construct public spheres. Given the uncertainties around big data, how reliable are the results of algorithmic calculations? The observation of algorithms also poses challenges in itself as algorithms introduce an element of recursiveness to public spheres. Various actors in the most diverse situation

react out of different interests to being presented with the results of algorithmic calculations. They may even try to influence algorithms in specific ways. The results of algorithmic analyses are thus presented to Internet users and then influence user behavior. As traces of user behavior are exactly what algorithmic analyses are based on, recursive feedback-loops of actions and reactions to algorithms emerge. How to address the challenges posed not only by the changeability and heterogeneity of big data but also by the interaction between big data and algorithms to future research of user behaviors remains an open challenge.

References

Ananny, M. (2015). Toward an Ethics of Algorithms: Convening, Observation, Probability, and Timeliness. *Science, Technology & Human Values*, 1–25. doi: 10.1177/0162243915606523

Author (2014). Corrupt personalization. *Social Media Collective Research Blog*. Retrieved from http://socialmediacollective.org/2014/06/26/corrupt-personalization/

Author, Hamilton, K., Karahalios, K., & Langbort, C. (2014). "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms." Paper presented to "Data and Discrimination," a pre-conference of the 64th annual meeting of the International Communication Association, Seattle, WA, USA

Author (in review). Twitter Friend Repertoires: Introducing a methodology to assess patterns of information management on Twitter.

Author (2015, October). *Twitter Friend Repertoires - Researching patterns of selective connectivity.* Paper presented at The Empiricist's Challenge: Asking Meaningful Questions in Political Science in the Age of Big Data, Mannheim, Germany.

Author, & Domeyer, H. (2012). Media repertoires as patterns of behaviour and as meaningful practices: A multimethod approach to media use in converging media environments. Participations. *Journal of Audience & Reception Studies*, 9 (2), 757–779

Author (2012). (Re-)Discovering the audience: The relationship between journalism and audience in networked digital media. *Information, Communication & Society*. doi:10.1080/1369118X.2012.665467

Author (2014). I always feel it must be great to be a hacker! The role of interdisciplinary work in social media research. In: WebSci'14. ACM Web Science Conference. Proceedings of the 2014 ACM conference on Web science, 91-98. New York: ACM.

Author (2014). 'I love thinking about ethics!' Perspectives on ethics in social media research. In: Selected Papers of Internet Research (SPIR). Proceedings of ir15 - Boundaries and Intersections.

Author (2015). Uncovering the Challenges in Collection, Sharing and Documentation: The Hidden Data of Social Media Research? In *Proceedings of the Ninth International AAAI Conference on Weblogs and Social Media (ICWSM)*, Standards and Practices Workshop.

Beam, M.I A. (2014). Automating the news. How personalized news recommender system design choices impact news reception. *Communication Research,* 41(8), 1019–1041. doi: 10.1177/0093650213497979

boyd, D.,& Crawford, K. (2012). Critical Questions for Big Data. *Information, Communication & Society,* 15(5), 662–679. doi:10.1080/1369118X.2012.678878

Bright, J., & Nicholls, T. (2013). The Life and Death of Political News: Measuring the Impact of the Audience Agenda Using Online Data. *Social Science Computer Review*, *32*(2), 170–181. doi: 10.1177/0894439313506845

Broersma. M., & Graham, T. (2013). Twitter as a news source. How Dutch and British newspapers used tweets in their news coverage, 2007–2011. *Journalism Practice*, 7(4), 446–464. doi: 10.1080/17512786.2013.802481

Bruns, A., & Moe, H. (2014). Structural layers of communication on Twitter. In: Weller, K., Bruns, A., Burgess, J., Mahrt, M. & Puschmann, C. (eds.) Twitter and Society. Peter Lang, New York, pp. 15-28.

Couldry, N., & Turow, J. (2014). Advertising, big data, and the clearance of the public realm: Marketers' new approaches to the content subsidy. *International Journal of Communication*, *8*, 1710–1726.

Datta, A., Tschantz, M. C., & Datta, A. (2015). "Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination." *Proc*. *Privacy Enhancing Technologies* (PoPETs).

Diakopoulos, N. (2014). Algorithmic accountability reporting. On the investigation of black boxes. Tow Center for Digital Journalism. Retrieved from http://towcenter.org/wp-content/uploads/2014/02/78524_Tow-Center-Report-WEB-1.pdf.

dos Reis, J. C., Benevenuto, F., Olmo, P., Prates, R. O., Kwak, H., & An, J. (2015). *Breaking the News: First Impressions Matter on Online News.* CoRR.

Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., & Author 1. (2015). I always assumed that I wasn't really that close to [her]: Reasoning about invisible algorithms in the news feed. *Proceedings of the 33rd Annual SIGCHI Conference on Human Factors in Computing Systems, Association for Computing Machinery (ACM):*153–162.

Eslami, M., Karahalios, K., Author 1, Vaccaro, K., Rickman, A., Hamilton, K., & Kirlik, A. (2016). First I "like" it, then I hide it: Folk Theories of Social Feeds. *Proceedings of the 34rd Annual SIGCHI Conference on Human Factors in Computing Systems, Association for Computing Machinery (ACM)*: page numbers TBD

Frické, M. (2014). Big Data and Its Epistemology. *Journal of the Association for Information Science and Technology*, 66 (4), 651–661.

Gillespie, T. (2012). Can an algorithm be wrong? *Limn*, 2. Retrieved from http://escholarship.org/uc/item/0jk9k4hj

Graham, S., & Marvin, S. (2001). *Splintering Urbanism: Networked Infrastructures, Technological Mobilities and the Urban Condition*. New York: Routledge.

Hutton, L., & Henderson, T. (2015). "I didn't sign up for this!": Informed consent in social network research. In *Proceedings of the Ninth International AAAI Conference on Weblogs and Social Media* (ICWSM), 178–187.

Kitchin, R. (2014). Big data, new epistemologies and paradigm shifts. *Big Data & Society,* April-June 2014: 1–12.

Kulshrestha, J., Zafar, M. B., Espin Noboa, L., Gummadi, K. P. & Ghosh, S. (2015, May). *Characterizing information diets of social media users.* Paper presented at the 9th international AAAI conference of web and social media. (ICWSM 15). Oxford: AAAI

Lecuyer, M., Ducoffe, G., Lan, F., Papancea, A., Petsios, T., Spahn, R., Chaintreau, A., & Geambasu, R. (2014). XRay: Increasing the Web's Transparency with Differential Correlation. *Proc. USENIX Security*.

Mager, A. (2012). Algorithmic ideology: how capitalist society shapes search engines. *Information, Communication and Society*, 15 (5), 769–787. doi:10.1080/1369118X.2012.676056

Mincy, R. (1993). The Urban Institute Audit Studies: Their Research and Policy Context. In Fix and Struyk, eds., *Clear and Convincing Evidence: Measurement of Discrimination in America.* Washington, DC: The Urban Institute Press, 165–86.

Pager, D. (2007). The Use of Field Experiments for Studies of Employment Discrimination: Contributions, Critiques, and Directions for the Future. T*he Annals of the American Academy of Political and Social Science, 609*(1), 104–33.

Parmelee, J. H. (2014). The agenda-building function of political tweets. New *Media & Society,* 16 (3), 434–450. doi: 10.1177/1461444813487955

Pariser, E. (2012). *The filter bubble: How the new personalized web is changing what we read and how we think*. New York: Penguin.

Pasquale, F. (2011). Restoring Transparency to Automated Authority. *Journal on Telecommunications and High Technology Law, 9*(235).

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Cambridge, MA: Harvard University Press.

Rohani, A. (2014). Upworthy And The Race For Your Attention. Retrieved from https://urbantimes.co/2014/04/upworthy-attention-news/

Soeller, G., Karahalios, K., Author 1, & Wilson, C. (2016). MapWatch: Detecting and Monitoring International Border Personalization on Online Maps. *Proceedings of the 25th Annual World Wide Web (WWW) Conference, Association for Computing Machinery (ACM)*: page numbers TBD.

Tufekci, Z. (2015). Algorithmic harms beyond Facebook and Google: Emergent challenges of computational agency. *Colorado Technology Law Journal*, 13(2), 203–217.

van Dalen, A. (2012). The algorithms behind the headlines: How machine-written news redefines the core skills of human journalists. *Journalism Practice*, 6(5-6), 648–658. doi: 10.1080/17512786.2012.667268

Zeng, A., Yeung, C. H., Shang, M., & Zhang, Y. (2012). The reinforcing influence of recommendations on global diversification. *EPL (Europhysics Letters)*, *97*, 18005–18011. doi:10.1209/0295-5075/97/18005

Zhou, T., Kuscsik, Z., Liu, J.-G., Medo, M., Wakeling, J. R., & Zhang, Y.-C. (2010). Solving the apparent diversity-accuracy dilemma of recommender systems. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(10), 4511–4515. doi:10.1073/pnas.1000488107

Ziewitz, M. (2016). Governing Algorithms: Myth, Mess, and Methods. *Science, Technology & Human Values 4(1), 3–16*.

Zimmer, M. (2010). But the data is already public: on the ethics of research in Facebook. *Ethics and Information Technology* 12(4),313–325.