

## Computational News Discovery: Towards Design Considerations for Editorial Orientation Algorithms in Journalism

Nicholas Diakopoulos

**To cite this article:** Nicholas Diakopoulos (2020): Computational News Discovery: Towards Design Considerations for Editorial Orientation Algorithms in Journalism, *Digital Journalism*, DOI: [10.1080/21670811.2020.1736946](https://doi.org/10.1080/21670811.2020.1736946)

**To link to this article:** <https://doi.org/10.1080/21670811.2020.1736946>



Published online: 09 Mar 2020.



Submit your article to this journal 



View related articles 



View Crossmark data 



# Computational News Discovery: Towards Design Considerations for Editorial Orientation Algorithms in Journalism

Nicholas Diakopoulos 

Department of Communication Studies, Northwestern University, Evanston, IL, United States

## ABSTRACT

Computational news discovery (CND) is a particular application area within computational journalism related to the use of algorithms to orient editorial attention to potentially newsworthy events or information prior to publication. Previous work in this area has been concentrated on prototyping CND tools, which can, for instance, send alerts and leads to journalists about social media events, documents of interest, or salient patterns in streams of data. This article describes a qualitative interview study of journalists as they incorporate CND tools into their practices. Findings provide insights into how CND tools interact with the internal attention economy and sociotechnical gatekeeping processes of the newsroom and how future CND tools might better align with necessary journalistic evaluations of newsworthiness and quality, while ensuring configurability, human agency, and flexible applicability to a wide range of use cases. These findings begin to outline a conceptual framework that can help guide the effective design of future CND tools.

## KEYWORDS

Algorithmic journalism; automated journalism; computational journalism; computational news discovery; data journalism; news production

## Introduction

There is perhaps no aspect of the news production pipeline that isn't increasingly impacted by the use of algorithms. Computational approaches are now broadly applied in journalistic work including in information gathering (Thurman et al. 2016), providing signals to assess the veracity of content or sources (Fletcher, Schifferes, Thurman 2017), automatically generating written articles (Graefe 2016), creating new interactive bot experiences (Lokot and Diakopoulos 2016; Ford and Hutchinson 2019; Jones and Jones 2019), and optimizing or otherwise influencing the distribution of content on homepages, apps, or platforms (Bucher 2016). Computational journalism considers how computing—defined as “the systematic study of algorithmic processes that describe and transform information” (Denning 2005)—is applied to support journalistic tasks and embrace journalistic values (Diakopoulos 2019a). Here “computational journalism” is emphasized rather than related terms such as “algorithmic journalism”, “automated journalism”, or “data journalism” (Thurman 2019;

Zamith 2019; Coddington 2015) as it broadly captures the idea of using algorithms to transform information and data for journalistic purposes. At the same time, computational approaches often entail the close collaboration of algorithms and people in hybrid systems that take advantage of the capabilities of algorithms for scale and speed, but also leverage the complementary capabilities of people (Brynjolfsson and McAfee 2014; Diakopoulos, 2019a). In particular this research examines this hybridization of algorithmic and human effort in the context of a specific application of computational journalism in news production: *computational news discovery*.

News discovery is described in some of the earliest ideation documents about computational journalism. Hamilton and Turner (2009) posit that, “a reporter could be alerted when a trend appears, an anomaly arises, or when a specific individual or entity or location is referred to in the data stream”. The premise is that with growing volumes of information, computing can offer a subsidy to public interest journalism, lowering costs and increasing the efficiency and scale at which new news stories can be identified (Hamilton 2016). The intervening years have brought forth a variety of system prototypes that have reified the use of computing for discovering news, which has recently been referred to as “computational news discovery” in a review of research on computational journalism (Thurman 2019). Monitoring systems can send alerts and craft leads that orient journalists’ attention to social media events, documents of interest, or anomalous patterns in streams of data that may reveal important news stories. For instance, in computational fact spotting scenarios algorithms sift through claims to identify those that can be checked (Graves 2018). In one particular implementation, a machine learned classifier is used to identify fact checkable statements in CNN transcripts, which are then sent to newsrooms as daily tipsheets that fact checkers may refer to in making coverage decisions (Hassan et al 2017; Adair et al 2019). With this prior work in mind, here I define computational news discovery (CND) as: *the use of algorithms to orient editorial attention to potentially newsworthy events or information prior to publication*.

Given the recent technical developments and demonstrations of CND systems, some research has begun to investigate how they are integrated into journalism practice, are applicable (or not) to different journalistic use-cases, and could be designed to enable more effective journalistic use (Diakopoulos 2019; Stray 2019). In furthering the technological lens to the study of computational journalism with a focus on the hybrid nature of newsroom technologies (Anderson 2013), here I adopt a human-centered and sociotechnical frame which considers how CND systems interact with and influence the experiences of journalists in undertaking their newswork. The focus is on individual perspectives on usage, including utility as well as social and normative expectations, rather than any wider organizational factors that may, of course, also play a role in adoption. In particular this work sets out to ask: What are the human-centered needs of journalists with respect to the effective use of computational news discovery systems in their sociotechnical gatekeeping practices of selecting and developing news items for publication? This research addresses this question by undertaking interviews with 18 stakeholders who have created, used, or both created and used computational news discovery systems in the context of journalism.

The results of this study offer insight into the experiences of practitioners with respect to computational news discovery systems. Based on these results this article contributes an initial elaboration of a conceptual design framework that can help inform and guide the future development of computational news discovery systems that are consistent with user needs and which begins to unpack the use of algorithms as part of sociotechnical pre-publication gatekeeping processes. In particular, the findings highlight the essential role of human effort and attention in developing leads into news items of publishable quality, elaborate factors related to the newsroom attention economy and how these tools fit within it, and suggest design opportunities for CND systems to better align with the various newsworthiness and quality evaluations journalists need to make across a wide range of journalistic scenarios.

## Related work

Here I consider two areas of related work which inform the conceptual approach taken in this research: (1) computational news discovery as an application area of computational journalism which speaks to the use of algorithms in sociotechnical gatekeeping processes, and (2) design-oriented and human-centered approaches to studying journalistic work enabled by computational tools that have the goal of identifying design implications.

### *Gatekeeping and computational news discovery*

The concept of gatekeeping captures the idea that information can be variously impeded or passed onward in the process of communication. Not all news information is published and made widely available. There is a matrix of forces at play which impact gatekeeping decisions, including individual cognitive differences or biases, work routines for news production, organizational characteristics, external social institutional actors such as advertisers or governments, and social systems such as culture or ideology (Shoemaker and Vos 2009). While some of the earliest work on gatekeeping focused heavily on the role of the individual (White 1950), contemporary models of gatekeeping consider not only human actors such as journalists, strategic professionals, and individual amateurs, but also technical actants such as algorithms and their role and interplay with other actors as part of broader sociotechnical gatekeeping practices (Wallace 2017; Thorson and Wells 2016; Lewis and Westlund 2015).

Oftentimes when algorithms are considered in gatekeeping processes the focus is on their role in distributing news to the public via feeds (DeVito 2017), aggregators (Nechushtai and Lewis 2019) and apps (Bandy and Diakopoulos 2020), including by examining how editorial values are embedded into the code of such curators (Weber and Kosterich 2018). A less closely studied aspect of algorithmic gatekeeping relates to the role that algorithms can play not only in distribution but also in news production (Heinderyckx and Vos 2016). In other words, algorithms used in the input and throughput stages of gatekeeping (Wallace 2017) to inform a sociotechnical process prior to wider publication. While some recent work has considered how algorithms used at these stages can introduce biases into the information suggested to journalists

in a sociotechnical gatekeeping process (Thurman et al 2016; Diakopoulos 2019), the current work aims to contribute more broadly to understanding how internally used algorithmic curators (i.e. CND systems) interact with human gatekeepers (i.e. journalists) in a sociotechnical news selection process.

CND systems can contribute to gatekeeping by allowing users to monitor the vast and overwhelming scale of content produced and published on social media platforms. They can help to detect newsworthy events, aggregate responses, and identify, track, and suggest useful sources and witnesses during breaking news or other types of scheduled events like speeches (Diakopoulos, DeChoudhury, and Naaman 2012). For instance, The City Beat tool was developed to detect and alert journalists to local events in New York City, and was deployed to several newsrooms on a trial basis (Schwartz et al 2015). More recently, the Tracer system was developed to monitor millions of daily tweets, cluster posts in order to detect events, and present those events to journalists in a sortable and searchable interface that has proven itself able to accelerate Reuters' news alerts in many cases (Liu et al. 2017; Nourbakhsh et al. 2017).

CND systems can also help monitor data sources such as numeric data streams or textual documents and identify items of interest to be brought to the attention of journalists. For instance, the BBC's Data Stringer prototype was developed to monitor data streams and trigger alerts when rules relating to trends or outliers were matched (Shearer, Simon, and Geiger 2014). The Marple system used statistical methods in order to send alerts to local journalists about anomalies, outliers, or trends in municipal data sets (Magnusson, Finnas, and Wallentin 2016). The Local News Engine scans data from courts, housing developments, and business licenses to detect the names of newsworthy people, places, or companies, which are sent to local media (Perrin 2017). The Tadam system ingests a variety of different documents from the Web, press releases, or document dumps in order to send alerts to reporters when a document shows up that matches their preset filters (Plattner, Orel, and Steiner 2016). The Washington Post's Lead Locator system mines a national voter file dataset in order to help national politics reporters identify interesting locations for their reporting based on demographic patterns and political relevance (Diakopoulos et al 2020). Interactive data-driven expert systems have been deployed to help journalists identify what might be meaningful and newsworthy patterns warranting story development (Broussard 2015). News discovery approaches have also been applied in fact checking workflows to help spot fact checkable claims mentioned in the media (Graves 2018; Hassan et al 2017; Adair et al 2019).

Journalism practice is also utilizing approaches from machine learning and data mining to expand the scope of individual investigations by filtering for known patterns that orient attention to entities or documents that suggest new lines of inquiry (Diakopoulos 2019a; Stray 2019). The current research considers both these bespoke approaches as well as more formalized systems (i.e. CND approaches and systems), with an emphasis on discovery from non-social media sources. These various approaches and tools represent a new journalistic source, offering access to information at the input stage of the gatekeeping process (Wallace 2017) by utilizing algorithms to orient editorial attention to items that journalists might not be aware of otherwise. In particular, this research seeks to understand the sociotechnical

gatekeeping processes around CND approaches and systems in order to inform the design of future CND tools.

### ***Design-oriented studies of newwork***

Human-centered design can be used to help develop new journalistic products, services, and experiences that de-center the role of technological affordances and instead focus on how to harness technology to meet user needs (Chaplin 2016). More broadly, design methods can facilitate a deeper understanding of a sociotechnical context and offer insights that can guide the creation of new technologies that support and align with journalistic goals. Recently, there have been calls for journalism studies to engage more deeply with human-computer interaction (HCI) research in order to better understand how journalists interact with algorithms and automation in news production and how journalistic values can come to be embedded into technologies (Aitamurto et al. 2019; Diakopoulos 2019). These in turn build on earlier observations of the need for scholarly attention towards the “imagined values and engineering design ... of journalistic artifacts” (Anderson 2013), and the importance of how technological actants are “inscribed and instructed by humans” (Lewis and Westlund 2015) so that such inscription is deliberate with respect to the goals of stakeholders and the intended contexts of use. This work addresses this conceptual space by examining the uses and practices surrounding current CND tools with an eye towards how the next generation of such tools could be designed to more effectively support journalistic contexts, work, and values. In particular, our findings are oriented towards identifying design implications that articulate various constraints, affordances, and social or normative expectations of the possible design space for future CND tools (Sas et al. 2014). A design orientation aligns with the conceptualization of digital journalism studies conveyed by Eldridge et al (2019) by enabling empirically informed opportunities for deliberate sociotechnical reconfiguration of journalistic practices.

A growing corpus of research looks at software or tools to support journalistic activity with an eye towards gaining insights that inform future designs and elaborate journalistic concepts. Recent studies in this vein have examined the uses and limitations of automated writing software (Thurman, Dörr, and Kunert 2017), and the utility of social surveillance tools (Thurman 2017) which suggest opportunities for how future computational tools could better support journalistic work. Another line of research has studied visual analytic tools for investigative journalists, revealing different news discovery use-cases depending on whether users already had a hypothesis to verify (Brehmer, Ingram, and Stray 2014) or were more interested in hypothesis generation (Felix et al. 2015). Systems have also been built to help sort social media information such as Tweets or online comments, and user evaluations have underscored the desire of journalists to interactively configure information filters according to domain or scenario specific criteria, to receive adequate context to evaluate or verify information, and to use such tools as a starting point to identify angles of interest for subsequent inquiry (Park et al 2016; Diakopoulos, DeChoudhury, and Naaman 2012). Specific design requirements for software to support the use of user-generated content in newwork have also been developed, including requirements for flagging or marking

leads for later, reflecting update frequencies, and supporting the ongoing nature of verification work (Tolmie et al 2017). The INJECT system was designed to support creativity amongst journalists by providing cues about people, background, or consequences of a news story that could trigger ideas for novel stories (Maiden et al 2018). Taken together, these human-centered and design-oriented studies begin to suggest a rich and deeply contextual design space for CND systems, which the current research seeks to empirically elaborate and refine.

## Study methods

In order to better understand the use of CND systems and approaches from a user-centered perspective a qualitative study consisting of 18 semi-structured interviews was undertaken. The goal of the interviews was to elicit the perspectives and experiences of practitioners, with a particular eye towards how CND systems and approaches might be designed to more effectively integrate into journalism practice.

## Participants

Participants for this study were selected using a purposive sampling strategy in order to deliberately reflect a range of systems and perspectives on those systems. Various editorial orientation systems and projects were identified by reviewing the literature as well as online articles and blogs. Systems were then selected to straddle both data-driven tools and products as well as one-off projects that utilized algorithms for directing editorial attention for the purposes of story finding, alerting, or document investigation. Specific systems studied include RADAR,<sup>1</sup> Newsworthy,<sup>2</sup> Klaxon,<sup>3</sup> several computational fact spotting tools including those of the Tech & Check Cooperative (Adair et al 2019), FullFact,<sup>4</sup> and Chequeado,<sup>5</sup> as well as a range of internal projects from outlets ranging from regional and national newspapers, to online news publications, and interest-specific digital outlets. Individuals associated with the identified systems were contacted via email for interview. Snowball sampling was then used to expand the sample by asking initial interviewees for referrals to other relevant potential participants. These recruiting processes yielded 18 individuals that were interviewed.

Most of the participants were initially interviewed for the purposes of the author's journalistic endeavor to report on the media industry's use of automation and algorithms in news production. Post hoc IRB approval was obtained to re-analyze the data collected in these interviews for the purposes of this research. Light disguise is used in the reporting of results in order to protect the privacy and confidentiality of interviewees (Bruckman 2006). This means that, while precautions have been taken to preserve anonymity and to not identify participants directly, active members of the communities where such projects are discussed may be able to guess identities. Any sensitive details are therefore omitted, as are the identities of the organizations where internal projects were selected for study since participants might be more easily re-identified if those outlets were named. Participants include individuals with a diverse range of perspectives on computational news discovery systems such as creators

( $N=4$ , labeled C1 ... C4), users ( $N=7$ ; U1 ... U7), or both creators and users ( $N=7$ ; CU1 ... CU7).

### ***Interview materials and procedure***

An interview guide was developed iteratively as interviews were undertaken and analyzed, allowing for theoretical sampling of concepts as they emerged from the data (Glaser and Strauss 2009). In its final form the guide included 20 questions including follow-up prompts, addressing topics typically covered in the following ordering: how the system's leads are used in newswork, the overall utility of the system's leads, the information interface of the system including how leads are presented, the newsworthiness of the leads provided, the volume of leads received and developed into news reports including time spent on leads, how the system fits into workflow, whether the leads were trustworthy, and any ethical reflections or other thoughts on the wider use of such tools in journalism (See Appendix A for more details). The semi-structured interview procedure allowed some latitude to focus each interview and additionally tailor questions and prompts as well as their ordering according to an interviewee's particular expertise, experience, and perspective.

Interviews were conducted over a two-year span, from early 2017 to late 2018. The median interview lasted 50.5 min (min = 27; max = 64) and was conducted via audio connection (e.g. phone or Skype) in English. All interviews were audio recorded with consent and were later transcribed. No monetary incentive was provided to participants.

### ***Analysis***

Interviews were fully transcribed and then analyzed using an iterative qualitative method involving open coding of key excerpts, constant comparison, typologizing, and memoing (Glaser and Strauss 2009; Lofland and Lofland 1994). This process was ongoing as interviews were undertaken, which helped inform follow-up questions in latter interviews based on the analytic results from earlier interviews. Analysis of interview materials was further augmented and grounded using document analysis of related materials from the various systems studied, including any extant blog posts, product descriptions, and video presentations where a system's functionality or design were discussed.

### ***Findings***

From the iterative analysis of the transcribed interviews several factors emerged with respect to the use of CND systems and approaches. Chief amongst these factors is the role that humans must still play in developing and evaluating the leads produced by these systems. Several sociotechnical factors that moderate the attention environment with respect to CND systems are elaborated, including the willingness of reporters to pursue leads in different contexts, external factors related to news cycles, the scope of monitoring offered by tools, and the user interface used to convey leads. The findings

further expand on how systems can serve to modulate important editorial evaluations related to newsworthiness and quality assessment that journalists undertake in their lead development work.

### ***The human role: evaluation***

Participants consistently reinforced the idea that people should be involved in *evaluating* the leads produced by CND systems. This held for the simplest of leads, such as alerts signaling a change to a web page, and for more complex leads where an algorithm might itself embed evaluative criteria to rank or draw attention to a subset of information more likely to pan out. While an algorithm permits a scale of monitoring that would otherwise be impossible, having people evaluate leads imbues the overall system with a degree of flexibility to suit different use-cases. Despite the varying degrees of evaluation an algorithm might itself encode, human evaluation of leads was seen as essential in at least three areas: (1) *lead development*, (2) *newsworthiness assessment*, and (3) *quality assurance*. These are briefly described next and further elaborated in subsequent subsections.

Additional editorial effort was often seen as needed in order to develop leads towards publication. In some cases, such as for the leads produced by RADAR, the leads could be published as-is, largely because they had already undergone substantial human editorial development before being distributed as leads. Publishers were able to directly excerpt and use snippets of the text from some of the Newseworthy leads as well. In such cases the editorial effort is mainly that of curation. But for most CND tools studied there was more substantive editorial attention needed, typically involving additional reporting to gather and assess related information as the lead became publication-worthy. These contrasts highlight the spectrum of human effort and attention that might be invested in leads.

Newsworthiness assessment was another important role that people were seen to play. Participants articulated the entire gamut of news values in their evaluation of leads, but the role of people was seen as particularly essential in evaluating newsworthiness dimensions such as *audience fit* and *actuality* (i.e. relevance to the current moment). For instance, FullFact intentionally built their claim spotting system to separate “checkability” (which they thought an algorithm was suited to recognize) from “checkworthiness” (which they thought should be left for a person to evaluate). “We thought that importance was actually something that is an editorial decision and will change over time. Claims, for example, about the EU ... two years ago wouldn’t have been as important as they are now” (C3). By deferring checkworthiness judgements to people the goal was to keep the sociotechnical system more responsive and flexible to a dynamic world where the importance of statements might change over time.

The third area of human evaluation relates to the desire to assure the journalistic quality of leads that are developed and eventually published. Several journalists acknowledged that CND leads could only be a starting point: “I only ever viewed it as a preliminary screen that needed a lot more reporting ... I would be very very nervous about reporting just from algorithmic output” (CU7). Another participant concurred, “Reporters are ultimately going to want to vet everything themselves by hand

to ensure that it's correct, to ensure that they understand it" (CU5). Having people check the leads supplied by an algorithm was seen as the most reliable way to ensure the highest quality standards were met.

### ***Lead development effort and attention***

The editorial effort and attention needed for a lead to mature into a publishable story varies a great deal across use cases. In some cases a lead might entail a few hours of reporter effort to make calls or find illustrative local examples: "Typically the people quoted in the RADAR stories are national people of limited relevance to our readers. So for us as a local paper it's better to get reaction from the people involved locally" (U5). But in other cases it could be a whole day's work to do an interview, get a photo, and work up the details of the lead. RADAR leads are already highly refined, but for less fully formed leads there might be even more work to do. One participant estimated it could take a day or two of effort to prepare a fact check for publication. Another participant remarked that to develop a Newsworthy lead at their radio station, "We have to put like one or two people to work with it for one week" (U1). Statistical aberrations and trends may demand substantial work to assess how interesting they are and begin to explain them: "If we see in some areas that the prices of housing are going up very fast ... Why is that? Then we can ask questions in that region: What's happening here and try to find cases and do journalism from that" (U1).

While it was accepted that CND leads would demand effort and attention to develop into meaningful journalistic contributions, some participants recognized there may be opportunities to utilize lower-skill labor to initially assess leads before passing them along to more seasoned journalists for further investigation. To determine if there's a real story being suggested, one participant remarked, "I'm hoping that basically we can make it so that the leads are self-explanatory enough that people can basically sit their interns down and sit slightly less experienced reporters down" (CU2). Lower skill workers would thus act as an initial screen of leads produced by the system. On the other hand, some leads may not get taken up if there's not someone with *enough* skill to interpret and evaluate them. As one tool creator explained, "Some newsrooms that are maybe a little more data savvy, they tend to use these stories more whereas others that are not they don't use them as much" (C1). So while there is a spectrum of human effort needed, there is also a spectrum of skill that is needed to pursue certain leads.

The availability of human attention is a key factor in determining whether any given news lead develops into a full story. One participant described a situation where automatically generated leads were sent to collaborators but took months before they were looked at and yielded a story. Also, sending too many leads might overwhelm the available attention of a newsroom and users may simply tune them out: "We couldn't take care of everything ... there was too much information for us to handle because we are like 20 people working here in our newsroom" (U1). A lead may look interesting but there may not be enough human capacity to further develop it.

Aside from the raw human attention available, four factors emerged as modulating the degree of effort and attention given to leads: reporter willingness, external contextual factors, monitoring scope, and user interface design.

### ***Reporter willingness***

In some scenarios journalists may give additional attention to individual leads or make time to look at more leads from a system. Typically there is higher willingness to invest time if the journalist already has anecdotal evidence that a story will be found, if that story is aligned with an angle or question they are already pursuing, or if there is a desire to be comprehensive in the assessment of leads because of the nature of the story. For instance, in investigative journalism there may be more of a desire to leave no stone unturned. But in other cases, there is no expectation of finding everything, “I don’t think for one minute that we found everything that was out there, but that wasn’t the point it was a quick screen to find things that looked interesting that we thought were going to give us stories.” (CU7). In some cases reporters are willing to miss things, but in others they’re not. System performance can also impact reporter willingness. If there are too many false positive leads that don’t pan out reporters may start to pay less attention to the leads over time. The nature and strength of the knowledge claims a journalist wants to make also impact how comprehensively a set of leads will be evaluated. So not only do leads perhaps intrinsically need a variable amount of effort to develop, but journalists have a variable willingness to put in that effort based on factors like system performance and the type of claim they might make.

### ***External contextual factors***

During slow news cycles, such as holiday periods, there may be more demand for CND leads to balance the lack of other newsworthy events. “I have grown to appreciate it over the last few months … it showed its worth over Christmas,” explained one RADAR user (U5). A creator explained that they typically see a higher uptake of leads during slow news cycles, “[lead uptake] depends quite a lot on … factors such as what day of the week are the leads sent, what time of year are we on, and so forth” (C1). By the same logic, if other channels of lead development are productive it may crowd out the need to attend to CND system leads: “We haven’t done that many of those fact checks off of the tip sheets in a while but that’s because we have such an abundance right now of statements that need to be fact checked thanks to the political campaigns, the Trump administration … sometimes you have to hunt for leads other times they just kind of flood in” (U3). The attention given to leads, from a CND system or otherwise, is a function of what else is competing for that attention.

### ***Monitoring scope***

The range of inputs monitored by a CND system represent an important form of operator control over the types of things they could have their attention drawn towards. The journalist chooses the scope of monitoring by articulating sources or targets of interest, such as data sources or names of people and other entities. The Klaxon system, for instance, enables users to explicitly target the part of a web page they want

to watch, with the goal of not triggering alerts on irrelevant site changes that might end up overwhelming or distracting the user. A user of Newsworthy leads explained that he didn't want every lead the system produced, but rather only the ones he was topically interested in. The Chequeado system is deliberately set up to monitor media that include at least one outlet from each administrative unit of the country. The scope of monitoring is thus an editorial decision that dictates how wide a net the system fishes with. This in turn impacts the volume, precision, and relevance of leads produced, and presumably increases the likelihood that leads match with interests and receive attention. The scope of monitoring can also be algorithmically widened based on an initial query, augmenting human ability by obviating the need to know exactly what to monitor. If the monitoring scope of a system is set too wide, either through user configuration or algorithmic expansion, control can be maintained by supporting the *filtering* of leads according to various relevance criteria. This came up in the context of a fact spotting system in which several users indicated they were not interested in checking the claims of pundits or other journalists and so wanted to filter those out to reduce distraction, while also recognizing that other outlets might still be interested in those leads. This underscores the variability in interests between different journalists and outlets and how monitoring scope and post-filtering can enable editorial control and support diverse uses.

### ***User interface design***

A final factor that impacts the flow of attention towards leads is the nature of the user interface (UI) and how it frames information for users. Some systems, such as Newsworthy or Tech & Check, send discrete chunks of information via email. By explicitly marking individual items or sets of items (e.g. a set of claims for fact checking) for attention this may, however, create expectations around the relevance or importance of those items. Other systems, such as RADAR or Chequeado, provide an interface that presents a ranked list of items that can be browsed. A ranking UI has the advantage of communicating some degree of relevance that corresponds to the ordering of leads, without definitively marking some subset as worthy of attention. A mixture of the two approaches involves sending discrete leads via email with a link to an interactive UI containing more details and a full ranking. Date-time order (i.e. most recent at top) is a common default for rankings, though leads can be sorted according to other criteria. Some interfaces provide keyword search functionality allowing users to find leads based on specific interests.

Finding the right presentation of the information from the lead was seen as a challenge: "I'm fairly optimistic at this moment about us being able to generate useful leads. But I'm fairly pessimistic about us being able to communicate that well enough to the reporters who are supposed to then make sense of them" (CU2). At Chequeado they've carefully crafted the information shown in the UI. Each lead shows the claim, the media outlet where it was published, a link to the claim (e.g. article where it was detected), and then shows the paragraph where the claim was found as additional context. Initially users see the last 3 days of checkable claims, but can continue scrolling down until they eventually start seeing claims that are less fact checkable. The rationale for including some of the not fact checkable claims in the interface was to

be able to collect feedback to further train the machine learning (i.e. by including negative cases). Newsworthy leads, on the other hand, consist of a chart that visualizes the trend or statistical anomaly behind the lead, a few sentences of generated text describing what it's about, and a link to the original spreadsheet with the data backing the lead. Linking to the data was found to be important: "Most or almost all reporters that get these leads and do something with them actually look at the data themselves" (C1). This approach was also taken with the Tech & Check leads: a link from each lead provides quick access to the source transcript which allows reporters to assess the context of a statement before further pursuing. Importantly, lead presentations included vital context to help reporters launch into follow-up activities.

### ***Newsworthiness assessment***

Participants were found to assess the leads produced by CND systems according to classic indicators of newsworthiness, such as reference to elites, proximity, and conflict (Harcup and O'Neil 2016), as well as to newer factors related to online media such as the potential for online traffic an item might produce (Vu 2014). Newsworthiness dimensions such as actuality and audience fit were thought to be important to reserve for humans in the loop. For instance, different outlets from a newspaper chain could treat the same lead differently depending on what editors thought was important to each local community. Other newsworthiness dimensions that participants mentioned with respect to lead evaluation included exclusivity, originality and repetition, organizational fit, degree of public interest, significance of impact, and unexpected changes or novelty.

A potential issue with working off of leads that are produced centrally by an algorithm and then distributed to different newsroom is that they lack exclusivity. Everyone potentially gets the same lead from the same algorithm. One user of the Newsworthy leads thought this was a major issue. A lack of exclusivity meant they were reduced to competing based only on speed: "The problem for us was we didn't get the information exclusively ... if we didn't do this stuff immediately other channels might or local stations local newspapers might have just printed those statistics so that it was no use for us" (U1). However, exclusivity wasn't mentioned by other participants as a factor. In fact checking, for instance, the premise is to check publicly available statements that are by definition not exclusive. As a fact checker explained, "We're certainly competitive, but it's a little bit different. It's not scoop driven, so the nature of the competition is a little more nuanced", adding that "you don't want to fact check something that no one's ever heard of" (U3). These comments underscore how the competitive landscape and reporting context can modulate whether exclusivity of leads is a driving factor in news selection.

Some participants re-framed the issue of exclusivity as one of originality, which people could control by developing and differentiating the leads from the system. Originality was also discussed in terms of its opposite: repetition. A publisher may have already covered a story provided by a CND and may choose not to follow the lead as a result. "One lead was that this August was unusually hot in this little town. But then we had already reported on that, on the extreme drought of the summer. So

it wasn't really newsworthy because it wasn't a new story," explained a user of Newsworthy (U2). The newsworthiness of a lead and whether it's pursued is modulated by what's already been covered by a news organization or its competitors, "We don't want to duplicate effort or something that's run lately in a big city daily two or three weeks previously," explained a RADAR user (U7).

Organizational fit and agenda can also impact the uptake and pursuit of leads. One participant described the mismatch between the statistically interesting leads provided by the Newsworthy system and the need to identify impacted individuals for the coverage she felt her outlet specialized in. Other participants described how the content mix for their publication was important and could impact how many leads they pursue, "An individual weekly paper will ... only probably want to run five data stories ... because of the mix of content" (C4). Some outlets look at the distribution of their coverage over some dimension of interest, like how many claims are being fact checked on each side of a political spectrum, "We decide not only based on what the robot shows us but also in the balance that we're trying to have that week we try to cover the party and the government but also the opposition" (C2).

Public interest value was also mentioned as a factor in the selection of leads. As one investigative journalist remarked, "I was sitting on that for more than a year and unclear whether we ever write a story about it because we were waiting for what we felt were genuine public interest stories to be flagged from it," (CU7). Leads had to be detected that could be developed into stories that had a real public interest angle.

Participants also talked about the application of newsworthiness criteria in a way that was contingent on other factors like topic, domain, or location. For instance, the interestingness of different types of statistical patterns might vary based on what the data is about: "We have certain definitions of what's a story in crime data for example, but we need to work with other definitions in the unemployment data" (C1). For example, a new record level for real-estate prices is not much of a news story since housing prices tend to go up every month. Rule-based logic for triggering alerts depends on the domain of interest, underscoring the need for domain expertise in defining those rules. Location contrasts can also modulate newsworthiness: the importance of a statistical anomaly could be amplified or diminished based on what's going on in a neighboring area.

Participants also described how measures of significance, impact, or magnitude are often domain specific, or even uniquely story-specific. For example, one system for monitoring arrests ranks individuals based on bail amount. In this particular domain, sorting by bail magnitude helps a reporter identify the most significant cases according to that metric. Other measures of magnitude come into play in other domains. A fact checking system uses the reach of a media outlet to assess the magnitude of impact of claims published there. An investigative journalist described how the significance of a lead can depend on who's connected to whom.

Leads reflecting changes, deviations, and novelties were also viewed as potentially newsworthy. The simplest of leads reflecting the message "something has changed" could at times be the most effective for drawing attention from human journalists who could then make a further assessment of the significance of that change. As one participant explained, "Oftentimes the things that people really want to be notified of

are things that don't require any sort of fancy algorithm ... like when did Donald Trump file his most recent return" (CU5). By embedding an evaluative component in a change detection algorithm leads can also be tuned to identify changes that are unexpected or novel. "What we're looking for ... mostly falls under the rubric of anomaly detection: When something isn't what maybe we expect it to be ... What we're looking for is things that are deviating from an established pattern in some way, or frankly are just unusual" (CU4).

Participants described two main approaches for embedding newsworthiness evaluations into CND systems: rule-based and machine-learning based. Rule-based systems encode logic in the form of "IF some combination of data variables detected, THEN send alert" whereas machine-learning systems are able to learn patterns from a set of examples and then detect those same patterns in new datasets. Oftentimes rule-based systems are adequate and the trick is having to figure out what rule will not be too noisy with false positives, but also not miss the pattern of interest. A side-effect of the rule-based approach is that it makes editorial decisions very explicit, which can then facilitate deliberation and critique of news values. On the other hand, machine learning approaches require the selection of specific examples to train the system on, which end up implicitly defining and biasing the evaluation of newsworthiness.

### ***Quality assurance***

Participants articulated factors which support the assessment of quality and the development of trust in a CND system, underscoring the idea that leads produced by algorithms should be treated with caution and skepticism. As one participant explained, "Before we would ... use it as the basis for our own reporting we would want to understand ... what it was doing," before elaborating, "The higher bar is taking somebody else's software and essentially using it as a key point ... or the crux or basis of a story ... without really understanding it in fairly specific detail" (CU4). When there is transparency provided with the lead it can help engender trust, or at least verifiability in the absence of trust, in the way the data and algorithm have produced the lead. One Newsworthy user explained, "You could pick the statistics up for yourself if you want to because it's open databases ... and it was also easy to check if we wanted to check ourselves" (U1).

Transparency information can take many different forms (Diakopoulos and Koliska 2017), but one aspect that was explicitly mentioned in interviews was to make the source materials for the lead, such as data or documents, available for inspection via a link. RADAR leads, for instance, sometimes include a "workbook" that shows any data transformations in arriving at the final numbers. In the fact checking context, journalists always need to go back to the transcript, both to see the context and verify who the speaker was, but also to ensure the transcript itself is correct against an original recording. "There's all kinds of other glitches and errors in the transcript but we know that so we double check them," explained a fact checker (U4). Links to the source material can serve to quickly confirm if there was an error. In general, making the original document, media, or data behind a lead available for inspection helped users more quickly assess the lead.

A more subtle aspect of assuring quality relates to whether the interpretation of a lead is accurate according to other known context. As one user of RADAR leads explained, "It's important to me to make sure that what we give to our readers is something that has a little bit of context around it ... it's a standard conversation to have is do we need to get more context do we need to get more info to present this or is it okay as is ... sometimes data can be misleading if it isn't given proper context ... it's not [about] trusting the data, I trust the data. It's about making sure that we present it in a sensible and responsible way as publisher" (U5). In order to assess quality, journalists needed to understand how their local context could alter the interpretation of a lead.

Developing trust in leads is also about the organization supplying the leads. As a Newsworthy user explained, "many newsrooms are kind of suspicious of other independent journalism people, or people outside their own newsroom" (CU4). Another user reinforced the importance of organizational trust: "It's as much about trusting the lead as it is trusting the organization supplying the lead" (U1). One way some CND systems seek to engender trust in leads is to develop error prevention or quality assurance (Q/A) processes that are then communicated to users. The need to develop trust in the organizational and procedural source of CND leads may mean that building and using in-house systems could ease adoption.

## Discussion

The findings of this study elaborate the workings and contingencies of the internal newsroom attention economy and its relationship to CND systems and their role in sociotechnical gatekeeping processes, showcasing how normative expectations around news values and verification remain firmly intact and might be supported more effectively in future designs (Aitamurto et al 2019; Diakopoulos 2019). The following subsections build on these findings to suggest opportunities for the effective design of future CND systems to better align with the evaluative work of journalists while supporting the desire for human agency and configurability to suit a high variance in journalistic scenarios.

### *The newsroom attention economy*

CND systems naturally shape the set of items a reporter might examine. Their ability to direct attention by shaping what journalists attend to should be incorporated more directly into contemporary gatekeeping models (Wallace 2017), including the role that the design of CND user interfaces plays in that process. Importantly, CND systems do not grow the attention budget available in a newsroom. Instead, algorithmically generated leads must also compete with all of the other information channels that journalists are monitoring, potentially contributing to the intensification of journalistic labor (Cohen 2018). Findings suggest that in some cases, such as exhaustive investigative scenarios, CNDs can save time by accurately filtering away documents that otherwise would have consumed attention. But in other scenarios CNDs may instead divert attention rather than save it, albeit with the promise of discovering newsworthy

stories that wouldn't otherwise have been found. While it is still true that "new ways of combining data and algorithms could lower the costs of discovering stories" (Hamilton 2016, pg 282), the results here indicate that the development of those stories into publishable journalistic contributions will accrue extensive labor costs. This finding is in line with recent observations about the need for follow-on editorial work when using data mining for investigative journalism (Stray 2019), as well as with broader understandings of how automation technologies oftentimes end up creating new tasks or even roles for people in hybridized workflows (Diakopoulos 2019a; Ford and Hutchinson 2019; Jones and Jones 2019). CND system designs should therefore be sensitive to the supply and demand of attention within the newsroom and to their role in that economy.

Findings furthermore indicate that CND systems should be designed for the particular use case of the leads, taking into account the expected investment of effort that journalists may be willing to provide in different scenarios (e.g. Daily reporting vs. investigative journalism). In some cases journalists will want to comprehensively assess every lead, while in others they may be content to receive any reasonable lead. Another factor to consider is the degree of editorial skill that may be required to evaluate leads and whether leads could be designed for a range of differently skilled people. Systems might also be designed such that lead timing and volume could be tuned (either manually or automatically) to the pace of the news cycle or other external contextual factors. Since CND systems have the potential to divert attention, they may also be designed to support user agency in how that attention is directed, an approach that other designers of interactive AI systems have found effective (Heer 2019). In particular, in order to offer more individual agency in the input and throughput phases of gatekeeping (Wallace 2017), control could be provided over the scope of inputs that are monitored as well as the filtering of leads based on potentially domain-specific or even user-defined relevance criteria. Allowing journalists to adaptively align tools with their specific goals, such as open questions or angles they are pursuing, may serve to increase their willingness to invest the necessary effort to further develop leads.

A vital aspect of how leads are treated in the newsroom attention economy and incorporated into internal sociotechnical gatekeeping processes is their user interface—how they are presented to users, potentially grab attention by conveying what's most compelling about the lead, and provide information that reduces the uncertainty of pursuing the lead to ease an initial evaluation. According to Hamilton, "Reporters do change coverage depending on the costs of assembling a story," (2016, pg 293). It's at the user interface where designers will be able to influence, to some extent at least, the perception of that cost.

The role lead presentation interfaces play in attention management parallels an earlier strand of research on information surrogates, such as document previews, which allow users to make decisions about the potential relevance of an item (Greene et al 2000). Web search results are a prime example of highly refined information surrogates, and offer potential avenues for designing lead presentation result sets (Wilson 2012). In the context of news, leads function as information surrogates with the potential to attract the attention of journalists who may then invest additional effort in

developing a publishable story. Ideally surrogates would be designed to enable journalists to make an informed decision about whether a given lead is worth their time and energy to pursue. They could be thought of as “pitches” which communicate what the story is about, what the interesting angle is and why it’s significant, what the sources are, and so on (Diakopoulos 2019a). In other words, the surrogate should be designed to enable the evaluative decisions of journalists, including both newsworthiness and verification assessments. The easier a lead makes it to see whether there’s a real story and get started with subsequent quality assessments, the more likely it will be transformed into a publishable story. Such evaluative assessments can be supported by designs which enable access to underlying data and data transformations, algorithmic transparency into how a lead was discovered, and further context that may modulate interpretation. An interesting ethical dimension to consider is how aggressive leads might be in attracting a journalist’s attention based on the confidence of the system in the lead panning out. This study makes clear that future research will be needed to design and evaluate effective information displays and user interfaces for leads that suit different journalistic contexts.

### ***Algorithmic newsworthiness***

What is selected to become a news item is contingent on a range of individual, organizational, social, normative, economic, and technical forces (Shoemaker and Vos 2009). This was reflected in the wide range of newsworthiness criteria mentioned by participants, including a variety of classic newsworthiness factors like exclusivity, conflict, surprise, reference to elites, magnitude and significance, proximity, audience expectations, and organizational agenda (Harcup and O’Neil 2016), which were important to different participants in different combinations. The implication for CND systems, and for sociotechnical gatekeeping processes more broadly, is that they might be designed to support the *configurability* of newsworthiness assessments so that leads can be adapted to suit a wide range of journalistic scenarios. This would further allow reporters to align the leads produced with their own willingness to pursue them, as well as dynamically account for external contextual factors not quantified by the system. For instance, a journalist might want to specify their particular interest in leads reflecting a conflict between elites and have those leads sorted at the top of their interface, or filtered into the display from a longer list. Or a lead presentation display could highlight metadata, like the name of an elite person mentioned, in order to make that aspect of newsworthiness more salient to a user. If journalists were able to articulate the most interesting dimensions of newsworthiness for their context, specific CND systems could be designed to offer greater efficiency, effectiveness, satisfaction, and alignment with editorial requirements and news values when evaluating the leads.

Supporting more configurability in the application of newsworthiness criteria in evaluating, sorting, or filtering leads would demand the development of computational operationalizations and implementations of newsworthiness criteria. While some interestingness and newsworthiness metrics have been discussed in prior work (Geng and Hamilton, 2006; Liu et al 2017; Magnusson et al 2016; Opdahl and Tessem 2019),

an ambitious line of interdisciplinary future work should systematically consider how to computationally specify and interactively configure the entire range of newsworthiness definitions found in the journalism studies literature. The technical and ethical challenges of encoding newsworthiness into algorithms may entail the explication of difficult-to-articulate rules, thresholds, or weightings from journalists, the development of statistical models using machine learning on appropriately labeled data, and the need for methods to evaluate how well computational operationalizations align or come into tension with normative or practical expectations of news values (Diakopoulos 2019; Milosavljević and Vobić 2019; Stray 2019).

The ability of CND systems to direct editorial attention suggests that gatekeeping models (Wallace 2017) might be usefully updated to include the role of algorithms not only in curating consumer attention, but also in suggested information to various types of human gatekeepers prior to publication. On a technical level, news coverage might be shaped according to the various quantifications of newsworthiness that might be developed. Redefining a key metric and encoding it differently could cause a system to locate different leads. Some measures of newsworthiness may be easier to encode algorithmically, making them more likely to be implemented and increasing the volume of those types of leads. If other measures are not technically feasible to algorithmically encode then simpler definitions may prevail. Some types of newsworthy stories may ultimately be harder to define in ways that can be written into code, which in the extreme could cause those types of stories to receive less coverage because algorithms can't pick them up (Diakopoulos 2019a). Some of the fact checkers interviewed expressed awareness and concern over the bias of the leads provided by their claim spotting systems, resulting in additional monitoring strategies to track coverage and ensure that the balance of claims checked adhered to editorial goals. The findings here aligns with observations in previous research that have noted algorithmic bias concerns in journalistic tools for source finding on social media (Thurman 2017; Thurman et al 2016; Schwartz, Naaman, Teodoro 2015). To address these types of bias concerns, CND systems might be explicitly designed to empower editors to track the shape of leads produced, including the impacts of how their reporters have configured the tools in terms of the scope of monitoring and application of newsworthiness criteria. Future editors might set bounds on an acceptable range of configurations in order to encourage a particular shape of coverage, further supporting journalistic autonomy in defining the overall outcomes related to these systems.

## Conclusions

"Any technological innovation, once adopted, offers routine paths for news organizations to select and shape the news," note Shoemaker and Vos in their articulation of gatekeeping theory (2009). This article delves into this idea in some detail, offering insight into the experiences of practitioners with respect to a range of computational news discovery systems and their role in shaping internal sociotechnical gatekeeping processes. The findings offer an elaboration of various constraints, affordances, and social and normative expectations for CND tools, suggesting design implications that can guide the development of future CND tools that are consistent with user needs.

These include the benefits of considering controllability and configurability of CND tools (e.g. monitoring scope, output filtering, newsworthiness definition) to suit the demands of varying reporting scenarios and contexts and increase the relevance of leads for journalists' interests and goals; the need to consider timing and volume of leads with respect to available journalistic attention; the need to support follow-up verification activity through transparency and context; and the importance of user-interface and information design in modulating journalists' uptake of leads. In general, future work should strive to build on and elaborate these findings through triangulation via alternative methods, such as embedded ethnographic observation, focus groups, and iterative human-centered design involving evaluations through user studies and field deployments.

## Notes

1. <https://radarai.com/>
2. <http://newsworthy.se/>
3. <https://github.com/themarshallproject/klaxon>
4. <https://fullfact.org/>
5. <https://chequeado.com/>

## Acknowledgements

I would like to thank colleagues at Northwestern who offered helpful feedback on this article in the Research in Progress (RIP) seminar, and to the anonymous reviewers who pushed me to greatly improve the initial manuscript. I am also indebted to my interview participants for the time they spent speaking with me, thank you!

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Funding

This work was supported by the U.S. National Science Foundation under grant IIS-1845460.

## ORCID

Nicholas Diakopoulos  <http://orcid.org/0000-0001-5005-6123>

## References

Adair, Bill, Mark Stencel, and Cathy Clabby, 2019. Chengkai Li. "The Human Touch in Automated Fact-Checking." *Proc. Computation + Journalism Symposium*.

Anderson, C. W. 2013. "Towards a Sociology of Computational and Algorithmic Journalism." *New Media & Society* 15 (7): 1005–1021.

Atamurto, Tanja, et al. 2019. "HCI for Accurate, Impartial and Transparent Journalism: Challenges and Solutions." Workshop at the CHI Conference on Human Factors in Computing Systems.

Bandy, J., and Nicholas Diakopoulos. 2020. "Auditing News Curation Systems: A Case Study Examining Algorithmic and Editorial Logic in Apple News." *Proc. International Conference on Web and Social Media (ICWSM)*.

Broussard, Meredith. 2015. "Artificial Intelligence for Investigative Reporting." *Digital Journalism* 3 (6): 814–831.

Bruckman, Amy. 2006. "Teaching Students to Study Online Communities Ethically." *Journal of Information Ethics* 82: 82–98.

Bucher, Taina. 2016. "Machines Don't Have Instincts': Articulating the Computational in Journalism." *New Media & Society* 19 (6): 918–933.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W. W. Norton & Company.

Chaplin, Heather. 2016. "Guide to Journalism and Design." *Tow Center for Digital Journalism*.

Coddington, Mark. 2015. "Clarifying Journalism's Quantitative Turn." *Digital Journalism* 3 (3): 331–348.

Cohen, Nicole. 2018. "At Work in the Digital Newsroom." *Digital Journalism* 4 (5): 571–591.

Denning, Peter J. 2005. "Is Computer Science Science?" *Communications of the ACM* 48 (4): 27–31.

DeVito, Michael A. 2017. "From Editors to Algorithms." *Digital Journalism* 5 (6): 753–773.

Diakopoulos, Nicholas, and Michael Koliska. 2017. "Algorithmic Transparency in the News Media." *Digital Journalism* 5 (7): 809–828.

Diakopoulos, Nicholas, Munmun De Choudhury and Mor Naaman. 2012. "Finding and Assessing Social Media Information Sources in the Context of Journalism." *Proc. Conference on Human Factors in Computing Systems (CHI)*.

Diakopoulos, Nicholas. 2019a. *Automating the News: How Algorithms Are Rewriting the Media*. Cambridge, MA: Harvard University Press.

Diakopoulos, Nicholas. 2019. "Towards a Design Orientation on Algorithms and Automation in News Production." *Digital Journalism* 7 (8): 1180–1184.

Diakopoulos, Nicholas, Madison Dong, Leonard Bronner, and Jeremy Bowers. 2020. "Generating Location-Based News Leads for National Politics Reporting." *Proc. Computation + Journalism Symposium*.

Eldridge, Scott A., II, Kristy Hess, Edson C. Tandoc, Jr., and Oscar Westlund. 2019. "Navigating the Scholarly Terrain: Introducing the Digital Journalism Studies Compass." *Digital Journalism* 7 (3): 386–403.

Felix, Cristian, Anshul Vikram Pandey, Enrico Bertini, Charles Ornstein, and Scott Klein. 2015. "RevEx: Visual Investigative Journalism with a Million Healthcare Reviews." *Proc. Computation + Journalism Symposium*.

Fletcher, Richard, Steve Schifferes, and Neil Thurman. 2017. "Building the 'Truthmeter': Training Algorithms to Help Journalists Assess the Credibility of Social Media Sources." *Convergence: The International Journal of Research into New Media Technologies* 26 (1): 19–34.

Ford, Heather, and Jonathon Hutchinson. 2019. "Newsbots That Mediate Journalist and Audience Relationships." *Digital Journalism* 7 (8): 1013–1031.

Geng, Liqiang, and Howard J. Hamilton. 2006. "Interestingness Measures for Data Mining: A Survey." *ACM Computing Surveys* 38 (3): 9–es.

Graefe, Andreas. 2016. *Guide to Automated Journalism*. Tow Center for Digital Journalism.

Graves, Lucas. 2018. *Understanding the Promise and Limits of Automated Fact-Checking*. Reuters Institute.

Glaser, Barney, and Anselm Strauss. 2009. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. New Brunswick: Aldine.

Greene, Stephen, Gary Marchionini, Catherine Plaisant, and Ben Shneiderman. 2000. "Previews and Overviews in Digital Libraries: Designing Surrogates to Support Visual Information Seeking." *Journal of the American Society for Information Science* 51 (4): 380–393.

Hamilton, James, and Fred Turner. 2009. *Accountability through Algorithm: Developing the Field of Computational Journalism*. Center for Advanced Study in the Behavioral Sciences.

Hamilton, James T. 2016. *Democracy's Detectives*. Cambridge, MA: Harvard University Press.

Harcup, Tony, and Deirdre O'Neill. 2016. "What is News? News Values Revisited (Again)." *Journalism Studies* 23 (1): 1470–1488.

Hassan, Naeemul, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. "Toward Automated Fact-Checking: Detecting Check-Worthy Factual Claims by Claimbuster." Proc. Conference on Knowledge Discovery and Data Mining (KDD).

Heer, Jeffrey. 2019. "Agency plus Automation: Designing Artificial Intelligence into Interactive Systems." *Proceedings of the National Academy of Sciences* 116 (6): 1844–1850.

Heinderyckx, F., and T. P. Vos. 2016. "Reformed Gatekeeping." *CM: Communication and Media* 11 (38): 29–46.

Jones, Bronwyn, and Rhianne Jones. 2019. "Public Service Chatbots: Automating Conversation with BBC News." *Digital Journalism* 7 (8): 1032–1053.

Lewis, Seth, and Oscar Westlund. 2015. "Actors, Actants, Audiences, and Activities in Cross-Media News Work." *Digital Journalism* 3 (1): 19–37.

Liu, Xiaomo, et al. 2017. "Reuters Tracer: Toward Automated News Production Using Large Scale Social Media Data." Proc. IEEE BigData. <https://ieeexplore.ieee.org/abstract/document/8258082>

Lofland, John, and LynH. Lofland. 1994. *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis*. 3rd ed. Belmont, CA: Wadsworth.

Lokot, Tanya, and Nicholas Diakopoulos. 2016. "News Bots: Automating News and Information Dissemination on Twitter." *Digital Journalism* 4 (6): 682–699.

Magnusson, Ma°ns, Jens Finnäs and Leonard Wallentin. 2016. "Finding the News Lead in the Data Haystack: Automated Local Data Journalism Using Crime Data." Proc. Computation + Journalism Symposium.

Maiden, Neil, et al. 2018. "Making the News: Digital Creativity Support for Journalist." Proc. Conference on Human Factors in Computing Systems (CHI).

Milosavljević, Marko, and Igor Vobić. 2019. "Human Still in the Loop." *Digital Journalism* 7 (8): 1098–1116.

Nechushtai, Efrat, and Seth C. Lewis. 2019. "What Kind of News Gatekeepers Do We Want Machines to Be? Filter Bubbles, Fragmentation, and the Normative Dimensions of Algorithmic Recommendations." *Computers in Human Behavior* 90: 298–307.

Nourbakhsh, Armineh, Quanzhi Li, Xiaomo Liu, and Sameena Shah. 2017. "Breaking" Disasters - Predicting and Characterizing the Global News Value of Natural and Man-Made Disasters." Proc. Data Science + Journalism Workshop

Opdahl, Andreasl and Bjørnar Tessem. 2019. "Towards Ontological Support for Journalistic Angles." In *Enterprise, Business-Process and Information Systems Modeling*, edited by Reinhartz-Berger I, Zdravkovic J, Gulden J, and Schmidt R, 352:279–294. Cham: Springer, Cham.

Park, Deokgun, Simranjit Sachar, Nicholas Diakopoulos, and Niklas Elmquist. 2016. "Supporting Comment Moderators in Identifying High Quality Online News Comments." Proc. Conference on Human Factors in Computing Systems.

Perrin, William, et al. 2017. "Local News Engine: Can the Machine Help Spot Diamonds in the Dust?." In J. Mair., eds. *Data Journalism past, Present, Future*. Bury St Edmunds, UK: Abramis academic publishing.

Plattner, Titus, Didier Orel, and Olivier Steiner. 2016. "Flexible Data Scraping, Multi-Language Indexing, Entity Extraction and Taxonomies: Tadam, a Swiss Tool to Deal with Huge Amounts of Unstructured Data." Proc. Computation + Journalism Symposium.

Sas, Corina, Steve Whittaker, Steven Dow, Jodi Forlizzi, and John Zimmerman. 2014. "Generating Implications for Design through Design Research." Proc. CHI Conference on Human Factors in Computing Systems.

Schwartz, Raz, Mor Naaman, and Rannie Teodoro. 2015. "Editorial Algorithms: Using Social Media to Discover and Report Local News." Proc. International Conference on Web and Social Media 9: 407–415.

Shearer, Matt, Basile Simon, and Clément Geiger. 2014. "Datastringer: Easy Dataset Monitoring for Journalists." Proc. Computation + Journalism Symposium.

Shoemaker, Pamela, and Tim Vos. 2009. *Gatekeeping Theory*. New York, NY: Routledge.

Stray, Jonathan. 2019. "Making Artificial Intelligence Work for Investigative Journalism." *Digital Journalism* 7 (8): 1076–1097.

Thorson, Kjerstin, and Chris Wells. 2016. "Curated Flows: A Framework for Mapping Media Exposure in the Digital Age." *Communication Theory* 26 (3): 309–328.

Thurman, Neil. 2019. "Computational Journalism." In K. Wahl-Jorgensen and T. Hanitzsch, eds. *The Handbook of Journalism Studies*, Second Edition. New York, NY: Routledge.

Thurman, Neil, Steve Schifferes, Richard Fletcher, Nic Newman, Stephen Hunt, and Aljosha Karim Schapals. 2016. "Giving Computers a Nose for News: Exploring the Limits of Story Detection and Verification." *Digital Journalism* 4 (7): 838–848.

Thurman, Neil, Konstantin Dörr, and Jessica Kunert. 2017. "When Reporters Get Hands-on with Robo-Writing: Professionals Consider Automated Journalism's Capabilities and Consequences." *Digital Journalism* 5 (10): 1240–1259.

Thurman, Neil. 2017. "Social Media, Surveillance, and News Work." *Digital Journalism* 6 (1): 76–97.

Vu, Hong Tien. 2014. "The Online Audience as Gatekeeper: The Influence of Reader Metrics on News Editorial Selection." *Journalism: Theory, Practice & Criticism* 15 (8): 1094–1110.

Tolmie, Peter, et al. 2017. "Supporting the Use of User Generated Content in Journalistic Practice." Proc. Conference on Human Factors in Computing Systems (CHI).

Weber, Matthew S., and Allie Kosterich. 2018. "Coding the News." *Digital Journalism* 6 (3): 310–329.

White, David Manning. 1950. "The 'Gate Keeper': A Case Study in the Selection of News." *Journalism Quarterly* 27 (4): 383–390.

Wilson, Max. 2012. *Search User Interface Design*. Morgan & Claypool. <https://www.morganclaypool.com/doi/abs/10.2200/S00371ED1V01Y201111ICR020>.

Zamith, Rodrigo. 2019. "Algorithms and Journalism." In *Oxford Research Encyclopedia of Communication*. Oxford University Press. <https://oxfordre.com/communication/view/10.1093/acrefore/9780190228613.001.0001/acrefore-9780190228613-e-779?rskey=hU25XZ&result=1>.

## Appendix A: Semi-Structured interview guide

- Could you give me an overview of how you use the news leads from <service> in your work?
- What do you think is the main utility of such leads in your work?
- What do the leads look like that you receive from <service>? What kind of information is provided with each lead?
- Does the interface provided by <service> work well for your needs? E.g. how you receive, search, or filter leads?
- Are there any features or improvements to the leads that you wish <service> had?
- Do you think the leads sent by <service> are newsworthy, or could be? Why or why not?
- What might make them more newsworthy for you?
- Is there any additional contextual info that you think would make the leads more compelling, interesting, valuable, or easy for you to pursue?
- Do news leads from <service> help you be more competitive in your news coverage? Why or why not?
- Do you get an appropriate number of leads from <service>? How often and when do they come? Are there too many, or not enough?
- How much time do you spend looking at the leads from <service>?
- How often do leads from <service> mature into full-fledged news items for you? How long does it take for a lead to mature into a news item? What might improve the signal to noise ratio?
- How do these leads fit into your broader workflow? How else do you scan for story leads?
- If you were able to filter or configure the leads, how would you do so?

- Do you rely on or trust the leads you get from <service>? Why or why not?
- Can you recall any cases where there was an error or mistake in a lead provided by <service>? How did you cope with that?
- What would be your advice to other newsrooms that are looking to adopt such a story discovery tool?
- Do you think such story discovery services should be adopted more widely in news production?
- Do you think such services introduce any ethical decisions into journalism practice?
- Who else do you think I should talk to about this topic?