GIVING COMPUTERS A NOSE FOR NEWS: Exploring the limits of story detection and verification

Neil Thurman, Steve Schifferes, Richard Fletcher, Nic Newman, Stephen Hunt, and Aljosha Karim Schapals

Acknowledgements

This work was supported by the European Union, under Grant 287975, and by the VolkswagenStiftung under a Freigeist Fellowship.

Main and Corresponding Author:

Professor Neil Thurman
Institute of Communication Studies and Media Research
Ludwig Maximilian University of Munich
Oettingenstr. 67, 80538 Munich, Germany

Telephone: +49 (0)892180 9449
Email: neil.thurman@ifkw.lmu.de

This is the authors' original manuscript version of an article that has been accepted for publication in Digital Journalism. The DOI of the article is: 10.1080/21670811.2016.1149436. It will be published online at the following permanent link: http://dx.doi.org/10.1080/21670811.2016.1149436

The use of social media as a source of news is entering a new phase as computer algorithms are developed and deployed to detect, rank, and verify news. The efficacy and ethics of such technology are the subject of this article, which examines the SocialSensor application, a tool developed by a multidisciplinary EU research project. The results suggest that computer software can be used successfully to identify trending news stories, allow journalists to search within a social media corpus, and help verify social media contributors and content. However, such software also raises questions about accountability as social media is algorithmically filtered for use by journalists and others. Our analysis of the inputs SocialSensor relies on shows biases towards those who are vocal and have an audience, many of whom are men in the media. We also reveal some of the technology's temporal and topic preferences. The conclusion discusses whether such biases are necessary for systems like SocialSensor to be effective. The article also suggests that academic research has failed to fully recognise the changes to journalists' sourcing practices brought about by social media, particularly Twitter, and provides some countervailing evidence and an explanation for this failure.

Keywords: algorithmic news, automation, computerisation, employment, journalism, social media, topic detection, verification

Introduction

The ubiquity of computing in contemporary culture has resulted in human decision-making being augmented, and even partially replaced, by computational processes or algorithms using artificial intelligence and information-retrieval techniques. Such augmentation and substitution is already common, and even predominates, in some industries, such as financial trading and legal research. Frey and Osborne (2013) have attempted to predict the extent to which a wide spectrum of jobs is susceptible to computerisation. Although journalists were not included in their analysis, some of the activities undertaken by journalists—for example those carried out by interviewers, proofreaders, and copy markers—were, and had a greater than 50 per cent probability of being computerised. It is that potential for the automation of journalistic work that is explored in this article.

Frey and Osborne remind us of how automation can be aggressively resisted by workers, giving the example of William Lee who, they say, was driven out of Britain by the guild of hosiers for inventing a machine that knitted stockings. Such resistance also exists in the context of journalistic automation. For example, the German Federation of Journalists have said they "don't think it is … desirable that journalism is done with algorithms" (Konstantin Dörr, personal communication, 6 February 2015). There is, however, also appreciation of the benefits automation can bring in assisting journalists with the management of the huge volumes of information—particularly from social media—they currently deal with (see, for example, Schifferes et al. [2014]).

Such algorithms, embodied in the SocialSensor mobile[1] and web applications, are the focus of this article. SocialSensor was developed by a multidisciplinary, international team, including some of this article's authors.[2] The project has been designing, building, and testing a single tool that aims to "quickly [surface] trusted material from social media … with context" (Diplaris et al. 2012, 1243). SocialSensor is only one example of such technology. Others include Newswhip's Spike, which

promises to enable journalists to “find big stories while they are still small” (Newswhip.com n.d.), and Geofeedia, which promises to allow journalists to “discover and break stories first” by providing a means to “spot, understand, and report trends in location-based social media” (Geofeedia.com 2015).

Of course such tools are not just of professional interest to journalists; they also raise important issues for those who study journalism, prompting reflection on:

- The social and professional contexts that have incubated the development of such tools.
- How they work.
- What biases might they have and should they be challenged?
- And how could an increase in their use change what journalists produce?

This article explores these questions, drawing on what we have learnt from the SocialSensor project.

Social and professional contexts

More and more news stories are becoming known for their social media provenance. Take, for example, the Hudson River plane crash, which was reported on Twitter about 15 minutes earlier than anywhere else (Beaumont 2009). Since then, journalists tell us, social media has grown in importance as a source. “It’s an incredibly important source,” says Jonathan Rugman of Channel 4 News (personal communication, 12 September 2014), a view echoed by Krishnan Guru-Murthy, also of Channel 4, who says “we use [social media] constantly in all stories … I don’t think you can overstate its importance” (personal communication, 21 September 2014).

There is some contradiction, then, in the fact that research to date has suggested that a relatively small proportion of newspaper and broadcast stories uses social media as a primary source. Broersma and Graham (2013) showed fewer than four articles per newspaper per day quoted Twitter. And the numbers were even smaller in a study of US newspapers and TV stations (Moon and Hadley 2014). One explanation for this contradiction is that both studies used data from 2011 and, since then, there appears to have been a step change in the way journalists use social media. Laura Roberts believes that “[2011] was a bit of a turning point in the way that the *Daily Telegraph* looked at using social media” (personal communication, 2 October 2014). Jonathan Rugman concurs, saying “it’s been a huge change, and it started in 2011”.

Nevertheless, even research that uses data from after 2011 continues to show relatively infrequent mentions of social media as sources of information (Paulussen and Harder 2014; Wallsten 2015). One explanation for this mismatch between journalists’ statements about their heavy reliance on social media and how infrequently those sources are actually cited is that it is often used as a tip-off mechanism, with journalists corroborating the information elsewhere. “[Twitter] can help inform you where things are kicking off,” Krishnan Guru-Murthy told us. BBC Middle East Correspondent Yolande Knell assessed Twitter’s utility in a similar way: “[In the Egyptian revolution] Twitter would often guide you to where an event was taking place” (personal communication, 29 September 2014).

In addition to social media’s utility as a tip-off service it has the potential to offer much more, for example highlighting trending topics, delivering multimedia

content on running stories, and as a searchable archive of contacts and information. These additional capacities have not—yet—been fully exploited.

There are, however, also challenges that come with an increasing reliance on social media as a news source. Perhaps the greatest of these is the veracity of the information it carries. Although Channel 4's Lindsey Hilsum says she follows the "normal system for verifying stuff" when she gets a tip-off from social media ("ringing people you trust, ringing the original source, going and seeing"), to do so can be problematic if the source is anonymous or the news event inaccessible. Furthermore, journalists are increasingly working under severe time pressures and having to deal with large volumes of information, some of which seeks to mislead. For example, Lindsey Hilsum says the information is "frequently contradictory" (personal communication, 15 September 2014). Krishnan Guru-Murthy believes "it can be very misleading, it can be manipulated". Jonathan Rugman agrees that the use of social media as a source has "raised huge issues of verification". BBC News Online's Joe Boyle sums up the problems journalists face by saying that "there are just so many sources out there [on social media] that it's hard to judge what's true and what's not" (personal communication, 16 September 2014).

There are, then, both opportunities to be exploited and challenges to be faced for journalists seeking to make the most of the potential of social media. These are starting to be addressed through the development of technology that seeks to partially automate the identification and verification of news in social media. Although progress has been made, results have been mixed because of the inherent difficulties involved, for example the proportion of "noise" in social media output (i.e. unimportant or malicious content) and fragmentation (people discussing the same topic in different ways). Callison-Burch (n.d.) reports that applying one technique—first story detection (FSD)—to the problem produces "a mass of false positives" with "less than 1 percent of events detected in Twitter" being "news related". He gives an example of one of these false positives, a tweet incorrectly detected as being "news" by this approach:

This wine is going down a lil to smoothly. Here comes trouble.

In terms of verification, although some tools and technologies do exist, journalists believe that they have not been "sufficiently granular to help [them] make judgements on authenticity in a fast-moving news story" (Schifferes et al. 2014, 409).

The SocialSensor approach

The SocialSensor project has made a number of innovations in automating news detection and verification. For example, it monitors a limited number of what it calls "newshounds". These are people—such as journalists, politicians, and bloggers—interested in or expert on particular subjects who share that information with others. Finding and following these newshounds and monitoring what they say was one method it was hoped that newsworthy information could be gathered in a rapid, flexible, and trustworthy way. In SocialSensor a newshounds database is "grown" from an initial "seed" list, for example a manually compiled list of journalists on Twitter.[3] The initial seed list is expanded by including the social media accounts that the "seeds" follow. So, for example, a BBC journalist on one of SocialSensor's seed lists, Lyse Doucet, follows over 2,000 people, so all her followers were added, and so on. As could be expected, any "fully grown" database ends up containing tens or

hundreds of thousands of social media accounts (the putative newshounds). Other research has indicated that there is minimal advantage in monitoring such large numbers of newshounds. For this reason, and because of technical limitations, the number of accounts followed was reduced by applying a scoring mechanism (see table 1), with the 5,000 highest-scoring accounts becoming active newshounds.[4]

Table 1: Scoring mechanism used to select active SocialSensor newshounds

Criteria	Score
Present on our initial "seed" list	**150**
Sends at least 10 tweets per day	**50**
Present on at least 50 Twitter lists	**25**
Verified with Twitter's blue tick	**25**
Score for each account on the initial "seed" list that follows them	**5**
Score for each account on the initial "seed" list that they follow	**2**

Accounts that scored highly but were mostly irrelevant in news terms—such as Justin Bieber—were excluded from the database. Our own testing also confirmed that following more than 5,000 newshounds did not improve significantly the effectiveness with which the software could detect news. By following 5,000 newshounds, we were able to capture between 63 and 90 per cent of the most popular stories taken from a database of tens of thousands of online news publications over a given period (for more information see Hunt et al. [2015]). This indicates how the dispersion of news on social media is likely to be skewed to a log-normal distribution. In other words relatively few, well-informed and well-followed individuals play a key role in the dissemination of news.

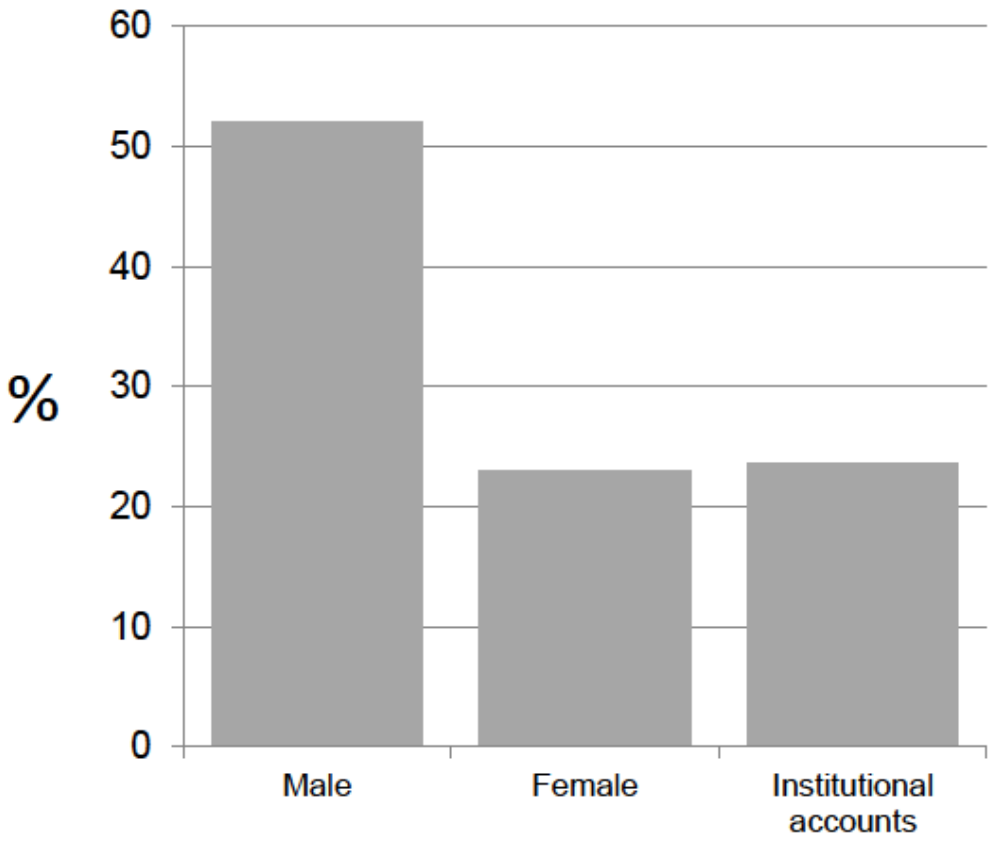
Who has influence?

But who are these active newshounds? If tools like SocialSensor become more widely used, such newshounds will have an important role as gatekeepers. The SocialSensor tool gives users the option of selecting different newshounds databases. The "UK Trends" newshounds database, for example, was seeded with

UK journalists on Twitter, many from the BBC; while the “US Trends” database was seeded with US journalists on Twitter, many from the *New York Times*. For this article the “UK Trends” database of 5,000 newshounds was analysed. A file containing the newshounds’ Twitter handles was ordered by the scoring system outlined in table 1. Three clusters of newshounds, selected stratificationally, were analysed: the top 100, middle 100, and bottom 100. Each newshound’s Twitter profile was examined and categorised by:

1. Type of account holder: e.g. “Female”.
2. Location: e.g. “UK”.
3. Affiliation: e.g. “Politician”.

Figure 1: Type of Twitter account holders in a sample (n=300) of SocialSensor’s “UK Trends” newshounds database



The results show there is a bias towards men (52 per cent of the accounts), with women and institutions (such as the UK Supreme Court) making up the rest with 23 and 24 per cent respectively (see figure 1). There is also a heavy bias towards mainstream media channels and those who work for them—they make up 63 per cent of our sample. Indeed if freelance journalists and bloggers are included, that figure goes up to 74 per cent. Politicians, government spokespeople, and government agencies make up 11 per cent. Experts/academics, celebrities, activists, and consultants make up 7 per cent. Lastly, corporations, NGOs, and public relations firms make up 5 per cent (see figure 2). The fact that the "UK Trends" database we analysed was seeded with journalists made little difference to its final composition. When an alternative "balanced" seed list was used (that included non-journalists such as experts and academics) the resulting database had a similar proportion of journalists. Journalists scored highly due to their activity and popularity, earning inclusion in the final database of active newshounds.

Our analysis of the location of the active newshounds showed that they were, overwhelmingly, UK-based (78 per cent), with a further 15 per cent from other developed countries and only 5 per cent from non-European developing countries (see figure 3).

Figure 2: Affiliation of Twitter account holders in a sample (n=300) of SocialSensor's "UK Trends" newshounds database

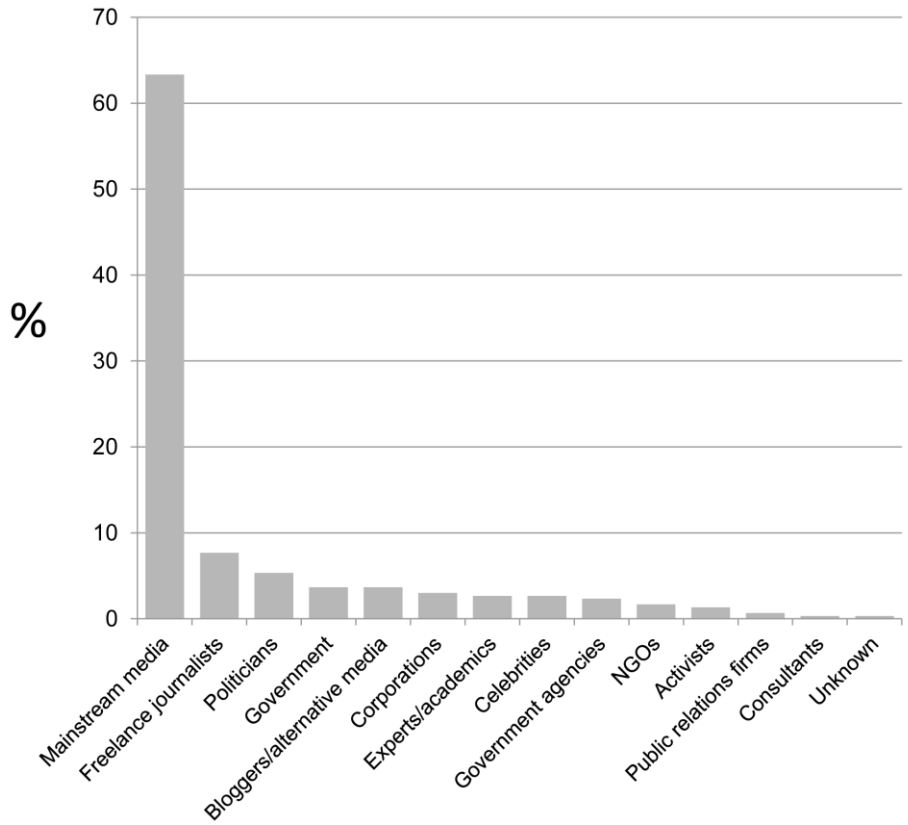
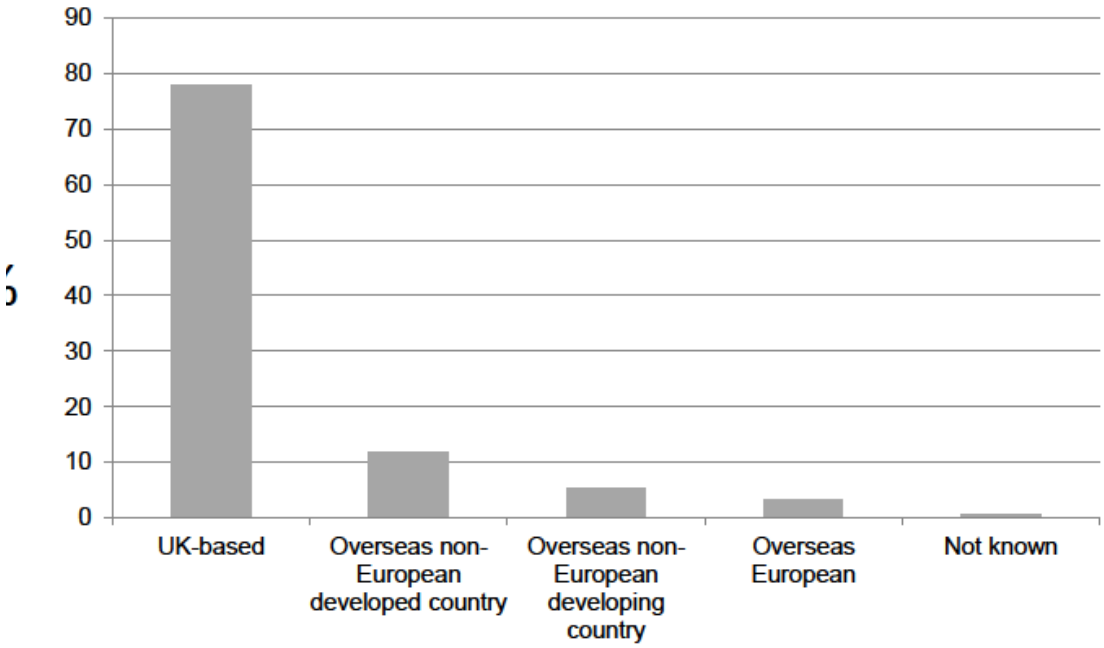


Figure 3: Location of Twitter account holders in a sample (n=300) of SocialSensor's "UK Trends" newshounds database



These results show how tools like SocialSensor rely on particular inputs and assumptions. Although here the selection of newshounds was determined by a scoring system that was blind to gender and, for at least 90 per cent of the newshounds, blind to affiliation and location, we see that there is still a strong bias towards men, towards those in the mainstream media, and towards those in the UK and other developed countries. This raises interesting questions about whether such tools should be tuned to listen, as was the case here, to those being talked about most, to those who do the most talking, and to those with the largest audience, or whether such tools could be agents for change, tuned to also pick up the stories and experiences of those under-represented in the public sphere. We will return to this question in the conclusion.

Our analysis also reveals how individuals can be more influential than institutions. For example, SocialSensor scored Tony Blair's retired media advisor, Alastair Campbell, higher than *The Economist*; and *The Guardian*'s director of digital strategy, Wolfgang Blau, higher than *Le Figaro*.

News discovery and clustering

Although there is no space to go into detail about how SocialSensor automatically identifies news stories, two of its techniques are of interest here. Broadly speaking, topics emerge because they are trending—that is, they are popular over a short period of time—or as a result of users' explicitly expressed (via search) or implicitly inferred (via analysis of their social media profile[5]) interests. Events that are trending (particularly over hours or a few days) are more likely to be detected (Aiello et al.

2013). It is interesting to consider whether the increasing use of tools that work, in part, by detecting such "bursty" events might reinforce relatively passive reporting of short-term events at the expense of longer-term issues like climate change and the economy. We should note here that social media is likely to become even more bursty with the launch of tools such as Thunderclap. Thunderclap deliberately concentrates particular social media messages into bursts so that users can, the company promises, "amplify your message with the power of the crowd" (Thunderclap n.d.). This prefigures a potential arms race between journalists using algorithms to find news and, on the other side, those with messages they want found trying to trick the technology.

Another characteristic of SocialSensor's news detection algorithm is the boost it gives to stories containing proper nouns such as people, places, and organisations (Aiello et al. 2013, 1274). Are such methods necessary to make these tools work? Might they exacerbate the trend towards human-interest stories and the coverage of "elite" people and celebrities? On the other hand, could news detection software like SocialSensor be instructed to consider alternative news values, for example focussing on "political, structural and natural root causes" and providing testimonies by "people concerned" and "positive images of women" (NGO–EC Liaison Committee quoted in Harcup [2015, 41])?

Verification

It is all very well for software to allow journalists to search or be alerted to clusters of social media posts representing news stories, but without context these clusters can be a trap as they are likely to contain a mixture of truth and lies. Twitter, for instance, carries significant amounts of misinformation, especially around breaking news events (see, for example, Burgess et al. [2012]). As a result journalists using social media need to exercise care with the material they are sourcing. This is, of course, already happening. There are some good guidelines (for example Silverman [2014]) and practice out there that focus on verification, often breaking the task down by looking at the content of the message, at the contributor, and the context around the message. An example of such content- or message-based analysis is given by Hermida (2015, 62) who suggests that tweets by baseball player Brett Lawrie from the scene of a shooting looked authentic because of "the language used in the tweets". User- or contributor-based analysis can also be useful. An investigation into a social media contributor purporting to be the Libyan prime minister showed he had previously tweeted that he would make all Libyans "tree hugging hippies". The account, of course, turned out to be fake, but not before it had been quoted by the media (Hermida 2015, 64). The context around social media posts can also be useful in verification. For example, Mendoza et al. (2010) showed that false rumours circulating on Twitter are much more likely to attract denials than confirmed truths.

However, the rising volume of information on social media makes it difficult to do such checking manually, and computers can assess the veracity of contributors and content in a way humans cannot. For example, they have the capacity to be trained using large volumes of historical material and once trained can work very quickly when a story breaks.

SocialSensor's experiments in the area of verification focussed on both content and contributors and looked at both textual and visual material. The focus on visual content was a result of the high priority the project's stakeholders gave to being able to better find relevant multimedia and the serious consequences to news

organisations of publishing fake images. The project started to develop techniques that could spot fake from real images that had been tweeted about actual news events by looking at both the images themselves and the textual content around the images. This was done by training the software using a corpus of fake and real tweets from Hurricane Sandy and the Boston Marathon bombings. Once trained the program was asked to assess a series of tweets it had not previously seen. It did a reasonably good job, achieving around a 75 per cent success rate, but only did so, however, in controlled conditions (Boididou et al. 2014). When more complexity was added (for example if the software was trained on one event and tested on another) its success rate was little better than random guesswork. With better training the success rate should improve; however, these early results go to show there are considerable challenges in automatically giving credibility scores to individual pieces of content on social media.

This is part of the reason why, at the time of writing, the SocialSensor software only gives credibility scores to contributors and not to individual pieces of content. With contributors there is more data to work with, for example analyses of their social media history and their network. In the project's experiments with verification, contributors' initial credibility scores were first computed by looking at their history, their popularity, and their influence, based on these metrics:

- Number of tweets.
- Frequency of tweets.
- Number of Twitter followers.
- Number of follows.
- Number of retweets achieved (Fletcher, Schifferes, and Thurman 2015).

Generally the higher the value for each of these metrics the higher the contributor's initial credibility score.[6]

In order to test this part of the software a different approach was taken than in the experiments with fake Twitter images. This is because contributors can post a mixture of trustworthy and less trustworthy information. It was decided to test the software by asking real—human—journalists to give a score to contributors and then comparing their scores with those given by the software. In the first round of testing the scoring came out pretty close, with mean scores of 5.67 (produced by journalists) and 5.71 (by the software). However, there were significant differences in some of the evaluations, indicated by the standard deviation of the scores produced by the software (2.45) against 2.10 by the journalists (ibid.). For example, there were some contributors who were relatively inactive on Twitter but who were—the human journalists thought—highly credible. In those cases the journalists were ignoring some signals (like inactivity) and paying attention to others (like the high quality of their followers). So the metrics were calibrated—adding new ones and weighting them (see table 2). This resulted in an improved set of scores, although the software remained reluctant to give any contributor eight or nine out of ten on our scale. For example, Harriet Harman, acting leader of the British Labour Party, has been scoring just five in the current version of the software.

Table 2: Weighted metrics used by SocialSensor to give an initial credibility score to social media contributors

Metric	Weight
Number of tweets	**1**
Frequency of tweets	**2**
Number of retweets	**2**
Ratio of followers to followings	**3**
Number of followers	**4**
Popularity*	**5**
Verified by Twitter	**5**

* Defined as: “the number of days since the account was created divided by the number of followers since then” (Fletcher, Schifferes, and Thurman, 2015).

Conclusions

Journalists are drawing on increasing volumes of social media content in their sourcing practices, and, at the same time, we are seeing the emergence of the “digital nose for news”. At least one of the tools that form part of this emergence, SocialSensor, relies on the input of journalists to power many of its processes, and it has been measuring its success by the “ground truth” of journalists’ perceptions and their current output. In that sense it is a tool that has been created in the media’s own image. This is not surprising as SocialSensor was designed explicitly to mimic journalistic judgements on newsworthiness in its discovery of trending news topics. It should be noted, however, that SocialSensor also allows journalists to search across social media for topics that they are particularly interested in or that they are already following closely. Whether they will use this capacity to develop unconventional sources and move beyond established journalism norms remains to be seen.

This article has revealed some of the inputs and instructions SocialSensor relies on. It prioritises stories that show a spike of interest in the short term, and stories about people, places, and organisations; and it listens to those who are vocal and have an audience—many of whom are men in the media. For reasons we have already put forward these characteristics are open to criticism. Is there, then, an

alternative? Well, yes and no. Changes to these characteristics are possible, yes, but may result in disadvantages that outweigh any benefits they bring. For example, the time window through which stories pass could be widened, but with the cost that the algorithm is compromised in its ability to detect breaking news events (SocialSensor 2013), something which is of great interest to news organisations. Stories mentioning people, places, and organisations could have their prioritisation rescinded, but to the detriment of the software's ability to detect effectively the stories being discussed in social media (Aiello et al. 2013). And a different set of newshounds could be monitored. How though would they be chosen and by whom? The transparency and accountability of SocialSensor's selection mechanism—based on social media contributors' popularity, productivity, and ability to engage—has considerable merit. More importantly, we know that the current selection mechanism does actually work, because, as our experiments have confirmed, there appears to be a log-normal distribution of people who disseminate news (see also Hindman [2008] and Lehmann et al. [2013]). The results of the SocialSensor project confirm some of the "contradictions of convergence" (Murdock and Golding 2002, 111): specifically, how social media amplifies and concentrates existing trends in news as much as it enlarges the discourse. The increasing use—whether automated or not—of social media as a news source is no guarantee of a correction to the gender inequalities and insularity of the mainstream media. Rather, we would emphasise the importance of changes to the demography of the journalism profession and to its practices.

In the SocialSensor project we found some of the greatest potential to be in the area of verification—to help combat the significant problem with fake social media content. To conclude we will return to William Lee, the early pioneer of industrial automation we mentioned in our introduction. When he showed his mechanised knitting machine to Queen Elizabeth I, she refused him a patent, saying that "to enjoy the privilege of making stockings for everyone is too important to grant to any individual" (Calvertonvillage.com n.d.). Given the centrality of verification to what we expect of our news media, we might agree that no individual entity—computational or otherwise—should have a monopoly on determining fact from fiction. For now at least, it seems as though software is unlikely to be able to take into account all the nuances of media verification. People, whether we call them journalists or not, will continue to be required to make that final judgement on truth and trust.

Notes

1. https://itunes.apple.com/gb/app/socialsensor/id906630117?mt=8.
2. http://www.socialsensor.eu.
3. Although Twitter was the first social media network mined by SocialSensor, subsequent iterations have incorporated Facebook, YouTube, and Flickr in order to accommodate journalists' desire for multimedia content from a wide range of platforms.
4. The newshounds database is not static but allows for additions and relegations.
5. Users of the software are able to link it to their Twitter account.
6. In the current version of the software, contributors' initial credibility scores are dynamically adjusted up or down depending on whether their contributions appear in the stories detected by SocialSensor.

References

Aiello, Luca Maria, Georgios Petkos, Carlos Martin, David Corney, Symeon Papadopoulos, Ryan Skraba, Ayse Göker, Ioannis Kompatsiaris, and Alejandro Jaimes. 2013. "Sensing Trending Topics in Twitter." *IEEE Transactions on Multimedia* 15 (6): 1268–1282.

Beaumont, Claudine. 2009. "New York Plane Crash: Twitter Breaks the News, Again." *The Telegraph*, 16 January. http://www.telegraph.co.uk/technology/twitter/4269765/New-York-plane-crash-Twitter-breaks-the-news-again.html.

Boididou, Christina, Symeon Papadopoulos, Yiannis Kompatsiaris, Steve Schifferes, and Nic Newman. 2014. "Challenges of Computational Verification in Social Multimedia." In *Proceedings of the 23rd International Conference on World Wide Web (WWW '14 Companion)*, 743–748. Geneva, Switzerland: International World Wide Web Conferences Steering Committee.

Broersma, Marcel, and Todd Graham. 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.

Burgess, Jean, Farida Vis, and Axel Bruns. 2012. "How Many Fake Sandy Pictures Were Really Shared on Social Media?" *The Guardian*, 6 November. http://www.theguardian.com/news/datablog/2012/nov/06/fake-sandy-pictures-social-media.

Callison-Burch, Chris. N.d. "Distilling Collective Intelligence from Twitter." Crowdsourcing-class.org. Accessed 6 August 2015. http://crowdsourcing-class.org/slides/twitter-first-story-detection.pdf.

Calvertonvillage.com. N.d. "William Lee." Accessed 22 August 2015. http://www.calvertonvillage.com/Willlee.html.

Diplaris, Sotiris, Symeon Papadopoulos, Ioannis Kompatsiaris, Nicolaus Heise, Jochen Spangenberg, Nic Newman, and Hakim Hacid. 2012. "'Making Sense of It All': An Attempt to Aid Journalists in Analysing and Filtering User Generated Content." In *Proceedings of the 21st International Conference on World Wide Web (WWW '12 Companion)*, 1241–1246. New York: ACM.

Fletcher, Richard, Steve Schifferes, and Neil Thurman. 2015. "Evaluating the Truthmeter: Improving Automated Social Media Contributor Credibility Assessments within the Context of Journalism." Paper presented at the *Artificial Intelligence, Robots and Media Conference*, University of Dubrovnik, Dubrovnik, Croatia, 30–31 October.

Frey, Carl Benedikt, and Michael A. Osborne. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerization?" Oxford Martin School, University of Oxford. http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.

Geofeedia.com. 2015. "Location-Based Media & Journalism Social Media Intelligence." https://geofeedia.com/solutions/media-journalism-social-media-tool/.

Harcup, Tony. 2015. *Journalism: Principles and Practice*. 3rd ed. London: Sage Publications.

Hermida, Alfred. 2015. "Filtering Fact from Fiction: A Verification Framework." In *Ethics for Digital Journalists: Emerging Best Practices*, edited by Lawrie Zion and David Craig, 59–73. New York: Routledge.

Hindman, Matthew. 2008. *The Myth of Digital Democracy*. Princeton: Princeton University Press.

Hunt, Stephen, Steve Schifferes, Neil Thurman, Nic Newman, Richard Fletcher, and David Corney. 2015. "Auto-Detection of News on Twitter: Tuning and Testing the SocialSensor App." Paper presented at the *Artificial Intelligence, Robots and Media Conference*, University of Dubrovnik, Dubrovnik, Croatia, 30–31 October.

Lehmann, Janette, Carlos Castillo, Mounia Lalmas, and Ethan Zuckerman. 2013. "Finding News Curators in Twitter." In *Companion Publication of the IW3C2 WWW 2013 Conference*, 863–869.

Mendoza, Marcelo, Barbara Poblete, and Carlos Castillo. 2010. "Twitter Under Crisis: Can We Trust What We RT?" In *Proceedings of the First Workshop on Social Media Analytics*, 71–79. New York: ACM.

Moon, Soo Jung, and Patrick Hadley. 2014. "Routinizing a New Technology in the Newsroom: Twitter as a News Source in Mainstream Media." *Journal of Broadcasting & Electronic Media* 58 (2): 289–305.

Murdock, Graham, and Peter Golding. 2002. “Digital Possibilities, Market Realities: The Contradictions of Communications Convergence.” *Socialist Register* 38: 111–129.

Newswhip.com. N.d. “Pro Tools.” http://www.newswhip.com/pro-tools#spike.

Paulussen, Steve, and Raymond A. Harder. 2014. “Social Media References in Newspapers: Facebook, Twitter and YouTube as Sources in Newspaper Journalism.” *Journalism Practice* 8 (5): 542–551.

Schifferes, Steve, Nic Newman, Neil Thurman, David Corney, Ayse Göker, and Carlos Martin. 2014. “Identifying and Verifying News through Social Media: Developing a User-Centred Tool for Professional Journalists.” *Digital Journalism* 2 (3): 406–418.

Silverman, Craig, ed. 2014. *Verification Handbook*. Maastricht: European Journalism Centre.

SocialSensor. 2013. “DySCO.” 23 August. http://www.socialsensor.eu/project/dyscos.

Thunderclap. N.d. “Homepage.” https://www.thunderclap.it/?locale=en.

Wallsten, Kevin. 2015. “Non-Elite Twitter Sources Rarely Cited in Coverage.” *Newspaper Research Journal* 36 (1): 24–41.