Research Challenges of Digital Misinformation: Toward a Trustworthy Web

Giovanni Luca Ciampaglia, Alexios Mantzarlis,
Gregory Maus, Filippo Menczer

■ *The deluge of online and offline misinformation is overloading the exchange of ideas upon which democracies depend. Fake news, conspiracy theories, and deceptive social bots proliferate, facilitating the manipulation of public opinion. Countering misinformation while protecting freedom of speech will require collaboration across industry, journalism, and academe. The Workshop on Digital Misinformation — held in May 2017, in conjunction with the International AAAI Conference on Web and Social Media in Montreal — was intended to foster these efforts. The meeting brought together more than 100 stakeholders from academia, media, and tech companies to discuss the research challenges implicit in building a trustworthy web. In this article, we outline the main findings from the discussion.*

The deluge of online misinformation is overloading the exchange of ideas upon which democracies depend. Many have argued that echo chambers are increasingly constricting the ability of alternative perspectives to provide a check on one's viewpoints (Sunstein 2009; Pariser 2011; Sunstein 2002). Suffering fragmentation and declining public trust, the Fourth Estate struggles to carry out its traditional editorial role of distinguishing fact from fiction (Pew Research Center 2017). Within this context, fake news, conspiracy theories, and deceptive social bots proliferate, facilitating the manipulation of public opinion.

Countering misinformation while protecting freedom of speech will require collaboration between stakeholders across the tech industry, journalism, and academia. To foster such collaboration, the Workshop on Digital Misinformation was held in conjunction with the International Conference on Web and Social Media (ICWSM) in Montreal, on May 15, 2017. The meeting brought together more than 100 stakeholders from academe, media, and tech companies to discuss research challenges toward a trustworthy web.

The workshop opened with a showcase of tools for studying digital misinformation, developed by the Center for Complex Networks and Systems Research and the IU Network Science Institute. These are part of a growing suite of publicly available tools called the Observatory on Social Media (Davis, Ciampaglia, et al. 2016). They include Hoaxy, a system for tracking competing claims and fact-checking that spread online (Shao et al. 2017) (figure 1); Botometer, an AI system for detecting social bots on Twitter (Davis, Varol, et al. 2016; Varol, Ferrara, Davis, et al. 2017) (figure 2); and unsupervised graph-mining methods for automatically estimating factual accuracy from DBpedia (Shiralkar et al. 2017) (figure 3). We then presented empirical data showing that, on social media, low-quality information often spreads more virally than high-quality information (Qiu et al. 2017). Factors that can explain this finding include the structural segregation and polarization in online social networks (Conover et al. 2011). The resulting echo chambers are exacerbated by algorithms that personalize online experiences and hinder exposure to ideologically diverse sources of information (Nikolov et al. 2015). Other factors include information overload (Qiu et al. 2017), limited attention (Weng et al. 2012), popularity bias (Nematzadeh et al. 2017), and manipulation through social bots (Ferrara et al. 2016). One of the key questions raised during discussion was how to empirically define quality of information in modern social media.

A lightning talk session was opened by BuzzFeed media editor Craig Silverman, who proposed a working definition of fake news as "fabricated news intended to deceive with financial motive." Silverman emphasized the global scope of the issue, describing similar problems in Germany, Japan, Italy, and Myanmar, even if the social networks used to spread such news vary. Journalists are now aware of the scope of the problem and are reporting on fake news stories — though not always well. Silverman called for more collaboration between journalists and academic researchers for the study of misinformation. In support of this call to action, his team at Buzzfeed has developed a curated list of partisan news sites that is openly shared (Silverman 2017).

Political scientist and communication scholar Leticia Bode (Georgetown University) reported on a number of findings about correcting social media misinformation. Her work focuses on health and science communication, where the distinction between factual and opinionated claims is more clear compared to other domains such as political communication (Bode and Vraga 2015). Bode and her collaborators found that certain topics (for example, GMOs) are easier to correct than others (for example, vaccines and autism). She also found that "social" fact-checking is more effective if it links to credible sources. Based on these findings, Bode recommended that news organizations should emphasize easily linked references and that corrections should be early and repeated. A new partnership model with social media platforms could satisfy these requirements, she concluded.

The perspective of a leading social media platform was given by Áine Kerr, leader of global journalism partnerships at Facebook. She started by sharing figures to appraise the scale at which Facebook's newsfeed operates, with hundreds of millions of links shared weekly. Kerr noted that the quality of those links varies dramatically, and quoted Mark Zuckerberg's call for amplifying the good effects of social media and mitigating the bad. Facebook is pursuing this goal with four approaches: (1) disrupting the financial incentives for fake news; (2) developing new products to curb the spread of fake news, such as allowing users or third-party fact-checkers to flag posted stories as untrue or unverified; (3) helping people make informed decisions by educating them on how to spot fake news; and (4) launching the News Integrity Initiative, a partnership between industry and nongovernmental organizations to promote media literacy (Mosseri 2016). Kerr noted that Facebook regularly engages with the research community via collaborative programs and grants, but acknowledged that there is a growing demand from third-party researchers for data to tackle the above problems. In its own bid to meet the problem head-on, the organization is constantly refining its best practices around data sharing. A lively discussion followed about the commitment of platforms to curbing the spread of misinformation. For example, it was pointed out that more should be done to deal with abuses that exploit social bots, fake accounts, Facebook pages, and "verified account" badges. It was also suggested that an API for access to public Facebook page data would be a great boon to the research community.

Computer scientist Paul Resnick (University of Michigan) argued that factual corrections are often ineffective and slow: they rarely reach the people originally influenced by the misinformation. Moreover, we are exposed to information filtered by sociotechnical mechanisms that largely prioritize popularity over accuracy, such as search engines, upvotes, and newsfeeds. To restore the balance in favor of accuracy, Resnick called for the development of reputation-based filtering mechanisms, and reported on
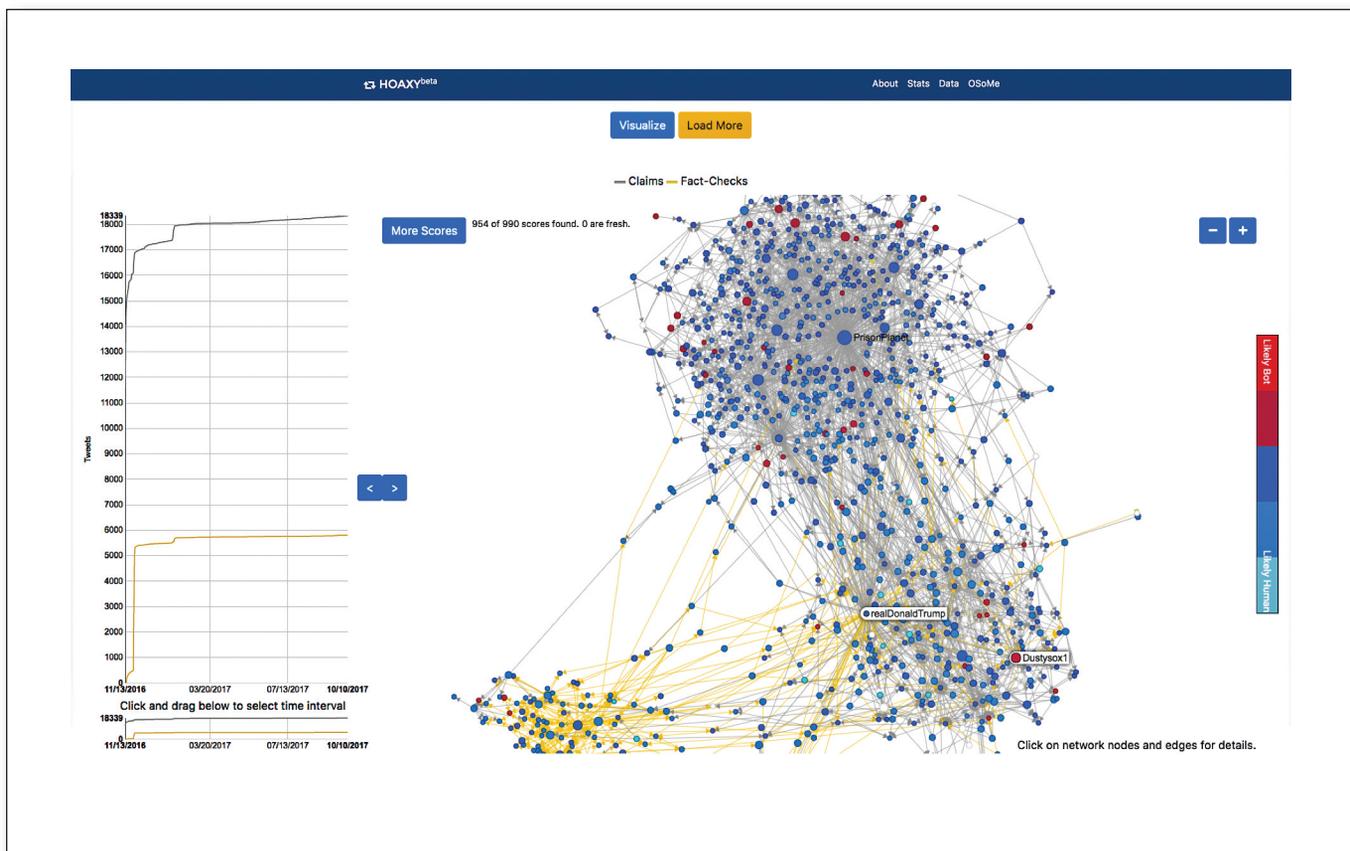
Figure 1. Screenshot of the Hoaxy System.

This is a visualization of the spread of fake news (purple) and fact-checking (orange) articles related to the query "Three Million Votes Illegal Aliens" on Twitter. The left panel plots the number of tweets with links to these articles over time. The right panel shows a diffusion network in which nodes represent Twitter accounts and edges represent the spread of the articles via retweets and mentions. (hoaxy.iuni.iu.edu).

his own work on gauging the trustworthiness of sources using a prediction market (Resnick and Sami 2008). In this scenario, third parties would verify the content shared by the sources, whose credibility would be automatically reevaluated in light of new verifications. Thus, those who consistently share accurate information would have greater influence than those who do not. Resnick argued that a more widespread adoption of reputation-based filtering technologies would lead to more accurate content being shared.

Data scientist Rolf Fredheim (NATO StratCom Centre of Excellence) discussed the recent spate of misinformation operations in Europe (especially Russian-sponsored ones) and the problems created by the decreasing credibility of governments and the media in the eyes of the public, as well as the increasing use of deceptive social bots (Fredheim 2017). He noted that while governments want a quick fix to the problem, comprehensive approaches are required, including adjusting the incentives for news production and consumption as well as providing education and awareness of misinformation. Fredheim insisted that social media need to be considered seriously as a tool of soft power and thus social media companies should be pressured to police their platforms.

How to Define and Detect Misinformation?

A first panel discussed the research challenge of defining and detecting misinformation in online settings. Computer scientist Jisun An (Qatar Computing Research Institute) demonstrated the economic threats of misinformation by discussing a 2013 incident in which hackers gained control of the Twitter feed of the Associated Press and used it to falsely announce an attack on President Obama, causing an immediate plunge in stock market prices. Computer scientist Meeyoung Cha (Korea Advanced Institute of Science and Technology) defined misinformation as information unverified at the time of circulation and later determined to be false, and proposed to employ AI algorithms for its detection (Zhao, Resnick, and
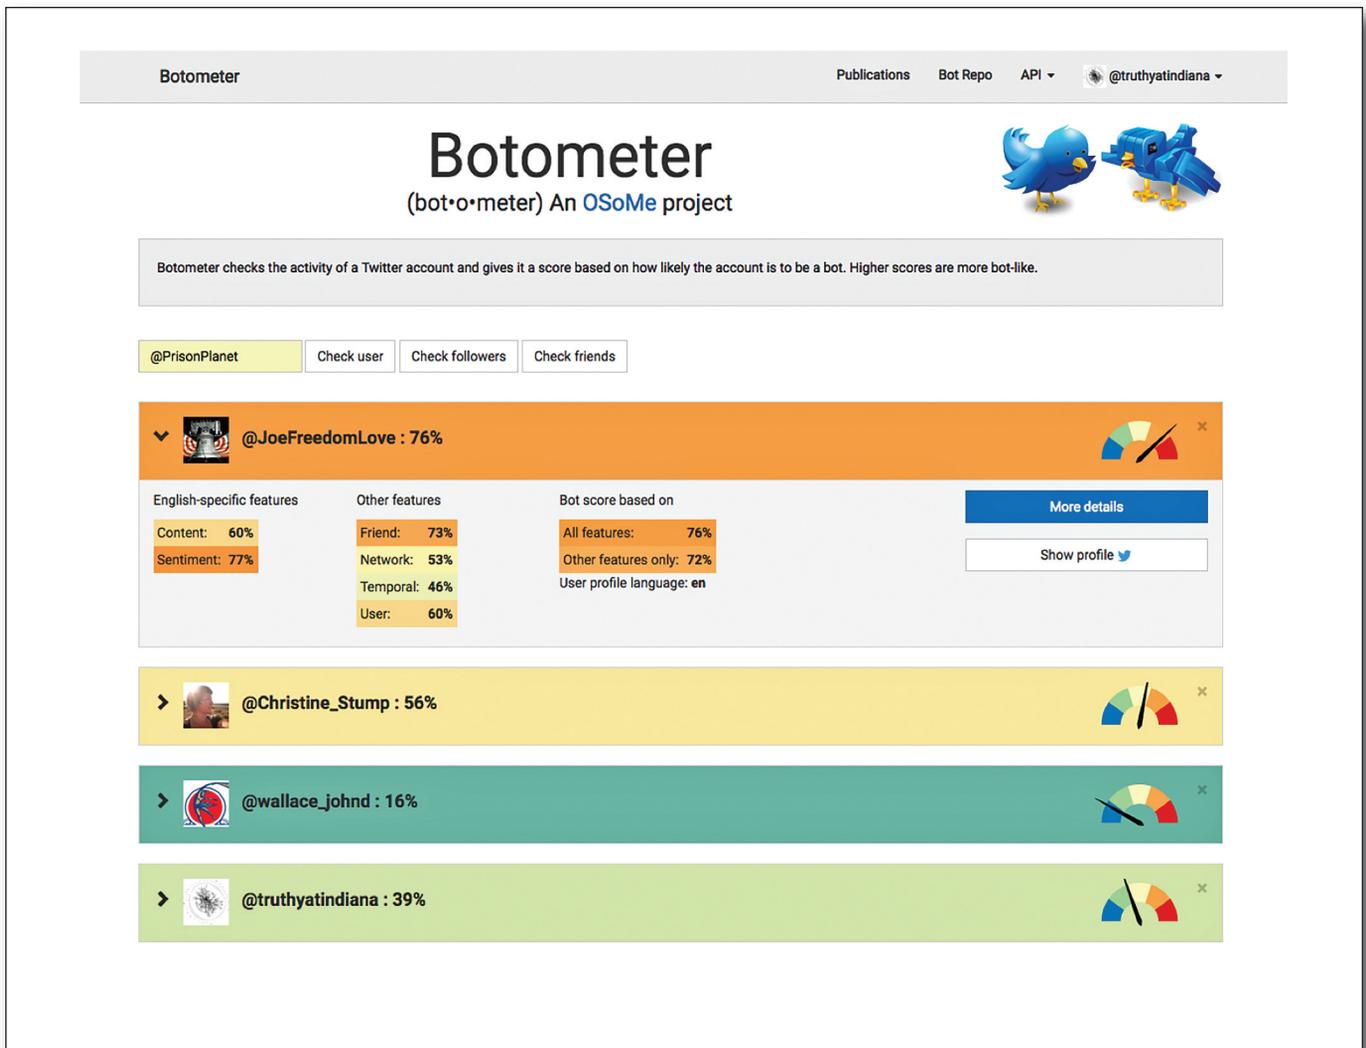
Figure 2. Screenshot of the Botometer System.

This is a supervised learning framework that calculates the likelihood that a
given Twitter account is controlled by software, that is, a social bot (botometer.iuni.iu.edu).

Mei 2015; Ma et al. 2016). More theoretical work is necessary to understand why people spread rumors, as this is less well understood than the incentives for creating rumors, particularly the financial ones.

Several statistical approaches for detecting misinformation require labeled instances to train machine learning models. Computer scientist James Caverlee (Texas A&M University) reported on the critical issue of producing such ground truth. He started by describing intuitive identification methods, such as the one according to which those users friending a fake account are themselves deserving of some scrutiny. He argued that while powerful, these simple heuristics can only address a fraction of the issue. When it comes to more complex claims, fact-checkers are key to identifying the ground truth. However, they cannot cope with the sheer volume and variety of misinformation. Caverlee proposed exploiting

aggregated signals to infer the reliability of a given piece of content, such as the reply-to-retweet ratio to flag controversy (Alfifi and Caverlee 2017; Kaghazgaran, Caverlee, and Alfifi 2017). Caverlee warned that crowdsourced information can be easily manipulated; for example, it is easy to recruit workers for astroturfing. He noted that the Chinese government often fabricates social media posts in an attempt to prevent the public from discussing civic issues. Caverlee called for more research on the problem of identifying the intent behind social media posts. Progress in this area could lead to tools for distinguishing organic conversations from covert coordinated campaigns on social media (Varol, Ferrara, Menczer, et al. 2017).

Computer scientist Qiaozhu Mei (University of Michigan) highlighted how hacked accounts can bypass reputation systems, challenging solutions
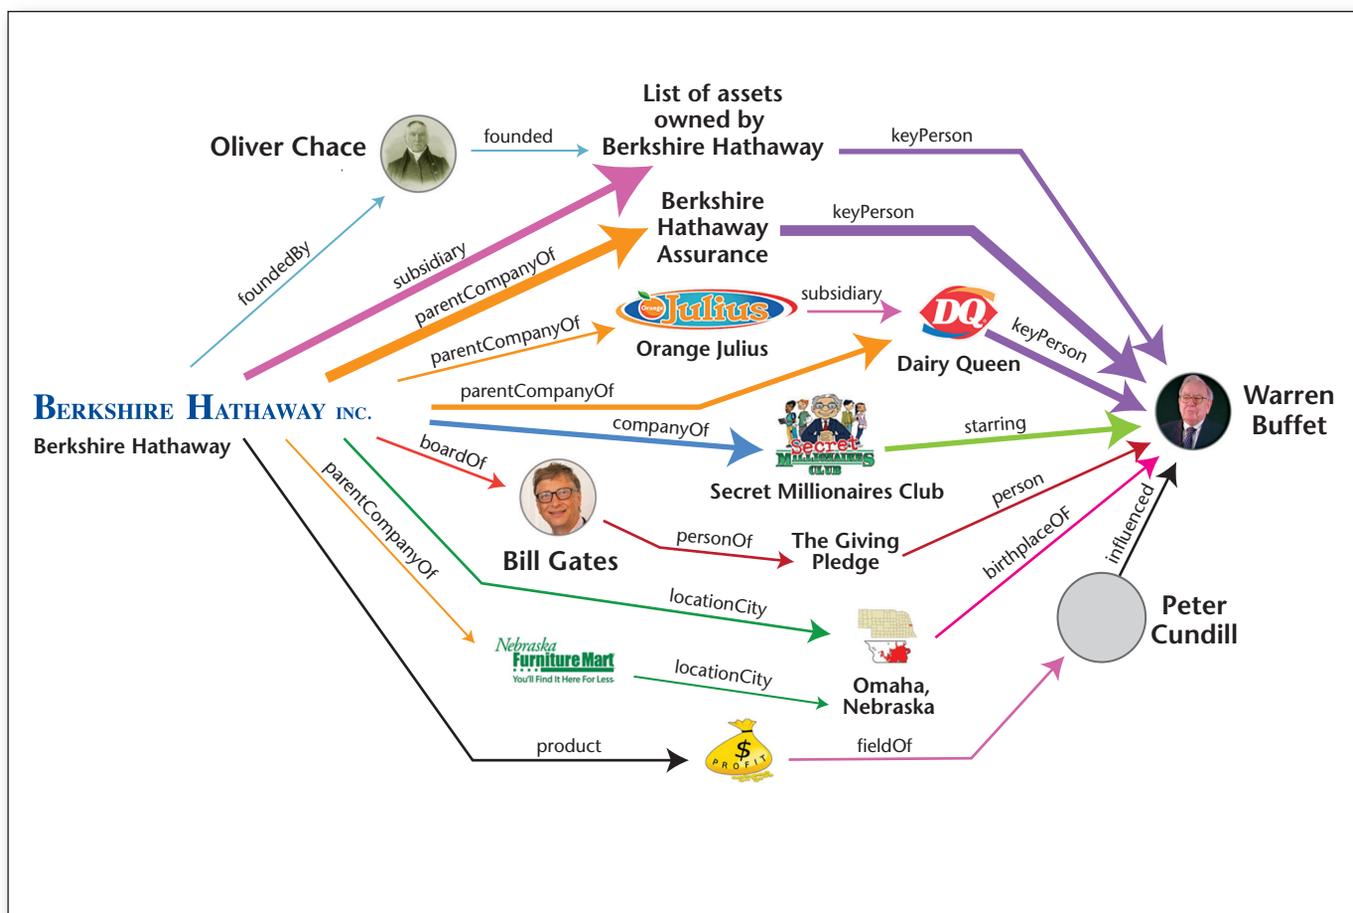
Figure 3. An Example of Connections Identified by a Computational Fact-Checking Algorithm to Verify a Claim about a Company and Its CEO.

A flow algorithm is used to identify relational paths in a knowledge network. The width of an edge is proportional to the flow.

suggested by previous speakers. He noted that while spotting known rumors is easy after the fact, identifying false rumors in their early stages is difficult. He proposed a four-part approach: (1) identifying emerging rumors early by mining search engine logs for questions about a statement's truthfulness; (2) retrieving all the posts about a rumor; (3) analyzing the impact of a rumor and its correction via visualization; and (4) predicting its future spread through deep learning (Zhao, Resnick, and Mei 2015; Li et al. 2017; Qazvinian et al. 2011; Resnick et al. 2014).

Computer scientist Eni Mustafaraj (Wellesley College) discussed an early example of fake news tactics, namely, a Twitter bomb against Massachusetts senate candidate Martha Coakley in 2010 (Mustafaraj and Metaxas 2010). A group of Twitter bots used hashtags targeted towards specific communities in a coordinated campaign to spread negative information about the candidate. Mustafaraj compared this attack with recent ones, leveraging fake Facebook accounts to target specific Facebook groups and spread links to fake news stories (Mustafaraj and Metaxas 2017). She

examined three motives for spreading such inaccurate information: financial, political, and ideological/cultural (including prejudices like sexism and xenophobia). She proposed that social media platforms should highlight provenance information about sources to help users determine their intents and trustworthiness. Finally, Mustafaraj urged platforms to provide researchers with data about how recipients of misinformation engage with it.

Information scientist Kate Starbird (University of Washington) delineated different types of rumor in a crisis situation, such as hoaxes, collective sense-making, and conspiracy theories (Arif et al. 2016). She categorized methods for detecting misinformation based on linguistic, network, account, URL domain, and crowdsourced features (Maddock et al. 2015; Huang et al. 2015; Dailey and Starbird 2014).

Workshop participants responded to an online wiki survey and identified five major research challenges related to the question of how to define and detect misinformation (table 1).

Survey of Participant Priorities by Panel
Panel 1: Defining and Detecting Misinformation
1. Identifying and promoting reliable information instead of focusing on disinformation
2. Tracking variants of debunked claims
3. Developing reputation scores for publishers
4. Creating an automated trustmark to promote journalistic integrity
5. Collecting reliable crowdsourcing signals
Panel 2: Cognitive, Social, and Technological Biases Creating Vulnerabilities to Misinformation
1. Investigating the use of language, images, and design in misinformation persuasiveness
2. Validating model predictions via field experiments
3. Studying the roles of algorithmic mechanisms in the spread of misinformation
4. Translating research findings into policy recommendations
5. Accessing behavioral data from social media platforms
Panel 3: Countermeasure Feasibility, Effectiveness, and Responsibility
1. Support and scaffold critical thinking
2. Increase prominence and availability of fact-checking information
3. Design trust & reputation standards for news sources and social media users
4. Build tools to track the provenance of digital content
5. Develop computational tools to support fact-checking

Table 1. Top Priorities for Each Topic.

After each panel, participants in the workshop proposed and chose top priorities
for each topic via a dynamic online head-to-head ranking system. This table records the top five results for each.

How to Best Study the Cognitive, Social, and Technological Biases that Make Us Vulnerable to Misinformation?

A second panel addressed research challenges in studying cognitive, social, and technological biases that create vulnerabilities to misinformation. Stanford University historian Joel Breakstone cited a study showing that students have difficulty distinguishing the reliability of news stories (McGrew et al. 2017). In this study, students' assessments of the veracity of the content they were presented were found to rely primarily on the appearance of the article rather than on consideration of the source. The same study also found that students were unable to distinguish native advertising from real news stories 80 percent of the time. Breakstone concluded that the fight against misinformation is one in which we all must take part, not just big tech companies.

Communications scholar R. Kelly Garrett (Ohio State University) sought to clarify several points in the discussion. First, he distinguished between holding a belief and being ignorant of the evidence against it, citing statistics on the number of people who know the scientific consensus on global warming but reject it (Funk and Kennedy 2016). However, he noted that online partisan news leads people to reject evidence, due to these outlets' emotional pull (Garrett, Weeks, and Neo 2016). While social media increase the profile of misinformation, it remains unclear how much this actually shifts public opinions.

Physicist Kristina Lerman (University of Southern California) emphasized that humans have limited information-processing capabilities, which makes it impossible to keep up with the growing volume of information, resulting in reliance on simple cognitive heuristics. These heuristics may in turn amplify certain cognitive biases. Lerman reported on her studies about popularity, engagement, and position bias (Kang and Lerman 2015; Lerman and Hogg 2014). The latter is the idea that people pay more attention to what is at the top of a list or screen. Experimental trials show that news stories at the top of a list are four to five times more likely to be shared. Lerman also discussed how social media reinforce network biases, creating echo chambers that distort our perceptions.

Economist David Rothschild (Microsoft Research) proposed to address misinformation as a market problem, framing it in terms of outcomes such as exposure to information and impact on opinion formation and decision making. Rothschild noted that research in the field tends to focus on what content people consume rather than the more difficult matter of how they actually absorb that information. He also questioned whether mass ignorance on a particular issue may be more harmful than the consumption of fake news about that issue. Research may also be distorted by overreliance on Twitter data, Rothschild suggested, which may not be as representative of news consumption by the general population as Facebook or television (Diaz et al. 2016).

Computer scientist Kazutoshi Sasahara (Nagoya University) presented work in progress on a simple model of online social network dynamics. The model demonstrates that online echo chambers are inevitable within the current social media mechanisms for content sharing, which tend to cluster individuals into segregated and polarized groups.

Workshop participants responded to an online wiki survey and identified five major research challenges related to the question of how to best study the cognitive, social, and technological biases that make us vulnerable to misinformation (see table 1).

What Countermeasures Are Most Feasible and Effective and Who Can Best Deliver Them?

The final panel discussed countermeasures against misinformation, as well as who can best deliver them. Computational journalist Nick Diakopoulos (University of Maryland) identified three relevant actors to be considered: tech platforms, individuals, and civil society. He discussed which combinations of the three groups could be most effective in combating fake news. Platforms are particularly powerful, but they raise the concern of influencing public discourse through their algorithms — although this could be mitigated through algorithmic transparency (Diakopoulos 2017; Diakopoulos and Koliska 2017). Civil society and individuals alone cannot fact-check everything. Diakopoulos concluded that the best partnership would be between civil society and platforms.

Fact-checker David Mikkelson (Snopes.com) argued that fake news is only as problematic as poor journalism: the very news outlets that are supposed to question and disprove misinformation often help spread it (Mikkelson 2016).

Computer scientist Tim Weninger (University of Notre Dame) addressed the lack of research on Reddit, a far larger platform than Twitter. He reported on findings that initial votes on posts have a strong impact on their final visibility, allowing coordinated attacks to game the system through a snowballing effect (Glenski and Weninger 2017). Weninger also found that a large portion of Reddit users merely scan headlines: most up or down votes are cast without even viewing the content (Glenski, Pennycuff, and Weninger 2017).

Computer scientist Cong Yu (Google Research) described the use of semantic web annotations such as the Schema.org ClaimReview markup to help surface fact-checks of popular claims on search engines (Kosslyn and Yu 2017). He argued that artificial intelligence can be a powerful tool to promote quality and trust in information. However, Yu recognized that users play a role in the spread of misinformation, which may be the most challenging problem to address.

Communication scientist Melissa Zimdars (Merrimack College) reported on efforts to collectively categorize news sources in an open-source fashion (opensources.co). She also recounted how, ironically, her research became the target of a fake news campaign.

Workshop participants responded to an online wiki survey and identified five major research challenges related to the question of what countermeasures are most feasible/effective and who can best deliver them (see table 1).

Conclusions

Unfortunately, AI is increasingly being exploited to manipulate public opinion. For example, sophisticated social bots can autonomously interact with social media users in an attempt to influence them or expose them to misinformation. Advances in machine generation of realistic video and voice have been already identified as the likely next-generation weapons in the digital misinformation arsenal (Suwajanakorn, Seitz, and Kemelmacher-Shlizerman 2017; Thies et al. 2016).

The good news is that AI can also play an important role in defending us from attacks against the integrity of the information space. In such an arms race, advances in supervised and unsupervised machine learning, representation learning, and natural language processing will be needed to help meet the above challenges.

Another area where more AI research is needed is the study of algorithmic bias. Social media platforms employ sophisticated ranking, filtering, and recommendation methods that are increasingly powered by cutting-edge AI algorithms. Unfortunately, these algorithms are also vulnerable to manipulation due to their focus on engagement and popularity, leading to echo chambers and selective exposure that amplify our own cognitive and social biases. A significant challenge will be to improve algorithms to take into account signals of trustworthiness and reliability.

Finally, reporters and fact-checking organizations are in great need of tools to help them manage the

volume of digital misinformation at scale. Semantic web technologies are reaching maturity and can be leveraged toward the development of computational fact-checking tools. Automatic extraction of semantic data from text provides large knowledge bases which, in combination with inference and network mining techniques, could accelerate the tasks of verification and news gathering, for example by connecting claims with relevant contextual information and previous analyses (see figure 3). Special emphasis should be placed upon interpretable results, which can aid both journalists and the general public make sense of the information to which they are exposed.

The ultimate goal of the workshop was to bootstrap a long-lasting initiative between various sectors (namely industry, academe, journalism, and civil society) with the aim of building a trustworthy web. Follow-up discussion and further collaborative activity is currently under way through online community spaces. We believe that support from both private foundations and federal agencies will be a key ingredient for the success of future collaborative activities, the scope of which must include research, education, and policymaking.

Acknowledgements

References

Alfifi, M., and Caverlee, J. 2017. Badly Evolved? Exploring Long-Surviving Suspicious Users on Twitter. In *Proceedings of SocInfo 2017, International Conference on Social Informatics.* Lecture Notes in Computer Science 10539, 218–233. Berlin: Springer. doi.org/10.1007/978-3-319-67217-5_14

Arif, A.; Shanahan, K.; Chou, F.-J.; Dosouto, Y.; Starbird, K.; and Spiro, E. 2016. How Information Snowballs: Exploring the Role of Exposure in Online Rumor Propagation. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (CSCW'16), 466–477. New York: Association for Computing Machinery. doi.org/10.1145/2818048.2819964

Bode, L., and Vraga, E. K. 2015. In Related News, That Was Wrong: The Correction of Misinformation Through Related Stories Functionality in Social Media. *Journal of Communication* 65(4):619–638. doi.org/10.1111/jcom.12166

Conover, M.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Menczer, F.; and Flammini, A. 2011. Political Polarization on Twitter. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media* (ICWSM), 89–96. Palo Alto, CA: AAAI Press. www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2847

Dailey, D., and Starbird, K. 2014. Journalists as Crowdsourcerers: Responding to Crisis by Reporting with a Crowd. *Computer Supported Cooperative Work* 23(4–6): 445–481. doi.org/10.1007/s10606-014-9208-z

Davis, C. A.; Ciampaglia, G. L.; Aiello, L. M.; Chung, K.;

Conover, M. D.; Ferrara, E.; Flammini, A. 2016. OSoMe: The IUNI Observatory on Social Media. *PeerJ Computer Science* 2(October):e87. doi.org/10.7717/peerj-cs.87

Davis, C. A.; Varol, O.; Ferrara, E.; Flammini, A.; and Menczer, F. 2016. BotOrNot: A System to Evaluate Social Bots. In *WWW '16 Companion: Proceedings of the 25th International Conference Companion on World Wide Web*, 273–274. Cambridge, MA: W3C. doi.org/10.1145/2872518.2889302

Diakopoulos, N. 2017. Enabling Accountability of Algorithmic Media: Transparency as a Constructive and Critical Lens. In *Transparent Data Mining for Big and Small Data*, eds. T. Cerquitelli, D. Quercia, F. Pasquale, 25–43. Studies in Big Data 32. Berlin: Springer. doi.org/10.1007/978-3-319-54024-5_2

Diakopoulos, N., and Koliska, M. 2017. Algorithmic Transparency in the News Media. *Digital Journalism* 5(7): 809–828. doi.org/10.1080/21670811.2016.1208053

Diaz, F.; Gamon, M.; Hofman, J. M.; Kıcıman, K.; and Rothschild, D. 2016. Online and Social Media Data as an Imperfect Continuous Panel Survey. *PLoS ONE* 11(1): e0145406. doi.org/10.1371/journal.pone.0145406

Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2016. The Rise of Social Bots. *Communications of the ACM* 59(7): 96–104. doi.org/10.1145/2818717

Fredheim, R. 2017. *Robotrolling*. Riga, Latvia: NATO Strategic Communications Centre of Excellence. (www.stratcom-coe.org/robotrolling-20171)

Funk, C., and Kennedy, B. 2016. *The Politics of Climate*. Washington, DC: Pew Research Center. (www.pewinternet.org/2016/10/04/the-politics-of-climate)

Garrett, R. K.; Weeks, B. E.; and Neo, R. L. 2016. Driving a Wedge Between Evidence and Beliefs: How Online Ideological News Exposure Promotes Political Misperceptions. *Journal of Computer-Mediated Communication* 21(5): 331–348. doi.org/10.1111/jcc4.12164

Glenski, M.; Pennycuff, C.; and Weninger, T. 2017. Consumers and Curators: Browsing and Voting Patterns on Reddit. *IEEE Transactions on Computational Social Systems* 4(4): 196–206. doi.org/10.1109/TCSS.2017.2742242

Glenski, M., and Weninger, T. 2017. Rating Effects on Social News Posts and Comments. *ACM Transactions on Intelligent Systems and Technology* 8(6): 78. doi.org/10.1145/2963104

Huang, Y. L.; Starbird, K.; Orand, M.; Stanek, S. A.; and Pedersen, H. T. 2015. Connected Through Crisis: Emotional Proximity and the Spread of Misinformation Online. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW'15). New York: Association for Computer Machinery. doi.org/10.1145/2675133.2675202

Kaghazgaran, P.; Caverlee, J.; Alfifi, M. 2017. Behavioral Analysis of Review Fraud: Linking Malicious Crowdsourcing to Amazon and Beyond. In *Proceedings of the 11th International AAAI Conference on Web and Social Media* (ICWSM 2017), 560–563. Palo Alto, CA: AAAI Press. (www.aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15659)

Kang, J. H., and Lerman, K. 2015. VIP: Incorporating Human Cognitive Biases in a Probabilistic Model of Retweeting. In *Proceedings of the International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction* (SBP 2015). Lecture Notes in Computer Science 9021, 101–110. Berlin: Springer. doi.org/10.1007/978-3-319-16268-3_11

Kosslyn, J., and Yu, C. 2017. Fact Check Now Available in

Google Search and News around the World. Google Blog, April 7. (www.blog.google/products/search/fact-check-now-available-google-search-and-news-around-world)

Lerman, K., and Hogg, T. 2014. Leveraging Position Bias to Improve Peer Recommendation. *PLoS ONE* 9(6): e98914. doi.org/10.1371/journal.pone.0098914

Li, C.; Ma, J.; Guo, X.; and Mei, Q. 2017. DeepCas: An End-to-End Predictor of Information Cascades. In *Proceedings of the 26th International Conference on World Wide Web* (WWW'17), 577–586. Cambridge, MA: W3C. doi.org/10.1145/3038912.3052643

Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B. J.; Wong, K.-F.; and Cha, M. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In P*roceedings of the 25th International Joint Conference on Artificial Intelligence* (IJCAI '16), 3818–3824. Palo Alto, CA: AAAI Press. www .ijcai.org/Abstract/16/537

Maddock, J.; Starbird, K.; Al-Hassani, H. J.; Sandoval, D. E.; Orand, M.; Mason, R. M. 2015. Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (CSCW'15), 228–241. New York: Association for Computing Machinery. doi.org/10.1145/2675133.2675280

McGrew, S.; Ortega, T.; Breakstone, J.; and Wineburg, S. 2017. The Challenge That's Bigger Than Fake News. *American Educator* (Fall): 4–9.

Mikkelson, D. 2016. We Have a Bad News Problem, Not a Fake News Problem. Snopes.com, November 17. (www.snopes.com/2016/11/17/we-have-a-bad-news-problem-not-a-fake-news-problem)

Mosseri, A. 2016. News Feed FYI: Addressing Hoaxes and Fake News. Menlo Park, CA: Facebook. (newsroom.fb.com/news/2016/12/news-feed-fyi-addressing-hoaxes-and-fake-news)

Mustafaraj, E., and Metaxas, P. T. 2010. From Obscurity to Prominence in Minutes: Political Speech and Real-Time Search. Paper presented at Wec Science 2010: Extending the Frontiers of Society On-Line (WebSci10), Raleigh, NC, April 26–27. (repository.wellesley.edu/computersciencefaculty/9)

Mustafaraj, E., and Metaxas, P. T. 2017. The Fake News Spreading Plague: Was It Preventable? In *Proceedings of the ACM Conference on Web Science (WebSci'17)*, 235–239. New York: Association for Computing Machinery. doi.org/10.1145/3091478.3091523

Nematzadeh, A.; Ciampaglia, G. L.; Menczer, F.; and Flammini, A. 2017. How Algorithmic Popularity Bias Hinders or Promotes Quality. arXiv Preprint. arXiv:1707.00574v2 [cs.CY]. Ithaca, NY: Cornell University Library.

Nikolov, D.; Oliveira, D. F. M.; Flammini, A.; Menczer, F. 2015. Measuring Online Social Bubbles. *PeerJ Computer Science* 1:e38. doi.org/10.7717/peerj-cs.38

Pariser, E. 2011. The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think. London: Penguin.

Pew Research Center. 2017. *State of the News Media*. Washington, DC: Pew Research Center. (www.pewresearch.org/topics/state-of-the-news-media)

Qazvinian, V.; Rosengren, E.; Radev, D. R.; and Mei, Q. 2011. Rumor Has It?: Identifying Misinformation in Microblogs. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'11), 1589–1599. doi:10.1038/s41562-017-0132

Qiu, X.; Oliveira, D. F. M.; Shirazi, A. S.; Flammini, A.; and Menczer, F. 2017. Limited Individual Attention and Online Virality of Low-Quality Information. *Nature Human Behaviour* 1(7): 0132. doi:10.1038/s41562-017-0132

Resnick, P.; Carton, S.; Park, S.; Shen, Y.; and Zeffer, N. 2014. RumorLens?: A System for Analyzing the Impact of Rumors and Corrections in Social Media. Paper presented at the Computation + Journalism Symposium, New York, October 24–25. (nicole.zeffer.com/cj2014.pdf)

Resnick, P., and Sami, R. 2008. Manipulation-Resistant Recommender Systems through Influence Limits. *ACM SIGecom Exchanges* 7(3): 10. doi.org/10.1145/1486877.1486887

Shao, C.; Ciampaglia, G. L.; Varol, O.; Flammini, A.; and Menczer, F. 2017. The Spread of Fake News by Social Bots. arXiv Preprint. arXiv:1707.07592v2 [cs.SI]. Ithaca, NY: Cornell University Library.

Shiralkar, P.; Flammini, A.; Menczer, F.; and Ciampaglia, G. L. 2017. Finding Streams in Knowledge Graphs to Support Fact Checking. In *Proceedings of the IEEE International Conference on Data Mining* (ICDM). Los Alamitos, CA: IEEE Computer Society.

Silverman, C. 2017. Inside The Partisan Fight For Your News Feed. BuzzFeedNews, August 8. (www.buzzfeed.com/craigsilverman/inside-the-partisan-fight-for-your-news-feed?utm_term=.beN9PdexE#.skpv67JG9)

Sunstein, C. R. 2002. The Law of Group Polarization. *The Journal of Political Philosophy* 10(2): 175–195. doi.org/10.1111/1467-9760.00148

Sunstein, C. R. 2009. *Republic.com 2.0.* Princeton, NJ: Princeton University Press.

Suwajanakorn, S; Seitz, S. M.; and Kemelmacher-Shlizerman, I. 2017. Synthesizing Obama: Learning Lip Sync from Audio. *ACM Transactions on Graphics* 36(4): 95. dx.doi.org/10.1145/3072959.3073640

Thies, J.; Zollhöfer, M.; Stamminger, M.; Christian T.; and Nießner, M. 2016. Face2Face: Real-Time Face Capture and Reenactment of RGB Videos. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (CVPR). Piscataway, NJ: Institute for Electrical and Electronics Engineers. IEEE. doi.org/10.1109/CVPR.2016.262

Varol, O.; Ferrara, E.; Davis, C. A.; Menczer, F.; and Flammini, A. 2017. Online Human-Bot Interactions: Detection, Estimation, and Characterization. In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media* (ICWSM). Palo Alto, CA: AAAI Press. (aaai.org/ocs/index.php/ICWSM/ICWSM17/paper/view/15587)

Varol, O.; Ferrara, E.; Menczer, F.; and Flammini, A. 2017. Early Detection of Promoted Campaigns on Social Media. *EPJ Data Science* 6:13. doi.org/10.1140/epjds/s13688-017-0111-y

Weng, L.; Flammini, A.; Vespignani, A.; and Menczer, F. 2012. Competition among Memes in a World with Limited Attention. *Scientific Reports* 2: 335. doi.org/10.1038/srep00335

Zhao, Z.; Resnick, P.; and Mei, Q. 2015. Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts. In *Proceedings of the 24th International Conference on World Wide Web* (WWW'15), 1395–1405. Cambridge, MA: W3C. doi.org/10.1145/2736277.2741637

Giovanni Luca Ciampaglia is an assistant research scientist at the Indiana University Network Science Institute

AIIDE-18 to be Held in Edmonton, Canada

Please join us for AIIDE-18, to be held in mid-November at the University of Alberta in Edmonton, Alberta, Canada. AIIDE-18 is the next in an annual series of conferences showcasing interdisciplinary research on modeling, developing, and evaluating intelligent systems in entertainment. AIIDE-18 provides a meeting place for academic AI researchers and professional software developers to discuss the latest advances in entertainment-focused AI. The conference has a long-standing history of featuring research on artificial intelligence in computer games. We also invite researchers, developers, and digital artists to share ideas on topics at the intersection of all forms of entertainment and artificial intelligence broadly. AIIDE-18 will feature invited speakers, paper sessions, workshops, tutorials, playable experiences, panels, posters, the Starcraft AI Competition, and a doctoral consortium.

Submissions for all programs are due May 25, 2018. For more information, please visit www.aiide.org, or write to aiide18@aaai.org.

(IUNI). Before joining Indiana University, he was an analyst for the Wikimedia Foundation, and a research associate at the Swiss Federal Institute of Technology in Zurich. He is interested in information quality in cyber-human systems, in particular the trustworthiness and reliability of information in intelligent systems. His work has been covered in major news outlets, including the *Wall Street Journal,* the *Economist, Wired, MIT Technology Review,* NPR, and CBS News.

Alexios Mantzarlis leads the International Fact-Checking Network at the Poynter Institute. Mantzarlis has helped draft the fact-checkers' code of principles, shepherded a partnership between third-party fact-checkers and Facebook, testified to the Italian Chamber of Deputies on the fake news phenomenon, and helped launch International Fact-Checking Day. He previously served as managing editor of P*agella Politica* and *FactCheckEU,* respectively Italy's main political fact-checking website and the EU's first multilingual crowd-checking project. While at *Pagella Politica,* he presented weekly fact-checking segments on Virus, a prime-time talk show airing on Italy's national broadcaster RAI 2. Before becoming a fact-checker, he worked for the United Nations and the Italian Institute for International Political Studies.

Gregory Maus is a PhD student in informatics at Indiana University, Bloomington. His work on social bots, personal data brokers, and algorithmic sociological influence has been featured in *IEEE Spectrum, Foreign Affairs,* and *Real Clear Technology.*

Filippo Menczer is a professor of informatics and computer science at Indiana University, Bloomington. He holds a Laurea in physics from the Sapienza University of Rome and a PhD in computer science and cognitive science from the University of California, San Diego. Dr. Menczer is an ACM Distinguished Scientist. He previously served as division chair and director of the Center for Complex Networks and Systems Research at the IUB School of Informatics and Computing. His research focuses on web and data science, social network analysis, social computation, web mining, and modeling of complex information networks. His work on the spread of information and misinformation in social media has been covered in many US and international news sources, including the *New York Times,* the *Wall Street Journal,* the *Washington Post,* NPR, PBS, CNN, BBC, the *Economist,* the *Guardian, the Atlantic, Reuters, Science,* and *Nature.* Menczer serves as associate editor of the *Network Science* journal and on the editorial boards of *EPJ Data Science* and *PeerJ Computer Science.*