IEEE Computer, **52**(1), March 2019, pp. 36-45

The Age of Social Sensing

Dong Wang, University of Notre Dame
Bolesław K. Szymański, Rensselaer Polytechnic Institute
Tarek Abdelzaher, University of Illinois
Heng Ji, Rensselaer Polytechnic Institute
Lance Kaplan, U.S. Army Research Laboratory

Online social media have democratized the broadcasting of information, encouraging users to view the world through the lens of social networks. The exploitation of this lens, termed social sensing, presents challenges for researchers at the intersection of computer science and the social sciences.

Technically speaking, social sensing predates the use of physical technological sensors and social media.[1] For thousands of years, words, either spoken or written, have been used to communicate for a variety of reasons. The messages in such communications can be both explicit and implicit. For instance, some scholars assert that secret messages were embedded within the works of Plato.[2] What is new today, however, is that technology makes it much easier to formulate and share thoughts with the entire world instantly, something that took many generations for Plato.

With the greater dissemination power comes a heightened danger of misuse. It is now not only significantly easier to share ideas but also increasingly possible to manipulate perceptions of reality at an unprecedented scale. A growing challenge today, therefore, is to find and understand the valuable and truthful messages

Digital Object Identifier 10.1109/MC.2018.2890173
Date of publication: 11 March 2019

 0018-9162/19©2019IEEE

in the much larger volume of social media content.

A NEW "MACROSCOPE"

The unprecedented connectivity afforded by social media also gives rise to revolutionary new perspectives on empowering individuals and societies to collectively generate value from information. Pierre Levy, who introduced the concept of *collective intelligence* in the 1990s, observes that human intelligence is derived from reflexive reasoning, with language being the semantic indexing scheme for making arguments and producing results.[13] He envisions social media as the enabler of the next leap in reflexive collective intelligence and describes a new metalanguage, which he calls the *information economy metalanguage*, that would empower reflexivity, facilitate the discovery of semantic ties, and document information provenance (using blockchain-like technologies) to keep track of the origins of ideas and preserve the collective reasoning behind them.

If posts on a social network collectively comprise a new way to index physical reality, social beliefs, concepts, biases, and ideas, a logical next question arises: Can one develop a new type of instrument—a new "macroscope"—to view the world state? The purpose of such a device would be to reliably observe physical and social phenomena at scale, as interpreted by the collective intelligence of social media users.

With this in mind, other questions must be considered:

- What would be the properties of such an instrument?
- How could these properties be influenced or optimized?
- How does the instrument distort the image of the world being observed?
- Do the posted observations themselves influence properties of the observed system, perhaps by engendering polarization and radicalization or, conversely, by leading to a better shared understanding, awareness, and agreement?
- Do they consequently affect future information propagation, thereby influencing the instrument itself?
- How susceptible is this instrument to the intentional manipulation of perceived reality, and what mitigation strategies are effective against such manipulation?
- Can one derive accuracy, error, and performance bounds for retrieved observations given the models of human behavior and biases?
- Conversely, can one infer human biases and trust relations from the reported observations?

An exciting aspect of these social-sensing challenges is that they open up a novel research field in which human-centered sciences meet research on physical, computing, and engineered systems to reach a better understanding of the new instrument and more fully characterize the properties of humans (posting on social media) as collective sensors.

This article focuses on one slice of this problem: the use of such a macroscope to reconstruct physical (as opposed to social) reality. The benefits of addressing this challenge are significant. Not only will better algorithms curb the misrepresentation of physical reality; they may also help build smarter urban services. Most systems, from urban transportation to disaster response, include humans as an integral part of the underlying sensing, management, and control loops. Acting as social sensors (who post on social media), humans are able to recognize, observe, describe, and report a much broader spectrum of events than do physical sensors. Especially important is the human ability to recognize when an activity represents an abnormal pattern or behavior. Examples include actions that call attention to suspicious individuals or suspected crime scenes.

Yet strong human sensing and interpretation skills are also liabilities. Unlike sensors that objectively report observed data, humans summarize observations to create their own interpretations of the observed state. These interpretations are often influenced by the opinions, biases, or beliefs of the observer. For example, a scene of street fighting between the police and demonstrators can be viewed by a government supporter as an example of brave police officers restoring law and order broken by unruly demonstrations; meanwhile, a government opponent might look at the same scene and see it as an example of police brutality against peaceful demonstrators with legitimate grievances against a corrupt government. How can one map such different "semantic indexing" representations of the same event back to a neutral reconstruction of what might have actually transpired? Figure 1 summarizes the goal of building the macroscope using observations from social-sensing platforms to understand the physical world in light of cyberphysical and linguistic challenges, as discussed in the following.

These challenges can be viewed from two important and deeply intertwined perspectives: challenges in the cyberphysical space and in the social and linguistic space.

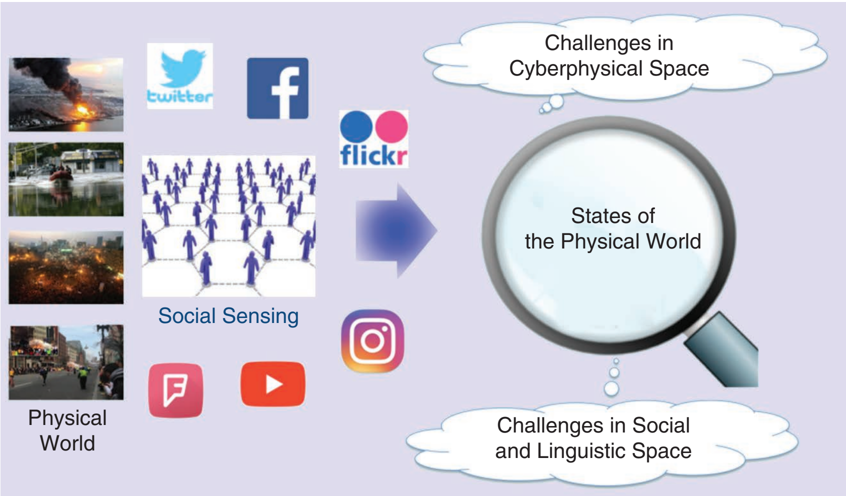

FIGURE 1. An illustration giving an overview of social sensing.

CYBERPHYSICAL CHALLENGES IN SOCIAL SPACES

Understanding the attributes of social-sensing systems involves modeling three interdependent components: 1) the humans in the loop and their cognitive models, 2) the algorithms required, and 3) the laws of nature that govern the underlying physical and engineered artifacts. For example, understanding how well a disaster-response team might be able to survey postdisaster damage within a given time will depend on

1) modeling the way survivors might respond to the disaster (including any information-sharing behaviors, such as reporting of actual damage and responding to potential rumors)
2) understanding the efficacy of software decision-support tools that distill the raw human response into more actionable information
3) accounting for physical-resource constraints that affect reporting and response.

Meeting these challenges requires interdisciplinary approaches that address the complex interactions among cyber, physical, and social components of the holistic system.

A recent issue of *Computer* summarized a few representative findings related to the challenges of human-in-the-loop systems.[3] In one article, Abowd defines the fourth generation of "collective computing," where technological forces (described as "the cloud," "the crowd," and "the shroud") tightly and interactively connect the physical and cyber worlds. In another article, Chang describes situation analytics, a new runtime computational model that explores users' mental states during the entire software lifecycle to better understand user intention. Additionally, Jiang et al. present their work in an emerging cyberphysical systems research area, safety-critical medical devices, and describe a state-of-the-art approach for validating closed-loop devices without jeopardizing human safety. It is informative to review the challenges brought about when cyberphysical systems operate and interact in social spaces.

Reliability is a key engineered property of cyberphysical systems. In physical sensing, measurement devices are typically well calibrated with well-understood error properties (documented in manufacturers' data sheets). By analogy, a key social-sensing challenge from a cyberphysical perspective is to understand the reliability properties of our new observational instrument: the social macroscope that converts social posts into an estimated state of observed phenomena. Recent history documents landmark events influenced by social media in the absence of tools to reliably assess information. We are reminded especially of the use of social media to spread fake news and terrorist propaganda as well as manipulate elections. These developments have generated interest in understanding signals and distortions on social media.

Modeling instrument distortion

One first needs to understand the nature of distortion. We can classify the distortion of information produced by human sources on social media into three broad categories.

The most challenging category is intentional disinformation sent to deceive. Done purposefully (e.g., changing small but important details and then multiplying them by the power of social networks, in which the malicious sources are hubs), this distortion can cause disinformation to be taken as truth for a long time.

The next category includes cases where a person makes assertions that are not sufficiently supported by data or observations but still might sound plausible to people with a certain bias and so be accepted as facts. This category includes biased interpretations by communities of people with similar opinions. In such cases, people recollect events in a manner consistent with their biases, beliefs, and preferences.

Finally, there are genuine random mistakes by people processing information, such as misspelling and typos. The massive number of sources, relative permanency of the social communities involved in collecting and propagating information, and our ability to partially track information provenance make it possible to build a rigorous and usable formal basis for reliable truth extraction from social sensors.[1]

An important consideration in social sensing is to account for human reporting, not of events but of the perception of events. These perceptions are governed by human cognition, which today can be simulated using such cognitive models as the Active Control of Thought–Rational (ACT-R) model.[4] Well-known limitations of cognition, such as limited attention span, the decay of unreinforced memory traces, or limited information-processing speed, affect human performance as sensors and have been adequately modeled based on the ACT-R. An interesting future direction is to enable simulated crowds of agents endowed with human cognitive models to reliably account for the impact of human cognition limitations on event reporting.

The tangled ties linking information distortion, bias, cognition, and human preferences mean that, unlike the case with physical sensors, distortions on social media carry a signal themselves. These distortions indicate human preferences and beliefs that offer great value in terms of tuning the product design, for example, in advertising campaigns, recommendation systems, and voter recruitment tools. In this article, however, we focus primarily on algorithms for removing, where possible, the biases to create a more objective representation of the physical reality.

Much work has been done on handling noisy sensors and unreliable (physical) sources. When humans act as sensors, however, a new model is needed to characterize the different types of distortion (discussed previously). From a system-reliability perspective, rumors, misinformation, exaggeration, and bias propagate along social channels, creating highly correlated reports (posts) and, thus, correlated errors at a large scale. It is well known that correlated errors decrease the reliability of systems. Social ties, trust relations, and homophily further contribute to correlated beliefs, opinions, and misrepresentations. This makes social networks a challenging instrument to use without a good understanding of their error properties and models. Recent research has used simplified models of human behavior to derive fundamental information-theoretic bounds for estimating accuracy from social media.[5]

Understanding the signal

If social media are abstracted as instruments that measure the physical world, what is the sensing modality of these instruments? Can social sensing be treated like acoustic, magnetic, or vibration sensing? For example, can Twitter be thought of as a medium, where a distribution of tokens of different types are emitted in response to a physical event, just as a distribution of acoustic waves of different frequencies may be emitted to reach an acoustic target (see Figure 2)? This analogy recently led to interesting approaches for event detection, localization, and tracking, where events reported on social media are algorithmically detected in a language-agnostic manner inspired by the target-tracking and data-fusion literature for physical signals.[6] More generally, we observe a rich set of data modalities generated by humans (e.g., text, sound, images, and video). Such rich data modalities lead to interesting research questions. For example, can we explore the interdependent relationship of sensing measurements with different data modalities to

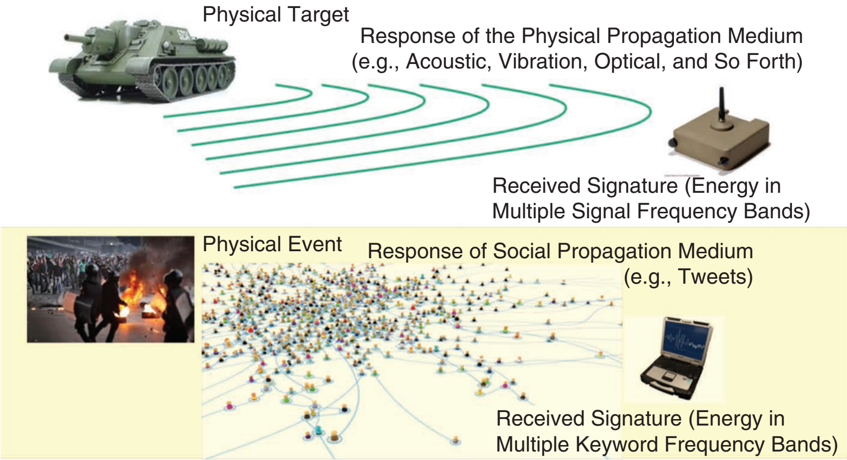

FIGURE 2. An Illustration showing an analogy for the social-sensing modality.

obtain more accurate sensing results? How can we build an appropriate model for versatile human sensors that can simultaneously generate sensing data of diversified modalities?

Researchers have recently identified a promising new direction for activity: the development of information-theoretic models of social channels. Information theory has made significant headway by abstracting information as a sequence of bits. This bit-stream abstraction offered ways to model noise, reason about error probabilities, derive capacity, and construct optimal decoders. Inspired by this approach, recent work on using Twitter as a data source viewed statements about physical reality as a binary signal (a bit). This is because the statement of the tweet is either true of the physical world or not, leading to the binary abstraction. Given this representation, Twitter becomes just another noisy channel where binary signals propagate. Conceptually, the source of this channel is the physical world itself. The output constitutes human observations of the world. Distortion is introduced, for example, when people make up rumors. This is equivalent to flipping a bit from zero (something that was not true of the physical world) to one (something claimed true in the output of the social medium).

Similarly, failures to report true events can be modeled as bit omissions, and claims that deny actual events are bit flips from one to zero. One can then reason about the probability of different bit flips and omissions and derive estimators of social-channel input (i.e., physical world state) given its output, using existing estimation-theoretic techniques. The approach has been successful for modeling statements about objectively observable realities (e.g., "it is raining"), where ground truth is unique and unambiguous. It needs extensions to cases where ground truth is undefined or subjective, such as for assessing the statement, "John is a great president!"

Quantifying data reliability and performance bounds

Recent maximum-likelihood estimation techniques adopting the stated binary model show much promise in both detecting the original state of the world from the distorted output of the social channel and estimating the distortion itself. This problem, known in classical data fusion as a *joint signal detection and channel estimation problem*, possesses the necessary analytic foundations for the solution approach. An optimal estimation-theoretic framework can jointly estimate both the reliability of data posted and the credibility of information sources involved without prior knowledge of either.[1]

Estimation theory offers solid analytical foundations for bounding the error of an optimal estimator. Representing social media as noisy binary communication channels, as described previously, enables the application of estimation-theoretic frameworks to the reliability analysis of data-cleaning systems operating on the outputs of social channels. A key recent contribution that leverages this insight exploits expressions of the error bounds of maximum-likelihood estimators to assess the quality of estimation results on social media.[1] This quality analysis is immensely important in practical settings where errors have consequences.

However, this analysis is complicated by two factors. First, it must account for correlated errors that result from rumor-spreading behaviors, such as when a person reports as his or her own observations received from others without verification. Second, the analysis must account for the expressed degree of vagueness or uncertainty in human observations, such as "the protest is possibly unsafe." In such cases, the estimator must consider the degree of confidence a source expresses in its messages to make proper assessments.

Recent work based on subjective logic (a type of uncertain probabilistic logic) has developed a framework to assess source reliability in such situations.[7] This framework presumes that the expressions of vagueness are quantified as specific probabilities. The conversion of natural-language expressions of vagueness into quantifiable numbers is a very difficult problem that remains to be fully solved (see the section "Challenges in the Linguistic Space").

Exploring the role of dependencies between sources

Dependent sources are common in social-sensing applications, resulting in uncertain data provenance. That is, it is not unusual for sources to report observations they received as if they were their own. This is common in instances where humans play the role of data sources connected through social networks (e.g., a follower–followee relationship on Twitter and a friend relationship on Facebook.). The rumor-spreading behavior of human sensors has no analogy in correctly functioning physical sensors. From a cyberphysical-system perspective, this means that errors in "measurements" across sources may be nonindependent, as one erroneous observation may be propagated by other sources without being verified. Recent research reports on attempts either to develop source-selection schemes for the careful selection of independent sources on social networks or to build reliable social-sensing models that explicitly model the source dependency into the social signal processing engine.[1] The complex and dynamic source dependency graphs on social networks deserve more investigation.

Understanding communities, social trust, and polarization

The influence that relations among sources have on results of the social macroscope is especially strong within human communities. Humans interact, operate, exchange, and propagate information much more frequently within communities than across them. Communities enable members to develop trust in one another. As a result of either the individual's selection of a community to join or the community members' influence on one another, communities tend to increase the level of homophily among their members.[8] This often results in members having similar opinions, attitudes, and beliefs as well as similar misconceptions, misrepresentations, and susceptibilities.

Communities serve the specific needs of their members, so the level of the homophily among members is often strongest for the traits related to these needs, such as skills, access to information, or interests including particular political movements, specific sports or hobbies, or a preference for a particular genre of music.[9] Typically, each person belongs to several communities and has varying levels of commitment to each, with a specific community satisfying different needs. The notion of communities is often associated with a sense of trust and shared opinions or biases. Understanding the importance and influence of communities is therefore key to understanding information reliability.

Returning to the Twitter example, the information-propagation characteristics of the social channel are significantly affected by trust relations among sources as well as by those sources' views. Individuals are much more likely to rebroadcast information they hear from a source they trust than from a source they mistrust or regard as neither especially reliable nor especially unreliable. Hence, when the same information is reported by multiple sources, the information's reliability cannot be accurately computed without understanding the underlying social-trust network. From a reliability perspective, trust relations increase correlations in reporting and, hence, the probability of correlated errors among individuals. Bias toward a particular viewpoint also affects information propagation. Individuals tend to propagate the claims that match their opinions. Accordingly, understanding community biases is important in assessing error correlations and, consequently, data reliability.

Equally interesting is the inverse problem. Because trust and biases modulate the propagation of information, observations of the propagation patterns themselves reveal the underlying trust relations and biases of sources.[7] This effect is demonstrated in Figure 3, which shows information-dissemination topologies among sources supporting different candidates in an Egyptian election. The problem is interdisciplinary: social science informs the development of generative models that predict how individuals modulate information propagation. Given such models, it is possible to estimate each individual's position on issues of interest by observing the collective information-propagation patterns on the social network.

A particularly interesting problem is to estimate the ground truth from social media outputs in the presence of polarization, a situation that occurs when two (or more) distinct camps in the social network propagate conflicting claims.

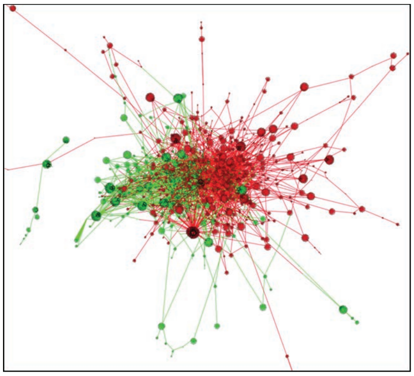

FIGURE 3. A diagram illustrating information-dissemination topologies among sources supporting different candidates in an Egyptian election. The dots represent information sources. The lines represent communication links. Green indicates sources and communications in support of one candidate. Red indicates sources and communications in support of another candidate.

Recent work addressed the challenge of detecting polarization and identifying the bias of individual sources in polarized communities.[10] Experiments showed that, when accounting for such bias, state reconstruction from social observations tends to align more closely with the ground truth in the physical world than when polarization is not considered. It is also possible to estimate the polarity-dependent propagation networks from tweet propagation patterns. Such an estimation enables the computation of correlations among sources to determine their error-dependence properties and, hence, inform fact-finder design.

Fusion of physical and social sensors

With error bounds computed on social-sensing observations, an interesting challenge arises in developing fusion engines that combine physical and social-sensing data within the sensing–control–actuation loops of cyberphysical systems. While social networks can be viewed as additional sensing sources that corroborate physical sensors, the fusion of social and physical sensing can also give rise to new kinds of capabilities. For example, such fusion can explain anomalies seen by sensors in view of data collected from social media.[11]

FROM A RELIABILITY PERSPECTIVE, TRUST RELATIONS INCREASE CORRELATIONS IN REPORTING AND, HENCE, THE PROBABILITY OF CORRELATED ERRORS AMONG INDIVIDUALS.

Interesting research questions remain to be answered.

- How is it possible to accurately correlate and fuse data from physical and social sources?
- How can different data modalities (e.g., numerical, text, images, and video) be handled effectively?
- How is it possible to automatically infer the causal relationships between identified events and quantify the accuracy of such inference?

Finally, sensing systems exist in the context of control loops. Current systems typically feature loops with only a few well-designed sensors. Can future applications use crowd sensing as a subsystem in their control loops? How can the closed-loop characteristics of such systems be analyzed? How can the impact of information errors on closed-loop stability and performance optimization be determined? In many such systems, individuals (besides being data sources) are part of system "control." Decision makers will act on the data in ways that affect the physical state being observed. Hence, an interesting challenge is to understand what information should be presented to the decision maker (and in what form) so as to offer the best decision support despite the inherent noise in the underlying social channel.

CHALLENGES IN THE LINGUISTIC SPACE

To complete our estimation of the state of the physical world, media content must be properly understood. This can be thought of as the challenge of "interfacing" to the human sensor. In addressing this obstacle, we need to seek dramatic advances in the rapid, low-cost development of information-extraction and text-mining technologies for understanding social media content.

Today's state-of-the-art language-processing technologies rely on either 1) supervised learning, which suffers from the high cost of large-scale manual annotation and the limited predefined fact types, or 2) unsupervised learning, which usually yields unsatisfactory performance due to the exclusion of language/task-specific knowledge. We need to systematically discover and unify latent and expressed knowledge from traditional symbolic semantics and modern distributional semantics through advanced machine-learning models. The idea of unified representations of semantics has been a focus of investigation both in social science (e.g., the seminal work of Lévy on metalanguages for expressing collective intelligence) and computing (e.g., the recent interest in language embedding as a way to abstract semantic spaces).

An information-extraction system should be able to simultaneously discover a domain-rich schema and extract information units with fine-grained types. It should have a "cold start" and be adaptable to any domain, genre, language, or data modality without any human-annotated data. Recent work on combining symbolic semantic and distributional semantics[12] has made it possible to discover semantic schema and extract facts simultaneously without relying on human-defined and -constructed ontologies for specific domains.[13] It should also be able to adapt within several hours to a new scenario using very few resources. The new framework should rapidly acquire, categorize, structure, and zoom in on incident-specific expectations from various nontraditional sources, including those with humans in the loop. The output of such a framework will constitute a structured set of data types and their instances/values derived from text (e.g., to describe the state of an observed environment). Such an output will be more amenable to integration within sensor data fusion systems than the original unstructured text generated by human observers. The process of collecting information introduces many challenges, which we describe in the following sections.

Ambiguity in sentences

A core challenge in processing natural language on social media is that important information is often presented implicitly. Such presentation contains varied instances of imprecision, ambiguity, vagueness, and implicit information. Natural-language processing technologies currently rely heavily on surface processing. This makes it difficult to exploit deep structure, background knowledge, and source information. Most posters on discussion forums assume the readers already know the on-topic entities and events: thus, they do not bother to elaborate the background for these target entities. They also tend to use short and informal references for efficient discussions. As a result, the local contexts in which an entity is mentioned are often not sufficient for disambiguation.

An entity linker needs to automatically construct a background knowledge base for disambiguation. Without understanding the global topic knowledge about an entity, entity-linking systems tend to mistakenly link the entity to a more popular one. Finally, entity disambiguation techniques are still weak in exploiting commonsense knowledge.

The importance of context

The current state of the art considers informal text from social media an impoverished alternative to formal text (e.g., newswire) and prone to producing erroneous and conflicting facts from informal and noisy data. Innovative techniques on implicit and morphed information extraction are required to handle imprecise language. For example, in some societies, the Internet is subject to censorship and surveillance. These actions include blocking websites, deleting posts, and filtering information about a specific entity or event.

With the growth of online social-networking services, people rapidly invent new ways to communicate sensitive ideas. We call this phenomenon *information morphing,*[14] which has also existed with traditional forms of communications. Motivations for morphing, besides dodging censorship, include expressing sarcasm/verbal irony or positive sentiment toward some entities or events, planning for secret actions, trading illegal products, masking politically unpopular views, or simply creating interesting conversations.

Morphing raises unique challenges for entity and event coreference resolution. Document-level profile and corpus-level temporal distribution-based features can narrow down the scope of referents, but deeper understanding is needed to discover concealed facts. While ambiguity, redundancy, and implied information can pose significant problems for systems designed to understand natural language, the impact of these errors can be mitigated by incorporating disparate deep sources of information.

Discourse ambiguity

Besides sentence-level and subsentence-level ambiguities, there is yet one more level of obscurity: supersentential or discourse-level ambiguity, which goes beyond sentence boundaries. This third level of ambiguity comes from two sources.

The first is coreference ambiguity, which occurs when pronouns refer to entities outside the current sentence and it is uncertain which entity is referenced. Failure to resolve this ambiguity undermines our ability to understand the original discourse. The second source of ambiguity is discourse structural ambiguity. This occurs when a discourse, like a sentence, has its own internal structure, often represented as a tree or graph. For deep understanding of the text, we need to know, for example, which sentence is the topic sentence, which sentences constitute elaboration on or contrast to another sentence, and the temporal structure between the sentences (in a story line).

Fuzzy and vague expressions

Individuals often use fuzzy language rather than precise terms when describing things or people (e.g., "she is very tall" instead of "she is six-feet, two-inches tall"). It is very difficult to convert fuzzy terms into numbers, as the use of such terms varies wildly between individuals and societies. It is much easier to calibrate physical sensors whose noise performance can be consistently measured. The problem becomes more difficult as humans add notions of vagueness in their reports (e.g., "she is quite tall"). The conversion of such terms into a distribution of possible height values is very much an open and difficult problem.

FUTURE DIRECTIONS

While research on various components of the social-sensing vision is underway, the field lacks a unified interdisciplinary problem formulation that takes a holistic approach to modeling humans as sensors and to modeling social media as noisy measuring instruments or channels. Several challenges remain. For example, we need to determine the following:

1) How can the language ambiguities discussed be incorporated into the channel bit error models considered by ground-truth estimators?
2) How can information theory be augmented to account for semantic errors and approximations?

ABOUT THE AUTHORS

DONG WANG is an assistant professor in the Department of Computer Science and Engineering at the University of Notre Dame, South Bend, Indiana. His research interests include developing analytic foundations for reliable sensing and information distillation systems as well as the foundations of data credibility analysis in the face of noise and conflicting observations. Wang received his Ph.D. in computer science from the University of Illinois at Urbana–Champaign. He has authored or coauthored more than 80 refereed publications in the area of big-data analytics, social sensing, cyberphysical computing, and real-time and embedded systems. He is a Member of the IEEE and ACM. Contact him at dwang5@nd.edu.

BOLESŁAW K. SZYMAŃSKI is the Claire and Roland Schmitt distinguished professor of computer science and the founding director of the Social Cognitive Networks Academic Research Center at Rensselaer Polytechnic Institute, Troy, New York. His research interests include network science, especially social networks and computer networks. Szymański received his Ph.D. in computer science from the Institute of Informatics in Warsaw, Poland. He is an author or coauthor of more than 300 publications and an editor of six books. He is a Fellow of the IEEE and a foreign member of the National Academy of Sciences in Poland. Contact him at szymab@rpi.edu.

TAREK ABDELZAHER is a professor and Willett faculty scholar in the Department of Computer Science at the University of Illinois Urbana–Champaign. His research interests include understanding and influencing the performance and temporal properties of networked embedded, social, and software systems in the face of increasing complexity, distribution, and degree of interaction with an external physical environment. Abdelzaher received his Ph.D. from the University of Michigan, Ann Arbor, on quality-of-service adaptation in real-time systems. He has authored or coauthored more than 200 refereed publications related to real-time computing, distributed systems, sensor networks, and control. He is a Member of the IEEE and ACM. Contact him at zaher@cs.uiuc.edu.

HENG JI is the Edward P. Hamilton development chair professor in the Computer Science Department of the Rensselaer Polytechnic Institute, Troy, New York. Her research interests include natural-language processing and its connections with data mining, social science, and vision. Ji received her Ph.D. in computer science from New York University. She has coordinated the National Institute of Standards and Technology Text Analysis Conference, Knowledge Base Population track, since 2010 and served as the program committee chair of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technology. She is a Member of the IEEE and ACM. Contact her at jih@rpi.edu.

LANCE KAPLAN is a researcher in the Networked Sensing and Fusion branch of the U.S. Army Research Laboratory, Adelphi, Maryland. His research interests include information/data fusion, reasoning under uncertainty, network science, resource management, and signal and image processing. Kaplan received his Ph.D. from the University of Southern California, Los Angeles, in electrical engineering. He serves on the Board of Governors for the IEEE Aerospace and Electronic Systems Society and as vice president of conferences for the International Society of Information Fusion. Previously, he served as editor-in-chief of *IEEE Transactions on Aerospace and Electronic Systems* and on the board of directors of the International Society of Information Fusion. He is a Fellow of the IEEE and of the U.S. Army Research Laboratory. Contact him at lance.m.kaplan.civ@mail.mil.

3) When individuals summarize information, how does the resulting lossy compression affect downstream error propagation?
4) Because sources often implicitly assume some shared background with the receiver, how can errors in interpretation that arise because of mismatch between the backgrounds of a message source and message recipients be predicted?
5) How should inference mechanisms that help recognize corroborating evidence be encoded?
6) How should these mechanisms identify and handle irreconcilable observations?
7) What is the right level of abstraction for modeling human behavior (in relaying information)?

Finally, individuals' actions with regard to information interpretation and processing before propagation are often correlated to the degree of the connectivity, familiarity, and trust among individuals. Hence, in addition to the actions of individuals it is necessary to understand the impact of relations among individuals and the influence of social groups and communities.

These questions call for the involvement of at least five synergistic disciplines. Specifically, the work needs

1) computational social scientists to model human behavior at both the individual and community scale and to quantify its susceptibility to errors, omissions, deceit, and other irregularities

2) linguists to model strengths and imperfections of human communication and their compounding effects on the reliability of information dissemination
3) information theorists to model social networks as imperfect communication media and derive fundamental capacity limits and uncertainty envelopes
4) data-mining experts to investigate the impact of the underlying error models on the reliability of knowledge extraction from imperfect information
5) cyberphysical experts to develop estimation-theoretic observability and control limits and tools that offer closed-loop robustness guarantees in the face of derived capacity limits, uncertainty envelopes, and knowledge errors.

The research challenges and future directions reviewed here call for the emergence of a new field that combines social and cognitive models, linguistics, estimation theory, information theory, and reliability analysis, with the goal of putting social media exploitation on well-understood analytic foundations, not unlike the fusion of hard data from physical sensors and signals. New interdisciplinary research is outlined to bring about novel solutions for a better theoretical and systematic understanding of emerging social-sensing systems in the sensor- and media-rich world of the future.

ACKNOWLEDGMENTS

Research for this article was sponsored by the Army Research Laboratory [Cooperative Agreements W911NF-09-2-0053 (the Army Research Laboratory Network Science Collaborative Technology Alliance) and W911NF-17-1-0409], the National Science Foundation (Grants CNS-1831669, CBET-1637251, CNS-1566465, and IIS-1447795), and the Office of Naval Research (grant N00014-15-1-2640). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. government. The U.S. government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation here.

REFERENCES

1. D. Wang, T. Abdelzaher, and L. Kaplan, *Social Sensing: Building Reliable Systems on Unreliable Data*. Burlington, MA: Morgan Kaufmann, 2015.
2. J. Kennedy, “Plato’s forms,” *Pythagorean Mathematics, Stichometry, Apeiron*, vol. 43, no. 1, pp. 1–32, 2010.
3. S. Helal, “2016 outlook: Humans in the technology loop,” *Computer*, vol. 49, no. 1, pp. 7–9, 2016.
4. J. R. Anderson, M. Matessa, and C. Lebiere, “ACT-R: A theory of higher level cognition and its relation to visual attention,” *Human-Comput. Interaction*, vol. 12, no. 4, pp. 439–462, 1997.
5. S. Yao et al., “On source dependency models for reliable social sensing: Algorithms and fundamental error bound,” in *Proc. 36th Int. Conf. Distributed Computing Systems (ICDCS)*, Osaka, Japan, June 2016. doi: 10.1109/ICDCS.2016.75.
6. T. Sakaki, M. Okazaki, and Y. Matsuo, “Earthquake shakes Twitter users: Real-time event detection by social sensors,” in *Proc. 19th Int. Conf. World Wide Web (WWW)*, 2010. doi: 10.1145/1772690.1772777.
7. L. Kaplan, M. Şensoy, S. Chakraborty, and G. de Mel, “Partial observable update for subjective logic and its application for trust estimation,” *Inform. Fusion*, vol. 26, pp. 66–83, 2015.
8. O. Lizardo et al., “Analysis of opinion evolution in a multi-cultural student social network,” in *Proc. 6th Int. Conf. Applied Human Factors and Ergonomics (AHFE)*, Las Vegas, NV, pp. 3974–3981, 2015.
9. P. Bogdanov, M. Busch, J. Moehlis, A. K. Singh, and B. K. Szymanski, “Modeling individual topic-specific behavior and influence backbone networks in social media,” *Social Network Anal. Mining*, vol. 4, no. 1, pp. 1–16, 2014.
10. M. T. Amin, C. C. Aggarwal, S. Yao, T. Abdelzaher, and L. Kaplan. “Unveiling polarization in social networks: A matrix factorization approach,” in *Proc. IEEE Infocom*, Atlanta, GA, May 2017. doi: 10.1109/INFOCOM.2017.8056959.
11. P. Giridhar et al., “ClariSense+: An enhanced traffic anomaly explanation service using social network feeds,” *Pervasive Mobile Computing*, vol. 41, pp. 381–396, 2017.
12. L. Huang et al., “Liberal event extraction and event schema induction,” in *Proc. 54th Annu. Meeting Association Computational Linguistics (ACL2016)*, pp. 258–268.
13. P. Lévy, 2013. *The Semantic Sphere 1: Computation, Cognition and Information Economy*. Hoboken, NJ: Wiley.
14. B. Zhang et al., “Be appropriate and funny: Automatic entity morph encoding,” in *Proc. 52nd Annu. Meeting Association Computational Linguistics (ACL2014)*, pp. 706–711.