Social media has become globally ubiquitous, transforming how people are networked and mobilized. This forum explores research and applications of these new networked publics at individual, organizational, and societal levels.
— Shelly Farnham, Editor

Antisocial Computing: Exploring Design Risks in Social Computing Systems

David W. McDonald, University of Washington, **David H. Ackley,** University of New Mexico, **Randal Bryant,** Carnegie-Mellon University, **Melissa Gedney,** Digital Promise, **Haym Hirsh,** Cornell University, **Lea Shanley,** University of Wisconsin-Madison

Social computing technologies offer many broad societal-level impacts—more effective crisis response, individually tailored education, effective workforce retraining, participatory governance, citizen journalism, entertainment and leisure activity, and improvements in individual wellness and healthcare, to name a few. But like many technologies, social computing is neither exclusively good nor bad (see sidebar). Technologies, methods, software infrastructures, and systems can be employed to achieve both socially desirable as well as potentially undesirable outcomes. An infrastructure designed to help us identify phishing attacks and malevolent computer viruses could just as easily identify individuals who might be good phishing targets or help hackers design a very effective virus.

A workshop jointly sponsored by the Computing Research Association, Computing Community Consortia (CCC), and the Commons Lab of the Woodrow Wilson International Center for Scholars recently convened to generate a national research agenda for human computation. Workshop participants explored research challenges in small groups. Our group focused on developing intellectual challenges that would address the risks of social computing.

The approach we used was to invert common pro-social goals and norms we often employ when designing. This design technique is useful because it allows us to unpack and explore spaces of a design that are rarely explored, similar to the way that some hackers work from inside an organization to understand its vulnerabilities. Designing systems to exploit users is work that is rarely undertaken, because few of us have the sociopathic tendencies of a social media grifter who would willingly take advantage of people's helpfulness, willingness to do good, and general tendency to please others.

We want to be clear that we are not advocating the use of social computing for criminal, abusive, or antisocial purposes. The goal of approaching the "anti-design" is to better understand the risk present in any complex social computing system.

DESIGNING TO EXPLOIT SOCIOTECHNICAL INTERACTIONS

We focused our design around three components but kept each at a relatively high level: first, a grand challenge with supporting motivation; second, the supporting social computing systems; and third, the essential supporting social computing techniques. We then used these three tiers to articulate a research agenda necessary to achieve the techniques, systems, and grand challenge.

The challenge. We chose the grand challenge of starting a regional armed conflict so that we could use the inside information of the conflict to more effectively arbitrage resources and use the economic turmoil to gain significant economic benefits. This might take the form of fomenting a war in Australia to be able to profit from trade in natural resources originating from that region. As part of brainstorming the grand challenge, we realized that an effective design could also be used to build extortion or "protection" rackets, or to more effectively hedge the trading of stocks. Like all good grand challenges, there are sets of key related activities from which, if the challenge is achieved, we would derive broader benefits.

Social computing systems. The key supporting social computing systems are those that effectively facilitate networks of users that can quickly spread information. These systems need to have relatively high connectivity (high edge density). Further, it is desirable that information have relatively robust properties once in the system. The system does not have to be uniform. That is, not every participant needs to trust and retain the specific information; we just need a significant number of actors to trust, value, and propagate that information. Further, the platform needs to enable participation from some form of *automated social actor* (ASA), or bot. That is, the platform cannot require known user validation. It turns out that many existing platforms currently have these properties—

Insights

- → It is time to take stock of the risks in social computing systems and design specific mitigations to those risks.
- → Exploring an "anti-design" can help designers better understand the risks present in any complex social computing system.

Facebook, Twitter, Wikipedia, Reddit, Pinterest, Instagram, and Flickr, to name just a few.

Social computing techniques. The essential computing technologies are a little more interesting. One of the technologies required is a special form of ASA called an *automated sociopathic actor*. This ASA does not need to be perfect or pass a Turing test, as many current ASAs do not. It just needs to do a reasonable job of echoing, mirroring, and trying to be like selected users to engender trust and an overall "nice guy" type of status in the social network. The ASA could be used to help propagate and amplify information as it moves through the network and to change the network structure in ways that help us achieve our desired goal. For example, the ASA could help isolate individuals who are inserting information that may counter our objectives.

A second key supporting technology is a *"sucker" network identification system*. The key here is to move beyond individuals who are gullible and easily influenced to well-connected networks. The system could not be very effective if it found only individuals who believed the information that might be injected. Most phishing scams are happy to spam email accounts and find individuals. Instead, the sucker network needs to be a well-connected network of people who are all equally gullible and influenceable. The sucker network facilitates the propagation of information, but it's not made up of just suckers. Once the sucker network is identified, it is essential to infiltrate it with ASAs to strengthen ties and enhance the spread of specific information.

A third supporting technology is an *information recasting tool*. Information has many different properties that can be manipulated. If we are to achieve the goal of spawning armed conflict, then we need to know how to recast information into subtly different forms that shift opinions in two or more networks to increase dissimilarity, distrust, and animosity. This is not simply political speech or advertising. Indeed, those types of messages are carefully crafted to achieve a goal—and there are likely to be cognitive and social science insights from related disciplines that study those types of message creation. But as yet, there is no generic tool that allows us to take a piece of information, a factoid, or an opinion and generate a message with a set of optimal information properties to achieve a specific goal.

Designing systems to exploit users is work that is rarely undertaken, because few of us have the sociopathic tendencies of a social media grifter.

Given the rough high-level design, we turned to brainstorming a set of research challenges that would be necessary to achieve our design. Our research challenges corresponded closely to the elements of our design.

DERIVED RESEARCH CHALLENGES

Our grand challenge opens an interesting research area in social networks. Finding the right sucker network is not a trivial task. Finding networks that are tightly connected is certainly possible, but rarely does social network analysis consider the veracity or accuracy of the information

IMAGE BY CHRIS LISTER

that crosses the ties that make up the connections among individuals. Further, the sucker network needs to be robust, which yields an interesting challenge in the area of structural holes. Most social network analyses that focus on structural holes work to identify them for the purposes of filling them so that the network itself is robust to flaws in information transmission. Instead, what we would need is to be able to create just the right structural holes so that disconfirming evidence could not be injected into our sucker network.

Establishing the right structural holes is an area where our social network research challenge relates to our sociopathic AI research challenge. One piece of work that the ASA might need to do is isolate participants to create the right type of structural hole and therefore make the information we need to inject into the sucker network more robust. The sociopathic AI research agenda needs to be able to effectively model differing tastes, personalities, and interaction styles across a network. One thing we know is that we tend to like people who are similar to ourselves. Effectively identifying and reflecting similarity is one key aspect of the sociopathic AI challenge that would allow us to build the right types of ASAs.

We identified a critical need for research in an area we termed *disinformation engineering*. The research challenge is focused on creating information messages that have optimum influence across specific constituencies in a social computing system. This should not be thought of as just propaganda or simply marketing. Disinformation engineering is an engineering problem with the goal of having an optimal disinformational impact and parameterizing how to achieve that goal for a set of participants within some specific tolerances.

Disinformation engineering has specific theory upon which to build. Claude Shannon is well known for characterizing the theoretic propagation of information to ensure robustness. The challenge for disinformation engineering is to be able to inject distortion or error to reliably generate a disinformation message with

IMAGE BY JENNIFER TOMALOFF

SOCIAL COMPUTING, BOTH GOOD AND BAD

This is Peking University in China, a place of those dreams of freedom and democracy. However, a young, 21-year-old student has become very sick and is dying. The illness is very rare. Though they have tried, doctors at the best hospitals in Beijing cannot cure her; many do not even know what illness it is. So now we are asking the world — can somebody help u

Excerpt from Usenet news post, April 1995

Via the worldwide computer Internet and other means of communication, physicians from coast to coast in the United States and at least 17 other countries have helped their mainland China colleagues treat a university student with a challenging array of signs and symptoms.

Journal of the American Medical Association 13 Dec 1995

We need more MAN then feds so Everyone run wild, all of london and others are invited! Pure terror and havoc & Free stuff.... just smash shop windows and cart out da stuff u want! Oxford Circus!!!!! 9pm

Excerpt from BlackBerry Messenger IM 8 August 2011

Over the weekend parts of London descended into chaos as riots and looting spread after a protest organised around the yet unexplained shooting of a man by Police. [...] But while Twitter has largely been the venue of spectators to violence and is a handy public venue for journalists to observe, it would appear the non-public BlackBerry BBM messaging network has been the method of choice for organising it.

Techcrunch.com 8 August 2011

specific content. Understanding how information is changed and distorted by people and how errors are introduced as information is propagated is a key theoretical challenge for disinformation engineering. Developing this theory would facilitate more effective design of informational content and structure, such that the information propagated through the network has our desired outcomes as it achieves its critical mass and most impactful state.

CONCLUSION

As we stated earlier, and must re-emphasize: We are not advocating the use of social computing for criminal, abusive, or antisocial purposes. In fact, all of us want social computing to be leveraged to achieve individual and societal good.

However, this exercise does reveal interesting aspects of social computing, and through the articulated research challenges we can see some things we might not otherwise have appreciated. There are facets in all of our research challenges that are very similar to existing streams of research. Researchers are actively studying how to model individualized aspects of human personality, taste, and interaction style. There is clearly also foundational theory and practice upon which to build disinformation engineering and make disinformation propagation a reality.

The concerns we raise go beyond our basic design exploration. Our working assumption was that our design and research would be a set of techniques and systems that would be external to the existing social computing systems we might want to change or influence. We did not explore the specific user risks or technical advantages that are provided to the entity that manages, runs, or owns a social computing platform. An owner or manager could more simply use interface design techniques to suppress or highlight specific information or to manipulate the social networks of users. These risks can be more profound when social computing systems have cross-platform federation through a common owner.

The risks for crowdsourcing, user-generated content, citizen science, citizen journalism, and the wider range of social computing systems are real. It's time for us as designers of these systems to think more carefully about how those systems might put users at risk. Further, it's time to understand these systems might put communities and society at risk and we should take care to design systems that mitigate those risks.

David W. McDonald is an associate professor and chair of the Department of Human Centered Design & Engineering at the University of Washington. His research spans areas of social computing from analysis of user activity through the design and evaluation of tools to facilitate mass interaction in social computing systems.
→ dwmc@uw.edu

David H. Ackley is an associate professor of computer science at the University of New Mexico. His prior research has involved neural networks and machine learning, evolutionary algorithms and artificial life, and biological approaches to security, architecture, and models of computation.
→ ackley@cs.unm.edu

Randal Bryant is dean and university professor in the School of Computer Science at Carnegie Mellon University. His research has covered a range of computing theory. Currently he focuses on how existing networks of computers can be repurposed to solve large-scale computing problems.
→ randy.bryant@cs.cmu.edu

Melissa Gedney is currently an associate at Digital Promise, and contributed to this article while working at the Wilson Center Science and Technology Innovation Program. She holds a B.A. in political science from the George Washington University.
→ melissa@digitalpromise.org

Haym Hirsh is dean and professor of computing and information science at Cornell University. He previously served as director of the Division of Information and Intelligent Systems at the National Science Foundation. His research has focused on machine learning, classification, and prediction techniques.
→ haym.hirsh@cornell.edu

Lea Shanley founded the Commons Lab within the Wilson Center's Science and Technology Innovation Program. She has also served as a fellow on the Mapping Science Committee of the National Academy of Sciences. Her research focused on community-based action research in geographic information science at the University of Wisconsin-Madison.
→ lshanley@wisc.edu

DOI: 10.1145/2675705 1072-5520/14/11 $15.00