

Global Manhunt Pushes the Limits of Social Mobilization

Iyad Rahwan, Sohan Dsouza, and Alex Rutherford, *Masdar Institute of Science & Technology, UAE*

Victor Naroditskiy, James McInerney, Matteo Venanzi, and Nicholas R. Jennings
University of Southampton, UK

Manuel Cebrian, *National ICT Australia and Massachusetts Institute of Technology*

Using social media and only the targets' mug shots, a team competing in the US State Department-sponsored Tag Challenge located three of five targeted people in five cities in the US and Europe in less than 12 hours.

To commemorate the Internet's 40th anniversary in 2009, the US Defense Advanced Research Projects Agency (DARPA) launched its "Red Balloon Challenge." The challenge aimed to test the limits of social mobilization and rapid information gathering using social media. It required competitors to locate 10 weather balloons tethered at random locations all over the United States.

The winning team, based at MIT, found all balloons within nine hours using social media.¹ Their strategy relied on an incentive scheme that rewarded people both for reporting balloon sightings and for recruiting their friends to look for balloons.² Further theoretical work proved that the strategy is optimal in terms of minimizing the investment to recover the information³ and is the most robust to misinformation.⁴

In March 2012, the "Tag Challenge" (www.tag-challenge.com), funded by the US State Department, raised the bar significantly. This challenge required teams to locate and photograph five people in cities across two continents—Bratislava, Slovakia; Washington, D.C.; London; New York

City; and Stockholm—within 12 hours. The winning team would receive \$5,000.

Each team received a single mug shot of each targeted person (or "suspect") at 8:00 a.m. local time on the day of the competition. As Figure 1 shows, each volunteer target wore a T-shirt with the competition logo (the appearance of which was also not known until the first mug shot was released) and was instructed to follow a 12-hour itinerary designed to reflect a normal day. For example, the New York City suspect started at Columbia University, had breakfast at a nearby cafe, took the subway to the World Trade Center site, went shopping, and so on.

The Tag Challenge was significantly harder than the Red Balloon Challenge because it required locating people in highly populated cities, where the pace of life can help people "hide in plain sight,"⁵ and making it much more difficult to spot the suspects than to find large red balloons. Furthermore, people are mobile, making it difficult to rule out locations that have already been visited. Although the suspects were not explicitly hiding or in disguise, finding a person moving about in a city such as New York, with




Figure 1. Target mug shots. Participating teams received a single mug shot of each “suspect” at 8:00 a.m. local time on the day of the competition.

a population of 8.2 million people, seemed nearly impossible. In particular, whether social media can help accomplish such a difficult task is an open question, as evidenced by the difficulty of finding suspects in police investigations.⁶

Despite these challenges, our team won the competition by locating three of the targets using a Web platform, a mobile application, and an incentive scheme. None of the team members were in any of the target cities.

Figure 2 shows the locations of the targets our team found and the approximate local time when we submitted their photos to the organizers. The London and Stockholm targets remained at large; pursuing them after the allotted 12 hours was not part of the competition.

The competition afforded a wide range of possible strategies, including search engine optimization, cash incentives, and deceptive measures. We based our approach on an understanding of the key challenges of rapid mobilization of crowdsourcing teams.

CROWDSOURCING CHALLENGES

Crowdsourcing rapid information gathering involves at least three distinct problems: mobilizing participants, aggregating information, and verifying information.

Mobilization

For search in social mobilization to succeed, individuals must be motivated to conduct the search and to participate in disseminating information. In an attempt to replicate Stanley Milgram’s “small world experiment,”⁷ Peter Dodds and colleagues found that most message-forwarding chains



Figure 2. Location of the five targets and local time when the three targets were found.

observed empirically terminate prematurely.⁸ Specifically, they concluded that “although global social networks are, in principle, searchable, actual success depends sensitively on individual incentives.” In other words, a key factor in social mobilization is the incentive challenge. Others have observed that the success of crowdsourcing mechanisms generally varies depending on the details of the financial incentive scheme in place.⁹

Although mobilization requires motivating people, in natural disaster situations, people seem to intrinsically want to contribute. Once they identify a focal point such as Ushahidi (www.ushahidi.com) or Cosm (www.cosm.com) where they can submit information, it becomes a hub, and everyone contributes and invites others to do so. Thus, providing motivation for participation and spreading information is less of a concern in such scenarios, and the efforts can focus on information aggregation and verification.

However, such intrinsic motivation often is not as strong or widespread. Consider the problem of finding a single missing person, as in the Tag Challenge, or a criminal. The sparsity of the information being sought requires motivating people to route the problem to others who are better placed to report on it—for example, to someone who lives in the same area as the target or who is likely to come across the target in the course of daily life. This is a much more

Even with the right incentives to motivate people to act, open questions remain regarding the crowdsourcing system's economics and efficiency.

difficult task because the person who eventually finds the target is only one element in a chain of participants. This is in contrast to more conventional crowdsourcing settings such as Ushahidi, where the bulk of each task begins and ends with an individual. In short, tasks on Ushahidi are independent, whereas each successful search task requires an uninterrupted chain of motivated participants.

Traditionally, finding missing persons or criminals requires a central actor to advertise a reward or ransom, respectively, as an incentive to route the information to others. In fact, someone who spots the criminal or has useful information about the criminal's whereabouts has an incentive to hide that information from others to avoid sharing the reward. These features act as barriers to large-scale mobilization in time- and resource-critical situations. Even with the right incentives to motivate people to act, open questions remain regarding the crowdsourcing system's economics and efficiency. Academic work in crowdsourcing is now exploring this problem, for example, to optimize the efficiency of microtask markets such as Amazon Mechanical Turk or CrowdFlower.⁹

Aggregation

Given enough motivated participants, the next challenge is to aggregate their responses. Even if we assume, for now, that all contributions are reliable, visualizing and synthesizing the information into an actionable form requires careful thought.

This is the main challenge that platforms such as Ushahidi and Cosm seek to address.¹⁰ If there is a forest fire or an earthquake, Ushahidi can aggregate geotagged reports posted by thousands of people, then visualize what is going on using the resulting crisis map. The system has proven efficient in rapid information aggregation through various deployments, including the aftermath of the devastating 2010 earthquake in Haiti.

However, aggregation is less trivial in some scenarios. For example, when participants contribute a high volume of information, it is often necessary to identify what is most relevant to avoid clutter in the visualized maps. This task is even more difficult when aggregating different types of data—for example, when combining spatial and temporal communication data to characterize the responses to massive emergencies.¹¹

Verification

Experience with Ushahidi also highlights the verification problem in crowdsourcing. Verification is important because people stand to benefit from submitting false reports—for example, to get more food or receive help sooner, often at the expense of others.

When conducting searches, false positives from well-intentioned participants and malicious reports from those seeking rewards are also likely. The Metropolitan Police in London recently released the photographs of 5,000 suspects believed to be involved in the London riots of 2011.⁶ In such events, cases of mistaken identity are inevitable, especially when matching against poor-quality CCTV images. This is known to happen with general crime appeals from television programs such as *Crimewatch* in the UK, and *America's Most Wanted* in the US.

In one of the first serious attempts to tackle automatic verification, the Ushahidi team developed the SwiftRiver set of algorithms, which use machine learning techniques to classify information and content such as Twitter accounts or individual tweets to filter the relevant sources. Nevertheless, verification remains a significant challenge, especially in contentious or competitive domains.¹²

OUR APPROACH

Our strategy addresses these challenges, focusing primarily on the mobilization aspect, as we believe that the most difficult factor is getting people to take action.

Strategy

Following the winning strategy for the Red Balloon Challenge, we used the incentive scheme depicted in Figure 3, which aimed to encourage people to both report if they found a target and help recruit other people to search for the target. We described the strategy as follows: “If we win, you will receive \$500 if you upload an image of a suspect that is accepted by the challenge organizers. If a friend you invited using your individualized referral link uploads an acceptable image of a suspect, you also get \$100. Furthermore, recruiters get \$1.00 for each recruit they refer to sign up with us through their individualized referral link, up to the first 2,000 recruits.”

The incentive to refer others had to be significant because attracting additional participants to the challenge would only decrease an individual's chances of getting

a reward. Thus, paying people for referring their friends, as well as for finding suspects, fundamentally changes the incentivized behavior.

Aggregation and verification were not major issues in this competition because we received relatively few submissions. In particular, we could manually verify that the target's photograph matched one of the mug shots the organizers provided. And, because we received a single submission from each city, aggregation was not required.

Each target suspect's T-shirt had a unique and previously unknown code, allowing organizers to determine whether the suspect had really been found. We had to make sure this code was legible in the photograph. In every case, we contacted the submitters directly; if part of their backstory had been inconsistent with expectations, we would not have trusted their submission. These activities were all possible because of the small number of submissions. Clearly, such a detailed approach would not be possible at larger scale, although we recently proposed a mechanism for crowdsourcing the verification tasks simultaneously with the search tasks.⁴

Implementation

To facilitate our strategy, we designed and built a Web application that lets people submit photographs of suspects online and recruit other people to join the effort. Where applicable, to reward participants for recruitment, we needed to be able to trace a new recruit to the referring recruit. We therefore provided an individualized referral link for participants to share on social networks, a process made easier through the site's template email and social media sharing links.

The heat map in Figure 4 shows the distribution of clicks to our team's website over the week leading up to the challenge.




Figure 3. Incentive scheme. The strategy encourages people to both report if they found a target and help recruit others to search for the target.

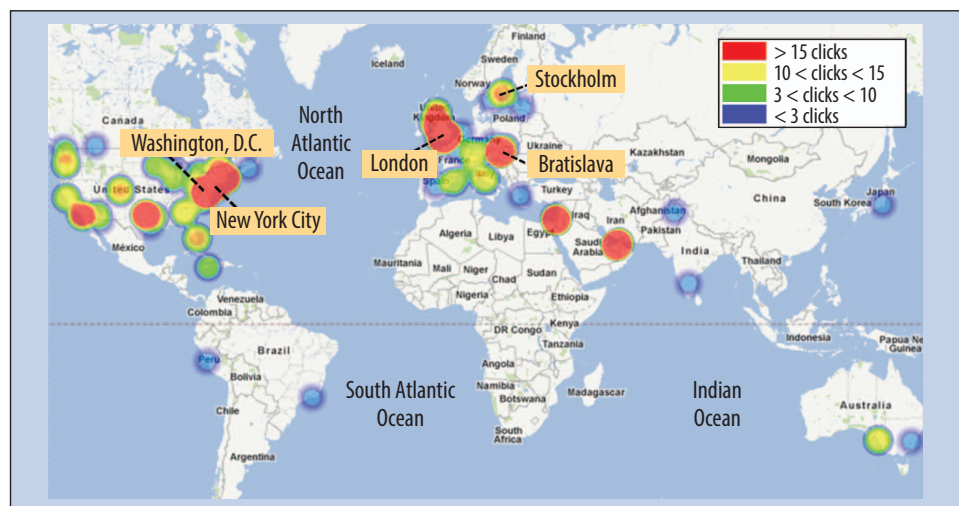


Figure 4. Heat map showing the distribution of visitors to the team's site. The areas of interest generally focused around the target cities: New York, Washington, D.C., Stockholm, London, and Bratislava.

In addition to the Web application, we built an Android mobile phone app that let people view the list of suspects as they were released and submit photos from their phones. We also created a Twitter feed and Facebook page for people to follow and "like," respectively, so they could receive updates.

Somewhat to our surprise, participants did not use either the Web platform or the mobile app much for submission during the competition. Successful participants in our team preferred to submit their photos via email, suggesting that when important information needs to be reported, people value direct communication rather than going through official channels. Perhaps this can be explained either on the grounds of convenience as each new platform involves

a learning curve or because of concern that highly valuable information such as the winning photos would get lost in a sea of noise. The website and app are both components of the crowdsourcing platform our team is developing, so addressing these issues is important for scalability.

Nevertheless, both the Web platform and mobile app enhanced our credibility by demonstrating commitment and Web presence. This helped significantly with the final element of our approach, which was to generate as much attention as possible for our team on blogs and news websites. We sent out press releases explaining our interest in the challenge, and posted notes about our team in related forums.

OTHER TEAMS

We are aware of five serious opposing teams. Some competitors, such as team @WeTagChallenge, used search engine optimization tactics to raise the profile of their websites above those of other teams. Another team with Twitter handle @TAG_challenge apparently attempted a “man-in-the-middle” attack using a name that was similar to the official @TAGchallenge account. We suspect they might have been attempting to pose as the actual Tag Challenge organizers to intercept submissions, even to the extent of duplicating the official tweets.

Yet another team, with the Twitter handle @TeamRave, ran its campaign primarily on Twitter and its website. They promised the entire \$1,000 share for each accepted suspect picture, leaving no incentive for people to recruit others.

Tag Team was the most serious and most aggressive competitor. They used an incentive scheme similar to ours (\$400 to finder, \$100 to referrer), but added a twist. Instead of using the remaining prize money to build a critical mass of recruits, they promised to give \$500 per suspect to charity. Their strategy for spreading awareness consisted predominantly of updating their Twitter account, surfing trending hashtags, and tweet-spamming individuals and social, governmental, and private organizations in the target cities, often with an explicit plea for a retweet. Most of these tweets were ignored and, we believe, reduced the team’s credibility. Interestingly, many of these tweets mentioned the @TAG_challenge handle, which was an imposter account.

In the days leading up to the challenge, Tag Team also tweeted messages targeting our team, including specific attacks on our team’s competence and members and tweets encouraging people not to support us. Tag Team also acquired a large number of fake Twitter followers, ostensibly purchased. Two days before the challenge began, their number of Twitter followers went from 37 to more than 12,000, literally overnight. Presumably, they sought to give the impression that they were the most promising team, thereby harnessing the bandwagon effect. This, however, does not seem to have helped, as their following did not subsequently increase much.

Most competitors seemed to focus purely on social media, using Twitter almost exclusively to spread their message. This narrow strategy was not sufficient, not least because several teams were perceived as spammers. As a result, none of the other teams successfully located more than one of the five suspects.

Our team was far more selective in its tweets and social media strategy, but we also made sure to generate as much attention as possible in blogs and online news. Even a small initial advantage in credibility and awareness can lead to eventual victory in this competitive domain.

The Tag Challenge aimed to test the limits of social mobilization in a time-critical competitive environment.

We can explain this phenomenon in the following way. A valid photograph of one of the targets has value because the organizers provide prize money. Hence, any participant in possession of such a photograph must decide carefully which team to send it to. The first barrier for a team in receiving this information is for the participant to simply be aware of it. But second, and just as importantly, the participant must trust the team leadership to give the promised reward. Finally, the participant must have confidence in the likelihood of the team’s victory—that is, an expectation that other participants would also submit their photos to that team.

Focusing on these qualities, and enhancing them through carefully considered incentives, gave our team an early lead in finding suspects. We then made sure to publicize these early successes in real time on social media, which contributed to even greater levels of credibility and awareness. This positive feedback loop ensured that we were always ahead of the other teams.

Nevertheless, other teams ranked higher than we did in the Klout score (www.klout.com), which uses data mining to measure influence and credibility across a user’s social network. As Table 1 shows, we use this score to quantify the influence of each team’s online presence.

Tag Team, which sent out numerous tweets and appeared to have purchased thousands of Twitter followers, ranked first by Klout score. This raises an interesting question: how important is this score for time-critical tasks and how well does it filter out attempts to engineer a higher score? It is possible that the score works for more conventional online behavior, but not for a time-critical task such as the Tag Challenge. Developing a time-critical Klout score is therefore an interesting open question. We were also able to classify the behaviors of each of the teams using

Table 1. Klout score, Reach, and Amplif for the five main teams participating in the Tag Challenge.

Team's Twitter Handle	Klout score	Reach*	Amplif*	Style
@WeTagChallenge	32	126	6	Socializer
@TagTeam_	26	54	4	Conversationalist
@CrowdScannerHQ (winner)	21	38	3	Explorer
@Tag_Challenge	16	20	3	Dabbler
@TeamRave	10	1	2	Observer

* Reach measures the number of followers and Amplif measures a team's ability to cause viral propagation.

Klout. Figure 5 shows the range of different interaction styles.

Another tactic used against our team seems to have been direct attack. Between GMT 16:57 and GMT 17:00, we received four submissions from the same user, each for a different suspect. The submitted files were copies of the same executable file and contained references to various critical operating system libraries. We took this as, at best, an attempt to annoy us, or, at worst, an attempt to make our systems vulnerable to an attack. We do not know the attacker's affiliation, but we cannot rule out malicious intent.

LESSONS LEARNED

The Tag Challenge aimed to test the limits of social mobilization in a time-critical competitive environment. We found only three of the five suspects within the allotted 12 hours, in contrast to the Network Challenge, where all of the balloons were found in less than nine hours. However, the Tag Challenge was arguably more difficult. Thus, the result of the competition—finding three of the five suspects—is a testimony to the power of social media, which pushes the boundaries of social mobilization.

Although members of our team were geographically close to one of the cities where suspects were present (Southampton is one hour away from London), we failed to find the London suspect. This could be a first lesson: geographical proximity does not guarantee anything, and ours is essentially a delocalized technique. Being in four countries, across two continents, and having to deal with three languages can also hinder recruitment. But a comprehensive understanding of the role of these factors requires more systematic investigation than a single run of the challenge affords.

Interestingly, social networks did not play an explicit




Figure 5. Representation of different Klout interaction styles.

role in our victory. Although Facebook and Twitter are the most natural tools for social mobilization, we found it extremely difficult to promote our Facebook page and Twitter accounts. Other teams seemed to face the same predicament. For example, all of our efforts amounted only to a few dozen Facebook “likes” and Twitter followers.

Nonetheless, these media were extremely important for at least two reasons. First, they added to our team's credibility by displaying its history in the form of wall posts and tweets. Second, the reach extended beyond our immediate fans and followers. Although fewer than 50 people “liked” our Facebook page, more than 300 unique users (as measured by Facebook statistics) were exposed to it before the competition, and we consequently received around 500

hits on our team's main website in the 24 hours leading up to and during the challenge. Furthermore, Facebook and Twitter enabled a broadcast-style communication to the most important people—those who explicitly expressed their interest. In particular, updates about our successes in finding targets were immediately posted to Facebook, Twitter, and our website, sending an encouraging signal that we were the strongest team.

Using crowdsourcing extensively and indiscriminately could result in information overload for participants, to the detriment of any one task.

Most visitors to our website came after reading a post about the challenge on Slashdot, and comments we left in forums discussing the challenge. More traditional forms of online media also helped make people aware of the challenge and our team: we were mentioned on CNET and ZDNet and by our respective universities' press teams.

It is also worth noting that two of the three people on our team who found a suspect had an existing interest or connection to crowdsourcing. A computer science graduate interested in crowdsourcing found the Bratislava suspect. David Alan Grier, president of the IEEE Computer Society, produces a regular podcast on crowdsourcing and found the suspect in Washington, D.C. In fact, most of our team's participants had an existing interest in crowdsourcing, an affinity that is important to motivation in social mobilization,¹³ and most postcompetition survey respondents indicated interest in social mobilization and crowdsourcing.

As in financial markets where investors copy successful strategies and lose profitability over time, query incentive cascades,¹⁴ although novel in the balloon challenge, were copied in this challenge. An open question is how this would play against the use of this strategy, if at all. Furthermore, using crowdsourcing extensively and indiscriminately could result in information overload for participants, to the detriment of any one task. In this case, it would be useful to study how to target tasks to specific individuals, depending on their location, interests, mobility, network structure, and other relevant features.

The influence of competition on search efficiency deserves some attention. Because we found every suspect that others found, we can assume that competition with the others did not harm our network's propagation. However, a conclusive answer requires further investigation into the role of competition.

We foresee a growing potential for applications that support time-critical social mobilization in the real world. In particular, one application uses social networks to find missing children or missing persons.

Our experience complements recent lessons learned about the use of crowdsourcing in processing large numbers of satellite images in the search for computer scientist Jim Gray, who was lost at sea—a search that unfortunately failed.¹⁵ We need more dedicated platforms to facilitate such efforts, allowing efficient sharing of useful information.

As is the case with the Tag Challenge and the Red Balloon Challenge, such applications will likely benefit from explicitly rewarding the routing of information and recruiting the right individuals for the search, such as those familiar with the geographies or urban areas of interest.¹⁶ People might also use these techniques to report environmental violations (for example, someone dumping pollutants in a river), or to map natural disasters in real time (such as the spread of a forest fire) or after they occur (such as damage after an earthquake).¹⁰ Again, social networks could provide useful information about volunteers who might have access to geographically relevant information. Crowdsourcing using social networks and human-based sensing can also help encourage data collection by citizen science volunteers.¹⁷ For example, the Cornell Lab of Ornithology uses volunteers to collect large-scale scientific data about urban living and wildlife.¹⁸

In short, this technology can help mobilize volunteers in faraway places to achieve highly distributed, possibly time-critical, tasks. We are working on a general “social mobilization app” that anyone can use to rapidly build and coordinate teams of volunteers to address difficult, geographically distributed challenges.¹⁷

Nearly 60 years ago, Stanley Milgram redefined the notion of social distance with his “six degrees of separation” experiment,⁷ showing that we are, on average, only six hops of friendship away from anyone on Earth. More recently, Facebook found the degree of separation to be only four in its network.¹⁹ Endeavors such as the Tag Challenge will redefine our conception of the temporal and spatial limits of technology-mediated social mobilization in the Internet age, showing that we can find any person who is not intentionally hiding in less than 12 hours. **■**

Acknowledgments

We thank the organizers of the Tag Challenge for running the experiment, as well as all of our volunteers, especially those who found the targets.

References

1. J. Tang et al., “Reflecting on the DARPA Red Balloon Challenge,” *Comm. ACM*, vol. 54, no. 4, 2011, pp. 78-85.

2. G. Pickard et al., "Time-Critical Social Mobilization," *Science*, vol. 334, no. 6055, 2011, pp. 509-512.
3. M. Cebrian et al., "Finding Red Balloons with Split Contracts: Robustness to Individuals' Selfishness," *Proc. 44th Symp. Theory of Computing (STOC 12)*, ACM, 2012, pp. 775-788.
4. V. Naroditskiy et al., "Verification in Referral-Based Crowdsourcing," *PLOS ONE*, vol. 7, no. 10, 2012, e45924.
5. A. Rutherford et al., "The Limits of Social Mobilization," *Proc. Nat'l Academy of Sciences*, 2013, in press.
6. D. Wilcock, "Police 'Facewatch' App Targets London Riot Suspects," *The Independent*, 26 June 2012; www.independent.co.uk/news/uk/crime/police-facewatch-app-targets-london-riot-suspects-7887778.html.
7. S. Milgram, "The Small World Problem," *Psychology Today*, vol. 1, no. 1, 1967, pp. 66-67.
8. P.S. Dodds, R. Muhamad, and D.J. Watts, "An Experimental Study of Search in Global Social Networks," *Science*, vol. 301, no. 5634, 2003, pp. 827-829.
9. W. Mason and D.J. Watts, "Financial Incentives and the 'Performance of Crowds,'" *Proc. ACM SIGKDD Workshop Human Computation*, ACM, 2009, pp. 77-85.
10. M. Hesse, "Crisis Mapping Brings Online Tool to Haitian Disaster Relief Effort," *The Washington Post*, 16 Jan. 2010; www.washingtonpost.com/wp-dyn/content/article/2010/01/15/AR2010011502650_Comments.html.
11. J. Bagrow, D. Wang, and A. Barabási, "Collective Response of Human Populations to Large-Scale Emergencies," *PLOS ONE*, vol. 6, no. 3, 2011, e17680.
12. H. Gao, G. Barbier, and R. Goolsby, "Harnessing the Crowdsourcing Power of Social Media for Disaster Relief," *IEEE Intelligent Systems*, vol. 26, no. 3, 2011, pp. 10-14.
13. J.L. Iribarren and E. Moro, "Affinity Paths and Information Diffusion in Social Networks," *Social Networks*, vol. 33, no. 2, pp. 134-142.
14. J. Kleinberg and P. Raghavan, "Query Incentive Networks," *Proc. 46th Ann. IEEE Symp. Foundations of Computer Science (FOCS 05)*, IEEE, 2005, 132-141.
15. J. Hellerstein and D. Tennenhouse, "Searching for Jim Gray: A Technical Overview," *Comm. ACM*, vol. 54, no. 7, 2011, pp. 77-87.
16. D. Horowitz and S.D. Kamvar, "Searching the Village: Models and Methods for Social Search," *Comm. ACM*, vol. 55, no. 4, 2012, pp. 111-118.
17. E. Hand, "Citizen Science: People Power," *Nature*, vol. 466, no. 7307, 2010, pp. 685-687.
18. R. Bonney et al., "Citizen Science: A Developing Tool for Expanding Science Knowledge and Scientific Literacy," *BioScience*, vol. 59, no. 11, 2009, pp. 977-984.
19. L. Backstrom, "Four Degrees of Separation," ArXiv:1111.4570, 2011.

Iyad Rahwan is an associate professor and head of the Computing and Information Science Program at Masdar Institute of Science & Technology, UAE, and Honorary Fellow at the University of Edinburgh, UK. His research interests include artificial intelligence and computational

social science. Rahwan received a PhD in artificial intelligence from the University of Melbourne, Australia. Contact him at irahwan@masdar.ac.ae.

Sohan Dsouza is a research engineer at Masdar Institute of Science & Technology, UAE. His research interests include social computing, machine learning, and cognitive agents. Dsouza received an MSc in informatics from the British University in Dubai. Contact him at sdsouza@masdar.ac.ae.

Alex Rutherford is a postdoctoral research fellow at Masdar Institute of Science & Technology, UAE. His research interests include spatial analysis, machine learning, and computational social science. Rutherford received a PhD in computational physics from University College London. Contact him at arutherford@masdar.ac.ae.

Victor Naroditskiy is a postdoctoral research fellow at the University of Southampton, UK. His research interests include incentive mechanism design, crowdsourcing, and social networks. Naroditskiy received a PhD in computer science from Brown University. Contact him at vn@ecs.soton.ac.uk.

James McInerney is a PhD candidate at the University of Southampton, UK. His research interests include human location behavior prediction and decision making under location uncertainty. McInerney received an MA in computer science from Oxford University, UK. Contact him at jem1c10@soton.ac.uk.

Matteo Venanzi is a PhD candidate at the University of Southampton, UK. His research interests include crowdsourcing, trust, and data fusion. Venanzi received an MEng in computer engineering from University of Rome La Sapienza, Italy. Contact him at mv1g10@ecs.soton.ac.uk.

Nicholas R. Jennings is a professor of computer science in the School of Electronics and Computer Science at the University of Southampton, UK, and a chief scientific advisor to the UK government. His research interests include multiagent systems, human-agent interactions, and artificial intelligence. Jennings received a PhD in artificial intelligence from the University of London. He is a fellow of IEEE and the Royal Academy of Engineering. Contact him at nrj@ecs.soton.ac.uk.

Manuel Cebrian is a senior scientist with National ICT Australia at the University of Melbourne and a senior research fellow with the MIT Media Laboratory. His research interests include crowdsourcing, urban economics, and behavioral game theory. Cebrian received a PhD in computer science from Universidad Autonoma de Madrid, Spain. Contact him at cebrian@mit.edu.



Selected CS articles and columns are available for free at <http://ComputingNow.computer.org>.