

Crowdsourcing Change

Hari Sundaram^{1,2}

¹School of Computing, Informatics and Decision Systems Engineering

²School of Arts, Media and Engineering

Arizona State University

Tempe, Arizona 85281.

hari.sundaram@asu.edu

Abstract

This paper discusses the role of computing in engendering cooperation in social dilemmas such as sustainability and public health. These cooperative dilemmas exist at a large scale, within heterogeneous populations. Motivated by analysis of cooperation from empirical field studies, we argue that an integrative computational framework that analyzes social signals and verifies behaviors through smartphone sensors can shape and mold individual decisions to cooperate. We discuss four interconnected technical challenges and example solutions. The challenges include community discovery algorithms for construction of small homogenous groups, persuasion of individuals in resource constrained networks, activity monitoring in the wild and detection of large scale social coordination. We briefly discuss new applications that arise from a computational infrastructure for cooperation, including fighting childhood obesity, cybersecurity and improving public safety.

1 Introduction

Sound, images and text have been long responsible for persuading individuals to act—Martin Luther King’s powerful speeches during the civil rights movement, for example, led to peaceful change. More recently in the Arab Spring, socio-technical networks (SMS services, Facebook, Twitter), have facilitated—through sharing of images, videos and social messaging—large scale protests in support of political change (Kirkpatrick, 2011). Socio-technical networks facilitated by multimedia communication also suggest an alternate outcome—the possibility of rich media networks to address some of the pressing challenges facing contemporary society including sustainability and public health. For example, how can sustainable behaviors—such as reducing individual energy consumption—be encouraged? How can participation in activities that reduce overall healthcare costs—such as compliance with preventive care routines and leading healthy lifestyles—be supported? The main technical question: how can we design a socio-technical network, with appropriate multimedia communication and representation to “nudge” individuals to cooperate towards these positive societal outcomes? That is, can we design an infrastructure to trigger meaningful activity in the crowd? While we are increasingly witnessing entrepreneurial activity in this space—JouleBug (<http://joulebug.com>) an iPhone application that encourages sustainable energy use, is but one example—we have barely scratched the surface of the research needed to ensure the success of such cooperative networks.

Technical components to solutions to address almost every important social problem, such as technologies that help the elderly, or that staying healthy, requires us to capture, interpret and distribute heterogeneous data types (images, videos, GPS, accelerometer, social communication data) from a population. These problems non-trivial because they are data sparse, key semantics appear at multiple scales, some semantics arise due to interaction between types, and finally because different data types are sampled at different time-scales. We can still address these challenges, because human activity is structured, constraining the relationships between different types of media. Multimedia research that focuses on large scale problems

at the level of the population is “big-data” research. The amount of data to be gathered, analyzed and distributed to address issues like sustainability or increasing the participation in voting at the level of the population of the United States is extraordinarily large, and we simply do not understand well how to design multimedia systems and analysis frameworks that trigger coordination (a new semantic process) among individuals in the population.

The rest of this paper is organized as follows. In the next section, we shall present some of the main ideas that animate the study of for cooperation in groups. In Section 3 we shall discuss key technical challenges to support cooperative behavior. In Section 4, we shall discuss three applications of a computational framework that enables cooperation. Finally in Section 5, we shall summarize the key ideas and present our conclusions.

2 The Scaling of Cooperation

In many cooperation dilemmas such as sustainability (recycling paper; using public transport instead of driving to work), public health (vaccinations; taking flu shots), information commons (contributing to wikipedia), the rational course of action is not to cooperate but to ‘free ride’ on the contributions of others. The main argument against cooperation was elegantly formulated by Hardin (1968) who said that in a group of N herders, one herder who overgrazes the meadow will keep all the profits while the losses due to overgrazing shall be distributed amongst all N herders. He argued that this scenario was simply an example of the well-known prisoners dilemma game and that the herders will converge to a Nash equilibrium where no one cooperates. He termed the scenario “the tragedy of the commons” and argued that only privatization of the resource or governmental regulation can help avoid the tragedy. Hardin’s paper was highly influential and set the tone for subsequent policy debates—we now view large scale cooperative dilemmas such as sustainability and public health within the purview of the state, with no role for ordinary citizens.

Ostrom (1990) and her colleagues developed a framework that helped explain the existence of robust cooperation¹ amongst ordinary citizens faced with cooperative dilemmas. A careful examination of empirical evidence revealed that successful cooperation occurred in small homogenous groups, where the individuals had the agency to collectively set the rules governing the resource, with delineated access rights, where the activities were verifiable through monitoring, with the ability to set graduated sanctions on the offenders.

How do we scale up the work of Elinor Ostrom? Dilemmas such as ‘global warming’ and ‘public health’ exist at large scales and exhibit heterogeneity—it is unclear if empirical observations that guided her work apply². There are two social trends that provide insight on the issue of scale. First, individuals are increasingly interconnected through large scale socio-technical networks—mobile networks, online social networks including Facebook and Twitter. Second, with a significant part of the population now accessing data through smartphones (Smith, 2011), we have a rich array of sensors in the phone to capture and monitor human behavior. We can develop an integrative computational framework that captures social signals (communication and participation in events) across networks of people and which senses their physical activity in the everyday through lightweight efficient sensors embedded in physical objects and smartphones. Analysis of social signals helps us understand the crucial role played by social ties when individuals take decisions to adopt behaviors. We can observe and analyze the context in which the behaviors occur through sensors in the smartphone. By coupling the analysis of communication with the analysis of actual behavior we derive insights into how to effectively shape and mold individual choices.

3 Technical Challenges in Engendering Cooperation

In this section, we shall identify some key technical challenges for engendering large-scale cooperation, including community discovery, engendering adoption in resource constrained networks, activity monitoring and rapidly identifying large scale change.

¹Irrigators in Valencia, Spain, have been cooperating since 1435 and herders in Törbel, Switzerland since 1483.

²Most of the successful examples of cooperation in the field (Ostrom, 1990) are of groups of several hundred people or smaller; one notable exception are the group of irrigators in Sri Lanka with several thousand individuals.

The first step towards facilitating cooperative behavior adoption within large heterogeneous populations is to unite similar individuals within the population into small homogenous groups or communities. Small homogenous groups will encourage the building of trusting relationships leading to greater cooperation. A community, in one operational definition, is a group of people interacting with each other in a consistent manner. Community detection algorithms (Newman and Girvan, 2004; Yang and Leskovec, 2013) identify the modular structure of a network, where nodes represent individuals and where links represent the interaction or similarity between individuals. These algorithms are closely related to the family of clustering algorithms. In recent work (Lin et al., 2012), we have developed a multi-relational soft-clustering objective function to focus on the constantly changing and co-evolving interaction contexts in online social media. The framework represents heterogeneous social contexts in social media—multi-relational and multidimensional social data—with tensors, which are a natural way to encode n -way relationship amongst entities. With this generalized objective function, different type of relations—such as user interaction, and media similarity—are considered simultaneously, allowing us to capture evolution of both user interaction and media use, within communities.

To persuade small homogenous groups to cooperate, we need to focus on three interrelated ideas: resource constrained individuals, identifying influentials, and personalized social signals. *First*, to model real-world behavioral adoption, we need to model individual with resource constraints—a feature missing in models of collective behavior (e.g. (Granovetter, 1978; Kempe et al., 2003)). Resource constraints influence individuals by affecting their ability enact a behavior (Ajzen, 1991) as well maintaining her social ties (Dunbar, 2008). A person interested in adopting a behavior (e.g. taking newspapers to a recycling station, going to cast a vote or taking a flu shot) may fail to do so, due to lack of resources—time, money or other tangible resources such as a car or a bicycle. *Second*, we need to identify influentials—individuals whose adoption of a behavior will persuade others to adopt (Watts and Dodds, 2007). We have developed an efficient algorithm to identify influentials as well the behaviors they need to adopt to maximize participation in resource constrained networks (Sarkar and Sundaram, 2013). In our framework, individuals adopt when they receive social signals of sufficient strength, are interested in the behavior and have the resources to adopt new behaviors. The main result is that individuals who are easily influenced and who connected to other easily influenced individuals are key to maximizing multiple behavior adoption. *Finally*, we need to design personalized social signals that persuade individuals to adopt. Two studies indicate that revealing hidden social norms is a powerful cue for cooperation. In a study of energy use Schultz et al. (2007) found that individuals were presented with affective symbols—a smiley face ‘☺’ when an individual household’s energy consumption was below the mean household consumption and a frownie face ‘☹’ when it was above the mean—in addition to mean neighborhood consumption, there was a net decrease in energy consumption over the entire population³. In another study Goldstein et al. (2008) demonstrated that presenting guests with the the descriptive norm “most guests reuse their towels” significantly improved towel reuse in comparison to the standard environmental message. We have developed a framework based on the idea of the “social proof” [Chapter 4, Cialdini (1993)], for generating personalized social signals for each resource constrained individual so that they are incentivized to adopt a set of behaviors. To persuade each individual to adopt a behavior, we identify relevant friends and other individuals with whom they share significant personal attributes, and then ‘re-wire’ their information network. The main result is that we can change behavior distribution within the group to achieve any target behavior distribution that satisfies certain constraints.

Robust verification of human activity in the wild is a major challenge to successful cooperation. Consider for example the case of an iPhone application such as JouleBug (<http://joulebug.com>) that encourages sustainable energy use; individuals receive points for participating in an sustainable activity (e.g. line drying clothes). However, since there is no verification of the activities, individuals can ‘free-ride’ and receive points for logging activities in which they do not actually participate. Absent a robust mechanism to verify activity, formation of trust within such networks is difficult. Without trust, these well-intentioned applications will end up in the Nash equilibrium where no one cooperates. Crowdsourcing the problem of labeling activity may be one mechanism. Recent work on crowdsourcing (Bigham et al., 2010; Lasecki et al., 2013) has concentrated on using paid workers at Mechanical Turk. However individuals in many social dilemmas such public health may be unwilling to pay to have their activity verified. Instead, crowdsourcing the verification

³OPower now provides this information to over a hundred power utility companies in the United States. <http://www.opower.com>.

amongst the members of the group attempting to cooperate is a more sustainable and scalable approach. The application “The Eatery” (<https://eatery.massivehealth.com>) attempts crowdsourcing of verification amongst its members. Each individual snaps a photo of their food and a set of strangers, chosen at random, rate their photo along a single dimension of “unhealthy–healthy.” Each individual thus receives a health score on their food eating habits. Notice that individuals are only scored on the foods they actually upload; it is easy to skip meals and only upload their most healthy meals.

Can we detect the onset of large scale, coordinated, behavior in large populations? Detecting onset of coordination action would lead us to identify viral behaviors. Viral behaviors, as well as rapid information diffusion can cause significant structural changes to networks within short time periods. Finding structural changes in massive social networks, in near real-time, is highly challenging. We examined two different approaches—one motivated by emergence of coordination in biological oscillators such as fireflies, and the other driven by the idea that changes to network structure are apparent in network spectra. In the first approach (De Choudhury et al., 2009) we developed a Bayesian user activity model, where the probability to act depends on the topic, prior activity history and behavioral coupling with neighbors. To predict coordination into the near future, we evolved the network by estimating sustained participation probability, and new user arrival rate. In the second approach (Lin et al., 2011) we introduced an innovative compressed sensing mechanism to encode the social data tensor streams as compact descriptors. Changes to the sensed coefficients allows us to quickly detect changes to the core tensor coefficients.

In this section we discussed several technical challenges to support large scale cooperation amongst individuals—detecting small homogenous groups, identifying influentials and designing social signaling schemes in resource constrained networks, detecting large scale change—and example mechanisms to address each challenge. Next we discuss novel applications arising from computational infrastructures for cooperation amongst individuals.

4 Novel Applications

A technical infrastructure that enables cooperation amongst individuals can lead to many novel applications beyond sustainability and public-health. We would like to highlight three different examples: childhood obesity, cybersecurity and public safety. In the United States and many parts of the world, childhood obesity is emerging as an important health issue (The Economist, 2012). Consider a mobile application that connects mothers of young children who are part of the same neighborhood and who are interested in providing more nutritious food for their children. Providing them with near real-time updates on what other mothers are providing their children for breakfast and what they are packing for lunch will reveal the social norm, and will encourage parents to provide similar nutritious food. In cyber security, a significant number of individuals do not patch up vulnerabilities in their software; appropriate social signals from their friends as well as individuals with whom they can relate can increase the adoption rate of important security patches. Finally, individuals can increase public safety by using their mobiles phones. An app that measures cell phone density can exchange information with nearby mobiles thus constructing a real-time spatio-temporal street map that reveals which streets are crowded and at what time. Such real-time safety information will be particularly important for women, children and minorities.

5 Conclusions

In this paper, we discussed the critical role of computing in engendering cooperation in support of large scale social dilemmas such as sustainability and public health. Unlike the small empirical studies of the cooperation, cooperation for large social dilemmas is challenging due the large scale and population heterogeneity. We argued that we can extract small homogenous groups by using algorithms for community discovery, thus fulfilling an important criteria. Within these groups, we need to identify influentials, as well as design signaling schemes so that individuals see signals from the relevant friends and strangers with whom they share important attributes, and are thus “nudged” into adopting the behavior. We discussed an important

unresolved question—monitoring of human activity in the wild; absence of effective, voluntary, mutual monitoring is a key barrier to the formation of trust in networks. Finally, we discussed several ways to address the challenge of discovering large-scale change. Beyond public health and sustainability, we discussed several interesting applications that benefit from a computational infrastructure for cooperation: fighting childhood obesity, cybersecurity and improving public safety.

While engendering cooperation is a multi-disciplinary effort, computing plays a critical role by analyzing social relations and the context in which successful behaviors occur. There are many important computing challenges that have attracted little attention. For a network of resource constrained individuals, the theoretical capacity (the fraction of available resources utilized) of the network of individuals willing to cooperate is yet unknown. Since the ultimate costs and benefits that accrue for example, from living healthy, are delayed, it is easy to abstain from participation. Developing automated, persuasive and personalized affective framing of information that takes advantage of our adherence to social norms is another open problem.

Bibliography

- Icek Ajzen. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211, 1991.
- Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and Tom Yeh. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, UIST '10, pages 333–342, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0271-5. doi: 10.1145/1866029.1866080. URL <http://doi.acm.org/10.1145/1866029.1866080>.
- Robert B Cialdini. Influence: The psychology of persuasion. *New York, Morrow*, 1993.
- Munmun De Choudhury, Hari Sundaram, Ajita John, and Dorée Duncan Seligmann. Social synchrony: Predicting mimicry of user actions in online social media. In *CSE'09, International Conference on Computational Science and Engineering, 2009.*, volume 4, pages 151–158. IEEE, 2009.
- R. I. M. Dunbar. Cognitive constraints on the structure and dynamics of social networks. *Group Dynamics: Theory, Research, and Practice*, 12(1):7 – 16, 2008. ISSN 1089-2699.
- Noah J. Goldstein, Robert B. Cialdini, and Vladas Griskevicius. A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35(3):472 – 482, 2008. ISSN 00935301.
- Mark Granovetter. Threshold models of collective behavior. *American journal of sociology*, pages 1420–1443, 1978.
- Garrett Hardin. The tragedy of the commons. *Science*, 162(3859):pp. 1243–1248, 1968. ISSN 00368075.
- David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.
- David D. Kirkpatrick. Wired and shrewd, young egyptians guide revolt, Feb 2011. URL <http://www.nytimes.com/2011/02/10/world/middleeast/10youth.html>.
- Walter S. Lasecki, Young Chol Song, Henry Kautz, and Jeffrey P. Bigham. Real-time crowd labeling for deployable activity recognition. In *Proceedings of the 2013 conference on Computer supported cooperative work, CSCW '13*, pages 1203–1212, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1331-5. doi: 10.1145/2441776.2441912. URL <http://doi.acm.org/10.1145/2441776.2441912>.

- Yu-Ru Lin, K. Selçuk Candan, Hari Sundaram, and Lexing Xie. Scent: Scalable compressed monitoring of evolving multirelational social networks. *ACM Trans. Multimedia Comput. Commun. Appl.*, 7S(1):29:1–29:22, November 2011. ISSN 1551-6857. doi: 10.1145/2037676.2037686. URL <http://doi.acm.org/10.1145/2037676.2037686>.
- Yu-Ru Lin, Hari Sundaram, Munmun De Choudhury, and Aisling Kelliher. Discovering multirelational structure in social media streams. *ACM Trans. Multimedia Comput. Commun. Appl.*, 8(1):4:1–4:28, February 2012. ISSN 1551-6857. doi: 10.1145/2071396.2071400. URL <http://doi.acm.org/10.1145/2071396.2071400>.
- M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113, Feb 2004. doi: 10.1103/PhysRevE.69.026113. URL <http://link.aps.org/doi/10.1103/PhysRevE.69.026113>.
- Elinor Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.
- K. Sarkar and H. Sundaram. How Do We Find Early Adopters Who Will Guide a Resource Constrained Network Towards a Desired Distribution of Behaviors? *ArXiv e-prints*, March 2013.
- Wesley P Schultz, Jessica M. Nolan, Robert B. Cialdini, Noah J. Goldstein, and Vladas Griskevicius. The constructive, destructive, and reconstructive power of social norms. *Psychological Science (Wiley-Blackwell)*, 18(5):429 – 434, 2007. ISSN 09567976.
- Aaron Smith. Smartphone adoption and usage. Technical report, Pew Internet & American Life Project, July 2011. URL <http://pewinternet.org/Reports/2011/Smartphones.aspx>.
- The Economist. The big picture, special report:obesity, December 2012. URL <http://econ.st/GzDTdZ>.
- Duncan J. Watts and Peter Sheridan Dodds. Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 34(4):441–458, 2007.
- Jaewon Yang and Jure Leskovec. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining, WSDM '13*, pages 587–596, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1869-3.