Computational Perspectives on Social Good and Access to Opportunity

Rediet Abebe
Department of Computer Science
Cornell University
red@cs.cornell.edu

Abstract

Computational techniques show immense promise to both deepen our understanding of socioeconomic inequality and inform interventions aimed at mitigating it. With the increasing collaborations across disciplines and the availability of large datasets, there is a wealth of areas where nuanced questions and novel techniques can reveal powerful observations and propose innovative solutions. My research focuses on this interface of algorithms, artificial intelligence, and applications to social good. In particular, I use algorithmic, computational, and network-based insights to study under-explored facets of inequality and provide solutions for effective interventions for improving access to opportunity.

Algorithms and artificial intelligence can be used to deepen our understanding of inequality and find interventions aimed at improving access to opportunity. None of the dimensions of inequality, such as economic, social, and cultural, individually conveys the whole picture, and many of these are complex and hard to measure (Tumin, Grusky, and Ku, 2012). This indicates immense potential for algorithmic and artificial intelligence techniques, in conjunction with the social sciences, to deepen our understanding of the structure of inequality. In turn, we can leverage these to find effective interventions to improve societal welfare. Thus far, computational applications have ranged from mechanism design for kidney exchange (Roth, Sonmez, and Unver, 2004), to machine learning for poverty mapping (Jean et al., 2016), and developing technology for under-resourced communities (Patel et al., 2010). While great strides have been made over the years, in many domains there are still major opportunities for exploration, and the prospect that we may be able to develop unified frameworks for applying computational insights.

To facilitate research at this interface, I co-founded and have been co-organizing a multi-institutional, interdisciplinary research group called *Mechanism Design for Social Good* with Kira Goldner (Abebe and Goldner, 2016). The goal of the group is to explore directions where algorithms, optimization, and mechanism design can be used to improve access to opportunity and resources across domains including housing, healthcare, and education. Since its inception in September 2016, the group has hosted numerous talks by researchers and domain experts from computer and information sciences, economics, global health, operations research, public policy, and sociology. We have forged research collaborations across disciplines and held the first *Workshop on Mechanism Design for Social Good* (MD4SG '17) at the ACM Conference on Economics and Computation (EC '17). The workshop highlighted research from our communities, invited domain experts from related fields, and held a discussion on best practices and future directions. To ensure that our work has impact within and beyond the academy, we are currently working with Thomas Kalil, Deputy Director for Policy for the White House Office of Science and Technology Policy under the Obama administration, to create collaborations with policy makers, non-government organizations, and other agencies with a shared objective.

In the next sections, I describe a sampling of my research, which uses algorithmic, computational, and network-based approaches to improve allocation of resources and targeted education efforts in various domains.

Allocation of Resources. Solutions aimed at mitigating inequality often involve fair and efficient resource allocations. This appears in many settings including housing, healthcare, and disaster relief. My research thus far has both established new notions of fairness and efficiency as well as introduced allocation mechanisms that satisfying various desirable criteria.

On the theoretical front, I have explored how to improve upon concepts of fairness. A popular definition of fairness in the literature is *envy-freeness*, which stipulates that no agent envies the allocation of another agent. Theoretical studies of fair division have treated this condition as a global constraint, which can be unnecessarily restrictive and lead to inefficient solutions. There is empirical evidence from social psychology indicating that local fairness notions, where agents compare their allocation to that of their friends' might be more natural. In Abebe, Kleinberg, and Parkes (2017), we use this insight to introduce new local definitions of fairness when the agents are on a graph (such as a social network). We show interesting relationships between these definitions and their global analogues, and give protocols for allocation of heterogeneous goods for a large class of networks.

A key challenge in resource allocation is that the individuals' valuations for the resources is not known to the

designer. As such, we would like to provide mechanisms that incentivize individuals to report their valuations truthfully. In Abebe et al. (2017a), we design a truthful mechanism for the allocation problem of Hylland and Zeckhauser (1979). The valuation of each agent for her allocation under this mechanism approximates her utility in the outcome that optimizes the Nash social welfare (a solution that combines natural fairness and efficiency requirements, but is cannot be obtained via a truthful mechanism in our setting). Through a set of experiments, we also show that our mechanism outperforms popular mechanisms in the literature. We are currently exploring performance of our mechanism for course allocations on a dataset obtained from the Wharton School of Business which contains students' reported cardinal valuations over courses.

Targeted Education and Interventions. Access to health care and health information is of major global concern. There are two prominent challenges that impede identification of major public health concerns and implementation of effective interventions; the first is that there are stark inequalities in the availability of health-related data by country, demographic groups, and socioeconomic status (Buvinic, Furst-Nichols, and Koolwal, 2014; UN, 2014). This data gap ranges from basic disease statistics, such as disease prevalence rates, to more nuanced information, such as public attitudes and information needs. The second is that, even when information about individuals is known, targeted education and campaigning efforts can use strategies that fail to take into account the heterogeneity in the population. Wakefield et al. (2003) discuss the need for segmenting of populations in anti-smoking campaigns to ensure that each group is exposed to the campaigns tailored to it. Exposure to ineffective campaigns can lead to inefficient and, at times, detrimental solutions. One example where there have been studies exploring the varying successes and limitations of a health campaign is Project D.A.R.E., aimed at preventing alcohol, tobacco, and illicit drug use among school-aged youths (West and ONeal, 2004).

Related to the data gap, a prominent challenge is that health organizations and policy-makers struggle to identify specific health information needs of individuals in developing nations. Such organizations face roadblocks in implementing effective targeted education efforts. To address some of this data gap, in Abebe et al. (2017b), we explore the role that search engines can play in understanding health information needs of individuals in Africa. We argue that search queries can and should be used to fill in gaps in our understanding of public health needs in data-sparse regions. We focus on HIV/AIDS and other infectious diseases; we demonstrate that searches related to these diseases show strong associations with the disease prevalence. Using topic modeling, we show that the themes that emerge cover standard information such as symptoms, testing, and medication as well as many hard-to-survey topics such as natural remedies, stigma, and concerns about gender inequality.

A rich line of research has studied the effectiveness of campaigning strategies based on person-to-person recommendations within a social network. This process, known as viral marketing or word-of-mouth recommendations, is popular in various domains including health campaigns. The goal here is to target an initial, small set of agents with a specific information with the hope of starting a cascade that reaches many members of the population. Theoretical models for this process have, thus far, assumed that the objective is to reach as many individuals as possible. However, studies such as those measuring the effectiveness of D.A.R.E., and research on viral marketing in commercial settings have shown that indiscriminate exposure in such a setting can lead to *overexposure*, whereby the campaign reaches people who evaluate it negatively. In Abebe, Adamic, and Kleinberg (2017), we ask how we should make use of social influence when there is risk of overexposure. We develop and analyze a theoretical model for this process and provide an efficient algorithm for finding the optimal target-set. We also experimentally show how it captures a number of the qualitative phenomena associated with overexposure.

References

Abebe, R.; Adamic, L.; and Kleinberg, J. 2017. Mitigating overexposure in viral marketing. *arXiv preprint arXiv:1709.04123*.

Abebe, R., and Goldner, K. 2016. Mechanism Design for Social Good. `http://md4sg.com`. Accessed: 2016-09-16.

Abebe, R.; Cole, R.; Gkatzelis, V.; and Hartline, J. 2017a. A truthful cardinal mechanism for one-sided matching markets. Under submission.

Abebe, R.; Hill, S.; Schwartz, H. A.; and Vaughan, J. W. 2017b. Using search queries to understand health information needs in africa. In *Machine Learning for the Developing World at NIPS*.

Abebe, R.; Kleinberg, J.; and Parkes, D. C. 2017. Fair division via social comparison. In *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '17.

Buvinic, M.; Furst-Nichols, R.; and Koolwal, G. 2014. Mapping gender data gaps. *Data2X report, United Nations Foundation, Washington, DC*.

Hylland, A., and Zeckhauser, R. 1979. The efficient allocation of individuals to positions. *Journal of Political economy* 87(2):293–314.

Jean, N.; Burke, M.; Xie, M.; Davis, W. M.; Lobell, D. B.; and Ermon, S. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353(6301):790–794.

Patel, N.; Chittamuru, D.; Jain, A.; Dave, P.; and Parikh, T. S. 2010. Avaaj otalo: a field study of an interactive voice forum for small farmers in rural india. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 733–742. ACM.

Roth, A. E.; Sonmez, T.; and Unver, M. U. 2004. Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.

Tumin, M. M.; Grusky, D.; and Ku, M. C. 2012. Gloom, doom, and inequality.

2014. A world that counts: Mobilising the data revolution for sustainable development. Technical report, United Nations Data Revolution Group.

Wakefield, M.; Flay, B.; Nichter, M.; and Giovino, G. 2003. Effects of anti-smoking advertising on youth smoking: a review. *Journal of health communication* 8(3):229–247.

West, S. L., and ONeal, K. K. 2004. Project dare outcome effectiveness revisited. *American journal of public health* 94(6):1027–1029.