The Design and Analysis of Algorithmic Vaccine Allocation Frameworks

Jeffrey Keithley

University of Iowa
jeffrey-keithley@uiowa.edu

The COVID-19 pandemic starkly demonstrated the importance of effective interventions for societal functioning, particularly through timely and strategic vaccine distribution. My research develops transparent algorithmic decision support tools to help public health officials make evidence-based resource allocation decisions during pandemic responses. These tools provide guidance informed by disease models, and their development focuses on improving our understanding of how vaccine allocation can be solved as a discrete optimization problem with approximation algorithms. This work revolves around three pillars: (1) Formulate vaccine allocation discrete optimization problems. (2) Design and analyze vaccine allocation approximation algorithms. (3) Develop network disease-spread models to inform the previous pillars.

My progress on each of these pillars contribute to the goal of designing fast and interpretable algorithm frameworks for vaccine allocation with high quality solutions. In particular, simulating diffusion processes over graphs holds considerable significance in the modern world, with a wide number of applications from marketing to social network analysis (Kempe, Kleinberg, and Tardos 2003; Soma et al. 2014; Domingos and Richardson 2001).

Overview

The vaccine allocation problem arises when decision-makers must determine how to distribute vaccines in such a way to minimize the impact of an infectious disease. These decisions depend on factors such as vaccine availability, disease impact on vulnerable populations, and social interaction patterns (Medlock and Galvani 2009; Bubar et al. 2021; Zhang and Prakash 2015). A key challenge in the decision making process is deciphering how those factors interact, so researchers and policymakers rely on disease-spread models to predict and evaluate the effects of different allocation strategies (Boey et al. 2025; Shayegh et al. 2023). The type of disease models I will discuss here are known as compartmental models, and the most basic is the mean-field SI-ODE model (Kermack and McKendrick 1927; Ross 1916). These models provide high-level insight into simple disease spread dynamics within large, well mixed populations, but more granular models must be used for more granular results. Agent-based models (ABMs) are often used to study disease spread at an individual level; they use graphs to represent contact patterns within a population, where each node represents an individual (or agent) and each edge represents a contact between pairs of agents. Simple ABMs also commonly use an underlying SIR model, where each individual is classified as susceptible, infected, or recovered, and a set of rules governs how individuals move between those states (Germann et al. 2006; Kerr et al. 2021). Formally, vaccine allocation informed by disease-spread models can be posed as a discrete optimization problem that seeks to minimize infection rates or maximize population health outcomes subject to supply constraints.

Vaccine Allocation to Regions

Focusing allocation strategies on subpopulations rather than individuals produces policy recommendations that are more feasible to implement in practice. Despite this, vaccine allocation over subpopulations has only received limited research interest, and the associated computational challenges are relatively unknown. To address this gap, we study vaccine allocation problems over geographically distinct subpopulations in this project. In particular, our vaccine allocation problems seek to find the best way to distribute a limited number of vaccines to a set of regions, such as the counties of a state. The distribution of vaccines can be evaluated using what is known as a metapopulation disease-spread model, where each county has its own mean-field compartmental model that takes the disease status of people moving between counties into account. The most obvious way to optimally solve the vaccine allocation problem is to test out every possible combination of how many vaccines each county receives. As it turns out, this approach is grossly infeasible (formally, the problem itself is known as NP-hard). For instance, if the disease model runs in 0.01 seconds, finding the optimal way to allocate just 10 vaccine shipments to the state of Iowa would take about 5000 years! This raises the question: *How can we allocate vaccines quickly and dependably?*

One approach is to use greedy algorithms, which are a well-known family of approximation algorithms often applied to resource allocation problems that provide high quality solutions (Bar-Noy et al. 2001; Jain, Mahdian, and Saberi

2002). In resource allocation problems, a greedy algorithm allocates resources one at a time, always selecting the option that provides the greatest marginal benefit at each step according to the objective function defined in the problem formulation. A significant benefit of using a greedy algorithm is that, if the objective function exhibits a property known as submodularity (*i.e.*, exhibits diminishing returns), the solution value found by a greedy algorithm is guaranteed to be above $\sim$63% of the optimal solution (Nemhauser, Wolsey, and Fisher 1978). Unfortunately, common objective functions for vaccine allocation, including ours, do not have this property. As a result, the previously discussed guarantee doesn't hold. The way we addressed this roadblock was to use approximate submodularity, which is a measure of how far a function is from being submodular. From this measure, derived performance guarantees in terms of this "distance to submodularity", often in the form of what is called a submodularity ratio (Das 2018; Qian et al. 2018). We then experimentally established the near-optimal performance of greedy algorithms in the regional vaccine allocation problem (Keithley et al. 2025). This project played a formative role in establishing the direction of my research, and each subsequent project can be traced back to questions raised during this project's exploration: (1) How can this problem be formulated and adapted for equitable vaccine allocation? (2) How do the disease model parameters and contact structure affect submodularity ratio values? (3) How can network-based disease simulations be made to run in less time to enable effective scaling?

Equitable Vaccine Allocation

The COVID-19 pandemic was a powerful reminder that existing societal inequalities get amplified during public health emergencies (Khan et al. 2023). In response to the pandemic, organizations such as the CDC, WHO, and public health departments developed frameworks for equitable allocation of vaccines, using well-established ethical principles as a foundation (NASEM 2020; World Health Organization 2020). The overall goal of this project was to translate these policy frameworks into a computational framework that can be used by public health departments to equitably allocate vaccines in a transparent manner during the initial stages of a pandemic, when vaccine demand far exceeds supply (Srivastava and Priyadarshini 2021; Liu and Lou 2022). The first step was to develop a mathematical model for disease-spread that accounts for social vulnerability, geographic barriers to healthcare access, and differences in work constraints (CDC 2021; Ye et al. 2024; Gaffney et al. 2023). On the basis of this model, we present two optimization formulations of a vaccine allocation problem that aims to reduce overall disease prevalence while also reducing disparity in outcomes for a given "protected class" relative to the general population. Experiments are focused on allocating vaccines at the census block group granularity in Johnson County, Iowa. The experimental test bed incorporates social vulnerability index, a hospital accessibility index, and essential worker status into CovaSim, a state-of-the-art agent-based COVID-19 model (Kerr et al. 2021). This leads to two main takeaways. First, it is possible to substantially reduce disparity in the outcomes of the protected class (for various choices of this class) with negligible worsening in overall disease-prevalence. Second, it is critical for disparity to be considered at all stages of the computational framework, *e.g.*, incorporating it into the optimization formulation without considering it in the modeling stage has very limited value. This project was primarily focused on modeling and experimentation. There are two natural immediate extensions to this work; deriving associated approximation guarantees and working with public health practitioners to adapt the framework for ease of use by policy experts (Keithley, Bonner, and Pemmaraju 2024, 2025).

Next Steps

The two projects outlined above form a foundation for my direction of research, and their planned and in-progress extensions can be summarized in three parts: (1) The vaccine allocation strategies in both projects assume preemptive distribution, which applies only to a limited range of scenarios. In practice, vaccine allocation is typically non-preemptive, meaning vaccines are distributed throughout an outbreak rather than allocated entirely at its onset. This realistic setting introduces two significant complications: the decision space for vaccine allocation becomes substantially larger, and deriving approximation guarantees for such dynamic allocation schemes proves considerably more challenging. (2) The effectiveness of the greedy algorithms used in the regional vaccine allocation project depend on the submodularity ratio, which is combinatorially defined and NP-hard to compute (Das 2018). Given this, my goal is to demonstrate how disease model properties (such as infectivity rates) affect the submodularity ratio, with early findings suggesting that less infectious diseases lead to better algorithm performance. The analysis focuses on limited graph structures including complete graphs and Erdos–Renyi graphs under independent cascade and SIR models, with potential for extending to spectral analysis. (3) ABM simulations offer high granularity for developing intervention methods but are computationally expensive compared to homogeneous-mixing models. This computational cost becomes problematic for intervention algorithms with high query complexity that require running ABMs many times. The speed of ABM simulations depends on contact network density, since each edge represents a potential infection that must be evaluated during simulation. One way to accelerate this process is through graph sparsification, which reduces the number of edges in the contact network (Benczúr and Karger 2015; Spielman and Teng 2011). Since disease spread behavior depends on the contact network structure, effective methods must approximately preserve this behavior while reducing computational complexity.

Additionally, I aim to enhance the practical applicability of this work through interdisciplinary collaboration with public health researchers and by developing more computationally efficient algorithms that can handle large-scale, realistic epidemic scenarios. These advances will strengthen both the theoretical foundations and real-world impact of submodular optimization approaches to equitable vaccine allocation.

Acknowledgments

Funding for this research was provided as part of the CDC MInD Healthcare group under cooperative agreement U01CK000594 and associated Covid19 supplemental funding, with additional funding provided by NSF Award Number 1955939. The author and his collaborators on this research acknowledge feedback from members of the Computational Epidemiology research group at the University of Iowa and the CDC MInD-Healthcare group.

References

Bar-Noy, A.; Bar-Yehuda, R.; Freund, A.; Naor, J.; and Schieber, B. 2001. A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48(5): 1069–1090.

Benczúr, A. A.; and Karger, D. R. 2015. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM Journal on Computing*, 44(2): 290–319.

Boey, L.; Baharmand, H.; Phillips, R. O.; Vandaele, N.; Balcik, B.; Kjøndal, J. O.; Birkeland, A.; Fossli, H.; Saeed, N.; and Decouttere, C. 2025. Developing an intuitive decision support system for equitable vaccine distribution during pandemics. *Scientific Reports*, 15(1): 16339.

Bubar, K. M.; Reinholt, K.; Kissler, S. M.; Lipsitch, M.; Cobey, S.; Grad, Y. H.; and Larremore, D. B. 2021. Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. *Science*, 371(6532): 916–921.

CDC. 2021. Social Vulnerability Index 2020 Database.

Das, A. 2018. Approximate Submodularity and Its Applications: Subset Selection, Sparse Approximation and Dictionary Selection. 34.

Domingos, P.; and Richardson, M. 2001. Mining the Network Value of Customers. In *Proceedings of the Seventh International Conference on Knowledge Discovery and Data Mining*, 57–66. New York, NY, USA: ACM.

Gaffney, A.; Himmelstein, D. U.; McCormick, D.; and Woolhandler, S. 2023. COVID-19 Risk by Workers' Occupation and Industry in the United States, 2020–2021. *American Journal of Public Health*, 113(6): 647–656.

Germann, T. C.; Kadau, K.; Longini, I. M.; and Macken, C. A. 2006. Mitigation Strategies for Pandemic Influenza in the United States. *Proceedings of the National Academy of Sciences*, 103(15): 5935–5940.

Jain, K.; Mahdian, M.; and Saberi, A. 2002. A new greedy approach for facility location problems. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, STOC '02, 731–740. New York, NY, USA: ACM. ISBN 1-58113-495-9.

Keithley, J.; Bonner, M.; and Pemmaraju, S. V. 2024. Balancing Efficiency and Equity in Iterative Agent-Based Vaccine Allocation Algorithms. Presentation at MIDAS Network Annual Meeting.

Keithley, J.; Bonner, M.; and Pemmaraju, S. V. 2025. Models and Algorithms for Balancing Efficiency and Equity in Vaccine Allocation. In *Proceedings of the 2025 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '25. Submitted.

Keithley, J.; Choudhuri, A.; Adhikari, B.; and Pemmaraju, S. V. 2025. Your PLOS Paper Title. *PLOS Computational Biology*. Under review (revision submitted).

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the Spread of Influence through a Social Network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, 137–146. New York, NY, USA: Association for Computing Machinery.

Kermack, W.; and McKendrick, A. 1927. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London*, 15(772): 700–721.

Kerr, C. C.; Stuart, R. M.; Mistry, D.; Abeysuriya, R. G.; Hart, G. R.; Rosenfeld, K.; Nunez, R. C.; Cohen, J. A.; Selvaraj, P.; Hagedorn, B.; George, L.; Jastrzebski, M.; Izzo, A. S.; Fowler, G.; Palmer, A.; Delport, D.; Scott, N.; Kelly, S. L.; Bennette, C. S.; and Klein, D. J. 2021. Covasim: An agent-based model of COVID-19 dynamics and interventions. *PLOS Computational Biology*, 17(7): e1009149.

Khan, N.; Mu, K.; Moharrami, M.; and Subramanian, V. 2023. Backward and Forward Inference in Interacting Independent-Cascade Processes: A Scalable and Convergent Message-Passing Approach. arXiv:2310.19138.

Liu, K.; and Lou, Y. 2022. Optimizing COVID-19 Vaccination Programs during Vaccine Shortages. *Infectious Disease Modelling*, 7(1): 286–98.

Medlock, J.; and Galvani, A. P. 2009. Optimizing Influenza Vaccine Distribution. *Science*, 325(5948): 1705–1708.

NASEM. 2020. *Framework for Equitable Allocation of COVID-19 Vaccines*. Washington, DC: The National Academies Press.

Nemhauser, G.; Wolsey, L.; and Fisher, M. 1978. An Analysis of Approximations for Maximizing Submodular Set Functions–I. *Mathematical Programming*, 14(1): 265–94.

Qian, C.; Zhang, Y.; Tang, K.; and Yao, X. 2018. On Multiset Selection With Size Constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32. AAAI Press.

Ross, R. 1916. An Application of the Theory of Probabilities to the Study of a Priori Pathometry.—Part I. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 92(638): 204–230.

Shayegh, S.; Andreu-Perez, J.; Akoth, C.; Bosch-Capblanch, X.; Dasgupta, S.; Falchetta, G.; Gregson, S.; Hammad, A. T.; Herringer, M.; Kapkea, F.; Labella, A.; Lisciotto, L.; Martínez, L.; Macharia, P. M.; Morales-Ruiz, P.; Murage, N.; Offeddu, V.; South, A.; Torbica, A.; and Melegaro, A. 2023. Prioritizing COVID-19 vaccine allocation in resource poor settings: Towards an Artificial Intelligence-enabled and Geospatial-assisted decision support framework. *PLOS ONE*, 18(8): e0275037.

Soma, T.; Kakimura, N.; Inaba, K.; and Kawarabayashi, K.-i. 2014. Optimal Budget Allocation: Theoretical Guarantee and Efficient Algorithm. 11.

Spielman, D. A.; and Teng, S.-H. 2011. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4): 981–1025.

Srivastava, V.; and Priyadarshini, S. 2021. Vaccine Shortage Dents India's Coronavirus Adult Immunisation Drive. *Nature India*.

World Health Organization. 2020. WHO Concept for Fair Access and Equitable Allocation of COVID-19 Health Products. Final working version, World Health Organization.

Ye, P.; Ye, Z.; Xia, J.; Zhong, L.; Zhang, M.; Lv, L.; Tu, W.; Yue, Y.; and Li, Q. 2024. National-scale 1-km maps of hospital travel time and hospital accessibility in China. *Scientific Data*, 11(1): 1130.

Zhang, Y.; and Prakash, B. A. 2015. Data-Aware Vaccine Allocation Over Large Networks. *ACM Transactions on Knowledge Discovery from Data*, 10(2): 20:1–20:32.