Research Statement – "Dynamic control of spreading processes on networks"

My fundamental goal is to understand the role of timing, network structure, and uncertainty in making better decisions for complex networked systems. My work characterizes **resource-allocation policies** for, and within, such heterogeneous **social, biological, and technological systems**. I integrate the fields of optimal control and dynamic programming, (nonlinear) optimization, games, and networks to simplify complex dynamic decisions and to characterize the resulting evolution of networked systems. This both allows the design of **better policies** (e.g., for curbing the spread of epidemics under resource constraints) and the design of **better incentives** for autonomous decision-makers (e.g., for the persistence of ethical behavior in an organization). The unifying theme across my work has been stressing that in dynamic processes, *when* one takes an action is as important as how it is targeted; and ignoring the temporal element of decision-making can lead to sub-optimal outcomes.

My work has impacted research on **epidemics, social influence,** and **energy-constrained networks**. In the future, I will develop new network decision tools for complex interventions by:

1- Improving interventions by incorporating network effects in resource allocation decisions,
e.g., creating network-based metrics to evaluate dynamic treatment allocations in an epidemic;
2- Creating analysis frameworks for incentive-driven behavior change in networks,
e.g., modeling the spread of unethical behavior in organizations; and
3- Developing improved methods to evaluate the effectiveness of networked interventions,
e.g., creating plausible counterfactuals for such policies.

Below, I outline my major past contributions and current plans.

Epidemics: Epidemics spread through contact; so, to combat them, we need to understand the local and networked aspects of contact mechanisms as well as the timing of interactions. My work has focused on epidemic control policies that are dynamic and heterogeneous, much like the underlying processes.

My work on epidemics started during my Ph.D. at Penn through epidemic-like message spreading protocols in wireless networks. I then studied more general heterogeneous networked epidemic models and showed that optimal coordinated curative actions involve offering vaccinations and treatments at maximum effort with no delay, regardless of the topology of the network [1]. I also proved this structure persists for curbing malware spread in computer networks, where the remedy (patches) can, unlike vaccines, spread immunity. I delved further into studying stealthy complex malware, like Stuxnet, that aimed to spread to reach a particular target while avoiding detection. I showed that the most damaging attack would be one that spreads the most virulent variant of the malware first, at the risk of discovery [2]. Surprisingly, this remains the most damaging attack even if the network implements various popular quarantine policies upon malware detection. In my thesis, I created an overarching taxonomy of heterogeneity for epidemic control and characterized the optimal control of epidemics in the presence of heterogeneity [3]. As a postdoctoral researcher at YSPH, I am designing efficient vaccine schedules and contact-tracing policies under real-world constraints, based on data from recent outbreaks.

I will continue developing networked models and optimal dynamic allocation mechanisms to mitigate the spread of epidemics, focusing on understanding the direct relationship between network structure and allocations (Aim 1) as well as creating valid counterfactuals for policy evaluation in the field (Aim 3).

Social Influence: Describing information spread and behavior adoption is a key question for several fields of inquiry, e.g., sociology, marketing. Modeling such spreads and designing network targeting strategies to maximize impact has driven my work in network science.

In my postdoctoral work at Cornell, I showed how political and marketing campaigns should optimally divide their budget between channels across time to maximize purchase decisions and votes [4]. I proved that an optimal campaign should initially prioritize reach, and only focus on targeting likely voters/purchasers late in the decision cycle, while using each channel in waves, i.e., cycling the use of channels. As a postdoctoral associate at YINS, I investigated social influence and targeting with complex behavioral models and limited network information. I showed how limited network information can be used in scalable algorithms to choose optimal seed sets to create cascades [5-6], as well as how the stability of a social group can be quantified, even without knowledge of group norms, using a targeting heuristic [7]. I also investigated differing organizational whistleblowing policies using game theory, showing that networked policies lead to better organizational outcomes and that, surprisingly, whistleblowing only thrives when punishments for unethical behavior are relatively light [8, 9].

I will continue studying the effect of incentives on network-based social decision-making, with a focus on validating models through experimentation (Aim 2).

Energy-Aware Design: Dynamic, distributed resource allocation is especially important when resource-availability varies in hard-to-predict ways, as is the case for electrical systems using renewable energy. My work has also addressed specific challenges in wireless communication and power networks.

As a Ph.D. student at Penn working on delay-tolerant wireless networks, where end-to-end connectivity is rare and messages must be relayed to reach their destination, I created a distributed energy-aware message-forwarding algorithm that guaranteed message delivery while maximizing network lifetime [10, 11] As an intern at NEC Labs, I proposed a patent-pending framework for the use of grid-scale batteries, showing how their optimal use in a microgrid can be mapped into a problem of dynamically (shadow-)pricing battery power and created a novel, principled way for obtaining the prices from price, demand, and generation data [12]. As a postdoctoral associate at YINS, I have developed distributed methods to create reliable energy-efficient wireless network backbones in hostile environments, such as battlefields [13-16].

I will apply such distributed decision-making tools to social systems (Aim 1) while considering the effect of resource constraints on interdependent social and technological networks, *e.g.,* electrical and transportation grids.

Future Research Directions

In the next 5 years, I will investigate 3 research aims that will propel my long-term research goal: creating an integrated framework to design and evaluate effective dynamic network interventions.

1 – How to design better dynamic public health intervention and testing policies

I will study how to design and evaluate public health interventions using an optimization framework. Public health interventions are costly to develop, and discarding effective interventions can lead to significant loss of social welfare. To evaluate the effectiveness of an intervention, we need to understand how it can best be implemented.

First, I will investigate the effect of network-related exposure on the outcomes of interventions, e.g., the effect of vaccine waning and refusal on the success of vaccination strategies. This will tie into my earlier work on resource allocation strategies to mitigate epidemic spread factoring in network-related heterogeneities [1-3, 9, 10].

Furthermore, I will investigate dynamic network-based targeting policies derived directly from the resource allocation framework [4]. This framework avoids substituting network structural centrality measures for true measures of network-related intervention effectiveness. I will then develop approximations to these metrics that only require local network information [13-16]. This is especially important in social systems where mapping the global network is prohibitively costly or there is significant measurement error.

Finally, I will incorporate the effect of uncertainty and information accumulation in the decision-making framework using robust optimization and reinforcement learning. I will develop adaptive resource allocation tools that focus on optimizing the *trajectory* of a spreading process while at the same time minimizing the uncertainty cone of possible realizations resulting from the interventions.

The ultimate goal of this aim is to incorporate incoming data into an adaptive simultaneous modeling and decision-making ecosystem, allowing decision-makers to tune trade-offs.

2- How to incentivize collective behavior change through network policies

I will study the effects of network structure on the spread of behavior. This is more complex than considering spillover from interventions, as it involves understanding and shaping incentives. Network and evolutionary games offer promising approaches to creating identifiable models for the incentive-driven spread of behavior.

First, I will consider the effect of complex social behaviors, such as providing care for sick acquaintances ("alters"), on overall networked outcomes. Understanding the interaction between behavior and public health outcomes, such as epidemic spread, is one of the goals of this step.

Furthermore, I will investigate the design of incentives and interventions to help spread beneficial actions. This involves taking the complex decision models generated in the first step and considering responses to policy interventions, e.g., my work on whistleblowing policies in organizations [8, 9].

Finally, I intend to validate networked models of decision-making, as well as the usefulness of interventions, through simple experiments on online experimentation platforms, e.g., Breadboard.

The ultimate goal of this aim is to understand the effect of network structure on population-level behavior change and learning of successful strategies and to use these insights in policy-design for public health interventions.

3- How to evaluate network interventions

I will create a framework to plausibly simulate the effect of alternative networked interventions from limited observations. Network data is costly to procure, measurements are noisy, and such experiments typically have few arms (subsets of participants assigned to a particular treatment) and observations due to budget constraints. Counterfactuals (which provide answers to "what-if" treatment assignment questions) can be used to establish causal relationships and to convince a policy-maker of the effectiveness of an intervention vis-à-vis alternatives.

The first step in this direction involves the creation of plausible counterfactuals at the individual level that balance matching (in feature space) and the accuracy of the counterfactual estimate. I seek to answer the question: How sure can we be that administering a different treatment to the same individual would have led to a different outcome?

Furthermore, I will leverage this individual model to incorporate time and network effects in the counterfactual, factoring exposure. These dynamic networked counterfactuals will allow us to simulate the effect of dynamic policy alternatives from observations, as a parallel to the model-based policy comparisons of the first two aims.

The final step involves calculating the value of network information in the intervention. Being able to simulate alternative policies allows us to estimate counterfactuals from policies that do not use network information. I will then compare the outcomes resulting from procured network information against the costs of network mapping.

The ultimate goal of this aim is to partner with public health researchers running network targeting policies to evaluate their ongoing interventions and to suggest dynamic improvements.

Research Impact

The research outlined above will provide the basis for a program that will develop **explainable resource allocation methods for effective networked interventions**. This program requires the development of tools to improve existing policies (Aim 1), incentives to promote beneficial behavior (Aim 2), and evaluation frameworks for such interventions (Aim 3). Ultimately, I will create a tighter connection between these three elements, using tools such as reinforcement learning to make real-time explainable policy recommendations.

For these goals, I will seek funding from external sources commensurate with the goal, with the National Science Foundation (NSF) being a target for more theoretical projects, e.g., Aim 1, and the National Institutes for Health (NIH) being a target for more applied ones, e.g., Aim 3. At Yale, I have contributed to multiple grant applications to the Army Research Office, the Office of Naval Research, NIH, and NSF. I have also been part of the multi-institution Network Science Collaborative Technology Alliance (NS-CTA) for a year and the International Technology Alliance in Distributed Analytics and Information Sciences (DAIS-ITA) for two years (both under the Army Research Labs), submitting white-papers and contributing to two funded project proposals

Network effects have been posited as great amplifiers for the impact of interventions, yet such interventions have had limited success in the field. My proposed work has the potential to improve their expected impact, as well as to create more realistic benchmarks for their evaluation. My work would most directly translate to improving public health outcomes, but can also apply to organizations and technological systems.

References

1- S. ESHGHI, M. H. R. Khouzani, S. Sarkar, S. S. Venkatesh, “Optimal patching in clustered malware epidemics,” in **IEEE/ACM Trans. on Networking**, 24(1), 283–298, 2016.

2- S. ESHGHI, S. Sarkar, S. S. Venkatesh, “Visibility-aware optimal contagion of malware epidemics,” in **IEEE Trans. on Automatic Control**, 62(10), 5205-5212, 2017.

3- S. ESHGHI, “Optimal control of epidemics in the presence of heterogeneity,” Ph.D. thesis, University of Pennsylvania, 2015.

4- S. ESHGHI, V. M. Preciado, S. Sarkar, S. S. Venkatesh, Q. Zhao, R. D'Souza, A. Swami, “Spread, then target, and advertise in waves: optimal budget allocation across advertising channels,” in **IEEE/ACM Trans. on Network Science and Engineering**, 2018.

5- S. ESHGHI, S. Maghsudi, V. Restocchi, L. Tassiulas, R. K. E. Bellamy, N. R. Jennings, S. Stein, “Influence maximisation in partially observable social networks,” in Proceedings of Conference on Complex Networks Workshop, 2017, 98-100.

6- S. ESHGHI, S. Maghsudi, V. Restocchi, S. Stein, L. Tassiulas, “Efficient influence maximization under partial network visibility,” in Proceedings of IEEE CAOS ’19 @ Infocom, 2019.

7- S. ESHGHI, G. R. Williams, G. B. Colombo, L. D. Turner, D. G. Rand, R. M. Whitaker, L. Tassiulas, “Social group stability and fracture,” in Proceedings of the 2017 IEEE Allerton Conference, 478-485.

8- S. ESHGHI, L. Tassiulas, “Innovation, cheating, and whistleblowing - a game theoretic perspective,”, in Proceedings of the 2018 IEEE CISS.

9- S. ESHGHI, L. Tassiulas, “Whistleblowing games in networks”, in Proceedings of the **2018 IEEE CDC**.

10- M. H. R. Khouzani, S. ESHGHI, S. Sarkar, N. B. Shroff, S. S. Venkatesh, “Optimal energy-aware epidemic routing in DTNs,” in Proceedings of **Mobihoc 2012**, 175–182, 2012.

11- S. ESHGHI, M. H. R. Khouzani, S. Sarkar, N. B. Shroff, S. S. Venkatesh “Optimal energy-aware epidemic routing in DTNs,” in **IEEE Trans. on Automatic Control**, 60(6), 1554–1569, 2015.

12- S. ESHGHI, R. M. Patil, “Optimal battery pricing and energy management for microgrids,” in Proceedings of the **2015 IEEE ACC**, 4994–5001, 2015.

13- D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, “Distributed algorithms for multi-layer connected edge-dominating sets”, in Proceedings of the **2018 IEEE CDC**.

14- D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, “Energy-aware distributed edge domination of multilayer networks”, in Proceedings of the **2019 IEEE ACC**.

15- D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, “Distributed algorithms for multi-layer connected edge-dominating sets”, in **IEEE Control Systems Letters (L-CSS)**, 3(1), 31-36, 2019.

16- D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, “Energy-efficient backbone formation in military multi-layer ad-hoc networks”, in **Ad Hoc Networks**, 81, 17-44, 2018.