

Clinical Trial of an AI-Augmented Intervention for HIV prevention in Youth Experiencing Homelessness

Bryan Wilder¹, Laura Onasch-Vera², Graham Diguseppi², Robin Petering³,
Chyna Hill², Amulya Yadav⁴, Eric Rice², Milind Tambe¹

¹ Center for Research on Computation and Society,

John A. Paulson School of Engineering and Applied Sciences, Harvard University

² Suzanne Dworak-Peck School of Social Work and Center for Artificial Intelligence in Society,
University of Southern California

³ Lens Co

⁴ College of Information Sciences and Technology, Pennsylvania State University

Abstract

Youth experiencing homelessness (YEH) are subject to substantially greater risk of HIV infection, compounded both by their lack of access to stable housing and the disproportionate representation of youth of marginalized racial, ethnic, and gender identity groups among YEH. A key goal for health equity is to improve adoption of protective behaviors in this population. One promising strategy for intervention is to recruit peer leaders from the population of YEH to promote behaviors such as condom usage and regular HIV testing to their social contacts. This raises a computational question: which youth should be selected as peer leaders to maximize the overall impact of the intervention? We developed an artificial intelligence system to optimize such social network interventions in a community health setting. We conducted a clinical trial enrolling 713 YEH at drop-in centers in a large US city. The clinical trial compared interventions planned with the algorithm to those where the highest-degree nodes in the youths' social network were recruited as peer leaders (the standard method in public health) and to an observation-only control group. Results from the clinical trial show that youth in the AI group experience statistically significant reductions in key risk behaviors for HIV transmission, while those in the other groups do not. This provides, to our knowledge, the first empirical validation of the usage of AI methods to optimize social network interventions for health. We conclude by discussing lessons learned over the course of the project which may inform future attempts to use AI in community-level interventions.

Introduction

Each year, approximately 4.2 million youth in the United States experience some form of homelessness (Morton et al. 2018). One of the key health challenges for this population is high HIV prevalence, with reported prevalence in the range of 2-11% (Young and Rice 2011), up to 10 times that for youth with stable housing (National HCH Council 2012).

One proposed mechanism for fostering behavior change in high-risk populations is the *peer change agent* model. The main idea is to recruit peer leaders from the population of

youth experiencing homelessness (YEH) to serve as advocates for HIV awareness and prevention. Use of peer leaders has been suggested in the public health and social science literature due to the central role that peers play in risk behaviors for YEH, including related to HIV spread (Green Jr et al. 2013; Rice et al. 2010, 2012). Indeed, peer change agent models have succeeded in past HIV prevention interventions in other contexts (Medley et al. 2009). However, there have also been notable failures (Group et al. 2010), and it has been argued such failures may be attributable to how peer leaders are selected (Schneider, Zhou, and Laumann 2015). The long-standing and most widely adopted method in the public health literature for selecting peer leaders is to identify the most popular individuals in the social network of the youth (Kelly et al. 1997) (formally, the highest degree nodes). This poses the question: are high-degree youth the best peer leaders to disseminate messages about HIV prevention? This question has relevance far beyond HIV prevention; analogous social network interventions are used widely across development, medicine, education, etc. (Kim et al. 2015; Paluck, Shepherd, and Aronow 2016; Banerjee et al. 2013; Valente and Pumpuang 2007).

Information dissemination on social networks is the focus of a long line of research in computer science. In particular, the *influence maximization* problem, formalized by (Kempe, Kleinberg, and Tardos 2003), asks how a limited number of seed nodes can be selected from a social network to maximize information diffusion. Influence maximization has been the subject of extensive work by the theoretical computer science and artificial intelligence communities (Chen, Wang, and Yang 2009; Chen, Wang, and Wang 2010; Goyal, Lu, and Lakshmanan 2011; Borgs et al. 2014; Tang, Xiao, and Shi 2014). However, to our knowledge, no work prior to this project had connected the computational literature on influence maximization to the use of network-driven interventions in public health and related fields. Computational work has mainly focused on developing highly efficient algorithms for use on large-scale social media networks (often motivated by advertising), while interventionists in health domains have not used explicitly algorithmic approaches to optimize the selection of peer leaders. Pre-

vious computational work assumed access to data (e.g., the full network structure and a model of information spread) which are simply not available in a public health context.

This paper reports the results of a project which bridges the gap between computation and health interventions. As a research team composed of computer scientists and social workers, we developed, implemented, and evaluated an intervention for HIV prevention in YEH where the peer leaders are algorithmically selected. This intervention was developed over the course of several years, alternating between algorithm design and smaller-scale pilot tests to evaluate feasibility. The final system, which we refer to as CHANGE (CompreHensive Adaptive Network samplinG for social influence), was evaluated in a large-scale clinical trial enrolling 713 youth across two years and three sites. The trial compared interventions planned with CHANGE to those using the standard public health methodology of selecting the youth with highest degree centrality (DC), as well as an observation-only control group (OBS). *Results from this clinical trial demonstrate that CHANGE was substantially more effective than the standard DC method at increasing adoption of behaviors protective against HIV spread.* To our knowledge, this is the first empirically validated success of using AI methods to improve social network interventions for health. It is critically important for “AI for Social Good” work to result in deployed and rigorously evaluated interventions, and this paper provides one such example.

The remainder of the paper is organized as follows. First, we survey related work from both a computational and application perspective. Second, we introduce a formalization of the problem of selecting peer leaders from a computational perspective. Third, we briefly review the design of the CHANGE system to address this problem (deferring most details to earlier technical publications (Wilder et al. 2018b; Wilder 2018; Wilder et al. 2018a)). Fourth, we present the design of the clinical trial. Fifth, we present and analyze results from the trial. Sixth, we discuss lessons learned over the course of the project which may help inform future attempts to design and implement AI-augmented public health interventions.

Related Work

A great deal of research in computer science has been devoted to the influence maximization problem. The majority of this has focused on computationally efficient algorithms for large networks (Chen, Wang, and Yang 2009; Chen, Wang, and Wang 2010; Goyal, Lu, and Lakshmanan 2011; Borgs et al. 2014; Tang, Xiao, and Shi 2014) and assumes that the underlying social network and model of information diffusion are perfectly known. There is also more recent literature on algorithms to learn or explore these properties. Predominantly though, such work requires many repeated interactions with the system. For example, algorithms to estimate the parameters of an unknown model of information diffusion (Du et al. 2014; Pouget-Abadie and Horel 2015; Narasimhan, Parkes, and Singer 2015; He et al. 2016; Kalimeris et al. 2018) typically require the observation of hundreds of cascades on the same network. Collecting this amount of data is intractable for public health interventions,

where a single round of the intervention takes months. Other work concerns the bandit setting, where the algorithm can repeatedly select sets of nodes and observe the resulting cascade (Wen et al. 2017; Chen, Wang, and Yuan 2013; Wang and Chen 2017). Similarly, these algorithms accept poor performance in early rounds as the price for improvement over the long run, but waiting tens or hundreds of rounds for improved performance is not an option in our domain. Such techniques are a much better fit for problems concerning on-line social networks (for example, in advertising domains) where repeated experiments and large datasets are possible.

The most closely related related computational work to ours concerns a robust version of the influence maximization problem (He and Kempe 2018; Chen et al. 2016; Lowalekar, Varakantham, and Kumar 2016), building on the earlier work of (Krause et al. 2008) on general robust submodular maximization problems. Our algorithm for robust submodular optimization, for which an overview is provided below, differs from these approaches mainly in that it solves a fractional relaxation of the problem instead of repeatedly calling a greedy algorithm for discrete submodular optimization, which helps improve computational performance.

There is a large literature on social network interventions in public health (Valente and Pumpuang 2007; Kim et al. 2015), clinical medicine (Young et al. 2003), international development (Cai, De Janvry, and Sadoulet 2015; Banerjee et al. 2013), education (Paluck, Shepherd, and Aronow 2016), etc. Common strategies involve selecting high degree nodes (as compared to in our trial), selecting nodes at random, or asking members of the population to nominate others as influencers. The empirical evidence for the relative effectiveness of different strategies is mixed; (Kim et al. 2015) reports no or marginal improvement for nominations vs random selections (depending on the outcome measure), while (Banerjee et al. 2019) report statistically significant improvements for a nomination-based selection mechanism. (Chin, Eckles, and Ugander 2018) introduce improved statistical methods to compare the effectiveness of seeding strategies and conclude that nomination-based strategies do not measurably improve performance. Indeed, (Akbarpour, Malladi, and Saberi 2018) show that in some theoretical network models it may be preferable to recruit a slightly larger number of influencers at random rather than carefully map the network. We contribute to this literature by developing and empirically evaluating an algorithmic framework which combines both features reminiscent of the nomination-based strategies proposed by others (for gathering information about network structure) as well as robust optimization techniques for jointly optimizing the entire set of influencers who are selected (not part of previous empirically evaluated strategies). Our clinical trial demonstrates statistically significant improvements from this strategy compared to the baseline of selecting high-degree nodes, providing (to our knowledge) the first real-world evidence that systematic optimization leads to improved results.

Problem Description

The population of youth are the nodes of a graph $G = (V, E)$. We seek to recruit a set of youth S to be peer leaders,

where $|S| \leq k$. In domain terms, this budget constraint reflects the fact that peer leaders are given a resource-intensive training and support process. The objective is to maximize the total expected number of youth who receive information about HIV prevention, given by the function $f(S)$. Here, f encapsulates the dynamics of a probabilistic model of information diffusion across the network (discussed below). The optimization problem $\max_{|S| \leq k} f(S)$ is the subject of the well-known influence maximization problem. When the objective function f is instantiated using common models for information diffusion, the resulting optimization problem is submodular (i.e., there are diminishing returns to selecting additional peer leaders). While finding an optimal solution is NP-hard, a simple greedy algorithm obtains a $(1 - 1/e)$ -approximation (Kempe, Kleinberg, and Tardos 2003).

The most common choice for the model of information diffusion is the independent cascade model. In this model, each node who receives information transmits it to each of their neighbors with probability p . All such events are independent. The process proceeds in discrete time steps where each newly informed node attempts to inform each of their neighbors, and concludes when there are no new activations. $f(S)$ calculates the number of nodes who receive information when the nodes S are informed at the start of the process, in expectation over the random propagation.

The standard influence maximization problem concludes here. However, while developing an algorithmic framework applicable to public health contexts, we came across challenges which must be solved before, during, and after the setting imagined in standard influence maximization. These challenges opened up new algorithmic questions, addressed in a series of publications in the AI literature (Wilder et al. 2018b; Wilder 2018; Wilder et al. 2018a). Here, we detail three steps for deploying an influence maximization intervention in the field.

First, information about the network structure G must be gathered. Previous work on influence maximization assumed that the network structure is known in advance. While this assumption may be reasonable for online social networks, we aim to disseminate information through the network consisting of real-world interactions between youth at a given center. Moreover, pilot studies revealed that information from an online social network (Facebook) was a poor proxy for actual connections at the center – not all youth used Facebook, and of those who did, many were not friends with their actual contacts at the drop-in center. Instead, network information must be gathered through in-person interviews where social workers ask youth to list those who they regularly interact with. Collecting data in this manner is time-consuming and expensive, often requiring a week or more of effort on the part of the social work team. Accordingly, the first stage of our algorithmic problem is to decide which nodes to query for network information. The algorithm is allowed to make M queries, where each query reveals the edges associated with the selected node. The queries can be adaptive, i.e., the choice of the i th node to be queried can depend on the answers given by nodes $1 \dots i - 1$.

Second, this network information is used to select an initial set of peer leaders. This stage more closely resembles the

standard influence maximization problem. However, there is an additional complication that the propagation probability p is not known. Indeed, there is no data source from which it could be inferred (as opposed to online platforms with abundant data; see related work). Instead, we formulate an uncertainty set \mathcal{U} containing a set of possible values for p which are consistent with prior knowledge (in CHANGE, we took \mathcal{U} to be a discretization of the interval $[0,1]$, reflecting limited prior knowledge). The aim is to find a set S which performs near-optimally for every scenario contained in \mathcal{U} . Formally, this corresponds to the robust optimization problem

$$\max_{|S| \leq k} \min_{p \in \mathcal{U}} \frac{f(S, p)}{OPT(p)}$$

where $OPT(p)$ denotes $\max_{|S| \leq k} f(S, p)$, i.e., the best achievable objective value if the propagation probability p were known. Normalizing by $OPT(p)$ encourages the algorithm to find a set S which simultaneously well-approximates the optimal value for each $p \in \mathcal{U}$ and avoids the trivial solution where solution to the inner min problem is always the smallest possible value of p . Note that since $OPT(p)$ is constant with respect to S , $\frac{f(S, p)}{OPT(p)}$ remains submodular with respect to S . Robust optimization of submodular functions is substantially more difficult than optimization of a single submodular function; in fact, it is provably inapproximable in general (Krause et al. 2008) and the aim is instead to approximate a tractable relaxation of the problem.

Third, after an initial set of peer leaders S is identified, recruitment proceeds in an adaptive manner. Not all youth invited to become peer leaders will actually attend the training session. A number of potential barriers exist, e.g., a given youth could have been arrested or not have had enough money for a bus ticket. Formally, we model that each youth who is invited will actually attend with probability q (based on experience in pilot studies, we took $q = 0.5$), where the attendance of each youth is independent of the others. For a given value of p , the resulting objective function is $f(S, p, q)$, which takes an expectation over both the randomness in which nodes are successfully influenced at the start of the process and in the subsequent diffusion. It is easy to show (Wilder et al. 2018b) that f remains submodular with this additional randomness. Because of this variation in attendance, as well as capacity limits for the initial training, peer leaders are recruited over multiple rounds, where the peer leaders selected in round t can depend on those who were successfully recruited in rounds $1 \dots t - 1$. In each round t , we select a set of peer leaders S_t with $|S_t| \leq k_t$ and observe which nodes are successfully recruited as peer leaders. The process continues for T rounds in total.

System Design

Our final proposed system for intervention planning is called CHANGE. CHANGE was originally introduced in (Wilder et al. 2018b). The final version of CHANGE summarized here is nearly the same as the original, with the exception of the algorithm used for robust optimization, which was separately developed and published in (Wilder 2018). We now provide an overview of CHANGE, mirroring the steps of the earlier problem formulation.

Network sampling CHANGE uses a simple but well-motivated heuristic to select a subset of nodes to be queried for network information (in the discussion section, we briefly review our earlier work on a more theoretically sophisticated solution, and the rationale for choosing a simpler method). The chosen method splits the query budget M into two halves. Each query in the first half is made to a node selected uniformly at random from the network. Each query in the second half follows a query in the first half, and selects a uniformly random neighbor of the first node. This design is motivated by the friendship paradox, the observation that high-degree nodes are overrepresented when we sample random neighbors (Feld 1991). Hence, the two stages of the query process balance between competing objectives: the first step encourages diversity, since random sampling ensures that we cover many different parts of the network, while the second step tends towards high-degree nodes who can reveal a great deal of network information.

Robust optimization We now provide an overview of how CHANGE handles parameter uncertainty within a single stage of the planning process, before considering the multi-stage problem (with uncertain attendance) below. As mentioned above, max-min submodular optimization is NP-hard to approximate (within any nonzero factor) (Krause et al. 2008). Accordingly, we need to somehow relax the problem to obtain meaningful guarantees. Let \mathcal{I} denote the set of all feasible solutions (sets S where $|S| \leq k$) and $\Delta(\mathcal{I})$ be the set of all distributions over \mathcal{I} (i.e., the $|\mathcal{I}|$ -dimensional simplex). We developed an algorithm for the problem

$$\max_{D \in \Delta(\mathcal{I})} \min_{p \in \mathcal{U}} \mathbb{E}_{S \sim D} \left[\frac{f(S, p)}{OPT(p)} \right] \quad (1)$$

which allows the algorithm to select a distribution over feasible sets and evaluates the worst case only in expectation over this distribution. In game theoretic terms, this allows the algorithm to select a mixed strategy instead of a pure strategy. At run-time, we sample from D ; the resulting set has guaranteed performance in expectation over the sampling, but strong guarantees cannot be obtained ex-post for the sampled set (as a result of the computational hardness of the original max-min problem). However, in practice we find that sampling several random sets and selecting the best one gives excellent empirical performance (i.e., closely matching or exceeding the expected value of the distribution).

Our algorithm for this problem, detailed in (Wilder 2018), uses a compact representation of the space of distributions (keeping track of only the marginal probability that each node is selected instead of each of the exponentially many potential subsets). It solves a fractional relaxation of the discrete max-min problem using this compact representation via a stochastic first-order method which is adapted to the particular properties of the objective. Then, we can use known rounding algorithms for submodular maximization to sample random sets from the distribution encoded by the solution to the fractional relaxation. This procedure guarantees a $(1 - 1/e)^2$ -approximation for Problem 1, which can be improved to $(1 - 1/e)$ with some additional steps (which we did not find empirically necessary).

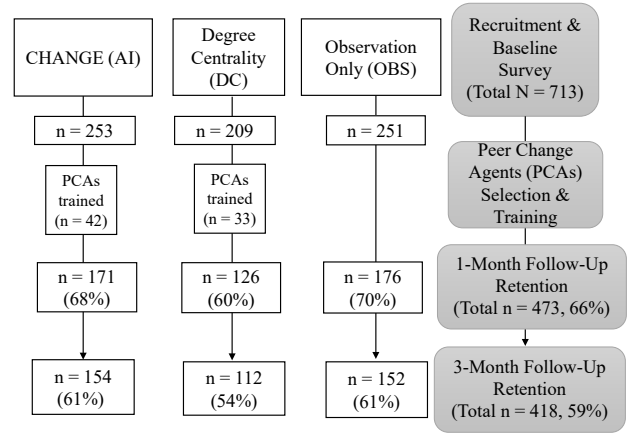


Figure 1: Number of participants recruited and retained in each arm of the study.

Multi-stage intervention with attendance uncertainty

We handle the multi-stage nature of the intervention by running the robust optimization problem at each stage, calculating the objective function in expectation over which peer leaders will attend and conditioning on the selection of those who have attended previous interventions. Formally, this means that at stage $t > 1$, we solve

$$\max_{D \in \Delta(\mathcal{I})} \min_{p \in \mathcal{U}} \mathbb{E}_{S_t \sim D} \left[\frac{f(S_t \cup S_1 \cup \dots \cup S_{t-1}, p, q)}{\max_{|S^*| \leq k} f(S^* \cup S_1 \cup \dots \cup S_{t-1}, p, q)} \right]$$

where S_1, \dots, S_{t-1} denote the sets of peer leaders who were successfully recruited in each previous stage. It is easy to show that the inner objective f remains submodular in S_t (see (Wilder et al. 2018b)), and so we retain the earlier guarantees on the quality of the solution obtained at each individual step. Moreover, in (Wilder et al. 2018b) we show that the multi-stage problem as a whole enjoys the property of *adaptive submodularity*, meaning that for any fixed parameter value p , solving

$$\max_{D \in \Delta(\mathcal{I})} \mathbb{E}_{S_t \sim D} \left[\frac{f(S_t \cup S_1 \cup \dots \cup S_{t-1}, p, q)}{\max_{|S^*| \leq k} f(S^* \cup S_1 \cup \dots \cup S_{t-1}, p, q)} \right]$$

at each step t and selecting the resulting set S_t enjoys an approximation guarantee relative to the optimal adaptive policy for selecting a sequence of sets S_1, \dots, S_t (again, with respect to a fixed p). More detailed discussion of the theoretical properties can be found in (Wilder et al. 2018b).

Study Design

We now move to the empirical portion of the project and provide an overview of the design of the clinical trial. All study procedures were approved by our institution's Institutional Review Board. The study was designed to compare the efficacy of two different means of selecting peer leaders: the CHANGE system described above and the standard DC approach in public health (selecting the highest-degree youth). We additionally included an observation-only control group (OBS), for three arms in total. The study was conducted at three drop-in centers for YEH in a large US

city. Drop-in centers provide basic services to YEH (e.g., food, clothing, case management, mobile HIV testing). Due to high transience in the YEH population, most clients at a given center leave within approximately six months. Accordingly, we tested each of the three methods at each of the three drop-in centers (giving nine deployments in total, each with a unique set of youth), ensuring that successive deployments at a given drop-in center were separated by six months. Youth were only allowed to enroll in the study once, so even the small number of youth who were present at the center across multiple deployments were included only on the first time they attempted to enroll. Testing each method at each drop-in center helps account for differences in the demographic and other characteristics of youth who tend to access services at each center.

Each of the nine deployments used the following procedure. Figure 1 shows the number of youth recruited and retained for each phase of the study in each arm.

First, youth were recruited at the drop-in center over the course of a week to participate in the study. All participants gave informed consent. Each participant completed a baseline survey which assessed demographic characteristics, sexual behaviors, and HIV knowledge. Demographic characteristics included age, birth sex, gender identity, race/ethnicity, and sexual orientation. Youth were also surveyed about their living situation and relationship status.

Second, peer leaders were selected and trained (for the CHANGE and DC arms of the study). Each individual training consisted of approximately 4 youth and there were 3-4 trainings per deployment (depending on exact attendance). In total, approximately 15% of survey participants in each deployment were trained as peer leaders. In the CHANGE arm of the study, network information was queried from approximately 20% of the participants (sampled according to the mechanism described above). In the DC arm, we used a full survey of the network to find high-degree nodes, in order to give the strongest possible implementation to compare to.

Third, peer leaders had three months to disseminate HIV prevention messages. Peer leaders were supported via 7 weeks of 30-minute check-in sessions with study researchers, which focused on positive reinforcement of their successes as well as problem-solving strategies and goals for the future. All peer leaders attended at least one check-in session, with modal attendance at five sessions. Peer leaders received \$60 in compensation for attending the initial training and \$20 for each check-in session.

Fourth, follow-up surveys were administered to the original study participants from the first step. Follow-up surveys assessed the same characteristics as the baseline survey. Differences in reported sexual behavior between baseline and follow-up were used as the primary metrics to evaluate the interventions. All such metrics were self-reported; we followed best practices in social science research to minimize bias in self-reported data (surveys were self-administered on a tablet and participants were guaranteed anonymity, each of which aim to reduce social desirability bias in reporting sensitive information). Additionally, any bias would be expected to influence each arm of the study equally, including the observation-only control group.

The training component of the peer change agent intervention was delivered by two or three facilitators from the social work research team. The training lasted approximately 4 hours (one half-day). Training was interactive and broken into six 45-minute modules on the mission of peer leaders (sexual health, HIV prevention, communication skills, leadership skills, and self-care). Peer leaders were asked to promote regular HIV testing and condom use through communication with their social ties at the drop-in center.

Study Results

We now present the results of the clinical trial, starting with an overview of the outcome variables and methodology for statistical analysis, and then giving the main results.

Outcome Variables

We compare two outcome variables across arms of the study. First, condomless anal sex (CAS), assessed via a survey question asking whether youth had anal sex without a condom at least once in the previous month. Second, condomless vaginal sex (CVS), assessed via a survey question asking whether youth had vaginal sex without a condom at least once in the previous month. CAS and CVS are both important behavioral risk factors for HIV transmission and so provide a direct assessment of the success of the intervention at producing a material health impact.

Statistical Methodology

We provide both the average value of each outcome variable at each time point for the three arms of the study as well as an analysis of statistical significance. The statistical analysis used a generalized estimating equations (GEE) model. GEE is an extension of generalized linear models which incorporates repeated measurements of data across a population. It is a standard choice for analysis of clinical data in this form (Zeger, Liang, and Albert 1988). We specified a linear model for each outcome variable which included terms for both the improvement caused by participating in a given arm of the study (our estimand of interest) as well as terms for a range of control variables which account for differences in demographics and the baseline rate of risk behaviors in each arm of the study. The demographic control variables were age, birth sex, transgender identity, LGBQ identity, the combination of male sex and LGBQ identity, race, committed relationship, housing status, and drop-in center. We also included a “time” variable to account for changes in the entire population over time regardless of participation in a particular arm of the study. This combination of control variables helps separate the impact of the intervention from pre-existing differences between arms of the study and intervention-independent trends.

The linear model combined contributions from each of these variables through a logistic link function. Since each outcome is binary, we present results in the form of the odds ratio (OR), which measures the ratio in the odds of the outcome in youth who are exposed to a given intervention vs youth in the observation-only group (after controlling for

	CAS		CVS	
	OR	CI	OR	CI
Baseline				
CHANGE	1.43	0.91, 2.28	0.77	0.52, 1.13
DC	1.49	0.89, 2.48	1.07	0.67, 1.68
Post-intervention				
CHANGE	0.69*	0.49, 0.98	0.78 [†]	0.57, 1.04
DC	0.80	0.55, 1.17	0.88	0.62, 1.23
Time	1.05	0.82, 1.33	0.87	0.71, 1.06

Table 1: Results of statistical analysis. Each column gives the effect size and confidence interval for one of the outcome variables. Each row gives the corresponding estimates for one of the variables included in the GEE model. The “baseline” category measures pre-existing differences between the groups (relative to the observation-only group) on enrollment in the study. The “post-intervention” category measures the estimated impact of participating in each arm of the intervention (relative to the observation-only group, and after controlling for both demographics and baseline behaviors). “Time” gives the estimated contribution of a trend over time independent of which arm of the study a participant was enrolled in. [†] $p < 0.1$; * $p < 0.05$.

demographics and baseline rate of risk behaviors). For all quantities, we also present 95% confidence intervals and indicate where significant p -values are obtained.

Results are known only for youth who completed the follow-up surveys, leading to missing data due to participant attrition (as is expected for a study enrolling YEH). Of the 713 participants who completed the baseline survey, 245 (34%) missed the 1-month follow-up, 300 (42%) missed the 3-month follow-up, and 180 (25%) missed both follow-ups. However, missingness had no statistically significant association with CAS or CVS, indicating that youth were not significantly over or under represented in the follow-up data based on their baseline level of risk behavior.

Results

We start by presenting the main results of the statistical analysis; the full results can be found in Table 1.

CAS We find that CAS reduced in the CHANGE group over time by a statistically significant amount (OR = 0.69, $p < 0.05$). The estimated OR of 0.69 indicates that, in the GEE estimates, a youth who is enrolled in the CHANGE arm of the study has 31% lower odds to engage in CAS than if they were enrolled in the observation-only group. That is, a youth who is enrolled in CHANGE has 31% lower odds to engage in CAS post-intervention than a youth with identical starting characteristics (including baseline rate of CAS) who did not receive the intervention. For the DC group, there was not a statistically significant change in CAS over time relative to the observation-only group.

CVS The GEE model estimated that CVS decreased by a marginally statistically significant amount in the CHANGE




Figure 2: Average value of each outcome variable at each point in time for the three arms. These plots show the results without any statistical processing, while the analysis above attempts to control for pre-existing differences between participants in each arm.

group (OR = 0.78, $p < 0.1$). For the DC group, there was no statistically significant change in CVS over time relative to the observation-only group.

We conclude from the analysis that only CHANGE provided a statistically significant improvement in HIV risk behaviors compared to the observation-only baseline.

Direct examination of the average values of the outcome variables for each arm at each point in time (Figure 2) shows another interesting trend. Improvements in the CHANGE group happen faster than the DC group: most of the improvement for CHANGE occurs by the one-month survey, while improvements in the DC group are not fully realized until month three. Fast results are important for two reasons. First, rapid adoption of protective behaviors helps to immediately curtail transmission in a high-risk population. Second, high transience among YEH means that a non-negligible portion of youth will have left the center by the time a three-month intervention is completed. We conclude that the AI-augmented intervention implemented with CHANGE has substantial advantages over an intervention where peer leaders are selected with the standard DC method.

Discussion

This project provides evidence that AI methods can be used to improve the effectiveness of social network interventions in public health: significant reductions in HIV risk behaviors were observed in groups where our CHANGE method was used to plan the intervention, with no significant changes in behavior when the status quo method (selecting high degree nodes) was employed. More broadly, we hope that our experiences over the course of the project can provide generalizable lessons about how AI research can be successfully employed for social good. There have been recent attempts by others to synthesize principles for AI for Social Good research (Floridi et al. 2020; Tomašev et al. 2020). We offer a complementary perspective shaped by the process of deploying a specific community-level intervention. In particular, existing discussions of best practice often focus in large part on ethics, data privacy, and building trust with stakeholders. While such considerations are indispensable, it is also important for the research community to investigate the on-the-ground components of developing and deploying an impactful intervention. We highlight five points.

First, the starting point was to listen to domain experts and

understand where in the problem domain AI could be most impactful. We did not approach this project with a preexisting intention to apply influence maximization to the choice of peer leaders. Rather, this emerged organically from discussions between the AI and social work sides of the research team as a topic where an AI-augmented intervention was both technically feasible and likely to improve outcomes. *Success is less likely when AI researchers start with a favored technique and search for an application.*

Second, data was overwhelmingly the bottleneck to the AI component of the intervention. Computational work on influence maximization to date had largely assumed a great deal of information would be known – the structure of the graph, the model for information diffusion, etc. None of this information was in fact available for YEH (or would likely be available in other public health settings). Moreover, gathering this data is itself time-consuming and costly, requiring unsustainable effort on the part of an agency wishing to deploy the intervention on their own. Much of the technical focus of the research consisted of finding ways to reduce the amount of data which needed to be gathered for the intervention to succeed. *Finding ways to reduce or eliminate data needs through improved algorithm design is an important part of producing deployable AI interventions in a community health context.*

Third, simplicity is valuable. As an example, prior to developing CHANGE, we designed a much more theoretically sophisticated algorithm for collecting network data which enjoyed provable guarantees for certain families of graphs (Wilder et al. 2018a). However, it quickly became apparent that this algorithm would be difficult to deploy in practice because it required a large number of sequential queries (the node which is queried on step 1 determines the node who is to be queried on step 2, and so on). This was impractical in the context of a program working with YEH where any given youth may be difficult to find, interrupting the entire process. More generally, if the algorithm requires tight coupling with the outside world (many steps where information is input, the algorithm recommends a very specific action, more information is input, and so on) then there are more things that can go wrong which are not captured in the computational formalization of the problem. This poses a contrast to the way that simplicity is often operationalized in AI for social good work as either *explainability* (Floridi et al. 2020) or as *methodological simplicity* (Tomašev et al. 2020) (e.g., using well-developed techniques instead of a new algorithm). Both explainability and methodological simplicity are of course valuable in many settings but in our experience neither was first-order requirement: the algorithm can solve a complex optimization problem internally so long as the way that it interacts with the outside world is simple and robust. *We believe that this operational simplicity is an under-emphasized design criterion for AI for Social Good.*

Fourth, smaller pilot tests were a valuable part of the project prior to embarking on a larger clinical trial. We conducted several such tests, each of which consisted of a deployment at a single drop-in center, in order to test earlier versions of our system (Wilder et al. 2018b; Yadav et al. 2016, 2017). This helped reveal key issues which needed

to be addressed. For example, we quickly discovered that a plan to collect network information via Facebook was not viable with this population and that manual collection of network data entailed a great deal of effort. We also quickly observed that peer leaders often did not attend the training, requiring on-the-fly adjustments over the course of the program. Addressing such issues was necessary to the success of the overall project (and turned out to provide much of the technical challenge involved). *It would have been very difficult to identify these challenges without piloting algorithms in the actual environment where they will be used.* It was also helpful for computer scientists on the research team to be regularly present onsite during the pilot deployments to learn more about the environment and help coordinate the initial attempts at using the algorithm.

Fifth, community engagement and trust was essential to the success of the project. Beyond the research team, a number of stakeholders needed to be involved in the process. For example, we needed buy-in from each of the drop-in centers to conduct the study at the center, enroll their clients, and use their facilities. We regularly convened a community advisory board with representatives from each of the drop-in centers along with members of the research team to provide information about the study progress, explain the methods being used, and share information which could be helpful to other center activities. Just as critical as the center leadership though, were the youth themselves. We asked youth to disclose sensitive information, including their HIV risk behaviors and social contacts. Especially for the YEH population, which is less inclined than most to engage with authority figures, building trust is essential. We found two factors to be especially important in establishing this trust. First, the social work portion of the research team had deep roots in the community, having regularly offered services at these drop-in centers for the past ten years. Second, transparency about why information was being collected was critical. We observed substantially increased willingness to disclose information related to social contacts when researchers explained how this information would be used in the study (i.e., that a computer program would be used to select some people as peer leaders based on their contacts) than when such an explanation was not proactively given. *A critical part of the peer change agent model is empowering youth to make a difference in their community, and this philosophy extends to the way that AI should be used in a community setting.*

Our hope is that this project provides one example towards a broader research agenda aiming at AI techniques which can be successfully used to improve health and equity within our communities. A great deal of work remains. Just within the context of social network intervention, future work should explore other intervention designs (e.g., interventions which attempt to modify network structure by fostering supportive relationships), methods for further reducing data requirements (e.g., by using administrative data to infer social connections), and more deeply investigate the relationship between information diffusion and behavioral change. However, the results from this clinical trial provide evidence that AI can substantially improve the quality of services offered to the most vulnerable among us.

Acknowledgments

This work was supported by the Army Research Office (MURI W911NF1810208) and the California HIV/AIDS Research Program.

References

- Akbarpour, M.; Malladi, S.; and Saberi, A. 2018. Diffusion, seeding, and the value of network information. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 641–641.
- Banerjee, A.; Chandrasekhar, A. G.; Duflo, E.; and Jackson, M. O. 2013. The diffusion of microfinance. *Science* 341(6144).
- Banerjee, A.; Chandrasekhar, A. G.; Duflo, E.; and Jackson, M. O. 2019. Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies* 86(6): 2453–2490.
- Borgs, C.; Brautbar, M.; Chayes, J.; and Lucier, B. 2014. Maximizing social influence in nearly optimal time. In *Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms*, 946–957. SIAM.
- Cai, J.; De Janvry, A.; and Sadoulet, E. 2015. Social networks and the decision to insure. *American Economic Journal: Applied Economics* 7(2): 81–108.
- Chen, W.; Lin, T.; Tan, Z.; Zhao, M.; and Zhou, X. 2016. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 795–804.
- Chen, W.; Wang, C.; and Wang, Y. 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge discovery and Data Mining*, 1029–1038.
- Chen, W.; Wang, Y.; and Yang, S. 2009. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 199–208.
- Chen, W.; Wang, Y.; and Yuan, Y. 2013. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, 151–159.
- Chin, A.; Eckles, D.; and Ugander, J. 2018. Evaluating stochastic seeding strategies in networks. *arXiv preprint arXiv:1809.09561*.
- Du, N.; Liang, Y.; Balcan, M.; and Song, L. 2014. Influence function learning in information diffusion networks. In *International Conference on Machine Learning*, 2016–2024.
- Feld, S. L. 1991. Why your friends have more friends than you do. *American Journal of Sociology* 96(6): 1464–1477.
- Floridi, L.; Cows, J.; King, T. C.; and Taddeo, M. 2020. How to Design AI for Social Good: Seven Essential Factors. *Sci Eng Ethics*.
- Goyal, A.; Lu, W.; and Lakshmanan, L. V. 2011. Celf++ optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th International Conference Companion on World Wide Web*, 47–48.
- Green Jr, H. D.; de la Haye, K.; Tucker, J. S.; and Golinelli, D. 2013. Shared risk: who engages in substance use with American homeless youth? *Addiction* 108(9): 1618–1624.
- Group, N. R.; et al. 2010. Results of the NIMH collaborative HIV/sexually transmitted disease prevention trial of a community popular opinion leader intervention. *Journal of Acquired Immune Deficiency Syndromes* 54(2): 204–214.
- He, X.; and Kempe, D. 2018. Stability and robustness in influence maximization. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12(6): 1–34.
- He, X.; Xu, K.; Kempe, D.; and Liu, Y. 2016. Learning influence functions from incomplete observations. In *Advances in Neural Information Processing Systems*, 2073–2081.
- Kalimeris, D.; Singer, Y.; Subbian, K.; and Weinsberg, U. 2018. Learning diffusion using hyperparameters. In *International Conference on Machine Learning*, 2420–2428.
- Kelly, J. A.; Murphy, D. A.; Sikkema, K. J.; McAuliffe, T. L.; Roffman, R. A.; Solomon, L. J.; Winett, R. A.; Kalichman, S. C.; and Collaborative, T. C. H. P. R. 1997. Randomised, controlled, community-level HIV-prevention intervention for sexual-risk behaviour among homosexual men in US cities. *The Lancet* 350(9090): 1500–1505.
- 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*, 137–146.
- Kim, D. A.; Hwong, A. R.; Stafford, D.; Hughes, D. A.; O’Malley, A. J.; Fowler, J. H.; and Christakis, N. A. 2015. Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet* 386(9989): 145–153.
- Krause, A.; McMahan, H. B.; Guestrin, C.; and Gupta, A. 2008. Robust submodular observation selection. *Journal of Machine Learning Research* 9(Dec): 2761–2801.
- Lowalekar, M.; Varakantham, P.; and Kumar, A. 2016. Robust Influence Maximization. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 1395–1396.
- Medley, A.; Kennedy, C.; O’Reilly, K.; and Sweat, M. 2009. Effectiveness of peer education interventions for HIV prevention in developing countries: a systematic review and meta-analysis. *AIDS Education and Prevention* 21(3): 181–206.
- Morton, M. H.; Dworsky, A.; Matjasko, J. L.; Curry, S. R.; Schlueter, D.; Chávez, R.; and Farrell, A. F. 2018. Prevalence and correlates of youth homelessness in the United States. *Journal of Adolescent Health* 62(1): 14–21.
- Narasimhan, H.; Parkes, D. C.; and Singer, Y. 2015. Learnability of influence in networks. In *Advances in Neural Information Processing Systems*, 3186–3194.
- National HCH Council. 2012. HIV/AIDS among Persons Experiencing Homelessness. In *Focus Quarterly Research Review*.

- Paluck, E. L.; Shepherd, H.; and Aronow, P. M. 2016. Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences* 113(3): 566–571.
- Pouget-Abadie, J.; and Horel, T. 2015. Inferring graphs from cascades: A sparse recovery framework. *arXiv preprint arXiv:1505.05663*.
- Rice, E.; Barman-Adhikari, A.; Milburn, N. G.; and Monro, W. 2012. Position-Specific HIV Risk in a Large Network of Homeless Youths. *American Journal of Public Health* 102(1): 141–147. doi:10.2105/AJPH.2011.300295. URL <https://doi.org/10.2105/AJPH.2011.300295>. PMID: 22095350.
- Rice, E.; Monro, W.; Barman-Adhikari, A.; and Young, S. D. 2010. Internet use, social networking, and HIV/AIDS risk for homeless adolescents. *Journal of Adolescent Health* 47(6): 610–613.
- Schneider, J. A.; Zhou, A. N.; and Laumann, E. O. 2015. A new HIV prevention network approach: sociometric peer change agent selection. *Social Science & Medicine* 125: 192–202.
- Tang, Y.; Xiao, X.; and Shi, Y. 2014. Influence maximization: Near-optimal time complexity meets practical efficiency. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, 75–86.
- Tomašev, N.; Cornebise, J.; Hutter, F.; Mohamed, S.; Picciarriello, A.; Connelly, B.; Belgrave, D. C.; Ezer, D.; van der Haert, F. C.; Mugisha, F.; et al. 2020. AI for social good: unlocking the opportunity for positive impact. *Nature Communications* 11(1): 1–6.
- Valente, T. W.; and Pumpuang, P. 2007. Identifying opinion leaders to promote behavior change. *Health Education & Behavior* 34(6): 881–896.
- Wang, Q.; and Chen, W. 2017. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Advances in Neural Information Processing Systems*, 1161–1171.
- Wen, Z.; Kveton, B.; Valko, M.; and Vaswani, S. 2017. Online influence maximization under independent cascade model with semi-bandit feedback. In *Advances in neural information processing systems*, 3022–3032.
- Wilder, B. 2018. Equilibrium computation and robust optimization in zero sum games with submodular structure. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Wilder, B.; Immerlica, N.; Rice, E.; and Tambe, M. 2018a. Maximizing influence in an unknown social network. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Wilder, B.; Onasch-Vera, L.; Hudson, J.; Luna, J.; Wilson, N.; Petering, R.; Woo, D.; Tambe, M.; and Rice, E. 2018b. End-to-End Influence Maximization in the Field. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 1414–1422.
- Yadav, A.; Chan, H.; Jiang, A. X.; Xu, H.; Rice, E.; and Tambe, M. 2016. Using Social Networks to Aid Homeless Shelters: Dynamic Influence Maximization under Uncertainty. In *International Conference on Autonomous Agents and Multiagent Systems*, volume 16, 740–748.
- Yadav, A.; Wilder, B.; Rice, E.; Petering, R.; Craddock, J.; Yoshioka-Maxwell, A.; Hemler, M.; Onasch-Vera, L.; Tambe, M.; and Woo, D. 2017. Influence Maximization in the Field: The Arduous Journey from Emerging to Deployed Application. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 150–158.
- Young, J. M.; Hollands, M. J.; Ward, J.; and Holman, C. J. 2003. Role for opinion leaders in promoting evidence-based surgery. *Archives of Surgery* 138(7): 785–791.
- Young, S. D.; and Rice, E. 2011. Online social networking technologies, HIV knowledge, and sexual risk and testing behaviors among homeless youth. *AIDS and Behavior* 15(2): 253–260.
- Zeger, S. L.; Liang, K.-Y.; and Albert, P. S. 1988. Models for Longitudinal Data: A Generalized Estimating Equation Approach. *Biometrics* 44(4): 1049–1060. ISSN 0006341X, 15410420. URL <http://www.jstor.org/stable/2531734>.