

Large-scale 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

¹Center for Research on Computation and Society,

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

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

³ Lens Co

⁴College 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 over 700 YEH at drop-in centers in Los Angeles. 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 [26]. One of the key health challenges for this population is high HIV prevalence, with reported prevalence in the range of 2-11% [47], up to 10 times the rate for youth with access to stable housing [11]. Youth experiencing homelessness (YEH) are a particularly vulnerable population, and these health disparities are compounded by the overrepresentation among YEH of racial and ethnic minority youth, sexual minority youth, and non-cisgender youth [26], all groups at greater risk for HIV transmission [35, 30, 34]. Improving the adoption of protective behaviors to limit HIV transmission among YEH is hence a crucial goal for public health and health equity.

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 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 [15, 32, 31]. Indeed, peer change agent models have proved successful in past HIV prevention interventions in other contexts [25]. However, there have also been notable failures [16], and it has been argued that failures

of the peer change agent model may be attributable to how peer leaders are selected for participation in the program [36]. 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 [20] (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 applications in development, medicine, education, and more [22, 28, 2, 38].

Information dissemination on social networks is the focus of a long line of research in the computer science literature. In particular, the *influence maximization* problem, formalized by Kempe, Kleinberg, and Tardos [21], asks how a limited number of seed nodes can be selected from a social network to maximize the extent of information diffusion. Influence maximization has been the subject of extensive work from the theoretical computer science and artificial intelligence communities [8, 7, 14, 4, 37]. 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 applications on advertising), while interventionists in health domains have not used explicitly algorithmic approaches to optimize the selection of peer leaders.

This paper reports the results of a project which aims to bridge 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 the feasibility of proposed methods. The final system, which we refer to as CHANGE (CompreHensive Adaptive Network samplinG for social influenceE), was evaluated in a large-scale clinical trial enrolling over 700 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.

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 [43, 41, 42]). 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 the development of computationally efficient algorithms for large networks [8, 7, 14, 4, 37] and assumes that the underlying social network and model of information diffusion are perfectly known. There is also a more recently developed literature on algorithms to learn or explore these properties. Predominantly though, such work requires many repeated interactions with the system. For example, a body of work exists on algorithms to estimate the parameters of an unknown model of information diffusion [12, 29, 27, 18, 19], but these algorithms 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 to deploy. Another line of work concerns algorithms for the bandit setting, where the algorithm can repeatedly select sets of nodes and observe the resulting cascade [40, 9, 39]. 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 online social networks (for example, in advertising domains) where repeated experiments are possible and large amounts of data can be collected.

The most closely related computational work to ours concerns a robust version of the influence maximization problem [17, 6, 24], building on the earlier work of [23] 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 around the use of social network interventions in public health [38, 22], clinical medicine [46], international development [5, 2], education [28], and related domains. 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 these strategies is mixed; [22] reports no or marginal improvement for nominations vs random selections (depending on the outcome measure), while [3] report statistically significant improvements for a nomination-based selection mechanism over random targeting. [10] introduce improved statistical methods to compare the effectiveness of seeding strategies by exploiting randomization to develop an off-policy estimator, and conclude that nomination-based strategies do not measurably improve performance. Indeed, [1] show that in some theoretical network models it may be preferable to invest resources that would have been spent carefully mapping the network in simply recruiting a slightly larger number of influencers at random. 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 (which has not been a feature 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 of the selection leads to improved results.

Problem description

The population of youth at whom the intervention is targeted are the nodes of a graph $G = (V, E)$. We seek to recruit a set of youth S to be peer leaders, where S obeys the budget constraint $|S| \leq k$. In domain terms, the budget constraint reflects the fact that peer leaders 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 (specific choices for this model are 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 $(1 - 1/e)$ -approximation can be obtained via a simple greedy algorithm [21].

The most common choice for the model of information diffusion is the independent cascade model. In this model, when a node receives information, they transmit the information 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 randomness in propagation.

The standard formulation for the influence maximization problem concludes at this point. However, in the process of 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 [43, 41, 42]. 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 communities, our intervention aims 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 the people who they actually reported interacting with at the center. Instead, network information must be gathered through in-person interviews with youth 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 to for a given center. Accordingly, the first stage of the 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, in the sense that 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 where algorithms have been developed to learn from repeated and easily observable cascades; 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 [23] 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 the intervention with probability q (based on experience in the 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 [43] 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 [43]. 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 [41]. We now provide an overview of the main components of CHANGE, mirroring the steps of the problem formulation in the previous section.

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 [13]. 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) [23]. 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 [41], 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).

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 [43]), and so we retain the earlier guarantees on the quality of the solution obtained at each individual step. Moreover, in [43] 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 [43].

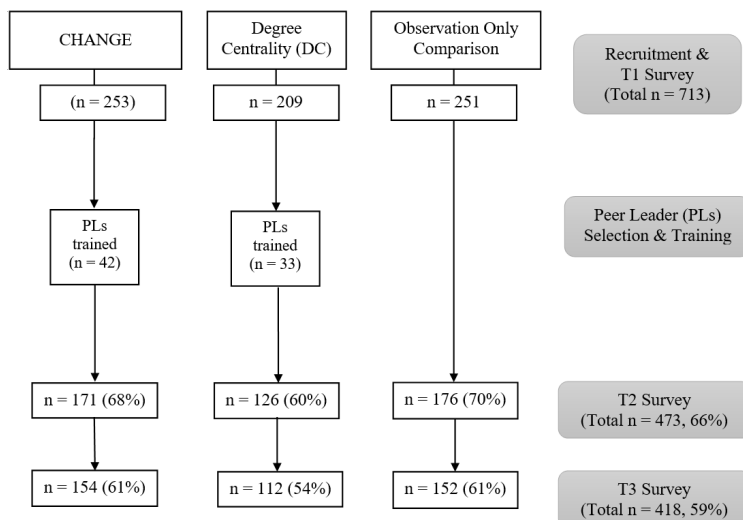


Figure 1: Number of participants recruited and retained in each arm of the study.

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 the University of Southern California 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 the Los Angeles area. 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 are replaced by new youth 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)¹, 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. On enrollment, 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 for details of 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 at each training). 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 used 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 rein-

¹Randomizing treatments at an individual level is clearly impossible for a social network intervention, so this is an example of a quasi-experimental design where entire populations of youth were assigned to one treatment or another.

forcement 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 recruited in the first step. Follow-up surveys assessed the same characteristics as the baseline survey, and differences in reported sexual behavior and HIV-related knowledge 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.

Details of the peer change intervention have been published in the social work literature [33]. Training was delivered by two or three facilitators from the social work research team. The primary intervention 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 an overview of the main results.

Outcome variables

We compared five outcome variables across arms of the study. First, condomless anal sex (CAS). CAS was 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). CVS was assessed via a survey question asking whether youth had vaginal sex without a condom at least once in the previous month. Third, HIV testing, assessed via a survey question asking whether youth had taken an HIV test in the previous six months. Fourth, HIV knowledge, assessed as the percentage of correct answers to a six-item questionnaire on HIV transmission, testing, and prevalence which was incorporated into the survey. Fifth, PrEP awareness. PrEP (pre-exposure prophylaxis) is a medication which people at high risk of contracting HIV can take to help avoid becoming infected. Awareness was assessed on a 4-point scale in the survey.

CAS and CVS are the two most important behavioral risk factors for HIV transmission which were included in the study. Hence, these are the most critical outcome variables to assess the success of intervention. HIV testing is also included because regular testing helps interrupt transmission chains by allowing people who know they are infected to adopt protective behaviors for their partners. HIV knowledge and PrEP awareness measure the contribution of the intervention towards general background knowledge of the youth, regardless of whether this led to changes in behavior.

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 [48]. 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

	CAS		CVS		HIV testing		HIV knowledge		PrEP knowledge		
	OR	CI	OR	CI	OR	CI	<i>b</i>	SE	IRR	CI	
Baseline											
CHANGE	1.43	0.91, 2.28	0.77	0.52, 1.13	1.17	0.75, 1.81	-0.01	0.02	1.49*	1.07, 2.07	
DC	1.49	0.89, 2.48	1.07	0.67, 1.68	0.92	0.56, 1.49	-0.03	0.02	0.69 [†]	0.47, 1.021	
Post-intervention											
CHANGE	0.69*	0.49, 0.98	0.78 [†]	0.57, 1.04	0.81	0.58, 1.12	0.04**	0.01	1.03	0.81, 1.31	
DC	0.80	0.55, 1.17	0.88	0.62, 1.23	0.99	0.64, 1.53	0.04**	0.01	1.30 [†]	0.99, 1.72	
Time	1.05	0.82, 1.33	0.87	0.71, 1.06	1.34*	1.05, 1.71	0.003	0.01	1.47	1.23, 1.76	

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.

[†] $p < 0.1$; * $p < 0.05$; ** $p < 0.01$

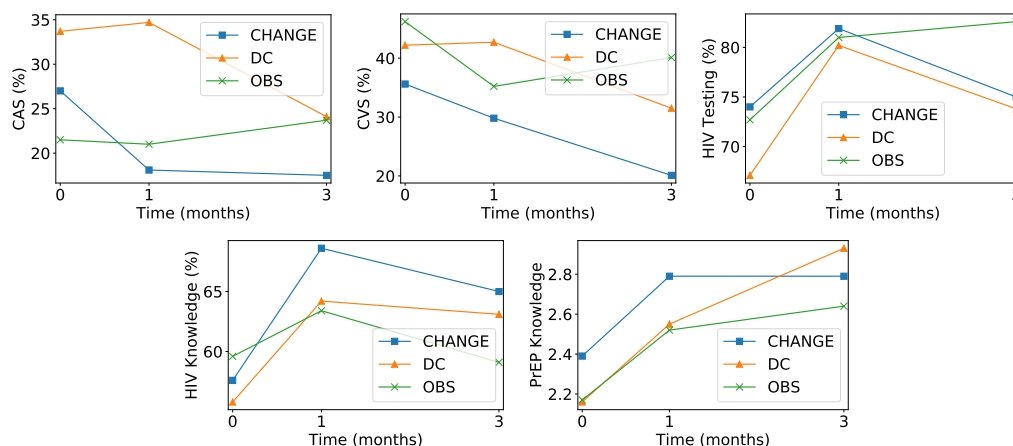


Figure 2: Average value of each outcome variable at each point in time for the three arms. These plots show the results without any statistical processing, while the analysis above attempts to control for pre-existing differences between participants in each arm.

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 link function (logistic for binary outcomes and Poisson for ordinal). For binary outcomes (CAS/CAV/HIV testing), this allows us to 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 demographics and baseline rate of risk behaviors). For PrEP awareness (ordinal), we present the incidence rate ratio (IRR), which gives the ratio in the estimated Poisson rate between the intervention and observation groups (again, after controlling for the other variables). For HIV knowledge (a real number), we present the estimated linear regression coefficient b for participating in the intervention. For all quantities, we also present 95% confidence intervals and indicate where significant p -values are obtained from the GEE model.

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 was had no statistically significant association with CAS, CVS, or HIV testing, 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. 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 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.

HIV testing There was no statistically significant change in HIV testing for either the CHANGE or DC groups relative to the observation-only group. The observation-only group experienced improvements in HIV testing rates over time, and there is insufficient evidence to conclude that either intervention improved HIV testing beyond this trend.

HIV knowledge HIV knowledge increased by a statistically significant amount in both the CHANGE group ($b = 0.04$, $p < 0.01$) and DC group ($b = 0.04$, $p < 0.01$) relative to the observation-only group.

PrEP knowledge PrEP knowledge increased by a marginally statistically significant amount in the DC group compared to observation-only (IRR = 1.3, $p < 0.1$) but not in the CHANGE group. Estimation for this outcome variable was complicated by the fact that PrEP knowledge was higher by a statistically significant margin in the CHANGE group at baseline (prior to the intervention) than the observation-only group, and lower in the DC group by a marginally significant amount.

We conclude from the statistical analysis that both intervention arms (CHANGE and DC) offered improvement compared to the observation-only baseline. However, the CHANGE group demonstrated significantly greater improvement in the most important outcome variables: CAS and CVS. The DC group did not experience a statistically significant improvement in either of these variables, which were the main HIV-related risk behaviors assessed in the study. Both interventions demonstrated some level of effectiveness at improving HIV-related knowledge.

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 in CAS, CVS, and HIV knowledge occur by the one-month survey, while improvements in the DC group are not fully realized until the three-month survey. Fast results are especially important in this domain 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, particularly in the context of community-level interventions. There are four points which we highlight here.

First, 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.

Second, 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 [42]. 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 an action, more information is input to produce the next action, and so on) then there are more things that can go wrong outside of the simplified version of the world which is captured in the computational formalization of the problem.

Third, 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 [43, 44, 45]. 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.

Fourth, 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 the names of their 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 conducted studies and 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 for the program based on their contacts) than when such an explanation was not proactively given. A critical part of the peer change agent model is empowering the youth themselves 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 research was supported by MURI Grant W911NF-11-1- 0332, the California HIV/AIDS Research Program, and a NSF Graduate Research Fellowship.

References

- [1] Mohammad Akbarpour, Suraj Malladi, and Amin Saberi. Diffusion, seeding, and the value of network information. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 641–641, 2018.
- [2] Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. The diffusion of microfinance. *Science*, 341(6144), 2013.
- [3] Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies*, 86(6):2453–2490, 2019.
- [4] Christian Borgs, Michael Brautbar, Jennifer Chayes, and Brendan Lucier. Maximizing social influence in nearly optimal time. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 946–957. SIAM, 2014.
- [5] Jing Cai, Alain De Janvry, and Elisabeth Sadoulet. Social networks and the decision to insure. *American Economic Journal: Applied Economics*, 7(2):81–108, 2015.
- [6] Wei Chen, Tian Lin, Zihan Tan, Mingfei Zhao, and Xuren Zhou. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 795–804, 2016.
- [7] Wei Chen, Chi Wang, and Yajun Wang. 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*, pages 1029–1038, 2010.
- [8] Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 199–208, 2009.
- [9] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159, 2013.
- [10] Alex Chin, Dean Eckles, and Johan Ugander. Evaluating stochastic seeding strategies in networks. *arXiv preprint arXiv:1809.09561*, 2018.
- [11] National HCH Council. HIV/AIDS among persons experiencing homelessness, 2012.
- [12] Nan Du, Yingyu Liang, Maria Balcan, and Le Song. Influence function learning in information diffusion networks. In *International Conference on Machine Learning*, pages 2016–2024, 2014.
- [13] Scott L Feld. Why your friends have more friends than you do. *American Journal of Sociology*, 96(6):1464–1477, 1991.
- [14] Amit Goyal, Wei Lu, and Laks VS Lakshmanan. Celf++ optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th international conference companion on World Wide Web*, pages 47–48, 2011.

- [15] Harold D Green Jr, Kayla de la Haye, Joan S Tucker, and Daniela Golinelli. Shared risk: who engages in substance use with American homeless youth? *Addiction*, 108(9):1618–1624, 2013.
- [16] NIMH Research Group et al. 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, 2010.
- [17] Xinran He and David Kempe. Stability and robustness in influence maximization. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 12(6):1–34, 2018.
- [18] Xinran He, Ke Xu, David Kempe, and Yan Liu. Learning influence functions from incomplete observations. In *Advances in Neural Information Processing Systems*, pages 2073–2081, 2016.
- [19] Dimitris Kalimeris, Yaron Singer, Karthik Subbian, and Udi Weinsberg. Learning diffusion using hyperparameters. In *International Conference on Machine Learning*, pages 2420–2428, 2018.
- [20] Jeffrey A Kelly, Debra A Murphy, Kathleen J Sikkema, Timothy L McAuliffe, Roger A Roffman, Laura J Solomon, Richard A Winett, Seth C Kalichman, and The Community HIV Prevention Research Collaborative. Randomised, controlled, community-level hiv-prevention intervention for sexual-risk behaviour among homosexual men in US cities. *The Lancet*, 350(9090):1500–1505, 1997.
- [21] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- [22] David A Kim, Alison R Hwang, Derek Stafford, D Alex Hughes, A James O’Malley, James H Fowler, and Nicholas A Christakis. Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet*, 386(9989):145–153, 2015.
- [23] Andreas Krause, H Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust submodular observation selection. *Journal of Machine Learning Research*, 9(Dec):2761–2801, 2008.
- [24] Meghna Lowalekar, Pradeep Varakantham, and Akshat Kumar. Robust influence maximization. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1395–1396, 2016.
- [25] Amy Medley, Caitlin Kennedy, Kevin O’Reilly, and Michael Sweat. 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, 2009.
- [26] Matthew H Morton, Amy Dworsky, Jennifer L Matjasko, Susanna R Curry, David Schlueter, Raúl Chávez, and Anne F Farrell. Prevalence and correlates of youth homelessness in the United States. *Journal of Adolescent Health*, 62(1):14–21, 2018.
- [27] Harikrishna Narasimhan, David C Parkes, and Yaron Singer. Learnability of influence in networks. In *Advances in Neural Information Processing Systems*, pages 3186–3194, 2015.
- [28] Elizabeth Levy Paluck, Hana Shepherd, and Peter M Aronow. Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences*, 113(3):566–571, 2016.
- [29] Jean Pouget-Abadie and Thibaut Horel. Inferring graphs from cascades: A sparse recovery framework. *arXiv preprint arXiv:1505.05663*, 2015.
- [30] María C Rangel, Loretta Gavin, Christie Reed, Mary G Fowler, and Lisa M Lee. Epidemiology of hiv and aids among adolescents and young adults in the United States. *Journal of adolescent health*, 39(2):156–163, 2006.

- [31] Eric Rice, Anamika Barman-Adhikari, Norweeta G. Milburn, and William Monro. Position-specific hiv risk in a large network of homeless youths. *American Journal of Public Health*, 102(1):141–147, 2012. PMID: 22095350.
- [32] Eric Rice, William Monro, Anamika Barman-Adhikari, and Sean D Young. Internet use, social networking, and hiv/aids risk for homeless adolescents. *Journal of Adolescent Health*, 47(6):610–613, 2010.
- [33] Eric Rice, Eve Tulbert, Julie Cederbaum, Anamika Barman Adhikari, and Norweeta G. Milburn. Mobilizing homeless youth for HIV prevention: a social network analysis of the acceptability of a face-to-face and online social networking intervention. *Health Education Research*, 27(2):226–236, 01 2012.
- [34] Mary Jane Rotheram-Borus, Emily Davis, and Roxana Rezai. Stopping the rise of hiv among adolescents globally. *Current opinion in pediatrics*, 30(1):131, 2018.
- [35] Mary Jane Rotheram-Borus and Marya Gwadz. Sexuality among youths at high risk. *Child and Adolescent Psychiatric Clinics of North America*, 2(3):415–430, 1993.
- [36] John A Schneider, A Ning Zhou, and Edward O Laumann. A new hiv prevention network approach: sociometric peer change agent selection. *Social science & medicine*, 125:192–202, 2015.
- [37] Youze Tang, Xiaokui Xiao, and Yanchen Shi. Influence maximization: Near-optimal time complexity meets practical efficiency. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 75–86, 2014.
- [38] Thomas W Valente and Patchareeya Pumpuang. Identifying opinion leaders to promote behavior change. *Health education & behavior*, 34(6):881–896, 2007.
- [39] Qinshi Wang and Wei Chen. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Advances in Neural Information Processing Systems*, pages 1161–1171, 2017.
- [40] Zheng Wen, Branislav Kveton, Michal Valko, and Sharan Vaswani. Online influence maximization under independent cascade model with semi-bandit feedback. In *Advances in neural information processing systems*, pages 3022–3032, 2017.
- [41] Bryan Wilder. Equilibrium computation and robust optimization in zero sum games with submodular structure. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [42] Bryan Wilder, Nicole Immerlica, Eric Rice, and Milind Tambe. Maximizing influence in an unknown social network. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [43] Bryan Wilder, Laura Onasch-Vera, Juliana Hudson, Jose Luna, Nicole Wilson, Robin Petering, Darlene Woo, Milind Tambe, and Eric Rice. End-to-end influence maximization in the field. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1414–1422, 2018.
- [44] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *AAMAS*, volume 16, pages 740–748, 2016.
- [45] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. 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*, pages 150–158, 2017.
- [46] Jane M Young, Michael J Hollands, Jeanette Ward, and CD’Arcy J Holman. Role for opinion leaders in promoting evidence-based surgery. *Archives of Surgery*, 138(7):785–791, 2003.
- [47] Sean D Young and Eric Rice. Online social networking technologies, HIV knowledge, and sexual risk and testing behaviors among homeless youth. *AIDS and Behavior*, 15(2):253–260, 2011.

- [48] Scott L. Zeger, Kung-Yee Liang, and Paul S. Albert. Models for longitudinal data: A generalized estimating equation approach. *Biometrics*, 44(4):1049–1060, 1988.