

Policy-Embedded Graph Expansion: Networked HIV Testing with Diffusion-Driven Network Samples

Akseli Kangaslahti¹, Davin Choo¹, Ling kai Kong¹, Milind Tambe¹,
Alastair van Heerden^{2,3} and Cheryl Johnson⁴

¹Harvard University

²University of Witwatersrand

³Wits Health Consortium

⁴World Health Organization

akselikangaslahti@g.harvard.edu, davinchoo@seas.harvard.edu, lingkai.kong@g.harvard.edu,
tambe@g.harvard.edu, alastair.vanheerden@wits.ac.za, johnsonc@who.int

Abstract

HIV is a retrovirus that attacks the human immune system and can lead to death without proper treatment. In collaboration with the WHO and the University of Witwatersrand, we study how to improve the efficiency of HIV testing with the goal of eventual deployment, directly supporting progress toward UN Sustainable Development Goal 3.3. While prior work has demonstrated the promise of intelligent algorithms for sequential, network-based HIV testing, existing approaches rely on assumptions that are impractical in our real-world implementations. Here, we study sequential testing on incrementally revealed disease networks and introduce Policy-Embedded Graph Expansion (PEGE), a novel framework that directly embeds a generative distribution over graph expansions into the decision-making policy rather than attempting explicit topological reconstruction. We further propose Dynamics-Driven Branching (DDB), a diffusion-based graph expansion model that supports decision making in PEGE and is designed for data-limited settings where forest structures arise naturally, as in our real-world referral process. Experiments on real HIV transmission networks show that the combined approach (PEGE + DDB) consistently outperforms baselines (e.g., 17.3% improvement in discounted reward and 15.4% more HIV detections with 25% of the population tested) and explore key tradeoffs that drive solution quality.

1 Introduction

Human immunodeficiency virus (HIV) attacks the human immune system and, without treatment, can lead to acquired immunodeficiency syndrome (AIDS). Although there is currently no cure available for HIV, proper medical treatment can help control the virus and its spread. UNAIDS has proposed the 95-95-95 targets for HIV that aim for 95% of individuals with HIV to be aware of their status, 95% of those aware of their positive status to receive treatment, and 95% of people

that are treated to reach viral suppression [UNAIDS, 2022]. However, a more recent UNAIDS report has shown that the first of these goals has proven to be the most difficult to reach, as about 1 in 7 individuals with HIV have not been diagnosed and 1.3 million new infections occur each year [UNAIDS, 2024]. Furthermore, with recent funding shifts, projections expect these gaps to grow further [Ten Brink *et al.*, 2025]. Therefore, in line with the UN Sustainable Development Goal 3.3 [United Nations, nd], we aim to help people with HIV get tested as quickly as possible so that they can begin receiving treatment and eventually achieve viral suppression. This also facilitates efficient delivery of HIV prevention such as pre-exposure prophylaxis (PrEP) to high-risk individuals with connections to people who test positive.

In some countries like South Africa with high HIV prevalence, the concept of Universal Test and Treat (UTT) is in place. In UTT, everyone in a community is offered HIV testing and anyone who tests positive is immediately started on antiretroviral therapy (ART) to suppress the virus, reduce transmission, and improve individual health. However, testing resources such as test kits and staff can be limited in practice, so we need to develop approaches for determining an efficient testing order. The World Health Organization (WHO) has recommended network-based testing (NBT) for HIV [World Health Organization, 2024], which has proven to be effective in South Africa [Jubilee *et al.*, 2019] and for infectious diseases beyond HIV (e.g., [Juher *et al.*, 2017; Monroe-Wise *et al.*, 2025]). The goal of NBT is to identify individuals at risk for STIs by engaging their social, sexual, or family networks. This strategy helps reach stigmatized populations that are often inaccessible via centralized outreach through peer referrals that leverage existing social trust between individuals. A prominent instantiation is the Enhanced Peer Outreach Approach (EPOA), where influential “seeds” from key populations use performance-based incentives to recruit high-risk contacts, effectively bridging the gap between hidden networks and clinical services [FHI 360, nd]. However, the complexity of NBT makes it difficult for humans to determine optimal decisions by hand, motivating the need for an intelligent decision-making framework. Prior work has proposed network-based testing strategies for HIV [Zhang

and Heath, 2023; Choo *et al.*, 2025] and other diseases like COVID-19 [Cui *et al.*, 2021]. However, these approaches require full network observability, which is prohibitively expensive to collect in practice. While some work on unknown networks exists, many efforts target different testing problems [McFall *et al.*, 2021] or base recommendations on time of infection data, which we do not have access to [Nikolopoulos *et al.*, 2016; Morgan *et al.*, 2019].

With these considerations in mind, we study the problem of allocating a limited number of HIV tests across individuals in a network to identify as many HIV-positive individuals as possible, as quickly as possible. Specifically, we consider a sequential frontier-based formulation (see Section 3) where the network is revealed incrementally as test recipients refer contacts to testing services. Furthermore, as in many health domains, we need to adapt to small and noisy training data, which is an especially difficult challenge when combined with the high-dimensional node covariates that we observe and the combinatorial explosions that are natural in graph-based problems.

1.1 Contributions

In this paper, we study how generative graph expansion can be embedded in the sequential decision-making loop to tackle the limited observability of the disease network and data limitations. Specifically, our contributions are:

- 1) We formalize our HIV testing setting into a new sequential decision-making problem on incrementally revealed graphs, which has applications beyond HIV testing.
- 2) We introduce a new approach called Policy-Embedded Graph Expansion (PEGE) for solving sequential decision-making problems on incrementally revealed graphs. PEGE is the first approach to this class of problems that directly embeds a generative graph expansion model into the sequential decision-making policy loop, enabling intelligent decisions that respect network uncertainty without requiring full topological accuracy on hidden network reconstruction. PEGE is well suited for problems with data limitations and/or high variance in the realization of the unobservable graph.
- 3) We further propose Dynamics Driven Branching (DDB), a new generative graph expansion model tailored to data-scarce settings where forest topologies occur naturally, e.g., our network referral system, that trains without graph-level supervision and supports decision-making in PEGE. DDB is a hybrid diffusion and Gaussian Process Regression (GPR) model that learns to capture the node-to-node dynamics that drive variation in the unobservable portion of the network.
- 4) We empirically evaluate our approach (PEGE + DDB), along with several baseline algorithms, showing that it outperforms all baselines on real-world HIV networks. We also provide additional experiments that highlight key tradeoffs that drive solution quality.

1.2 Interdisciplinary Collaboration and Impact

Our interdisciplinary research team consists of both AI researchers and domain experts from the WHO and the University of Witwatersrand. This work on network-based testing strategies for intelligent HIV diagnosis and prevention is a part of a larger ongoing collaboration to study how AI can be



Figure 1: HIV self-test kit distribution in South Africa.

used to support HIV prevention and treatment efforts in South Africa. The domain experts on our team have played a crucial and iterative role in designing the problem formulation to make it as practical as possible, assessing real-world applicability of proposed approaches, and connecting our research to avenues for deployment.

We are currently aiming to deploy the work presented in this paper to have a real-world impact in improving our HIV testing efforts in South Africa. Figure 1 shows the team distributing HIV self-test kits in the region. Specifically, we want to integrate our methods to determine testing sequences in clinics with limited testing resources like test kits and staff.

2 Related Work

2.1 Network-Based Disease Testing

Network-based modeling of HIV and other infectious diseases is a well-studied field [Mattie *et al.*, 2022]. One area is in algorithms for selecting groups from known networks for pooled testing to maximize resource efficiency [Silva and Singh, 2021; Sewell, 2022], with some efforts incorporating dynamically updating networks [Srinivasavaradhan *et al.*, 2024]. Another effort has proposed a Gittins index-based policy for sequential frontier-based testing that we leverage in this work [Choo *et al.*, 2025]. However, these efforts assume full network observability, which is often unrealistic in practice. Other studies on unknown networks have either focused on other types of testing-based problem formulations [McFall *et al.*, 2021; Tan *et al.*, 2025] or on higher-level strategies reliant on temporal infection data [Nikolopoulos *et al.*, 2016; Morgan *et al.*, 2019], which are not applicable to our setting.

2.2 Sequential Decision-Making on Incrementally Revealed Graphs

Many types of sequential decision-making problems on incrementally revealed graphs have close ties to the problem that we study in this work (see Section 3). For example, in active search on graphs, the goal is to query as many nodes with a hidden target label as possible given a budget and graph topology [Wang *et al.*, 2013; Tsui *et al.*, 2024]. Active Learning is a common approach to such problems. Other areas include

online graph exploration, where the goal is to discover all nodes in an initially unknown graph with minimal path cost [Megow *et al.*, 2012], for which learning-augmented and reinforcement learning approaches have been proposed [Eberle *et al.*, 2022; Chiotellis and Cremers, 2020]. Another relevant problem class is influence maximization (IM) with partial or myopic feedback, where the goal is to sequentially select a small set of root/seed nodes in a network that maximizes influence (e.g., spread of information) in a network [Yuan and Tang, 2017; Peng and Chen, 2019]. In fact, IM has been used in real network-based HIV interventions [Wilder *et al.*, 2021], although these algorithms do not apply to our settings since we cannot query the network to learn more information. One closely related effort in IM shows how partial graph observability can be aided via maximum likelihood link prediction conditioned on known node metadata [Tran *et al.*, 2024], solving an intermediate learning task to address partial observability similarly to PEGE. Similar approaches have been proposed in online graph exploration [Tan *et al.*, 2025]. Other efforts in robotic navigation have employed intermediate learning to supplement information about uncertain frontier nodes in a graph [Cui *et al.*, 2024]. Decision-focused learning is an alternative strategy where intermediate learning tasks such as link prediction can be directly optimized with respect to downstream decision-making outcomes [Wilder *et al.*, 2019]. In contrast, in PEGE, the generated graph expansions represent a distribution rather than a single prediction. The model is not trained on intermediate topological accuracy or downstream decision quality; rather, its goal is to iteratively capture variation in decision-relevant areas of a graph, which change as the graph grows incrementally. This makes PEGE more suitable for high-dimensional and data-scarce settings with high variance in the realization of the unobservable graph, where topology prediction is challenging and learned decision-focused dynamics from training may not transfer well to test graphs.

3 Problem Definition

We first define a more abstract problem class and solution approach for generalization beyond HIV testing, and in Section 5 show how it applies to our HIV testing application. Specifically, we define *Sequential Acting on Partially Observed Graphs (SAPOG)* as a class of sequential decision-making problems in which an agent acts on vertices of a fixed but initially partially observed graph, with actions revealing additional structure and node information over time.

Throughout this section, we will use $\subseteq$ to denote subsets and subgraphs, $|\cdot|$ to denote set cardinality, $N(v)$ to refer to the neighbors of a node v , and $\setminus$ to denote set subtraction. We will also refer to node attributes as covariates.

Definition 1 (Sequential Acting on Partially Observed Graphs (SAPOG)). A SAPOG instance is specified by the tuple $\mathcal{P} = (\mathcal{G}, \mathbf{X}, \mathbf{Y}, s^{(0)}, R, \gamma, T)$, where the fixed ground truth graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ is revealed incrementally over $T \in \mathbb{N}_{\geq 1}$ rounds with respect to discount factor $\gamma \in (0, 1]$. Functions $\mathbf{X} : \mathbf{V} \rightarrow \mathbb{R}^{d_c}$ and $\mathbf{Y} : \mathbf{V} \rightarrow \mathbb{R}^{d_\ell}$ map each node to covariates and labels (e.g., binary HIV status) respectively. At any timestep $t \in [T]$, the state $s^{(t)} = (\mathbf{V}^{(t)}, \mathbf{V}_Y^{(t)})$ is a




Figure 2: An illustration of a SAPOG instance at time t . Solid nodes and edges represent the induced subgraph $\mathcal{G}^{(t)} = \mathcal{G}[\mathbf{V}^{(t)}]$. Grey nodes represent nodes that have already been acted on, i.e., $\mathbf{V}_Y^{(t)}$. Green nodes represent frontier nodes, which are nodes that have been discovered but not acted on, and thus have an unrevealed status. These frontier nodes make up the action space $\mathbf{V}^{(t)} \setminus \mathbf{V}_Y^{(t)}$. Dashed nodes and edges represent unrevealed nodes, i.e., $\mathbf{V} \setminus \mathbf{V}^{(t)}$ and their unrevealed incident edges. (1) Agent evaluates the state $s^{(t)}$. (2) Agent acts on the red node $a \in \mathbf{V}^{(t)} \setminus \mathbf{V}_Y^{(t)}$. (3) Agent observes the label of node a and reveals its previously undiscovered neighbors in the next state.

pair of vertex subsets with $\mathbf{V}_Y^{(t)} \subseteq \mathbf{V}^{(t)} \subseteq \mathbf{V}$, where $\mathbf{V}^{(t)}$ is the set of nodes discovered up to time t and $\mathbf{V}_Y^{(t)}$ is the subset with revealed labels. In other words, given state $s^{(t)}$, we observe the induced subgraph $\mathcal{G}^{(t)} = \mathcal{G}[\mathbf{V}^{(t)}]$, along with covariates $\mathbf{X}^{(t)} = \{\mathbf{X}(v) \in \mathbb{R}^{d_c} : v \in \mathbf{V}^{(t)}\}$ and labels $\mathbf{Y}^{(t)} = \{\mathbf{Y}(v) \in \mathbb{R}^{d_\ell} : v \in \mathbf{V}_Y^{(t)}\}$. Upon picking an action from the frontier $a \in \mathbf{V}^{(t)} \setminus \mathbf{V}_Y^{(t)}$, we reveal the label of node a and its previously undiscovered neighbors, i.e., $\mathbf{V}^{(t+1)} = \mathbf{V}^{(t)} \cup N(a)$ and $\mathbf{V}_Y^{(t+1)} = \mathbf{V}_Y^{(t)} \cup \{a\}$. Given some initial state $s^{(0)} = (\mathbf{V}^{(0)}, \emptyset)$, the goal is to compute a policy π that maps each state $s^{(t)}$ to a node in $\mathbf{V}^{(t)}$, maximizing the expected total discounted reward:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{T-1} \gamma^t \cdot R(s^{(t)}, \pi(s^{(t)})) \right] \quad (1)$$

where $\pi(s^{(t)}) \in \mathbf{V}^{(t)} \setminus \mathbf{V}_Y^{(t)}$ is the frontier node selected by policy π and $R(s^{(t)}, v)$ is the (possibly label-dependent) reward for acting on node $v = \pi(s^{(t)})$ at round t .

Figure 2 shows an example timestep in a SAPOG instance. The SAPOG class captures several other existing problem formulations. For example, Adaptive Frontier Exploration on Graphs [Choo *et al.*, 2025] is a special case where $\mathbf{V}^{(0)} = \mathbf{V}$. Additionally, active search on graphs [Wang *et al.*, 2013] and selective harvesting [Murai *et al.*, 2018] are example undiscounted cases where $\gamma = 1$ and $T < |\mathbf{V}|$ and, in active search, $\mathbf{V}^{(0)} = \mathbf{V}$. The frontier-based decision-making in SAPOG is also analogous to other fields where local actions reveal information about the neighboring environment, such as robotic exploration [Keidar and Kaminka, 2014].

4 Methods

4.1 SAPOG Approach Analysis

There are several challenges in solving SAPOG instances. First, since $(\mathcal{G}, \mathbf{X}, \mathbf{Y}) \setminus (\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$ is unknown, the agent needs to make decisions without knowledge of their long-term impact or even their short-term rewards. Second, it is common that in practice, there are multiple different possible realizations of the completed graph, i.e., the agent should view $(\mathcal{G}, \mathbf{X}, \mathbf{Y})$ as a distribution conditioned on $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$. While algorithms for the case where $\mathbf{V}^{(0)} = \mathbf{V}$ can be formulated, they are difficult to apply to general SAPOG instances where the graph and covariates are only partially observable and incrementally revealed.

Prior work has shown that even in the fully observable case where $\mathbf{V}^{(0)} = \mathbf{V}$, directly learning the Q function of nodes is not an effective approach [Choo *et al.*, 2025], motivating some intermediate learning step to address the uncertainty in the unobservable portion of the graph. One natural approach is used in prior work for similar decision-making problems on partially observable graphs [Tran *et al.*, 2024]. In the context of SAPOG, the analogous approach would be to learn to predict the maximum-likelihood graph and covariate completion conditioned on the observable $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$, then iteratively plug the learned completion into an oracle that solves for optimal actions under the assumption that the learned completion is correct. More formally, we define such a decision-making oracle $\mathcal{O}$ as an algorithm that returns a real-valued evaluation (e.g., Q-value) of any action v given $\mathbf{Y}^{(t)}$ and any $(\mathcal{G}' \supseteq \mathcal{G}^{(t)}, \mathbf{X}' \supseteq \mathbf{X}^{(t)})$, i.e., any possible completion of the observed $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)})$ that is assumed by $\mathcal{O}$ to be the ground truth. However, we show in Proposition 1 that this strategy is not necessarily effective in SAPOG.

Proposition 1 (Non-Optimality of Maximum-Likelihood Graph Completion). *Assume that, during any timestep t of a SAPOG instance, an agent:*

1. *Has access to a decision-making oracle $\mathcal{O}$ that returns the Q-value of given actions.*
2. *Has access to $(\mathcal{G}^* \supseteq \mathcal{G}^{(t)}, \mathbf{X}^* \supseteq \mathbf{X}^{(t)})$, the maximum-likelihood completion of the observed $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)})$.*

Then, the action $v^ = \arg \max_{v \in \mathbf{V}^{(t)} \setminus \mathbf{V}_{\mathbf{Y}}^{(t)}} \mathcal{O}(v \mid \mathbf{Y}^{(t)}, \mathcal{G}^*, \mathbf{X}^*)$, is not necessarily optimal in expectation.*

Proof. Suppose that $(\mathcal{G}_1, \mathbf{X}_1), (\mathcal{G}_2, \mathbf{X}_2)$ are the only two possible completions of $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)})$, with $\Pr(\mathcal{G}_1, \mathbf{X}_1 \mid \mathbf{Y}^{(t)}, \mathcal{G}^{(t)}, \mathbf{X}^{(t)}) = \frac{3}{5}$ and $\Pr(\mathcal{G}_2, \mathbf{X}_2 \mid \mathbf{Y}^{(t)}, \mathcal{G}^{(t)}, \mathbf{X}^{(t)}) = \frac{2}{5}$. Thus, $\mathcal{G}^* = \mathcal{G}_1$ and $\mathbf{X}^* = \mathbf{X}_1$. Now, for some constant $c > 0$, suppose that there exist two possible actions $v_1, v_2 \in \mathbf{V}^{(t)} \setminus \mathbf{V}_{\mathbf{Y}}^{(t)}$ such that $\mathcal{O}(v_1 \mid \mathbf{Y}^{(t)}, \mathcal{G}_1, \mathbf{X}_1) = 2c$, $\mathcal{O}(v_2 \mid \mathbf{Y}^{(t)}, \mathcal{G}_1, \mathbf{X}_1) = c$, $\mathcal{O}(v_1 \mid \mathbf{Y}^{(t)}, \mathcal{G}_2, \mathbf{X}_2) = c$, and $\mathcal{O}(v_2 \mid \mathbf{Y}^{(t)}, \mathcal{G}_2, \mathbf{X}_2) = 10c$. Then, $v^* = v_1$, which is sub-optimal since the expected Q-value v_1 is $\frac{8c}{5}$ while v_2 has a larger expected Q-value of $\frac{23c}{5}$. $\square$

Algorithm 1 Policy-Embedded Graph Expansion (PEGE)

Input: SAPOG instance $\mathcal{P} = (\mathcal{G}, \mathbf{X}, \mathbf{Y}, s^{(0)}, R, \gamma, T)$, sampling parameter $k \in \mathbb{N}$, generative graph expansion model $\mathcal{M}$, decision-making oracle $\mathcal{O}$

Output: Actions $a^{(1)}, \dots, a^{(T)}$ for each time step $t \in [T]$

- 1: **for** $t = 0, 1, \dots, T - 1$ **do**
- 2: Observe state $s^{(t)} = (\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$
- 3: Using $\mathcal{M}(s^{(t)})$, sample k graph completions $\bar{s}_1^{(t)}, \dots, \bar{s}_k^{(t)}$ where $\bar{s}_j^{(t)} = (\mathcal{G}_j^{(t)}, \mathbf{X}_j^{(t)})$
- 4: For all actions $a \in \mathbf{V}^{(t)} \setminus \mathbf{V}_{\mathbf{Y}}^{(t)}$, define $f(a)$ as the empirical estimate of $\mathbb{E}_{\bar{s}^{(j)} \sim \mathcal{M}(s^{(t)})} [\mathcal{O}(a \mid \mathbf{Y}^{(t)}, \bar{s}^{(j)})]$:

$$f(a) = \frac{1}{k} \sum_{j=1}^k \mathcal{O}(a \mid \mathbf{Y}^{(t)}, \bar{s}_j^{(t)})$$

- 5: Take action $a^{(t+1)} \in \arg \max_{a \in \mathbf{V}^{(t)} \setminus \mathbf{V}_{\mathbf{Y}}^{(t)}} f(a)$
 - 6: **end for**
 - 7: **return** $a^{(1)}, \dots, a^{(T)}$
-

Let $\mathcal{D}^{(t)}$ be the true distribution of possible graph completions given $(\mathbf{Y}^{(t)}, \mathcal{G}, \mathbf{X}^{(t)})$. Since we do not know the true future, we wish to evaluate an action v by maximizing

$$\mathbb{E}_{(\mathcal{G}', \mathbf{X}') \sim \mathcal{D}^{(t)}} [\mathcal{O}(v \mid \mathbf{Y}^{(t)}, \mathcal{G}', \mathbf{X}')] \quad (2)$$

As shown in Proposition 1, we should not decide on the action via $\mathcal{O}(v \mid \mathbf{Y}^{(t)}, \text{mode}(\mathcal{D}^{(t)}))$, where $\text{mode}(\mathcal{D}^{(t)})$ denotes the maximum-likelihood completion in $\mathcal{D}^{(t)}$, but instead learn the distribution $\mathcal{D}^{(t)}$ itself. This would then allow us to approximate Eq. (2) via samples. For this task, we consider a generative graph expansion model $\mathcal{M}$ that, given any state $s^{(t)} = (\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$, can be sampled to produce a graph expansion $\bar{s}^{(t)} = (\mathcal{G}', \mathbf{X}')$ such that $\mathcal{G}^{(t)} \subseteq \mathcal{G}'$ and $\mathbf{X}^{(t)} \subseteq \mathbf{X}'$. We will see in Section 4.2 that $\mathcal{M}$ does not need to accurately predict the fully topologies in $\mathcal{D}^{(t)}$ in order to support decision-making in PEGE.

4.2 Policy-Embedded Graph Expansion (PEGE)

With the challenges of SAPOG and our analysis in mind, we propose Policy-Embedded Graph Expansion (PEGE), a new approach to sequential decision-making problems over incrementally revealed graphs. To the best of our knowledge, PEGE is the first attempt at leveraging a generative graph expansion model to address limited graph observability in the sequential decision-making loop for an incrementally revealed graph. PEGE solves SAPOG instances as described in Algorithm 1 and illustrated in Figure 3.

The generative graph expansion model $\mathcal{M}$ and the decision-making oracle $\mathcal{O}$ (defined in Section 4.1) have a closely intertwined relationship: the graph expansions are plugged into $\mathcal{O}$ to decide which node to test on timestep t , which influences which nodes $\mathcal{M}$ expands from during timestep $t+1$ after neighboring true graph nodes are revealed. Further, there can be less of a burden on $\mathcal{M}$ of learning the full distribution $\mathcal{D}^{(t)}$ conditioned on the partially observed




Figure 3: Workflow Diagram for PEGE. Grey nodes represent $\mathbf{V}_Y^{(t)}$ while green nodes represent frontier nodes, i.e., $\mathbf{V}^{(t)} \setminus \mathbf{V}_Y^{(t)}$. The red node represents an action. Blue nodes and edges represent graph expansions generated by $\mathcal{M}$ that are treated as true observed nodes and covariates by the oracle. 1) The state $s^{(t)}$ is observed. 2) We sample k graph expansions up to depth d from frontier nodes. In our case, we use a diffusion-driven model for $\mathcal{M}$, as described in Section 5. In this example, $k = 3$ and $d = 1$. 3) The oracle solves each of the k instances separately and the resulting action evaluations are aggregated. 4) The agent takes an action v . 5) Reward, the label of v , and new nodes adjacent to v and their covariates are observed as we transition to the next state.

graph since its output is only used for downstream discounted decision-making. Normally, learning $\mathcal{D}^{(t)}$ is very challenging, as $|\mathbf{V}|$ is unknown, so the model has no sense of how large $\mathcal{G}$ is. However, with a discount factor $\gamma < 1$ in place, generated nodes become less relevant to decision-making as they get further from the frontier nodes. Thus, we can limit the graph expansion to a maximum depth d from frontier nodes to help address the unknown size of the complete $\mathcal{G}$ while ensuring that $\mathcal{M}$ can still capture the variation in possible graph completions that most strongly impacts decision-making. $\mathcal{M}$ must be able to adapt when these decision-relevant areas of the graph (i.e., the areas near the frontier) change as the graph is incrementally revealed and the frontier updates. The Monte-Carlo nature of PEGE helps produce actions that are more robust to variation in this decision-relevant portion of $(\mathcal{G}, \mathbf{X}, \mathbf{Y})$ conditioned on $(\mathcal{G}^{(t)}, \mathbf{X}^{(t)}, \mathbf{Y}^{(t)})$.

5 Application to Network-Based HIV Testing

In our sequential network-based HIV testing setting (see Section 1.2), we are interested in producing an efficient testing order that identifies as many HIV-positive individuals as possible as quickly as possible. With budgeted testing kits and staff, this makes efficient use of limited testing resources, helps people with HIV and associated high-risk contacts get started with treatment and prevention as early as possible, and limits further spread of the virus. In this application, $\mathcal{G}$ represents a sexual contact network that grows incrementally from a set of root individuals $\mathbf{V}^{(0)}$ with initially unknown statuses ($\mathbf{V}_Y^{(0)} = \emptyset$) as we allocate tests and test recipients refer contacts. One additional and important property in our application is that $\mathcal{G}$ is known to be a forest (see Section 6.1). $\mathbf{X}$ represents binary covariates that are collected about each individual in the network, e.g., gender and occupation, ($d_c = 72$) and $\mathbf{Y}$ represents binary HIV status ($d_\ell = 1$). Our reward

function is simply $R(a) = \mathbf{Y}(a)$. For the oracle $\mathcal{O}$ used in PEGE, we use a Gittins index-based policy that produces real-valued evaluations of nodes and has proven performance guarantees when $\mathcal{G}$ is a forest [Choo *et al.*, 2025]. This oracle learns a joint distribution of infection statuses, modeled via potential functions of individual covariates; see Appendix B of [Choo *et al.*, 2025] for details. At test time, the oracle uses the learned function parameters θ to compute scores via a discrete version of the branching bandit formulation of [Keller and Oldale, 2003], and selects the node with the largest score.

However, one significant challenge remains with applying PEGE to our application: the choice of the generative graph expansion model $\mathcal{M}$. While off-the-shelf models may be more applicable in different problem instances, we have quite specific requirements for $\mathcal{M}$ given the PEGE setting, our application of the SAPOG problem, and our small, noisy dataset (see Section 6.1):

- 1) $\mathcal{M}$ must be a generative model that can produce a distribution of possible graph expansions.
- 2) $\mathcal{M}$ must be able to generate both edges and high-dimensional node covariates.
- 3) $\mathcal{M}$ must be adaptable to various sizes of known partial graphs and tunable expansion depth limits (d).
- 4) $\mathcal{M}$ must preserve the forest property in the graphs that it generates (see Section 6.1).
- 5) $\mathcal{M}$ must be suited to data-scarce settings.

No current model that we are aware of meets all of these criteria. Thus, we design Dynamics-Driven Branching (DDB), a new generative graph expansion model that trains only on low-level node-to-node dynamics with no graph or sequence-level supervision. DDB is a hybrid model that fuses denoising diffusion with Gaussian Process Regression (GPR) and learns to recursively expand branches off of the observed frontier nodes. The forest prior of $\mathcal{G}$ and the recursive property of DDB allow us to tabularize data from training graphs, reducing training of DDB to only node-to-node dynamics. Specifically, let v_{parent} and v_{child} be a pair of parent and child nodes with covariates $\mathbf{x}_{\text{parent}}$ and $\mathbf{x}_{\text{child}}$, respectively. Then:

- 1) We train a GPR model [Pedregosa *et al.*, 2011] to produce a Gaussian distribution of $N(v_{\text{parent}})$ conditioned on $\mathbf{x}_{\text{parent}}$.
- 2) We train a denoising diffusion model to learn a distribution of $\mathbf{x}_{\text{child}}$ conditioned on $\mathbf{x}_{\text{parent}}$. Our diffusion model employs a sinusoidal time embedding followed by four linear layers with GELU activations, trained via MSE loss.

At inference, DDB starts at each frontier node and recursively samples the GPR model to generate the number of successors at each expanded node, then samples the diffusion model to generate a vector of covariates for each of those successors. This way, DDB generates the stochastic node-to-node dynamics that drive variation in possible graph completions near frontier nodes, rather than directly learning specific structures and sequences present in the limited training data.

The task of generating neighboring node covariates is especially challenging since the covariates have 72 dimensions and it is unclear which specific features in the parent node covariates are the most relevant. The diffusion model is thus a critical element of DDB, as its mode covering property helps learn an expressive distribution of the high-dimensional $\mathbf{x}_{\text{child}}$ conditioned on $\mathbf{x}_{\text{parent}}$. Additionally, targeting this task rather

than learning to generate the entire graph in one shot helps limit training time of the diffusion model to only a few minutes and reduce overfitting or mode collapse.

Since DDB only expands new child nodes from parent nodes, it always preserves the forest structure of the input graph. Expansion ends when either the branching naturally terminates or the maximum node depth of d from frontier nodes is reached. DDB is sampled to produce the k graph expansions used in PEGE. Since DDB is conditioned on frontier nodes only and has a tunable expansion depth d , it is indifferent to the size of the known partial graph and can generate graph expansions of any specified depth.

The Gittins index-based oracle operates in $O(m^2)$ time on trees with m nodes [Choo *et al.*, 2025]. In the worst-case, a frontier of n nodes each generating b^d descendants results in $m = nb^d$, yielding a per-sample complexity of $O(n^2b^{2d})$. However this is a very loose upper bound. In practice, the frontier is a small subset of the n nodes in the known graph, and the average branching factor is much lower than the maximum branching factor b . Furthermore, our method is embarrassingly parallel: disjoint connected components (i.e., “communities” that make up a larger population) and each graph expansion can be processed independently across parallel threads before being their results are combined. On an M4 Pro MacBook using 12 threads, a sample run for $k = 24$ and $d = 3$ takes approximately 7, 11, 16 seconds when the observed graph has 50, 100, 150 nodes, respectively.

6 Evaluation

6.1 Dataset

We build our experiments on a publicly available real-world HIV sexual contact network dataset with anonymized entries provided by the Inter-university Consortium for Political and Social Research (ICPSR) [Morris and Rothenberg, 2011]¹. This dataset contains 778 nodes comprising 188 connected components and has an 11.3% HIV prevalence rate. This small size is important for our setting since we expect similar data constraints as we deploy (see Section 1.2). Figure 4 shows an overview of the HIV sexual contact network.

Each node is associated with a 72-dimensional vector of binary covariates that are one-hot encoded from 17 categorical covariates related to gender, sexual orientation, occupation, etc. Additionally, to mirror the real-world data-collection process from our HIV prevention efforts in South Africa (see Section 1.2), we project each connected component in the test graph to a tree, forming a forest. In practice, edges are recorded as people refer their contacts to testing clinics and services. However, people are unlikely to (or may not be allowed to) accept a second referral to the same service, especially when referrals are limited. As such, the forest structure arises naturally and noise is introduced in the form of randomly unrecorded edges. These forest structures have been observed in other studies on peer-to-peer network recruitment




Figure 4: Overview of the ICPSR HIV sexual contact network dataset after forest projection (root nodes shown in red). Nodes in the graph represent people and edges represent sexual contacts.

as well (e.g., [Stein *et al.*, 2018]). Furthermore, prior research has observed that actual STI transmission networks are often sparse and tree-like, and that the Gittins index-based policy remains robust to non-tree graphs [Choo *et al.*, 2025].

We split the ICPSR HIV sexual contact network data into 5 partitions with an equal number of connected components such that no connected components are split up. We also ensure that partitions roughly match the distribution of connected component sizes that is present in the overall dataset. For each experiment in this section, we run 5 SAPOG instances with static test graphs and average the normalized results, weighting each instance by how many nodes are in the test graph. Each of the 5 data partitions serves as the test graph once and serves as part of the training data four times.

6.2 Baseline Algorithms

In addition to our method (PEGE + DDB), we evaluate five baseline algorithms for reference, which we describe below:

- 1) **Random (R)**: Always select a random frontier node.
- 2) **Greedy Neighbor (GN)**: During each timestep, select a frontier node that has a parent node that tested positive. If no such node exists, return a random frontier node. This is an intuitive strategy that a human might apply.
- 3) **Greedy Classifier (GC)**: Use the training data to train a linear support vector classifier (SVC) with balanced class weights [Pedregosa *et al.*, 2011] to predict the probability that a given node tests positive conditioned on its covariates. A linear SVC was selected for the greedy baseline as it outperformed logistic regression, RBF-SVC, and gradient boosting. At test time, we always select the frontier node with the greatest probability of testing positive. This is a heuristic approach that is indifferent to any network-based data and is representative of an individual risk-based algorithm that has been proposed for real-world deployment (see Section 1.2).
- 4) **Deep Q-Network (DQN)**: We implement a DQN baseline [Mnih *et al.*, 2015] with the Pytorch NNConv neural architecture [Fey and Lenssen, 2019] that uses a message-passing GNN with edge-conditioned weights [Gilmer *et al.*, 2017]. We allow the DQN to train on $\mathcal{G}$ and $\mathbf{X}$ (but not $\mathbf{Y}$), which serves as a conservative buffer against other learned policies.
- 5) **Fully Observable Gittins (FOG)**: Apply the Gittins

¹Our codebase is publicly available at <https://github.com/akanga-harvard/policy-embedded-graph-expansion>. The ICPSR dataset is publicly available after completing the user agreement in the ICPSR portal [Morris and Rothenberg, 2011]. Our full paper with additional experiments is available at arxiv.org/abs/2601.16233

index-based policy to $\mathcal{G}$, $\mathbf{X}$, $\mathbf{Y}^{(t)}$ on each timestep and select the frontier node with the highest Gittins index. The feature mapping θ used in the Gittins index policy is allowed to train on the test set, meaning the oracle itself is also stronger in this case. Since this algorithm can observe the entire ground truth network, FOG serves as an unattainable upper bound for any policy with the same information described in our model.

6.3 Experiments and Results

We evaluate methods based on how much normalized discounted reward they accrue throughout the experiments (see Eq. (1) for reward model), which captures our real-world goal of identifying as many HIV-positive individuals as possible as quickly as possible. First, we compare the average performance of our method (PEGE + DDB) to that of each baseline. For this experiment, we fix the expansion depth in DDB at $d = 3$, the number of graph expansions in PEGE at $k = 24$, the discount factor at $\gamma = 0.99$, and the number of tests at $T = |\mathbf{V}|$. We also cap the maximum expansion width at $b = 25$ to guarantee bounded graph expansions. The left side of Figure 5 shows the empirical results. Table 1 also summarizes some performance statistics, including area under the cumulative discounted reward curve (AUC) and cumulative discounted reward at different testing budgets.

On average, our PEGE + DDB method outperforms all baselines. The total AUC for PEGE + DDB is about 9.42% greater than that of the closest baseline (Greedy Classifier). Furthermore, PEGE + DDB remains advantageous under different testing budgets, outperforming the strongest baseline by 17.3%, 8.06%, and 6.83% under testing budgets of 25%, 50%, and 75% of the population, respectively. PEGE also captures about 78.4% of the AUC of FOG, and 76.3%, 76.6%, and 80.6% of the cumulative discounted reward accrued by FOG under testing budgets of 25%, 50%, and 75% of the population, respectively. The remaining gap is likely due to the observability gap, the stronger oracle that FOG accesses, and error in DDB. These boosts that we observe in discounted reward also indicate more positive detections early on. For example, when testing only 25% of the population, PEGE + DDB discovers about 15.4% more positive cases than the strongest baseline. If network-based testing algorithms are eventually deployed with tens or hundreds of thousands of participants, improvements of this scale can translate to hundreds or thousands of additional positive HIV identifications, assuming the same population-wide positivity rate of 11.3%.

Next, to isolate contributions of each component of our method, we present an ablation study comparing PEGE + DDB against Naive Gittins (no graph expansion), PEGE + DDB (Full Theta), which uses the θ trained on the full (train and test) network in the oracle, and the FOG upper bound. The results are shown in the right side of Figure 5. We observe: (i) *Data Richness*: Training θ on the full network increases PEGE + DDB’s AUC by 9.87%, highlighting possible gains from larger datasets. (ii) *Framework Value*: PEGE + DDB outperforms the AUC of Naive Gittins by about 10.2% (similar to improvements over other heuristics), isolating the framework’s benefit from the oracle’s baseline performance. (iii) *Early Impact*: DDB expansions are most useful in early rounds when graph uncertainty is highest. This makes the




Figure 5: Left: Performance of our method (PEGE + DDB) compared to several baselines and the FOG upper bound, as measured by normalized cumulative discounted reward. Right: Performance of our PEGE + DDB method compared to Naive Gittins with no graph expansion, PEGE + DDB with θ trained on the full network in the oracle, and the FOG upper bound.

Algorithm	AUC	25%	50%	75%
Random	0.506	0.389	0.598	0.670
Greedy Neighbor	0.561	0.473	0.622	0.715
Greedy Classifier	0.616	0.526	0.682	0.737
DQN	0.603	0.504	0.665	0.747
PEGE + DDB (Ours)	0.674	0.617	0.737	0.798
<i>Fully Observable Gittins</i>	<i>0.860</i>	<i>0.809</i>	<i>0.962</i>	<i>0.990</i>

Table 1: Total AUC and normalized discounted reward at different testing budgets (25%, 50%, 75% of population) for each algorithm. In each column, the best is **bolded** and the upper bound is *italicized*.

framework especially fit for resource-critical problems (e.g., at 10% budget we outperform Naive Gittins by 66.7%).

7 Conclusion

In this work, we explore a sequential HIV testing problem over an incrementally revealed graph that addresses gaps in prior network-based HIV testing work. We propose a new approach (PEGE + DDB) and perform experiments on real-world HIV networks that show that it consistently outperforms all other baselines. In future work, we are studying related problems in HIV testing, such as referral and incentive distribution algorithms for peer-to-peer network recruitment [Pan *et al.*, 2026] and batched HIV testing policies [Kong *et al.*, 2026]. We are also working to explore how PEGE can be applied to other applications of SAPOG.

Our interdisciplinary team is motivated by our real-world efforts in South Africa, where we are working to deploy the algorithms that we demonstrate in this paper. The improvements that our approach provides can lead to hundreds or thousands of additional HIV detections when deployed at a larger scale of tens or hundreds of thousands of testing participants, which is crucial for getting people with HIV started with treatment, getting their high-risk contacts started with prevention like PrEP, and limiting further spread of the virus. This represents direct progress towards the first of the 95-95-95 targets for HIV [UNAIDS, 2022] and the UN Sustainable Development Goal 3.3 [United Nations, nd].

Acknowledgements

This work was supported by ONR MURI N00014-24-1-2742. The findings and conclusions in this report are those of the authors and do not necessarily represent the official position of the WHO.

References

- [Chiotellis and Cremers, 2020] Ioannis Chiotellis and Daniel Cremers. Neural online graph exploration. *arXiv preprint arXiv:2012.03345*, 2020.
- [Choo *et al.*, 2025] Davin Choo, Yuqi Pan, Tonghan Wang, Milind Tambe, Alastair van Heerden, and Cheryl Johnson. Adaptive Frontier Exploration on Graphs with Applications to Network-Based Disease Testing. *Conference on Neural Information Processing Systems (NeurIPS)*, 2025.
- [Cui *et al.*, 2021] Yapeng Cui, Shunjiang Ni, and Shifei Shen. A network-based model to explore the role of testing in the epidemiological control of the COVID-19 pandemic. *BMC Infectious Diseases*, 21(1):58, 2021.
- [Cui *et al.*, 2024] Xinru Cui, Qiming Liu, Zhe Liu, and Hesheng Wang. Frontier-enhanced topological memory with improved exploration awareness for embodied visual navigation. In *European Conference on Computer Vision*, pages 296–313. Springer, 2024.
- [Eberle *et al.*, 2022] Franziska Eberle, Alexander Lindermayr, Nicole Megow, Lukas Nölke, and Jens Schlöter. Robustification of online graph exploration methods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9732–9740, 2022.
- [Fey and Lenssen, 2019] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with PyTorch Geometric. *Workshop on Representation Learning on Graphs and Manifolds, ICLR*, 2019.
- [FHI 360, nd] FHI 360. LINKAGES Enhanced Peer Outreach Approach, n.d. Available at <https://www.fhi360.org/wp-content/uploads/drupal/documents/resource-linkages-accelerating-implementation.PDF>. Accessed June 12th, 2026.
- [Gilmer *et al.*, 2017] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, pages 1263–1272. Pmlr, 2017.
- [Jubilee *et al.*, 2019] Makhahliso Jubilee, Faith Jiyeon Park, Knowledge Chipango, Kenoakae Pule, Albert Machinda, and Noah Tarubekera. HIV index testing to improve HIV positivity rate and linkage to care and treatment of sexual partners, adolescents and children of PL-HIV in Lesotho. *PLoS One*, 14(3):e0212762, 2019.
- [Juher *et al.*, 2017] David Juher, Joan Saldaña, Robert Kohn, Kyle Bernstein, and Caterina Scoglio. Network-centric interventions to contain the syphilis epidemic in San Francisco. *Scientific Reports*, 7(1):6464, 2017.
- [Keidar and Kaminka, 2014] Matan Keidar and Gal A Kaminka. Efficient frontier detection for robot exploration. *The International Journal of Robotics Research*, 33(2):215–236, 2014.
- [Keller and Oldale, 2003] Godfrey Keller and Alison Oldale. Branching bandits: a sequential search process with correlated pay-offs. *Journal of Economic Theory*, 113(2):302–315, 2003.
- [Kong *et al.*, 2026] Lingkai Kong, Anagha Satish, Hezi Jiang, Akseli Kangaslahti, Andrew Ma, Wenbo Chen, Mingxiao Song, Lily Xu, and Milind Tambe. Latent Spherical Flow Policy for Reinforcement Learning with Combinatorial Actions. *International Conference on Machine Learning (ICML)*, 2026.
- [Mattie *et al.*, 2022] Heather Mattie, Ravi Goyal, Victor De Gruttola, and Jukka-Pekka Onnela. A review of network models for HIV spread. *JAIDS Journal of Acquired Immune Deficiency Syndromes*, pages 10–1097, 2022.
- [McFall *et al.*, 2021] Allison M McFall, Bryan Lau, Carl Latkin, Aylur K Srikrishnan, Santhanam Anand, Canjeevaram K Vasudevan, Shruti H Mehta, and Sunil S Solomon. Optimizing respondent-driven sampling to find undiagnosed HIV-infected people who inject drugs. *AIDS*, 35(3):485–494, 2021.
- [Megow *et al.*, 2012] Nicole Megow, Kurt Mehlhorn, and Pascal Schweitzer. Online graph exploration: New results on old and new algorithms. *Theoretical Computer Science*, 463:62–72, 2012.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [Monroe-Wise *et al.*, 2025] Aliza Monroe-Wise, Magdalena Barr-DiChiara, Antons Mozalevskis, Busisiwe Msimanga, Maeve Brito de Mello, Kafui Senya, Niklas Luhmann, Cheryl Case Johnson, and Rachel Baggaley. Can network-based testing services have an impact beyond testing for HIV? *Sexual Health*, 22(2):SH24027, 2025.
- [Morgan *et al.*, 2019] Ethan Morgan, Britt Skaathun, Georgios K Nikolopoulos, Dimitrios Paraskevis, Leslie D Williams, Pavlo Smyrnov, Samuel R Friedman, and John A Schneider. A network intervention to locate newly HIV infected persons within MSM networks in Chicago. *AIDS and Behavior*, 23(1):15–20, 2019.
- [Morris and Rothenberg, 2011] Martina Morris and Richard Rothenberg. HIV Transmission Network Metastudy Project: An Archive of Data From Eight Network Studies, 1988–2001 (ICPSR 22140). ICPSR Data Holdings, 2011. Available at doi.org/10.3886/ICPSR22140.v1.
- [Murai *et al.*, 2018] Fabricio Murai, Diogo Rennó, Bruno Ribeiro, Gisele L Pappa, Don Towsley, and Krista Gile. Selective harvesting over networks. *Data Mining and Knowledge Discovery*, 32(1):187–217, 2018.

- [Nikolopoulos *et al.*, 2016] Georgios K Nikolopoulos, Eirini Pavlitina, Stephen Q Muth, John Schneider, Mina Psichogiou, Leslie D Williams, Dimitrios Paraskevis, Vana Sypsa, Gkikas Magiorkinis, Pavlo Smyrnov, *et al.* A network intervention that locates and intervenes with recently HIV-infected persons: The Transmission Reduction Intervention Project (TRIP). *Scientific Reports*, 6(1):38100, 2016.
- [Pan *et al.*, 2026] Yuqi Pan, Davin Choo, Haichuan Wang, Milind Tambe, Alastair van Heerden, and Cheryl Johnson. Adaptive Multi-Round Allocation with Stochastic Arrivals. *International Conference on Machine Learning (ICML)*, 2026.
- [Pedregosa *et al.*, 2011] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [Peng and Chen, 2019] Binghui Peng and Wei Chen. Adaptive influence maximization with myopic feedback. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- [Sewell, 2022] Daniel K Sewell. Leveraging network structure to improve pooled testing efficiency. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 71(5):1648–1662, 2022.
- [Silva and Singh, 2021] Arlei Silva and Ambuj Singh. Group testing on a network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4348–4356, 2021.
- [Srinivasavaradhan *et al.*, 2024] Sundara Rajan, Srinivasavaradhan, Pavlos Nikolopoulos, Christina Fragouli, and Suhas Diggavi. Dynamic group testing to control and monitor disease progression in a population. *IEEE Journal on Selected Areas in Information Theory*, 2024.
- [Stein *et al.*, 2018] Mart L Stein, Vincent Buskens, Peter GM Van Der Heijden, Jim E Van Steenberghe, Albert Wong, Martin CJ Bootsma, and Mirjam EE Kretzschmar. A stochastic simulation model to study respondent-driven recruitment. *PLOS One*, 13(11):e0207507, 2018.
- [Tan *et al.*, 2025] Chee Wei Tan, Pei-Duo Yu, Siya Chen, and H Vincent Poor. DeepTrace: Learning to optimize contact tracing in epidemic networks with graph neural networks. *IEEE Transactions on Signal and Information Processing over Networks*, 11:97–113, 2025.
- [Ten Brink *et al.*, 2025] Debra Ten Brink, Rowan Martin-Hughes, Anna L Bowering, Nisaa Wulan, Kelvin Burke, Tom Tidhar, Shona Dalal, and Nick Scott. Impact of an international HIV funding crisis on HIV infections and mortality in low-income and middle-income countries: a modelling study. *The Lancet HIV*, 12(5):e346–e354, 2025.
- [Tran *et al.*, 2024] Cong Tran, Won-Yong Shin, and Andreas Spitz. IM-META: Influence maximization using node metadata in networks with unknown topology. *IEEE Transactions on Network Science and Engineering*, 11(3):3148–3160, 2024.
- [Tsui *et al.*, 2024] Joseph L.-H. Tsui, Mengyan Zhang, Prathyush Sambaturu, Simon Busch-Moreno, Marc A. Suchard, Oliver G. Pybus, Seth Flaxman, Elizaveta Semenova, and Moritz U. G. Kraemer. Optimal disease surveillance with graph-based Active Learning. *medRxiv*, 2024.
- [UNAIDS, 2022] UNAIDS. Political Declaration on HIV and AIDS: Summary of 10 Targets, 2022. Available at <https://www.unaids.org/en/resources/documents/2022/political-declaration-summary-10-targets>. Accessed June 12th, 2026.
- [UNAIDS, 2024] UNAIDS. 2024 global AIDS report — The Urgency of Now: AIDS at a Crossroads, 2024. Available at www.unaids.org/en/resources/documents/2024/global-aids-update-2024. Accessed June 12th, 2026.
- [United Nations, nd] United Nations. Goal 3: Good Health and Well-being – Targets and Indicators, n.d. Available at https://sdgs.un.org/goals/goal3#targets_and_indicators. Accessed June 12th, 2026.
- [Wang *et al.*, 2013] Xuezhi Wang, Roman Garnett, and Jeff Schneider. Active search on graphs. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 731–738, 2013.
- [Wilder *et al.*, 2019] Bryan Wilder, Eric Ewing, Bistra Dilkina, and Milind Tambe. End to end learning and optimization on graphs. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.
- [Wilder *et al.*, 2021] Bryan Wilder, Laura Onasch-Vera, Graham Diguiseppi, Robin Petering, Chyna Hill, Amulya Yadav, Eric Rice, and Milind Tambe. Clinical trial of an AI-augmented intervention for HIV prevention in youth experiencing homelessness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14948–14956, 2021.
- [World Health Organization, 2024] World Health Organization. Consolidated guidelines on differentiated HIV testing services, 2024. Available at <https://www.who.int/publications/i/item/9789240096394>. Accessed June 12th, 2026.
- [Yuan and Tang, 2017] Jing Yuan and Shaojie Tang. No time to observe: Adaptive influence maximization with partial feedback. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, 2017.
- [Zhang and Heath, 2023] Jingyi Zhang and Lenwood S Heath. Adaptive group testing strategy for infectious diseases using social contact graph partitions. *Scientific Reports*, 13(1):12102, 2023.