Adaptive Frontier Exploration on Graphs with Applications to Network-Based Disease Testing

Davin Choo[*]
Harvard University
davinchoo@seas.harvard.edu

Yuqi Pan[*]
Harvard University
yuqipan@g.harvard.edu

Tonghan Wang
Harvard University
twang1@g.harvard.edu

Milind Tambe
Harvard University
tambe@seas.harvard.edu

Alastair van Heerden
University of Witwatersrand
Wits Health Consortium
alastair.vanheerden@wits.ac.za

Cheryl Johnson
World Health Organization
johnsonc@who.int

Abstract

We study a sequential decision-making problem on a n-node graph $\mathcal{G}$ where each node has an unknown label from a finite set $\mathbf{\Omega}$, drawn from a joint distribution $\mathcal{P}$ that is Markov with respect to $\mathcal{G}$. At each step, selecting a node reveals its label and yields a label-dependent reward. The goal is to adaptively choose nodes to maximize expected accumulated discounted rewards. We impose a frontier exploration constraint, where actions are limited to neighbors of previously selected nodes, reflecting practical constraints in settings such as contact tracing and robotic exploration. We design a Gittins index-based policy that applies to general graphs and is provably optimal when $\mathcal{G}$ is a forest. Our implementation runs in $\mathcal{O}(n^2 \cdot |\mathbf{\Omega}|^2)$ time while using $\mathcal{O}(n \cdot |\mathbf{\Omega}|^2)$ oracle calls to $\mathcal{P}$ and $\mathcal{O}(n^2 \cdot |\mathbf{\Omega}|)$ space. Experiments on synthetic and real-world graphs show that our method consistently outperforms natural baselines, including in non-tree, budget-limited, and undiscounted settings. For example, in HIV testing simulations on real-world sexual interaction networks, our policy detects nearly all positive cases with only half the population tested, substantially outperforming other baselines.[2]

1 Introduction

We study a sequential decision-making problem on a graph $\mathcal{G}$, where each node has an unknown discrete label from $\mathbf{\Omega}$. The labels follow a joint distribution $\mathcal{P}$, which we assume is specified by a Markov random field (MRF) defined over $\mathcal{G}$ [KF09]. When we act on a node, its label is revealed and we receive a label-dependent reward. Crucially, the entire process is *history-sensitive*: label realizations are stochastic and depend on previously observed labels, a setting that naturally arises in Bayesian adaptive planning [GK11]. In this paper, we study a setting where actions are subject to a *frontier exploration constraint*: the first node in each connected component is selected based on a pre-defined priority rule, and subsequent actions are restricted to neighbors of previously selected nodes. This constraint reflects realistic settings where local neighborhood information becomes

[*]Equal contribution
[2]Please see https://arxiv.org/abs/2505.21671 for the full version.

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

accessible only through exploration, as in active search on graphs [GKX+12], robotic exploration [KK14], and cybersecurity applications [LCH+25]. The objective is then to maximize the expected accumulated discounted reward over time by sequentially selecting nodes to act upon.

Definition 1 (The *Adaptive Frontier Exploration on Graphs* (AFEG) problem)**.** An AFEG instance is defined by a triple $(\mathcal{G}, \mathcal{P}, \beta)$, where $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ is a graph, $\mathcal{P}$ is a joint distribution over node labels that is Markov with respect to $\mathcal{G}$, and $\beta \in (0, 1)$ is a discount factor. The process unfolds over $n = |\mathbf{X}|$ time steps, with the state $\mathcal{S}_t$ at time t consisting of the current frontier and the revealed labels. Acting on a frontier node reveals its label, grants a label-dependent reward, and updates beliefs about other nodes via Bayesian inference under $\mathcal{P}$. The goal is to compute a policy π that maps each state to a frontier node, maximizing the expected total discounted reward:

$$\pi^* = \arg\max_{\pi} \sum_{t=1}^{n} \beta^{t-1} \sum_{v \in \mathbf{\Omega}} \mathcal{P}(X_{\pi(\mathcal{S}_{t-1})} = v \mid \mathcal{S}_{t-1}) \cdot r(X_{\pi(\mathcal{S}_{t-1})}, v),$$

where $X_{\pi(\mathcal{S}_{t-1})}$ is the node selected by policy π at time t, and $r(\cdot, \cdot)$ is the label-dependent reward.

While the optimal policy can be computed via dynamic programming, it is intractable for general graphs due to the exponential state space. A natural strategy is to leverage adaptive submodularity, which guarantees that greedy policies achieve a $(1 - 1/e)$-approximation [GK11]. Unfortunately, the objective in AFEG is not adaptively submodular in general: for instance, in disease detection, observing an infected neighbor can *increase* the marginal benefit of testing a node, violating the diminishing returns property of adaptive submodularity.

Our problem is closely related to the setting of active search on graphs [GKX+12, WGS13, JMC+17, JMA+18], where the goal is to identify as many target-labeled nodes as possible under a fixed budget, without exploration constraints. Since exact optimization is intractable, these works focused on practical heuristics such as search space pruning. AFEG differs in two key respects: (i) we impose a frontier constraint, and (ii) we consider an infinite-horizon objective with discounting, rather than a fixed budget. These differences are not merely technical but they enable provable optimality in meaningful special cases, particularly when the input graph $\mathcal{G}$ is a forest. Forest structures naturally arise in several relevant domains, including transmission trees in contact tracing [KFH06] and recruitment trees in respondent-driven sampling [Hec97, GS09]. Moreover, algorithms with guarantees on forests can be efficiently applied to sparse real-world interaction graphs, such as sexual contact graphs, which tend to be tree-like in practice.

1.1 Motivating application: network-based disease testing

A key motivating example of AFEG is network-based infectious disease testing where the goal is to identify infected individuals as early as possible. In particular, we focus on diseases that are transmitted through person-to-person contact[3], e.g., sex, exposure of blood through injecting drug use, or birth, where interaction information can be collected through interviews. In this context, frontier testing is both natural and operationally motivated: test outcomes substantially alter beliefs about neighboring individuals, making sequential expansion along the frontier an efficient strategy.

Public health motivation. The 95-95-95 HIV[4] targets proposed by UNAIDS [UNA22] aim for 95% of people with HIV to know their status, 95% of those to receive treatment, and 95% of treated individuals to achieve viral suppression — aligned with UN Sustainable Development Goal 3.3 [Nat]. Yet, the 2024 UNAIDS report [UNA24] reveals that the "first 95" remains the most elusive, with roughly one in seven people living with HIV still undiagnosed, and there continues to be 1.3 million new infections every year. Studies have shown that virally suppressed individuals will not infect others [CCM+11, RCB+16, BPP+18], leading to the U=U (undetectable = untransmittable) campaign [oAD19, OG20]. Thus, the faster we can detect infected individuals, the faster they can be enrolled onto treatment and limit the spread of the disease. To address this gap, the WHO recommends network-based testing strategies to reach underserved populations [Org24a]. These include partners and biological children of people with HIV, as well as those with high ongoing HIV risk. Network-based interventions have shown effectiveness in South Africa [JPC+19] and

[3]This is in contrast to illnesses like flu where transmission can occur to a room full of strangers.

[4]The human immunodeficiency virus (HIV) attacks the immune system and can lead to AIDS. It remains a major global health issue, having claimed over 42 million lives to date [Org24b].

have also been explored for other infectious diseases beyond HIV [JSK+17, MWBDM+25]; see also [CLJ+24] for a WHO-commissioned systemic review on social network-based HIV testing.

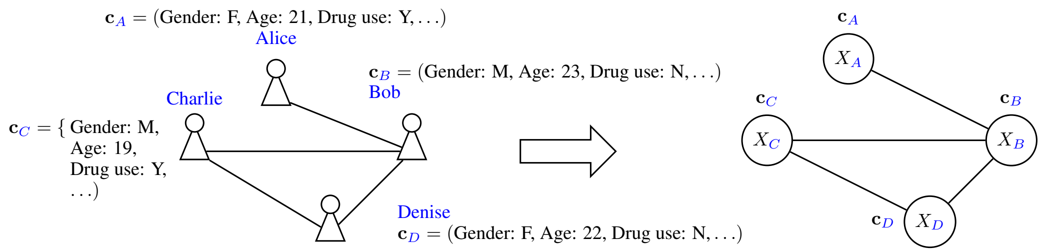

Figure 1: Illustration of how a real-world transmission graph (left) can be framed as an AFEG instance (right). Here, the joint distribution $\mathcal{P}$ over the labels $X_A, X_B, X_C, X_D \in \{+, -\}$ may depend on the covariates $\mathbf{c}_A, \mathbf{c}_B, \mathbf{c}_C, \mathbf{c}_D \in \mathbb{R}^d$ and underlying interaction graph structure.

Fig. 1 illustrates how we can model the network-based disease testing problem into a AFEG instance. Firstly, we use the network $\mathcal{G}$ as is, where nodes represent individuals and edges represent sexual interactions. Each node has a binary infection status (infected or not) that is drawn from some underlying joint distribution $\mathcal{P}$ on $\mathbf{X}$ over the labels $\mathbf{\Omega} = \{+, -\}$, where $\mathcal{P}$ may depend on the individual covariates and graph structure. The reward for testing individual X and revealing status $b \in \{0, 1\}$ is then $r(X, b) = b$. See Fig. 1 for an illustration. The goal is of trying to identify infected individuals as early as possible is implicitly enforced by the presence of *any* discount factor $\beta < 1$. Importantly, discounting reflects both practical constraints – such as sudden funding cuts [UNA25] – and clinical importance of early diagnosis, which improves patient outcomes and limits transmission [CCM+11]. See also [RN21] for other natural justifications for using discount factors β in modeling long-term policy rewards. While transmission graphs of sexually transmitted diseases are not truly forests and may have high-degree nodes (e.g., sex workers), empirical studies have also shown that such transmission graphs are often sparse and exhibit tree-like structure [BMS04, YJM+13, WKPF+17]. Finally, to apply the infinite horizon framework of AFEG in our finite testing setting, we give zero subsequent rewards after every individual has been already tested.

2 Our contributions

We state and discuss our contributions here.[5]

2.1 Gittins index-based policy for AFEG and new results for branching bandits

We show that when $\mathcal{G}$ is a forest, AFEG reduces exactly to the branching bandit framework [Wei88, Tsi94, KO03], for which Gittins index policies are known to be optimal [KO03]; see Fig. 2.

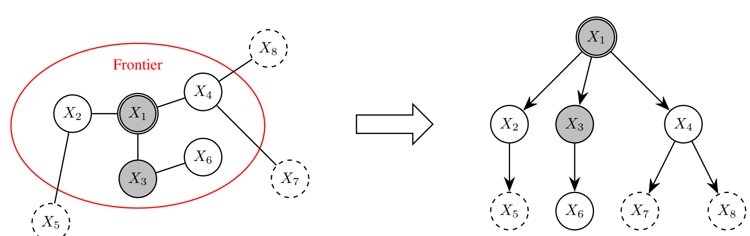

Figure 2: Reduction to a branching bandit on 8 nodes with root X_1. After acting on $\{X_1, X_3\}$, the frontier is $\{X_2, X_4, X_6\}$. Note that we have $\mathcal{P}(x_2 \mid x_1, x_3) = \mathcal{P}(x_2 \mid x_1)$ by the Markov property.

While [KO03] established the existence of an optimal Gittins index policy, they did not characterize the index explicitly nor provide an efficient method for computing it. We provide a novel characterization of Gittins indices for discrete branching bandits using piecewise linear functions, and develop a practical implementation that runs in $\mathcal{O}(n^2 \cdot |\mathbf{\Omega}|^2)$ time while using $O(n \cdot |\mathbf{\Omega}|^2)$ oracle calls to $\mathcal{P}$

[5] Please see https://arxiv.org/abs/2505.21671 for the full version.

and $O(n^2 \cdot |\mathbf{\Omega}|)$ space. Our policy also works for general non-tree AFEG instances, but without optimality guarantees, by first projecting $\mathcal{G}$ into a forest using breadth-first search from the root nodes of each connected component. Despite this, our proposed method still demonstrates strong empirical performance in experimental evaluations. Full derivation details are given in our full version.

2.2 Formalizing network-based disease testing as an AFEG instance

As shown in Section 1.1, network-based infectious disease testing can be cast as an instance of AFEG. To our knowledge, this is the first formal framework to model frontier-based testing as sequential decision-making on a probabilistic graph model for principled exploitation of network effects in diseases such as HIV. In the appendix of our full version, we propose a method to learn parameters from past disease data to define a joint distribution $\mathcal{P}$ on new interaction networks so as to define new AFEG instances.

2.3 Empirical evaluation

We benchmark our proposed GITTINS policy against several natural baselines — RANDOM, GREEDY, DQN, and OPTIMAL — on both synthetic and real-world graphs to evaluate performance on AFEG. To reflect the network-based disease testing application discussed in Section 1.1, we consider binary node labels, and define the immediate reward to be 1 if and only if the revealed label is positive. As such, it is natural to define the first node in every connected component as the node with the highest marginal probability of being positive amongst all nodes in that connected component.

Benchmarked policies. Given a problem instance $(\mathcal{G}, \mathcal{P}, \beta)$, a state in AFEG consists of the current set of frontier nodes and the revealed labels of previously tested nodes.

- RANDOM: Selects a random node from the frontier without using any state information.
- GREEDY: Selects the frontier node with the highest posterior probability of being positive, conditioned on the currently observed labels.
- DQN: Implements a deep Q-network baseline [MKS+15], using the NNConv architecture from PyTorch Geometric [FL19]. This model applies a message-passing GNN with edge-conditioned weights [GSR+17] to capture graph structure and node covariates.
- OPTIMAL: Computes the action that maximizes the expected total discounted reward for each possible state via brute-force dynamic programming. This method is tractable only on small graphs due to the combinatorial explosion of the state space.
- GITTINS: Our proposed method, which is provably optimal when $\mathcal{G}$ is a forest.

We present our full experimental details and results in our full version. On synthetic experiments, our Gittins index-based policy performs strongly even in settings where it is not provably optimal, including non-trees and finite-horizon scenarios. It also outperforms other baselines on public-use real-world sex interaction graphs on 5 sexually transmitted diseases (Gonorrhea, Chlamydia, Syphilis, HIV, and Hepatitis) from ICPSR [MR11].

3 Conclusion and discussion

We introduced and studied the adaptive frontier exploration on graphs problem (Definition 1), a framework for sequential decision-making with label-dependent rewards under a frontier exploration constraint. Our Gittins index-based policy is provably optimal on trees, runs in polynomial time, and demonstrates strong empirical performance on general graphs.

Broader impact and fairness. This work is motivated by public health challenges, where limited resources and reduced funding [UNA25] highlight the need for more efficient testing strategies. The AFEG framework supports targeted, adaptive exploration of interaction networks, guided by a joint distribution $\mathcal{P}$ can incorporate domain knowledge. It also enables fairness-aware interventions through reward shaping, allowing practitioners to prioritize specific subpopulations within the same decision-making framework. Our proposed Gittins index-based policy operates within this flexible setup, making it suitable for responsible and context-aware deployment. Additional discussion of limitations and fairness considerations is provided in our full version.

Acknowledgments and Disclosure of Funding

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. Davin would like to thank Bryan Wilder, Amulya Yadav, and Chun Kai Ling for their thought-provoking technical discussions, and Eric Rice, Geoff Garnett, Samuel R. Friedman, and Ashley Buchanan for generously sharing their domain expertise on HIV testing and transmission.

References

[BMS04] Peter S. Bearman, James Moody, and Katherine Stovel. Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks. *American journal of sociology*, 110(1):44–91, 2004.

[BPP+18] Benjamin R Bavinton, Angie N Pinto, Nittaya Phanuphak, Beatriz Grinsztejn, Garrett P Prestage, Iryna B Zablotska-Manos, Fengyi Jin, Christopher K Fairley, Richard Moore, Norman Roth, and et al. Viral suppression and HIV transmission in serodiscordant male couples: an international, prospective, observational, cohort study. *The Lancet HIV*, 5(8):e438–e447, 2018.

[CCM+11] Myron S. Cohen, Ying Q. Chen, Marybeth McCauley, Theresa Gamble, Mina C. Hosseinipour, Nagalingeswaran Kumarasamy, James G. Hakim, Johnstone Kumwenda, Beatriz Grinsztejn, Jose H. S. Pilotto, and et al. Prevention of HIV-1 Infection with Early Antiretroviral Therapy. *New England journal of medicine*, 365(6):493–505, 2011.

[CLJ+24] Annabelle Choong, Yi Ming Lyu, Cheryl C. Johnson, Rachel Baggaley, Magdalena Barr-DiChiara, Muhammad S. Jamil, Nandi L. Siegfried, Christopher K. Fairley, Eric P. F. Chow, Virginia Macdonald, and Jason. Ong. Social network-based approaches to HIV testing: a systematic review and meta-analysis. *Journal of the International AIDS Society (JIAS)*, 27(9):e26353, 2024.

[FL19] Matthias Fey and Jan E. Lenssen. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

[GK11] Daniel Golovin and Andreas Krause. Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization. *Journal of Artificial Intelligence Research (JAIR)*, 42:427–486, 2011.

[GKX+12] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian Optimal Active Search and Surveying. In *International Conference on Machine Learning (ICML)*, page 843–850, 2012.

[GS09] Sharad Goel and Matthew J. Salganik. Respondent-driven sampling as Markov chain Monte Carlo. *Statistics in Medicine*, 28(17):2202–2229, 2009.

[GSR+17] 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, 2017.

[Hec97] Douglas D. Heckathorn. Respondent-Driven Sampling: A New Approach to the Study of Hidden Populations. *Social Problems*, 44(2):174–199, 1997.

[JMA+18] Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic batch active search. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1099–1109, 2018.

[JMC+17] Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. Efficient Nonmyopic Active Search. In *International Conference on Machine Learning (ICML)*, pages 1714–1723, 2017.

[JPC+19] Makhahliso Jubilee, Faith Jiyeong Park, Knowledge Chipango, Kenoakae Pule, Albert Machinda, and Noah Taruberekera. HIV index testing to improve HIV positivity rate and linkage to care and treatment of sexual partners, adolescents and children of PLHIV in Lesotho. *PLoS One*, 14(3):e0212762, 2019.

[JSK+17] 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.

[KF09] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.

[KFH06] Don Klinkenberg, Christophe Fraser, and Hans Heesterbeek. The Effectiveness of Contact Tracing in Emerging Epidemics. *PloS One*, 1(1):e12, 2006.

[KK14] Matan Keidar and Gal A. Kaminka. Efficient frontier detection for robot exploration. *The International Journal of Robotics Research (IJRR)*, 33(2):215–236, 2014.

[KO03] Godfrey Keller and Alison Oldale. Branching bandits: a sequential search process with correlated pay-offs. *Journal of Economic Theory*, 113(2):302–315, 2003.

[LCH+25] Chun Kai Ling, Jakub Cerny, Chin Hui Han, Christian Kroer, and Garud Iyengar. How Deep Is Your Defense-in-Depth? Hardening Cybersecurity Network Control Against Adaptive Attackers. *AAAI 2025 Workshop on Artificial Intelligence for Cyber Security (AICS)*, 2025.

[MKS+15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, and et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[MR11] 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 https://www.icpsr.umich.edu/web/ICPSR/studies/22140/summary.

[MWBDM+25] 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), 2025.

[Nat] United Nations. Goal 3: Good Health and Well-being – Targets and Indicators. Available at https://sdgs.un.org/goals/goal3#targets_and_indicators.

[oAD19] National Institute of Allergy and Infectious Diseases. HIV Undetectable=Untransmittable (U=U), or Treatment as Prevention, 2019. Available at https://www.niaid.nih.gov/diseases-conditions/treatment-prevention.

[OG20] Rose McKeon Olson and Robert Goldstein. U=U: Ending stigma and empowering people living with HIV, 2020. Available at https://www.health.harvard.edu/blog/uu-ending-stigma-and-empowering-people-living-with-hiv-2020042219583.

[Org24a] World Health Organization. *Consolidated guidelines on differentiated HIV testing services*. World Health Organization, 2024. Available at https://www.who.int/publications/i/item/9789240096394.

[Org24b] World Health Organization. HIV and AIDS: Fact Sheet, 2024. Available at https://www.who.int/news-room/fact-sheets/detail/hiv-aids.

[RCB+16] Alison J. Rodger, Valentina Cambiano, Tina Bruun, Pietro Vernazza, Simon Collins, Jan van Lunzen, Giulio Maria Corbelli, Vicente Estrada, and et al. Sexual Activity Without Condoms and Risk of HIV Transmission in Serodifferent Couples When the HIV-Positive Partner Is Using Suppressive Antiretroviral Therapy. *JAMA*, 316(2):171–181, 2016.

[RN21] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 4th edition, 2021.

[Tsi94] John N. Tsitsiklis. A Short Proof of the Gittins Index Theorem. *The Annals of Applied Probability*, pages 194–199, 1994.

[UNA22] 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.

[UNA24] UNAIDS. 2024 global AIDS report — The Urgency of Now: AIDS at a Crossroads, 2024. Available at https://www.unaids.org/en/resources/documents/2024/global-aids-update-2024.

[UNA25] UNAIDS. Impact of US funding cuts on the global AIDS response – 28 March 2025 update, 2025. Accessed: 28 April 2025.

[Wei88] Gideon Weiss. Branching Bandit Processes. *Probability in the Engineering and Informational Sciences*, 2(3):269–278, 1988.

[WGS13] Xuezhi Wang, Roman Garnett, and Jeff Schneider. Active Search on Graphs. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 731–738, 2013.

[WKPF+17] Joel O. Wertheim, Sergei L. Kosakovsky Pond, Lisa A. Forgione, Sanjay R. Mehta, Ben Murrell, Sharmila Shah, Davey M. Smith, Konrad Scheffler, and Lucia V. Torian. Social and Genetic Networks of HIV-1 Transmission in New York City. *PLoS Pathogens*, 13(1):e1006000, 2017.

[YJM+13] A. M. Young, A. B. Jonas, U. L. Mullins, D. S. Halgin, and J. R. Havens. Network Structure and the Risk for HIV Transmission Among Rural Drug Users. *AIDS and Behavior*, 17(7):2341–2351, 2013.