Recommendations on the
Ethical Use of Novel HIV Data & Analytics

Working Group on the Ethical Use of Novel HIV Data and Analytics

May 2024

Funded by the Bill and Melinda Gates Foundation

ethics lab

GEORGETOWN UNIVERSITY
Center for Innovation in Global Health

COOPER/SMITH

BILL & MELINDA GATES foundation

WORKING GROUP MEMBERS

Janet Tatenda Bhila
Zimbabwe Young Positives (ZY+)

Aleny Couto, MD
Ministry of Health, STI and HIV/AIDS, Mozambique

Shona Dalal, PhD
WHO, Dept. of Global HIV, Hepatitis, and STIs

Will Fleisher, PhD
Georgetown University, Center for Digital Ethics

Jen Gennai
Google, Responsible Innovation

Rayid Ghani
Carnegie Mellon University, Machine Learning Dept. and College of Information Systems and Public Policy

Peter Godfrey-Faussett
London School of Hygiene and Tropical Medicine

Melissa Goldstein, JD
George Washington University, Milken Institute School of Public Health

Kathy Hageman, PhD, MPH
CDC, Division of Global HIV, AIDS, TB

Nina Hasen, PhD
Global Health Consultant, ACHIEVE Innovations

Micheal Ighodaro
Prevention Access Campaign, Washington, D.C.

Thoko Kalua, MBBS, MSc
Center for International Health, Education, and Biosecurity (Ciheb), University of Maryland

Jennifer Miller, PhD
Yale School of Medicine

Yogan Pillay, PhD
Department of Health, South Africa

Anton Pozniak, MD
Chelsea and Westminster Hospital, HIV Medicine; London School of Hygiene and Tropical Medicine

Miriam Rabkin, MD, MPH
Columbia University, ICAP Global Health, Systems Strategies

David Ribes, PhD
University of Washington, Dept. of Human Centered Design & Engineering, Data Ecologies Lab

Lisa Singh, PhD
Georgetown University, Massive Data Institute, Dept. of Computer Science

PROJECT TEAM

Charles Holmes, M.D., MPH, Co-PI
Georgetown University, Center for Innovation in Global Health

Maggie Little, Ph.D., Co-PI
Georgetown University, Ethics Lab

Heidi Weimer, J.D., MPH, LL.M, Project Manager
Georgetown University, O'Neill Institute for National & Global Health Law

Dylan Green, MPH, Subject Matter and Technical Expert
Cooper/Smith, Research and Science

Alicia Patterson, Ph.D.
Georgetown University, Ethics Lab

Jonathan Healey
Georgetown University, Ethics Lab

Sydney Luken
Georgetown University, Ethics Lab

INTRODUCTION

Key priorities of HIV programmes include improvements in the ability to identify those most at risk of acquiring HIV, experiencing treatment interruption, or in need of more support to remain on treatment—all in service of providing access to appropriate services. In support of these goals, there is increased interest by HIV researchers and programme innovators in deploying the new capabilities offered by Big Data, including data generated by individuals' use of digital services and devices, such as social media platforms, apps, and mobile phones. In order to leverage the power of Big Data, HIV programmes are using Machine Learning (ML) and Artificial Intelligence (AI) methods to develop models that can better predict which clients or patients might be living with HIV, at risk for treatment interruption, or otherwise underserved. Advances in Natural Language Processing (NLP) models has also led to interest in extracting information about patients' personal life experiences that may be contained as clinician notes in their Electronic Medical Records (EMRs).

While these strategies have strong potential for improving predictive models, they also bring with them critically important ethical questions. These include issues around privacy and justified surveillance, risks to individuals and groups should the data that is harvested and analytically generated become known, and the potential for algorithmic bias. It is thus critical to appreciate the distinct ethical issues posed by these new modes of surveillance and analysis, and consider frameworks for their responsible use.

International guidelines provide important ethical principles relevant to these questions, including guidance on surveillance in public health ethics,[1] the use of AI in health contexts,[2] and data privacy and protection.[3] The purpose of this guidance is to move from broad principles to pragmatically identifying and confronting considerations, complexities, and factors in the context of HIV. Its goals are to 1) help guide researchers and programme innovators in building designs that are ethically, as well as technically, feasible; 2) assist those charged with assessing proposals for funding and implementation to effectively assess the proposed use of Big Data and ML; and 3) identify actions HIV funders and multilateral organizations can take to advance responsible approaches to the use of novel data and ML models in HIV research and programmatic innovation.

This guidance is the product of an 18-member international, interdisciplinary, and intersectoral Working Group, supported by the Bill and Melinda Gates Foundation. The Working Group convened remotely for three workshops. Each workshop extended across three days; with half-day sessions. Members included HIV researchers and programme innovators with experience in novel data and machine learning; technical experts from computer, information, and data sciences; experts in global public health policy and health law; experts in data ethics and ethics of AI; and members from the community of people living with HIV who have been leaders in advocating for their community. Countries represented were Malawi, Mozambique, South Africa, the United Kingdom, the United States, and Zimbabwe. A series of briefings, extensive virtual table-top exercises, break-out groups, and anchoring case studies informed extensive and iterated plenary discussions.

---

[1] World Health Organization, "WHO Guidelines on Ethical Issues in Public Health Surveillance," Geneva: World Health Organization; 2017. License: CC BY-NC-SA 3.0 IGO.

[2] World Health Organization, "Ethics and Governance of Artificial Intelligence for Health: WHO Guidance," Geneva: World Health Organization; 2021. License: CC BY-NC-SA 3.0 IGO

[3] United Nations Development Group, "Data Privacy, Ethics and Protection: Guidance Note on Big Data for Achievement of the 2030 Agenda," resolution 45/95, November 2017, https://unsdg.un.org/resources/data-privacy-ethics-and-protection-guidance-note-big-data-achievement-2030-agenda. Digital Medicine Society's Toolkit for Inclusive Digital Health Measurement Product Development: https://datacc.dimesociety.org/development/

KEY ETHICAL CHALLENGES

Proposals to use data from individuals' use of digital services and devices, or highly personal information from EMRs, and deploy ML methods for their analysis, raise distinctive ethical issues, both in general and in the specific context of HIV. The following are especially critical.

Expansion of Digital Surveillance into Private & Communal Spaces

The ongoing collection of personal data from activities on browsers, platforms, apps, and use of mobile phones can impinge on what have become important sites of personal and social life. These raise important privacy concerns, and can have chilling effects on activities important to democratic activities such as political activism. Information included in EMR clinician notes can contain stigmatized information, such as reports of domestic violence and of suicidal ideation, which also raise privacy concerns, and may have chilling effects on what patients share with their health providers.

Gathering and Generating Sensitive Information

Information gathered from individuals' use of digital services and devices, from web pages accessed to patterns of movement, can include highly personal or potentially compromising information beyond sensitive HIV-related issues. The predictive models based on these data, in turn, generate HIV-related risk scores, and attach them to users' individual accounts. Each of these represents an expansion of informational risk to broad populations.

Risks of De-Anonymization

The use of big data and their associated analytics can substantially increase the ability to infer the identity of individuals in anonymized datasets. Identity can often be recovered when multiple datasets are merged together, as is common practice in big data analytics. The ongoing collection of mobile phone location data also carries especially strong concerns about de-anonymization. Because people's patterns of movement across time are unique, the detailed information on times and locations of the mobile phone's use can translate into the ability to identify the person behind the phone's use.

Fewer Surrounding Protections on Access

Personal data sourced from the use of digital platforms and apps is subject to far less regulation than is health data. Access and analysis of novel data also often involves commercial partners, such as telecommunications, app, or platform companies, which have interests that can be at odds with public health interests, including interests in monetizing shared data, or combining the analytic output with their own, often extensive, data on users.

Potential for Civil Rights Violations

Information on citizens' locations can be used to target them; HIV risk scores attributed to accounts can be used to persecute vulnerable populations. These issues are of particular concern in contexts of high political volatility, and contexts in which certain HIV risk-associated behaviors are criminalized.

Algorithmic Bias & Complexity in ML Models

Data used for training and validation can reflect ethnic, socioeconomic, differential access, or other structural disparities. Using that data to guide future decisions can reinforce those disparities. Many ML techniques, in turn, are sufficiently complex that even those who designed them do not fully know how their predictions are made. These forms of complexity make it more difficult to check for bias in the datasets used to train the models, and impede the ability of communities to appeal consequential decisions made on their basis.

CONSENSUS FACTORS FOR ETHICALLY RESPONSIBLE USE OF BIG DATA & AI IN HIV RESEARCH & PROGRAMMATIC INNOVATION

In navigating the ethical values at stake in the use of novel data and ML methods, research and programmatic innovators, as well as those who serve as funding or adoption gatekeepers, should address the following processes, standards, and expectations.

1. Embed Ethics Into the Concept Development Process

Concept development is a time of option scanning, early scoping, and preliminary feasibility assessments. While the technical, methodological, and institutional feasibility of options are critical, so too is their ethical feasibility.

Options raised should be probed for potential ethical risks from the start. Questions include how intrusive are the contemplated means of monitoring; how consequential would it be if the data collected or analysis generated were accessed; how readily might the data be de-anonymized, and who might have interest in doing so; whether the concept carries substantial risks of public misunderstanding or worsening of trust in public health efforts; whether the country's political context raise concerns for the possibility that the dataset or ML model might be misused against its citizens.

The intersection of technical, methodological, infrastructural, and ethical considerations is both complex and specialized. Project teams are strongly encouraged to convene a joint advisory group that combines technical experts in data science and ML methods, an ethicist fluent in the technical and public health issues, along with community members briefed on the technical and ethical stakes. Early identification of such a team can help at the concept development phase and beyond, as projects inevitably face recalibrations, refinements, and iterations.

2. Incorporate Robust Community Involvement

Community involvement in consideration of public health programmes - always critical - assumes particular importance programmes proposed involve expansion into new and less tested areas of digital surveillance.

Those who would be most affected by initiatives should have opportunities, not simply to comment on whether or not to accept a given proposal, but to influence the shape and direction of a proposal, to give feedback to ideas and offer their own—to ensure that priorities and goals, and ideas for how to meet them, are shaped with and by the community's perspectives.

Proper community involvement thus includes partnership at early stages in the conceptualization and design of a programme; in evaluations of updated information about post-launch efficacy and risks; and in decisions about potential changes, mitigations, or the decision to end a programme.

Project teams are encouraged to consult and adapt existing models that have successfully developed community-led approaches. These include building in community-partnered reassessment after community co-design to ensure the scope and changes will meet the community requirements; and giving grants to community organizations to develop their own proposals, design feedback mechanisms, and community governance.

3. Design Value-Centered ML

Programmatic innovations should aim to use interpretable methods wherever possible. If and where such methods are unviable, researchers should consider sacrificing some accuracy for the sake of increasing interpretability. Proposals for using uninterpretable models should be assessed, not just by the accuracy advantage they may bring, but whether that advantage is worth the added risks of bias and the social opacity they bring.

Those developing ML models should measure bias using fairness measures and reduce any model bias that may arise. Active areas of computer science research are focusing on developing privacy- and fairness-preserving techniques that should be explored. The addition of fairness in optimization metrics should be considered.

Projects should not use ML models that they cannot audit—both for accuracy and the values internal to the system. Audits should include understanding the training dataset used, optimization metrics deployed, and be assessed for alternatives that may represent a better overall balance of values. Audits should also be conducted to make sure performance is consistent over time after models are deployed in real world settings, and that they do not diminish over time to the point of limited utility (in which case ethical risks obviously outweigh benefits).

4. Ensure Robust Data Protection & Governance Structures

Given the sensitivity of the data involved and its linkage to individual accounts, data protection and governance structures must be especially robust.

Data governance boards should be designed in ways that protect the independence of the board and its ability to resist potential pressures for access. This includes attention to appointment policies, accountability audits, and checks and balances in the governing structure. It may also involve sequestering data on private servers, and/or use of neutral data intermediaries that can serve as gatekeepers to the data.

Systems must have robust structures for monitoring, oversight, and capacities to respond with mitigation should unintended consequences or risks emerge.[4] Plans should include regular assessment and re-evaluation points, with processes, responsibilities, and authorities demarcated for response. Policies should think ahead to conditions under which the data structure should be destroyed, and develop mechanisms and protocols for doing so before crises emerge.

Systems should retain data and analytical engines only as long as they are required for immediate programmatic needs. To drive down the ethical risk around the long-term life of these data sets, researchers and programmatic innovators should consider amassing data sets that would be designed to be destroyed after their intended use. Programs should also assess whether it is essential that an analytical engine persists indefinitely, or whether it should be a time-limited tool that self-destructs or is decommissioned by a certain date.

5. Require ERC/IRB Review for Research

Research that harvests personal data from extant sources should be subject to IRB/ERC approval when it will be linked with other, formally gathered research data, or when it will be linked to other sensitive data, including sensitive data that is analytically generated.

Informed consent for research participation requires the provision of clear explanations of what the research involves, its burdens and risks; and its purpose. When describing the study and potential risks, care must be taken to go beyond description of the immediate data that will be harvested, and include inferences that can and will be drawn from it, risk scores that will be linked to their individual records, and the risk of de-anonymization.

---

[4] World Health Organization, "WHO Guidelines on Ethical Issues in Public Health Surveillance;" United Nations Development Group, "Data Privacy, Ethics and Protection: Guidance Note on Big Data for Achievement of the 2030 Agenda."

Assessment of the research proposal should look ahead to the ethical feasibility of the programmatic proposal to which it is in service. If the proposed program—the ongoing harvesting of data it envisions, the the actor(s) that would house or have access to the data, etc.—is ethically unjustified, then the social value of the research, and the significant investment of resources it will require, is likely unjustified in turn.

6. Require Ethical Evaluation of Proposals for Programmatic Innovation

Proposals for programmatic innovations should be subject to ethical evaluation before adoption. Considerations to be addressed include the topics outlined in Part III of this guidance: autonomy burdens and informational risks to both individuals and groups; implications for trust and adoption of needed services; elements of bias, transparency, and accuracy for the ML model.

Attention must also be paid to the political context in which the research and programmatic innovation would take place. Tools justified in one context can be used for oppression and violations of civil rights in another. While those who build analytic tools cannot always foresee or control who may make use of their tools, concrete risks of oppressive uses and civil rights violations must always raise questions about the responsibility of pursuing the tool's development.

Evaluation of proposals should also address the comparative advantage of the proposed approach for achieving expected health outcomes relative to other, potentially less ethically risky or intrusive options. Considerations should be given to the anticipated data protection, governance, and monitoring plans. Particular focus should be given to the choice and design of ownership and access structure; details of any share-back arrangements or vulnerabilities they possess should be disclosed. Consideration should also be given to what community involvement was incorporated into the design and development of the programmatic proposal.

Evaluation of the proposed innovation should involve advisors or consultants knowledgeable about the technical details, programmatic experts, as well as community representatives—all of whom should be treated as equal partners.

7. Ensure Public Disclosure & Accountability

A core tenet of public health ethics is the need for public transparency, justification, and accountability of public health surveillance programmes.[5] [6] [7] This obligation applies both to governmental and large non-governmental entities working for public health goals. The public has a right to know what personal data is being monitored, collection, and analysis of personal data. accountability for effects, harms, and responsible reviews; and to give input into discussions of benefits and harms.

Public health authorities and non-governmental organizations engaged in the use of novel data collection and analytic methods must give meaningful public disclosure of ongoing data harvesting, including what data is being collected, the use to which it will be put, the reasoning and policies that lie behind its decisions, the results of monitoring and mediation activities; and provide meaningful opportunity for public input and objection, with special attention to those most directly impacted by the programmes.

Channels for transparently sharing program details with the public will vary between public health authorities and NGOs. The level of detail provided will vary according to the scale and import of a program.

---

[5] World Health Organization, "WHO Guidelines on Ethical Issues in Public Health Surveillance," Geneva: World Health Organization; 2017. License: CC BY-NC-SA 3.0 IGO.

[6] World Health Organization, "Ethics and Governance of Artificial Intelligence for Health: WHO Guidance,"

[7] United Nations Development Group, "Data Privacy, Ethics and Protection: Guidance Note on Big Data for Achievement of the 2030 Agenda,"2017