

National Artificial Intelligence Advisory Committee

Meeting Minutes

May 2, 2024

The National Artificial Intelligence Advisory Committee (NAIAC) held a hybrid public meeting on Thursday, May 2, 2024. The meeting was recorded and [is available online](#). Members attending virtually are marked with an (*) in the roster below.

NAIAC Members

- David Danks*
- Paula Goldman
- Susan Gonzales*
- Janet Haven
- Dan Ho
- Jon Kleinberg
- Ramayya Krishnan
- Ashley Llorens
- Haniyeh Mahmoudian
- James Manyika*
- Christina Montgomery
- Liz O’Sullivan*
- Fred Oswald
- Trooper Sanders
- Navrina Singh*
- Swami Sivasubramanian
- Reggie Townsend
- Miriam Vogel (Chair)

NAIAC-LE Members

- Anthony Bak*
- Jane Bambauer (Chair)*
- Jennifer Eberhardt*
- Armando Aguilar*

NIST Staff Members

- Melissa Taylor, NAIAC Program Manager
- Cheryl Gendron, Designated Federal Officer (DFO)

Meeting Minutes

Welcome Remarks

- Gendron called the meeting to order at 10:04 AM Eastern Time and confirmed the committee is operating under the Federal Advisory Committee Act and accessible to the public both in-person and via livestream. Time was reserved at the conclusion of the meeting for public questions.
- Taylor thanked members of the public for their participation, encouraged them to visit the NAIAC [website](#) to learn more about the Committee’s work, and highlighted the opportunity for public participation via email at naiac@nist.gov and the NAIAC [mailing list](#). Taylor then introduced Vogel, Chair of NAIAC, for opening remarks.

- Vogel thanked NIST staff for their support and NAIAC members for their work. Vogel applauded NAIAC members for the number of findings and recommendations they developed for Committee review and the many public discussions they have held with experts in health care, workforce, AI safety, and other topics.

NAIAC-LE Subcommittee Update

- Bambauer updated the Committee on changes to the Subcommittee's WG structure. The Subcommittee is transitioning from a fixed, theme-based WG structure to a dynamic, project-based one. When the Subcommittee identifies a new workstream, it will form a new WG dedicated to it, and the WG will dissolve once it has completed or tabled its project.
- Currently, the Subcommittee has five active WGs:
 - **Three original WGs:** (1) Identification and Surveillance, led by Aguilar, (2) Performance and Bias, led by Bambauer, and (3) Process, led by Farhang Heydari.
 - **Two new, project-based WGs:** (1) Accountability AI, led by Eberhardt and (2) Officer Training, led by Benji Hutchinson.

Deliberation on Subcommittee Draft Recommendations

Encourage the Creation of Statewide Repositories for Police Body-worn Camera Footage

Recommendation

Accountability AI WG

- **Overview:** The recommendation calls for the federal government to invest in the development of statewide repositories for body-worn camera (BWC) footage so academic researchers can access and analyze the data. The recommendation argues for a systematic, computationally driven approach to footage review to reveal large-scale patterns in police interaction with the public and calls for measures to ensure footage data security.
- The Committee considered potential benefits of BWC footage analysis. Academic research has found that police BWC are associated with decreased use of force, decreased citizen complaints, and decreased time to resolve complaints. However, the full accountability benefits of BWC depend on analysis of footage.
- Academic analyses of BWC footage have been shown to validate community perceptions of racial disparities in police communication with the public and has led to training that reduced disparities.
- The need to address data privacy risks of BWC footage repositories was discussed. Possible solutions include restricted access protocols, tiered data access levels (transcript/audio/video), and adoption of established privacy measures, such as those established by the National AI Research Resource. An individual opt-out was discussed but would introduce significant selection bias into repositories.

- Aspects of the recommendation can be clarified by foregrounding benefits of researcher access to police BWC footage, calling for the exclusion of footage in ongoing cases or appeals, and identifying mechanisms for community-engaged BWC footage research and for researchers to share datasets associated with their analyses.
- Some additional areas for consideration are incentives for increased research on privacy techniques for video data and participatory processes by which community members might engage in BWC footage repositories’ establishment and use.
- Following deliberation, the recommendation was tabled pending revision.

Require Public Summary Reporting on Use of High-Risk AI

Recommendation

Process WG

- **Overview:** The recommendation calls for the Office of Management and Budget (OMB) or another Executive Branch entity to require law enforcement agencies to annually publish summary usage reports for safety- or rights-impacting AI (as defined by OMB), for inclusion in the agency’s AI Use Case Inventory. The recommendation calls for summary statistics for which vendors can automate data collection and reporting and for determining which statistics are appropriate for different AI tools and use cases.
- After a brief discussion, Llorens motioned a vote on the recommendation, and Townsend seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>18</u>	<u>5</u>
Vogel	Ballantyne
Manyika	Clark
Danks	Espinel
Goldman	Howard
Gonzales	Strier
Haven	
Ho	
Kleinberg	
Krishnan	
Llorens	
Mahmoudian	
Montgomery	
O’Sullivan	
Oswald	
Sanders	
Singh	

Sivasubramanian
Townsend

Field Testing of Law Enforcement AI Tools

Findings and Recommendations

Performance and Bias WG

- The recommendation was tabled due to scheduling constraints.

Overview of NAIAC Department of Commerce Ethics Training

- Zach Prager, an attorney in the Ethics Law & Program Office at the Department of Commerce, provided a brief overview of the ethics training and conflict of interest guidance NAIAC members have received both individually and as a group.
- The presentation outlined the recusal requirement for a conflict of interest. Prager noted that NAIAC members are required to recuse themselves if and only if work on a matter meets *all* criteria outlined in the requirement. However, NAIAC members may choose not to participate in work, deliberation, or voting that does not require recusal if they seek to avoid even *appearing* to lose impartiality.

Deliberation on Committee Draft Deliverables

Enhancing AI's Positive Impact on Science and Medicine

Findings

AI Futures WG

- **Overview:** These findings on the use of AI to promote scientific and medical discovery synthesize expert testimony provided at a public briefing hosted by the WG in April 2024. A recording and summary of the briefing will be accessible at <https://ai.gov/naiac>.
- Based on the deliberation, the findings will be edited to (1) more clearly signal that the findings synthesize the views of invited expert panelists, (2) bolster claims by citing research, and (3) expand on the discussion of a systems-level approach to AI.
- Committee members called for future work that addresses the sociotechnical considerations inherent in scientific advances, such as individual data sovereignty and systemic barriers that hinder equal access to the benefits of scientific progress.
- Future deliverables summarizing expert testimony at briefings and workshops might be labelled “proceedings” rather than “findings,” adopting National Academy of Sciences terminology for documents that capture the perspectives of individuals rather than validated research results.

- Ho motioned a vote on the findings with the proposed edits, and Sanders seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>18</u>	<u>5</u>
Vogel	Ballantyne
Manyika	Clark
Danks	Espinel
Goldman	Howard
Gonzales	Strier
Haven	
Ho	
Kleinberg	
Krishnan	
Llorens	
Mahmoudian	
Montgomery	
O'Sullivan	
Oswald	
Sanders	
Singh	
Sivasubramanian	
Townsend	

Towards Standards for Data Transparency for AI Models

Proceeding

Rights, Trust, and Safety WG

- **Overview:** This proceeding synthesizes expert testimony provided at a February 2024 public briefing on the benefits and risks of establishing baseline standards for data transparency for model creators. A recording and summary of the briefing will be accessible at <https://ai.gov/naiac>.
- Based on the deliberation, the proceeding will be edited to (1) emphasize the diverse set of stakeholders to be included in the development of data transparency standards and (2) reference a NAIAC [recommendation](#) on AI literacy and underscore the role of AI literacy in meaningful stakeholder engagement.
- Committee members identified ways the WG could continue addressing data transparency and provenance in future work, such as by assessing the benefits and shortcomings of current data transparency measures (e.g., data cards) and elaborating on the connection between data transparency and model interpretability.

- Ho motioned a vote on the proceeding with the proposed edits, and Krishnan seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>17</u>	<u>6</u>
Vogel	Ballantyne
Manyika	Clark
Danks	Espinel
Goldman	Gonzales
Haven	Howard
Ho	Strier
Kleinberg	
Krishnan	
Llorens	
Mahmoudian	
Montgomery	
O'Sullivan	
Oswald	
Sanders	
Singh	
Sivasubramanian	
Townsend	

Harnessing AI for Scientific Progress

Recommendations

AI Futures WG

- **Overview:** The recommendation calls for sustained government funding and resources for AI in science, as well as support for AI education and training in scientific communities, and an assessment of infrastructural investments needed to leverage AI for scientific discovery.
- Based on the deliberation, the recommendations will be edited to (1) emphasize human effort by changing “AI-led scientific discovery” to “AI-enabled scientific discovery, (2) emphasize the need to pursue science in the public as well as commercial interest, (3) discuss data availability and access in the recommendation on infrastructure, and (4) call for multidisciplinary collaboration in the recommendation on AI training and education.
- The Committee suggested developing another recommendation focusing on a few areas of scientific inquiry in which AI might facilitate significant progress; this recommendation could identify specific infrastructure and resources that could accelerate progress on those topics.

- Ho motioned a vote on the recommendation with the proposed edits, and Sivasubramanian and Goldman both seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>17</u>	<u>6</u>
Vogel	Ballantyne
Manyika	Clark
Danks	Espinel
Goldman	Gonzales
Haven	Howard
Ho	Strier
Kleinberg	
Krishnan	
Llorens	
Mahmoudian	
Montgomery	
O’Sullivan	
Oswald	
Sanders	
Singh	
Sivasubramanian	
Townsend	

Provide Authority and Resources to Promote Responsible Procurement Innovation for AI at Government Agencies

Recommendation

Procurement WG

- **Overview:** The recommendation encourages use of flexible procurement authorities beyond the Federal Acquisition Regulations (FAR) by suggesting that three or more civilian federal agencies be provided Other Transaction Authority (OTA) on the condition that they develop, document, and disseminate procurement best practices through a procurement or acquisition innovation lab.
- Goldman, Llorens, Mahmoudian, Manyika, Montgomery, O’Sullivan, Singh, and Sivasubramanian declined to participate in the deliberation and vote on the recommendation.
- Based on the deliberation, the recommendation will be edited to (1) include in the synopsis the expectation that agencies participating in the pilot will share best practices and (2) highlight the potential to promote equitable government procurement through increased use of flexible procurement authorities. OTA is associated with a greater diversity of

awardees, including non-incumbents and small enterprises, whereas standard procurement practices can be opaque and prolonged, favoring larger, more established vendors. The recommendation emphasizes, however, that use of OTA should be paired with oversight to ensure the development of appropriate norms and guardrails.

- Krishnan motioned a vote on the recommendation with the proposed edits and Townsend seconded the motion. In the presence of quorum, Members of the Committee were polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>9</u>	<u>6</u>
Vogel	Ballantyne
Danks	Clark
Haven	Espinel
Ho	Gonzales
Kleinberg	Howard
Krishnan	Strier
Oswald	
Sanders	
Townsend	

Declined to Participate: 8

Data Challenges and Privacy Protections for Safeguarding Civil Rights in Government

Recommendation

Rights, Trust, and Safety WG

- **Overview:** The recommendation calls for federal agencies to address the tension between federal mandates to conduct demographic disparity assessments and federal laws codifying data minimization principles that limit the collection of demographic data and shares various approaches for doing so. The recommendation also urges agencies to consider and mitigate risks (e.g., data misuse, privacy concerns) entailed in demographic data collection.
- A public comment on the draft recommendation [\[INSERT LINK\]](#) was submitted by Miranda Bogen on behalf of the Center for Democracy and Technology (CDT). The comment applauds the recommendation and notes that CDT will soon release a report that outlines methods of managing the risks of demographic data collection.
- Based on the deliberation, the recommendation will be edited to (1) clarify that it focuses on privacy risks for tabular data rather than other data modalities (e.g., text), (2) highlight the need for effective and verifiable approaches to data collection challenges, (3) reference the principle of least privilege (which calls for computing system users to have the minimal access needed for an assigned task, (4) incorporate into the recommendation text some of

the language in the discussion section following it, and (5) change “mechanisms” for addressing data collection challenges to “approaches” or “methods.”

- Montgomery motioned a vote on the recommendation with the proposed edits and Krishnan seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>13</u>	<u>10</u>
Vogel	Ballantyne
Manyika	Clark
Haven	Danks
Ho	Espinel
Krishnan	Goldman
Mahmoudian	Gonzales
Montgomery	Howard
O’Sullivan	Kleinberg
Oswald	Llorens
Sanders	Strier
Singh	
Sivasubramanian	
Townsend	

AI Safety

Findings and Recommendations

Rights, Trust, and Safety WG

- **Overview:** The findings derive from two multidisciplinary, multisectoral WG panels that addressed both the conceptual definition of AI safety and methodologies to evaluate the safety of AI systems. Recordings and summaries of these panels will be accessible at <https://ai.gov/naiac>. The recommendations call for the U.S. AI Safety Institute (AIS) to address a range of dimensions of AI safety rather than focus narrowly on one and for the federal government to provide funding and resources to advance the science of AI safety.
- Krishnan motioned a vote on the findings and recommendations and Montgomery seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>13</u>	<u>10</u>
Vogel	Ballantyne
Manyika	Clark
Haven	Danks
Ho	

Krishnan	Espinel
Mahmoudian	Goldman
Montgomery	Gonzales
O’Sullivan	Howard
Oswald	Kleinberg
Sanders	Llorens
Singh	Strier
Sivasubramanian	
Townsend	

Year-Two Insights Report

NAIAC

- **Overview:** The NAIAC *Year-Two Insights Report* contextualizes the Committee’s efforts over the preceding year and compiles the work products it has released on a nearly monthly cadence since the publication of its *Year One Report* and lists Biden Administration and federal agency actions that align with NAIAC guidance. The Report also outlines NAIAC WG agendas for the remainder of 2024.
- The Committee suggested returning to an earlier discussion of potential moonshot projects to address in its upcoming work.
- Townsend motioned a vote on the *Year-Two Insights Report* and Oswald and Mahmoudian seconded the motion. In the presence of quorum, each Member of the Committee was polled and decided by majority vote to advance the recommendation.

Approve	Not Present
<u>13</u>	<u>10</u>
Vogel	Ballantyne
Manyika	Clark
Haven	Danks
Ho	Espinel
Krishnan	Goldman
Mahmoudian	Gonzales
Montgomery	Howard
O’Sullivan	Kleinberg
Oswald	Llorens
Sanders	Strier
Singh	
Sivasubramanian	
Townsend	

Briefing: UN High-Level Advisory Body on AI

- Presenters from the UN High-Level Advisory Body on AI were invited to give prepared remarks on its *Governing AI for Humanity* interim report.
 - **Amandeep Singh Gill**, UN Secretary-General's Envoy on Technology, explained the Advisory Body's goal of promoting internationally coordinated AI governance to solve global problems (e.g., climate change) and ensure consistent enforcement of standards.
 - **Alondra Nelson**, Harold F. Linder Chair in the School of Social Science, Institute for Advanced Study, highlighted the efficient work of the Advisory Body, and identified the five guiding principles and seven AI governance functions delineated in the interim report. Nelson emphasized that the goal of international governance is to *support* rather than *supplant* national and industry regulation.
 - **James Manyika**, Co-Chair, UN High-Level Advisory Body on AI; Senior Vice President, Google-Alphabet, outlined the Advisory Body's proposed steps over the next 12 to 18 months to meet the functional gaps identified in the interim report (e.g., establishing a forum for international AI standards exchange to promote global interoperability of these standards).

- Singh thanked the presenters for their insights and invited members to ask follow-up questions. Members invited speakers to discuss the following items. The discussion and presentation may be viewed in full [in the accompanying recording of the meeting](#):
 - How NAIAC and the Advisory Body could collaborate.
 - How international political differences might impact the desire for scientific consensus around AI.
 - How the Advisory Body discerns whether topics require global or localized governance.
 - Whether the Advisory Body is considering what kind of infrastructure is required to promote AI adoption globally.
 - How the Advisory Body's inclusive international dialogue regarding AI governance might complement its support for harmonizing AI standards.
 - To what extent the Advisory Body will consider the environmental impacts of AI development and use.
 - How to develop a compelling and ambitious vision of the future of AI to share with the public.

Public Comments

- One public comment was received prior to the meeting and was shared with the Committee and the public during the meeting.
- Vogel solicited comments from members of the public attending the meeting. None were forthcoming.

Closing Remarks

- Vogel expressed enthusiasm as the Committee begins its third year of work, thanked Committee members for sharing their expertise, effort, and time. Vogel thanked all attendees for their participation.
- Members of the public are encouraged to share comments and questions with NAIAC and NAIAC-LE by emailing naiac@nist.gov. They can visit ai.gov/naiac to subscribe for Committee updates. A recording and summary of the meeting will also be posted on ai.gov/naiac.
- Gendron adjourned the meeting at 3:51 PM Eastern Time.

National Artificial Intelligence Advisory Committee

Public Comments

May 2, 2024

The National Artificial Intelligence Advisory Committee (NAIAC) held a virtual public meeting on Thursday, May 2, 2024. The meeting was recorded and is available online. The following pages contain public comments received related to this meeting.

From: [Miranda Bogen](#)
To: [naiac](#)
Cc: [Samir Jain](#)
Subject: May 2, 2024, NAIAC Public Meeting
Date: Tuesday, April 30, 2024 12:49:09 PM
Attachments: [EXECUTIVE SUMMARY - Demographic Data for Fairness Report \(Forthcoming\).pdf](#)

Members of the National Artificial Intelligence Advisory Committee,

Thank you for the opportunity to provide public comments on the NAIAC's upcoming public meeting and the recommendations to be discussed. As the Director of the Center for Democracy & Technology's AI Governance Lab, I was pleased to serve as a panelist at NAIAC's March 5, 2024 panel on AI Safety and I write today with comments regarding the recommendation on **Data Challenges and Privacy Protections for Safeguarding Civil Rights in Government** that is to be discussed at this week's public meeting.

CDT is a nonprofit 501(c)(3) organization that works to advance civil rights and civil liberties in the digital age. Among our priorities, CDT advocates for the responsible and equitable use of technology by government agencies, and promotes the adoption of robust, technically-informed solutions for the effective regulation and governance of AI systems — including approaches related to advancing work related to fairness, bias, and nondiscrimination.

We commend the Administration's efforts to ensure that AI systems are consistent with civil rights and antidiscrimination laws and its emphasis on proactive assessments to detect algorithmic bias. We also appreciate the NAIAC calling attention to challenges that government agencies may face when assessing AI-powered systems for demographic disparities, especially when demographic data may not be reliably available. NAIAC's recommendation acknowledges that federal agencies should proactively explore methods that may facilitate proactive assessment of disparities, while being attentive to privacy and other risks related to broadening data collection. To inform the continued exploration and implementation of this recommendation, please find attached to this email an executive summary of CDT's forthcoming report, **Navigating Demographic Measurement for Fairness and Equity**. The full report will be published in May 2024, and will discuss at length an array of measurement methods that practitioners have used to navigate this challenge, detail a variety of institutional safeguards that can be combined to guard against potential misuse and privacy harms, and make specific recommendations related to navigating this important issue.

We look forward to sharing the final report and engaging with interested stakeholders in the federal government on this matter.

Respectfully submitted,

Miranda Bogen | Director, AI Governance Lab
Center for Democracy & Technology | [cdt.org](#)
E: mbogen@cdt.org | **D:** 202.407.8827

EXECUTIVE SUMMARY: Navigating Demographic Measurement for Fairness and Equity

Center for Democracy & Technology

Forthcoming, May 2024

Governments and policymakers increasingly expect practitioners developing and using AI systems in both consumer and public sector settings to proactively identify and address bias or discrimination that those AI systems may reflect or amplify. Central to this effort is the complex and sensitive task of obtaining demographic data to measure fairness and bias within and surrounding these systems. This report provides methodologies, guidance, and case studies for those undertaking fairness and equity assessments — from approaches that involve more direct access to data to ones that don't expand data collection.

Practitioners are guided through the first phases of demographic *measurement* efforts, including determining the relevant lens of analysis, selecting what demographic characteristics to consider, and navigating how to hone in on relevant sub-communities. The report then delves into several approaches to uncover demographic patterns:

Approaches for Measuring Demographic Characteristics for Fairness Measurement
Measuring disparities related to real people
Collection: Directly asking individuals to self-report their demographic information
Observation and inference: Assigning perceived demographic characteristics based on observable features or predicting them using statistical methods or machine learning
Proxies and surrogate characteristics: Using signals that correlate with demographic characteristics to detect patterns or disparities without directly inferring individual demographics
Auxiliary datasets: Combining existing datasets containing demographic information with the data of interest
Cohort discovery: Using pattern detection techniques to identify groups experiencing negative outcomes, without explicitly naming demographic characteristics
Measuring disparities related to representations
Keywords and terms: Manually or automatically constructing lists of words and topics that relate to demographic characteristics and using them to probe systems
Observation and labeling of content: Automatically or manually assigning labels of apparent traits to unidentified people represented in audiovisual or text content
Measuring disparities across contexts
Synthetic data: Using artificially generated data that simulates the structure and distribution of real-world examples or populations
Exploratory analysis: Reviewing a system to reason about how its design, behavior, or other characteristics might lead to negative impact for certain communities
Qualitative research: Directly engaging with people using and affected by systems to capture more nuanced insights about people's lived experience

Given long histories of demographic data being misused to the detriment of vulnerable communities, the report emphasizes that responsibly *handling* demographic data is just as critical as the measurement methods themselves. Many of the approaches described have the potential to be mixed and matched with one another to strengthen protections against potential harms while helping to enable critical work.

Approaches for Handling Demographic Characteristics for Fairness Measurement
Data and infrastructure controls
Pseudonymization: Replaces personal identifiers with placeholder information or otherwise breaks the link between identifying data and other data about an individual
Infrastructure controls: Data and system architecture choices that limit how and by whom data and measurement methods can be accessed or used
Encryption: Scrambling data so it can't be easily deciphered without a mathematical key
Retention and ephemerality: Preventing data from being created or stored longer than needed
Privacy enhancing methods
Aggregation: Combining and summarizing data to reduce identifiability of individual data points
Differential privacy: Adding a specific amount of random statistical noise to datasets to realize particular privacy constraints
Secure multi-party computation: A cryptographic protocol that allows parties to conduct analyses across multiple datasets without sharing data with one another
Procedural controls
User controls: Providing people with the opportunity to decide whether to share their data and to request data be corrected or deleted
Organizational oversight: Processes to review proposed uses of data or measurement methods to ensure they comply with policies and follow necessary procedures
Separate teams: Assigning a specific team to be responsible for oversight and compliance with laws that implicate demographic measurement
Privacy impact assessments: Structured impact assessments to evaluate whether proposed use of data sufficiently mitigate against privacy risks
Transparency: Disclosure about data measurement and handling practices and other relevant information relevant to bias and equity measurement efforts

As policymakers and practitioners build regulatory and technical infrastructure to make progress in this domain, we highlight several recommendations to ensure that the balance remains tipped toward beneficial measurement efforts.

Practitioners should:

- Establish ongoing relationships with communities affected by measurement activities to co-design data collection and handling strategies, discuss potential risks and benefits, and collaboratively define fairness goals.
- Where possible, consider methods that avoid generating or storing sensitive demographic information in a way that can be easily connected to individuals.
- Take great care before using observation and inference methods to identify characteristics, especially those lacking precedent or that resist observation.
- Clearly differentiate between perceived or implied characteristics and actual ones
- Employ a robust combination of approaches to handling data and measurement methods to ensure appropriate use.
- Communicate openly about demographic measurement efforts, as well as how data is handled

Government agencies and regulators should:

- Recognize that a variety of approaches are available for companies to identify and measure disparities, even in the absence of comprehensive demographic data collection.
- Clarify criteria and expectations about acceptable measurement methods when it comes to civil rights compliance, and articulate minimum expectations for how data and methods should be handled.
- Explore how more measurement methods can be used to monitor compliance with Federal civil rights laws, including to conduct investigations and enforcement actions.
- Facilitate collaboration between NGOs, research institutes, and government data agencies to explore creative ways that existing administrative data can be used to conduct measurements in a privacy-respecting manner.
- Encourage continued research to explore how unsupervised, synthetic, privacy-enhancing, and content-related methods can be used to further the detection and remediation of bias and discrimination.

While there is no one-size-fits-all solution, this report makes clear that the lack of obvious access to raw demographic data should not be considered an insurmountable barrier to assessing AI systems for fairness, nor should it provide a blanket justification for widespread or incautious data collection efforts. From exploring privacy-preserving techniques to pursuing measurement of content-related bias when disparities affecting people are hard to measure directly, practitioners have a range of tools at their disposal. As practitioners navigate this complex but important landscape, they should engage early and often with impacted communities, clearly document and communicate their practices, and embed strong technical and institutional safeguards to prevent misuse. Ultimately, responsible demographic measurement demands extraordinary care — for technical choices and their implications, but even more for the people and communities this work must ultimately serve.

"Respecting people's self-determination and autonomy when it comes to sensitive data about who we are is complex and hard to do well. But ignoring that kind of data is also not an option." – [*Who Do They Think You Are? Categories, Classification, and Profiling*](#)