PALANTIR RESPONSE TO THE NATIONAL INSTITUTES OF HEALTH (NIH) ARTIFICIAL INTELLIGENCE /MACHINE LEARNING CONSORTIUM TO ADVANCE HEALTH EQUITY AND RESEARCH DIVERSITY (AIM-AHEAD) PROGRAM RFI

Copyright © 2023 Palantir Technologies Inc. and affiliates ("Palantir"). All rights reserved. The content provided herein is provided for informational purposes only and shall not create a warranty of any kind. Any data contained herein is notional. Actual results and experiences may vary.

AIM-AHEAD RFI Response

Use of AI/ML for health disparities and inequities research

Machine learning is being used at the individual and population levels to support risk stratification, predictive modeling, and decision support for diagnoses and treatment. Because these models are developed from real-world data, algorithms reflect current and historical biases in ways that can harden inequities by race and ethnicity, gender identity, sexual orientation, disability, age, social class, and geography. This response will focus on the data and systems issues NIH needs to address to support AI/ML equity research.

To address equity, NIH requires diverse data sources, covering a wide range of demographic, socioeconomic, and health-related data. These should offer social and economic context as well as psycho-social risk factors such as race, age, gender, economic status, pre-existing medical conditions, housing status, patient experience within and outside of clinical encounters, and data that adequately represent diversity within populations.

Because data outside the clinical encounter are not captured routinely or in a standardized way, there are significant challenges with ingesting and curating these highly complex and sensitive data. A robust informatics approach involves:

1) Understanding the current landscape of reported data, including assessing if reported data represents equity-related efforts and is sufficient for identifying bias in algorithms trained on these data.

2) Identifying and implementing strategies to improve existing data quality and completeness. These should include: a) continuous data quality identification, resolution and validation to address gaps, inaccuracies and bias; b) preemptive identification of data quality issues for rapid triage to the reporting entity; c) statistical imputations measures to fill in the gaps.

3) Tracking the effectiveness of those strategies to improve existing data.

4) Identifying novel sources of data. These sources could include screening for social needs at the clinic level, large public health cohorts such as the National Health and Nutrition Examination Survey (NHANES), and surveillance systems such as the Behavioral Risk Factor Surveillance System (BRFSS), large research cohort studies such as the National Institutes of Health's Precision Medicine Initiative, *All of Us* Research Program and the National COVID Cohort Collaborative (N3C)–which integrate participant-collected data sources from diverse groups. Other potential sources include mobile or remote sensing devices and online geocoded data, which may provide valuable insights on equity.

5) Model evaluation. While there are evolving analytic techniques to assess and account for algorithmic bias that may reflect social bias, researchers developing, testing, and using these models need to be sensitive to unexpected results and recognize bias within the data and within the experts who train the model. Community engagement and including the perspective of those who are underrepresented (those whom bias will most affect) is an essential step.

The informatics challenges of collecting and curating data to support machine learning research on equity requires numerous technical considerations. These range from unifying data assets and improving data quality to supporting data-driven decision-making and expanding public engagement. Use of a commercial platform requiring minimal configuration and tooling for such an effort enables consistent and flexible data incorporation and analysis while tracking progress toward achieving health equity goals.

Infrastructure and resources for AI/ML application and research.

Deploying useful and accurate AI/ML requires a secure, trustworthy data management platform that can provide transparent, generalizable data with a lack of bias. Such a platform is needed not only for developing but also deploying models on production data and continuously improving them based on user feedback.

When organizations fail to achieve consequential AI/ML research, it is often because they deploy models in a piecemeal fashion using legacy or disparate systems that are not designed to reliably manage data or collect feedback necessary for model improvement. Moving AI models from development into production is also one of the biggest pain points in traditional systems, as models are built and trained on static data in siloed environments separate from production data and systems. This disconnect makes it difficult to ensure models will operate effectively or continue to be relevant once deployed in an operational environment. Algorithms must be built with and learn from data that incorporates all relevant context, including the outputs of real-time operational decisions from across an institution. Furthermore, without integrated, clean data, model outputs will be incomplete and inaccurate, as will resulting applications of AI/ML. As such, proper data management is critical for performance modeling and accurate decision-making. These are just some of the considerations that should drive AIM-AHEAD to use a mature, existing data platform that is proven and successful within NIH and the broader research community.

One example of such a platform is the National COVID Cohort Collaborative (N3C) Data Enclave, which leverages Palantir Technologies Inc.‘s (“Palantir”) Palantir Gotham Platform configured with Foundry (“Foundry”) and demonstrates an existing resource that NIH grantees are actively using to implement and manage AI/ML models for research. Foundry is a FedRAMP Moderate-authorized Software-as-a-Service that provides organizations with fundamental data infrastructure capabilities and user interfaces to power any number of AI/ML workflows. The platform combines AI/ML and statistical models into its ecosystem and is backed by key components of data management such as versioning, branching, reproducibility, and lineage—therefore promoting best practices in model training and evaluation workflows. N3C users include an ML Domain Team that collaborates on AI/ML efforts in the N3C Data Enclave, which offers a variety of diverse and high-quality data sources including EHR data from more than 50 Clinical and Translational Science Awards (CTSA) Program sites and other relevant biomedical and clinical data (e.g., social determinants of health, genomics, and imaging data).

To enable high-quality modeling and other data-driven workflows, Foundry also offers critical data quality, provenance, and cleaning tools, including data health checks that protect against

faulty input data that could otherwise impact AI/ML models. Foundry's open architecture is built on distributed microservices and leverages leading open-source tools and open APIs, enabling interoperability with third-party applications and existing technology investments. The platform includes versioning, pipeline orchestration, purpose-based access controls, full auditing, and governance capabilities—helping ensure the N3C Data Enclave delivers high-quality, accurate data feeds to AI/ML models while protecting data privacy.

Partnerships approaches for AI/ML application.

AI/ML applications are most effective, consequential, and operational when approached at the systems level, rather than when treated as discrete/isolated technology components. An integrated systems approach to AI/ML to advance health equity and research diversity outcomes requires an expansive assortment of partnerships spanning disciplines and expertise. But integrating these constituent partnerships is as much an infrastructure challenge as it is an organizational coordination and governance challenge.

On an infrastructure level, a framework for effective partnerships should focus on a common platform to enable various partners and stakeholders to interoperate, build, test, deploy, and oversee the use of AI/ML technologies. One such example would involve a software platform layer with accessible APIs for specialized AI/ML vendors to be able to "plug in" their models or services into a modeling framework that tracks such critical performance, evaluation, and oversight features as model and training data versioning, performance metrics sliced by compulsory and elective population segments, confidence scores, model explainability details, applied confidence scores, etc. The platform could also be established to track certification standards attributable to ecosystem partners based on prior experiences, external evaluations, internal accreditations, and other formalized benchmarks. Overall, this technical infrastructure would be intended to provide a technical scaffolding for implementing partnership standards of due diligence, trust, accountability, reliability, and efficacy.

The infrastructure layer is, however, one part of a broader socio-technical implementation framework for partnerships. Coordination and governance must also be married with the technical platform capabilities to promulgate standards and ensure partners adhere to them. The role of an appropriately empowered Partnerships Governance Group should be *inter alia* to:

- Establish standards for partnership evaluation, accreditation, and ongoing involvement
- Determine baseline evaluation metrics and features to which AI/ML partners should be responsible
- Assess partner fitness for involvement and adjudicate disputes
- Evaluate instances of malfeasance or other violations of partnership community standards and obligations
- Create and promulgate AI/ML Ethics guidelines and assess partners' adherence
- Establish protocols and standards for data collection, management, and sharing
- Administer access to common resources, including the aforementioned infrastructure, as well as data, models, etc.
- Appoint appropriate governance officials (e.g., Data Protection and/or Privacy Officers) to oversee questions of formal legal and ethical responsibility

- Liaise with institutional review boards and other stakeholders with responsibilities for examining and assessing efforts undertaken within the AIM-AHEAD Consortium's auspices
- Support the drafting and promulgation of relevant legibility, accountability, and oversight artifacts, such as privacy impact assessments, model explainability conformity assessments, public transparency measures and initiatives, etc.
- Periodically reassess overarching Governance framework to address changing and emerging risks and needs

The confluence of institutional and technical practices and measures are essential to ensuring a healthy partnership infrastructure that supports the overarching AIM-AHEAD Consortium's initiatives.