Artificial Intelligence Working Group Update

119th Meeting of the Advisory Committee to the Director (ACD)
December 13, 2019

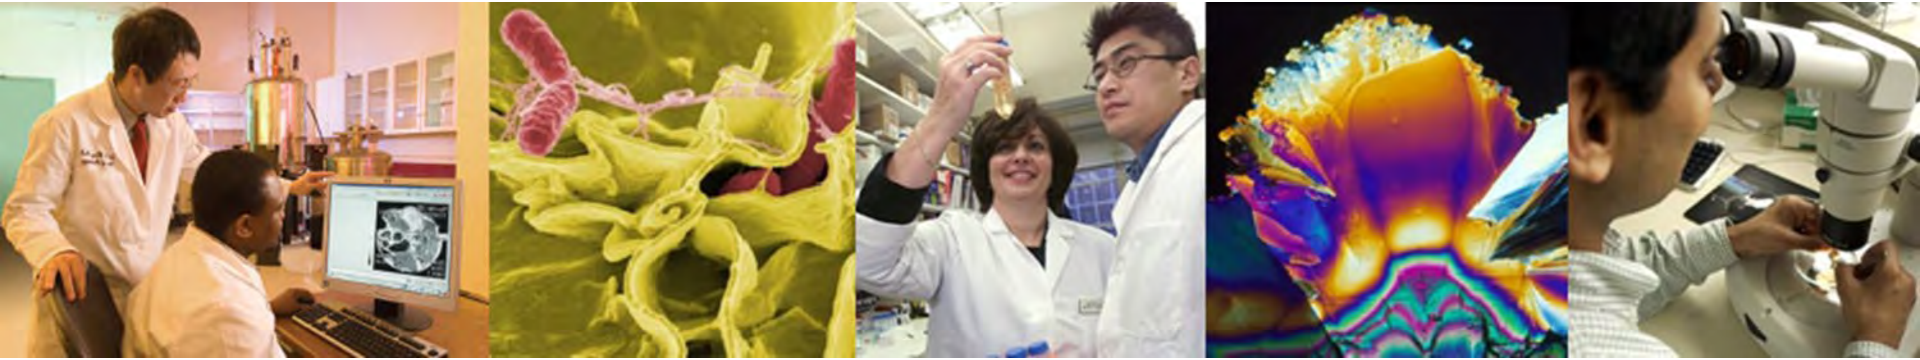

David Glazer
Engineering Director, Verily

Lawrence A. Tabak, DDS, PhD
Principal Deputy Director, NIH
Department of Health and Human Services

ACD Artificial Intelligence Working Group Members


Rediet Abebe
PhD
Harvard


Greg Corrado, PhD
Google


Kate Crawford, PhD
AI Now Institute


Barbara Engelhardt, PhD
Princeton

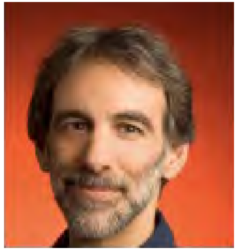
David Glazer
Verily (Co-Chair)


David Haussler, PhD
USCS


Dina Katabi, PhD
MIT Computer Science & AI Lab


Daphne Koller, PhD
insitro


Anshul Kundaje, PhD
Stanford University


Eric Lander, PhD
Broad Institute


Jennifer Listgarten, PhD
Berkeley


Michael McManus, PhD
Intel


Lawrence Tabak, DDS, PhD
NIH (Co-Chair)


Serena Yeung, PhD
Stanford

Agenda

- **our charge**
- big picture
- opportunities & challenges
- recommendations

We Generate Enormous Volumes of Data Daily

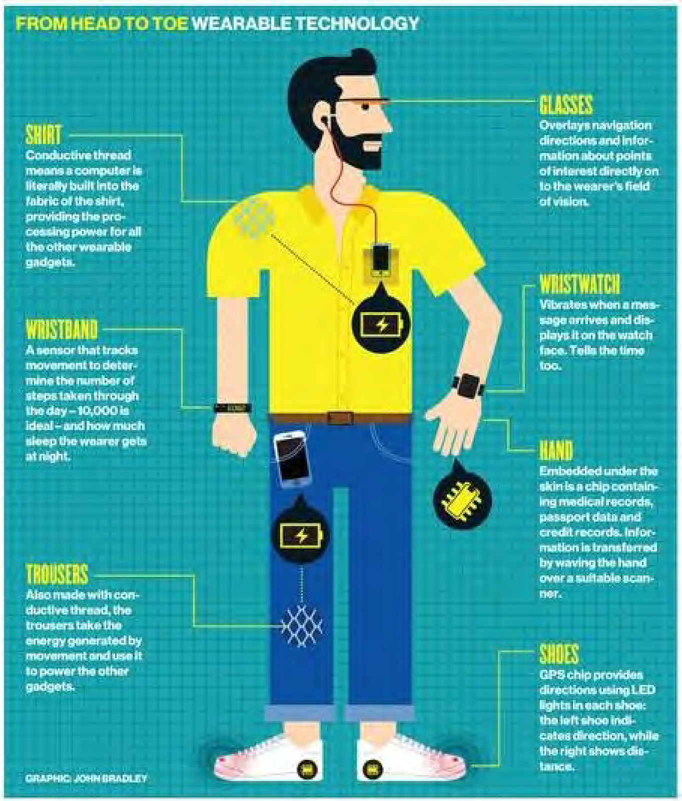

https://people.rit.edu/sml2565/iimproject/wearables/index.html

Every Day Artificial Intelligence Applications

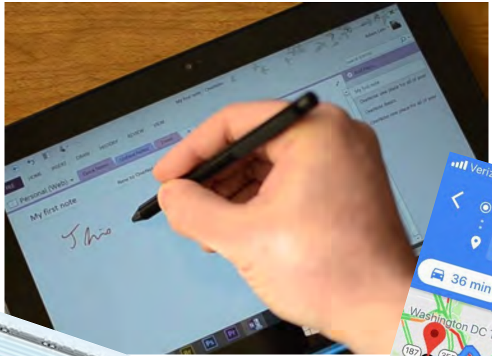

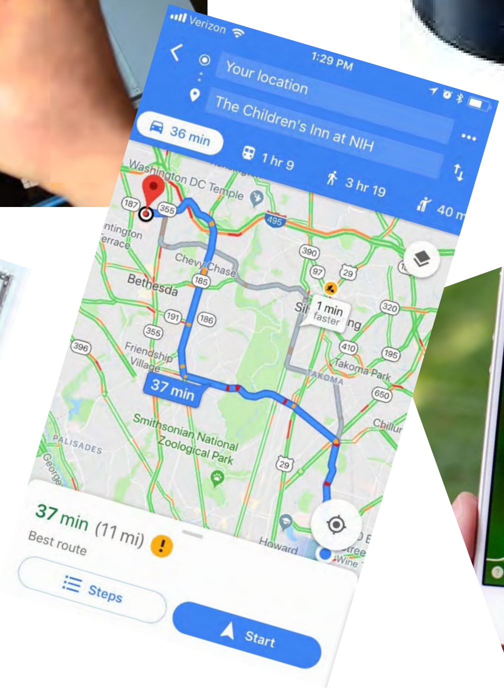

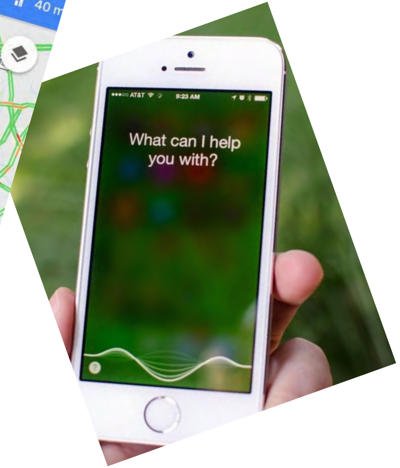

Charge to the AI Working Group (December 14, 2018)

- Are there opportunities for cross-NIH effort in AI? How could these efforts reach broadly across biomedical topics and have positive effects across many diverse fields?
- How can NIH help build a bridge between the computer science community and the biomedical community?
- What can NIH do to facilitate training that marries biomedical research with computer science?
    - Computational and biomedical expertise are both necessary, but careers may not look like traditional tenure track positions that follow the path from PhD to post-doc to faculty
- Identify the major ethical considerations as they relate to biomedical research and using AI/ML/DL for health-related research and care, and suggest ways that NIH can build these considerations into its AI-related programs and activities

Agenda

- our charge
- **big picture**
- opportunities & challenges
- recommendations

Parallel Revolutions:
Fusing Biomedicine and Machine Learning

Data Generation

more data about the biology and health of more individuals than ever before

Data Analysis

machine learning, other forms of artificial intelligence, cloud computing

ML-BioMed

- biomedical experiments that are designed for ML
- ML that's designed for biomedical experiments

Three Themes

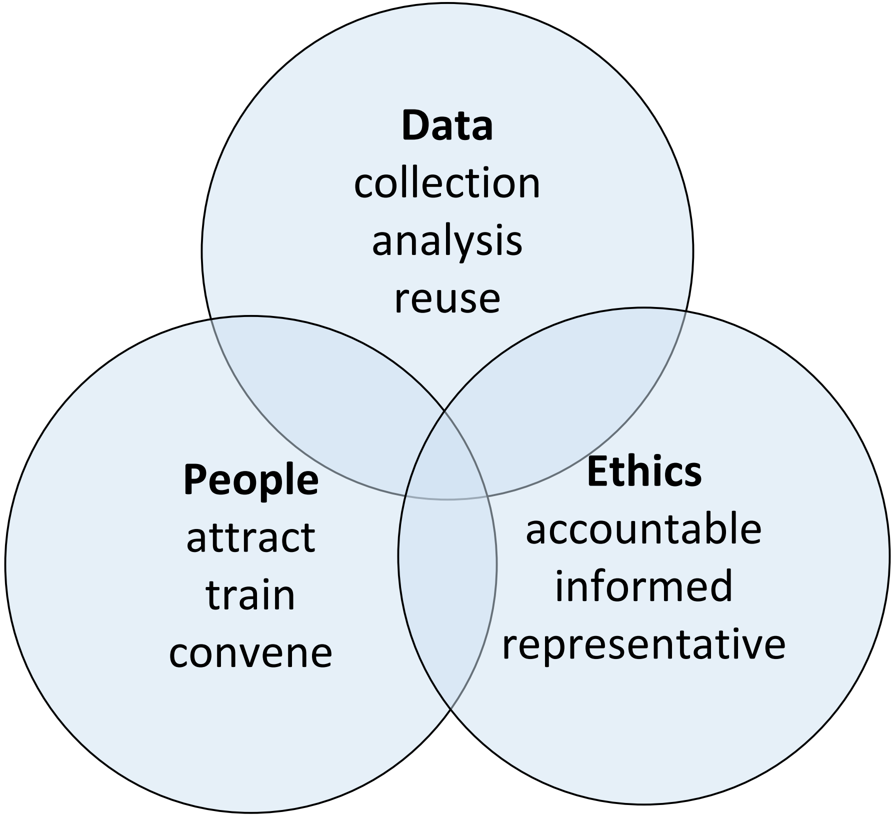

Our Starting Point

We feel strongly that:

1. the **opportunity** is huge
   - including to drive discontinuous change
2. we need new **data generation** projects
   - NOT business-as-usual
3. the single best way to attract the **right people** is with the right data
   - “show me the data”
4. the time to invest in **ethics** is now
   - before we dig a deeper hole

Agenda

- our charge
- big picture
- **opportunities** & challenges
- recommendations

Opportunities
Biomedical Science: Active Learning

- Active learning uses ML to search through high-dimensional spaces, by creating intelligent feedback loops.
  - use the results of existing measurements to prioritize the gathering of new measurements
- Biomedicine has an extraordinary range of measurement and perturbation tools.
- Examples of potential spaces to search:
  - protein design
  - drug regimens matched to the genomic state of cells
  - whole genome CRISPR multiplexed screens to dissect cellular circuitry

Opportunities
Biomedical Science: Understanding Mechanisms

- Use machine learning to ‘fill in’ pictures, at many levels, from partial information.
  - e.g. impute large-scale gene expression into histological images, based on partial data
  - e.g. predict effects of coding and non-coding mutations, based on large but incomplete data
  - e.g. infer the effects of combinatorial perturbations too vast to ever interrogate experimentally
- Go beyond using ML to make predictions; use ML to infer mechanisms.
  - Systematically learn how regulatory networks, cellular programs, and tissue-level interactions work, and how dysfunction of mechanisms leads to disease.

Opportunities
Clinical Care

- Develop more accurate markers for more conditions.
  - measure health “inputs” (molecular, behavioral, physical)
  - measure health “outputs” (interventions and outcomes from medical records)
  - train ML models to identify the most useful signals and to predict health outcomes
- Deploy personalized monitoring to individual patients. For example:
  - risk of developing metabolic syndrome and recommended interventions
  - early warnings of autoimmune disease risk and flare-ups
  - early signs of neurological diseases, based on wearable and contactless sensors
  - detection of heart disease and sleep disorders, based on wearable and contactless cardiac and respiration monitors and new blood-based marker signatures
  - cancer vaccine effectiveness monitoring via immune cell, cfDNA, and exosome analysis
  - cancer recurrence risk via initial tumor genome analysis followed by regular blood tests
- Scale to large and diverse populations.
  - sensor advances ⇒ inexpensive and comprehensive detection technology
  - computational advances ⇒ models optimized for ubiquitous devices (e.g. phones)

Opportunities
Case Study: Fundoscopy

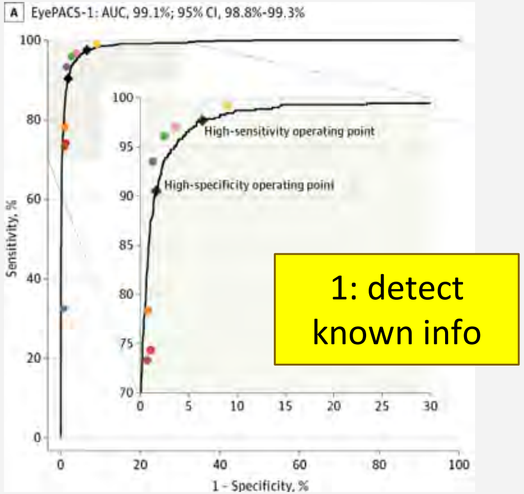

Validation Set Performance for Referable Diabetic Retinopathy. Performance of algorithm (black curve) and ophthalmologists (colored circles) for the presence of referable diabetic retinopathy on EyePACS-1 (8788 fully gradable images)

Gulshan et al., Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA. 2016;316(22):2402–2410.

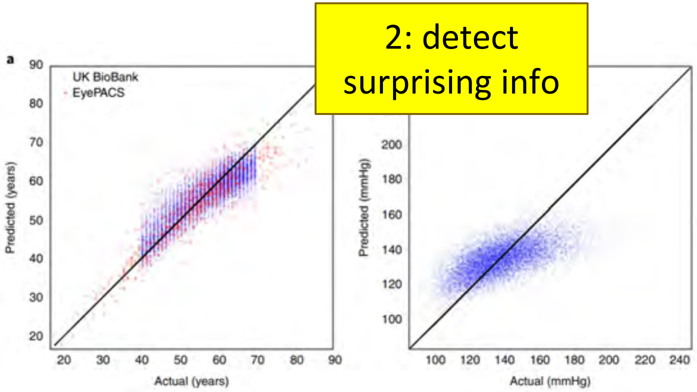

Predictions of age and SBP. a, Predicted and actual age in the two validation datasets. b, Predicted and actual SBP in the UK Biobank validation dataset.

Poplin et al. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. Nat Biomed Eng 2, 158–164 (2018)

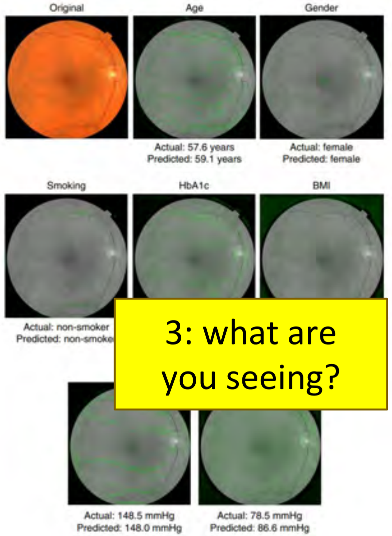

Attention maps. The top left image is a sample retinal image. The remaining images show the same retinal image in black and white, with the "soft attention heat map" overlaid in green, indicating the areas the neural network model is using to make its prediction..

Opportunities
Social Understanding of Health

- Build a future that reduces structural injustices and social harms, rather than amplifying and reinforcing them.
  - include social scientists and humanists from the beginning of projects, to complement technological and biomedical experts
  - incorporate understanding of systemic health inequities, organizational practices of healthcare, and diverse cultural approaches to health
- Examples of potential benefits:
  - earlier detection and resolution of data bias issues before systems are widely deployed
  - consideration of non-technical determinants of health
  - mechanisms for measuring and modulating patient-doctor communication
  - data-driven health programs that improve patient health w/o compromising patient privacy

Opportunities
Ethics and Data Sharing

Combining the ethos of experimentation from the ML field with the traditions of responsible data practices in the biomedical space will allow for research breakthroughs that improve the lives of all people, while preserving the privacy, agency, and respect of patients and their data.

We have an opportunity to drive new best practices across biomedicine, leading to the creation and widespread adoption of tools such as:

- novel ways to build informed consent thru the data lifecycle of a clinical study
- accountability and auditing mechanisms to record what data is used and where
- methods to assess efficacy of models trained on one population for use on others
- processes to include patient groups in the design of health data collection and use

Agenda

- our charge
- big picture
- opportunities & **challenges**
- recommendations

Challenges: data

Many current datasets:

- lack crucial metadata (e.g. demographics, collection methods)
  - limits ability to detect and correct for batch effects
- were originally collected for a narrow purpose
  - e.g. healthcare delivery data is often colored by the need for reimbursement
  - e.g. contain a limited number of modalities and time points
- include biases and a lack of representation
- have small sample sizes
- have prohibitively restrictive access policies

Challenges: consent

Guidelines for participant-facing consent and researcher-facing data access have not kept up with the opportunity for ML-BioMed.

- gaps between biomedical research norms and ML research norms
  - e.g. scraping the internet of public information, without explicit consent for reuse
- need to responsibly enable wide use (and reuse) of data
  - protect all populations, *and* ensure research benefits flow to all populations
- lack of clear guidelines on:
  - handling of data from historically marginalized groups (e.g. sovereign AI/AN tribal nations)
  - when data reuse is allowed, and how to inform participants
  - who are “bona fide researchers”, and what data they are allowed to use
  - how to streamline data access mechanisms, potentially including a “data passport” model

Challenges: ethics

- Oversight and investigation into the use of ML tools has not caught up to the proliferation of use, leading to unanswered questions around:
  - fairness and equity
  - privacy and consent
  - reliability, safety, and security
  - accountability and governance
  - education
- Without new coordinated efforts, machine learning can reinforce existing blind spots and biases in medicine, adding a new unjustified veneer of technical credibility.
  - cautionary tale: Obermeyer et al., Dissecting racial bias in an algorithm used to manage the health of populations

Challenges: people

- ML and biomed experts often struggle to collaborate effectively.
  - Realizing the full value of ML-BioMed requires collaborative teams, with experts from multiple domains.
- The groups speak different languages, and have very limited understanding of each others’ fields.
  - Computational scientists often aren’t aware of the most useful problems to solve.
  - Biomedical scientists often aren’t aware of what tools are useful for what problems.
- Educational systems (from high school through grad school and beyond) and professional conference landscapes are disjoint and do not provide sufficient opportunity for cross-pollination.

Agenda

- our charge
- big picture
- opportunities & challenges
- **recommendations**

Recommendations

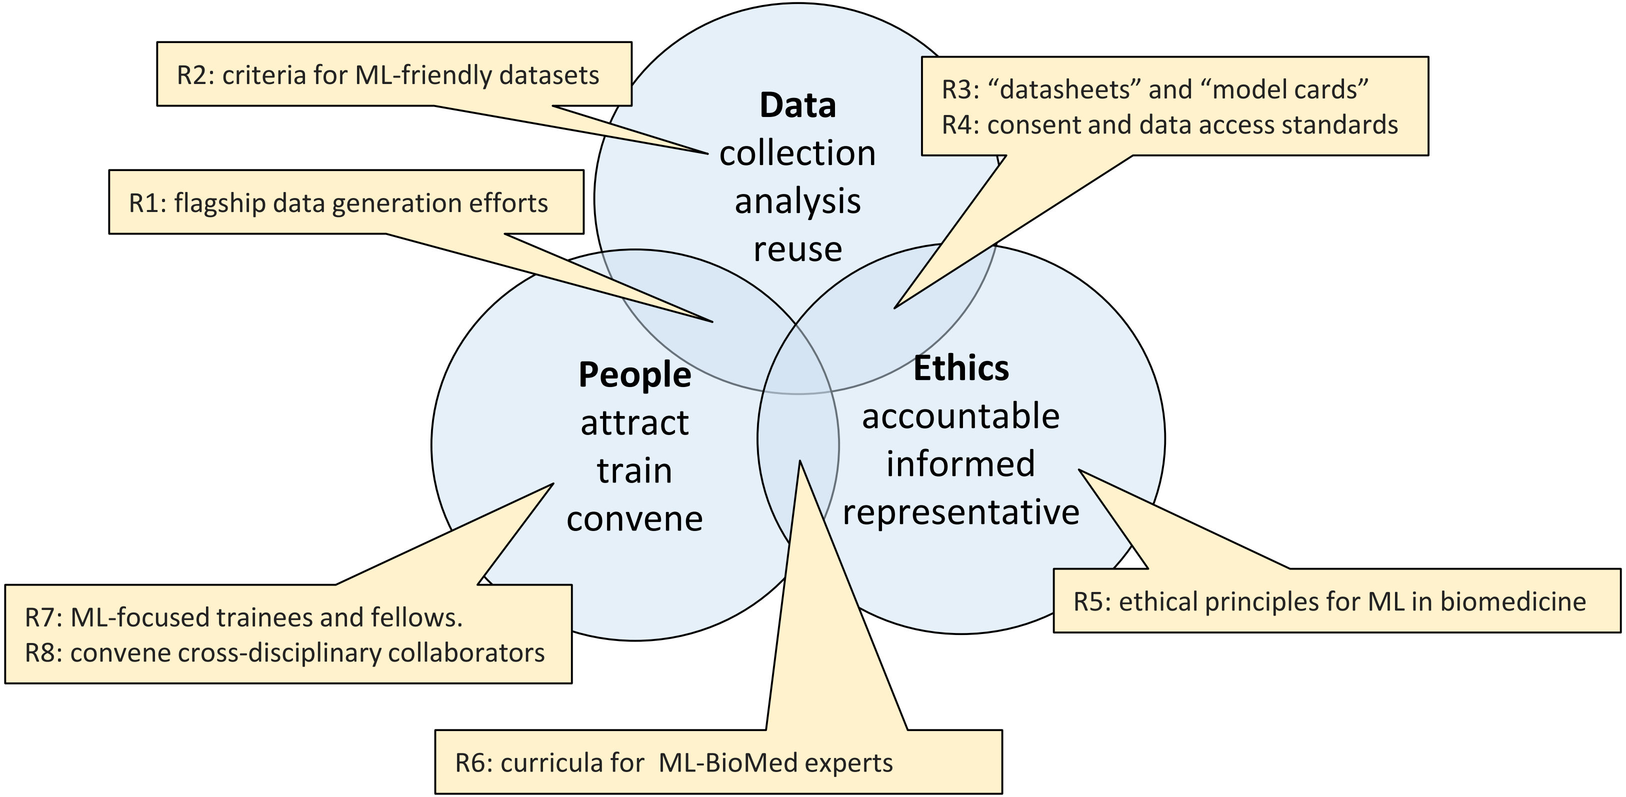

Recommendations

1. Support flagship **data generation** efforts to propel progress by the scientific community.
2. Develop and publish criteria for **ML-friendly** datasets.
3. Design and apply **“datasheets” and “model cards”** for biomedical ML.
4. Develop and publish **consent and data access standards** for biomedical ML.
5. Publish **ethical principles** for the use of ML in biomedicine.
6. Develop **curricula** to attract and train ML-BioMed experts.
7. Expand the pilot for ML-focused **trainees and fellows**.
8. **Convene** cross-disciplinary collaborators.

Non-recommendations

We discussed, but are *not* recommending:

- NIH investment in improving general-purpose ML techniques
  - other venues will do that better
- additional focus on continued use of existing ML tools on existing data
  - that will continue to happen naturally and productively
- investment in scalable secure cloud infrastructure for biomedical data
  - that’s an essential prereq, but not ML-specific, and is already happening

(1) Support flagship data generation efforts to propel progress by the scientific community.

Support flagship efforts that generate large-scale experimental data, with billions of data points designed to:

i. be well-suited for ML analysis and inference
ii. address key biomedical challenges
iii. stimulate new approaches in machine learning

And that implement processes designed to:

i. develop improved criteria and technical mechanisms for data access
ii. strengthen ethical criteria for dataset use (consent, privacy, accountability, ...)

Projects should:

- address key biomedical challenges using ML methods
- advance ML methods for future use in biomedicine
- produce transformative data sets, designed with ML in mind
- propel new ways to gather massive data in biomedicine
- **involve strong engagement from leading ML researchers**

Project review should:

- **incorporate expertise in ML as well as traditional biomedical domains**

(1) Flagship data generation: examples of potential topics

- Cellular pathways: Inferring cellular pathways based on large-scale gene perturbation data
- Genetic variants: Inferring the role of non-coding variants based on observational and perturbational data
- Disparities in healthcare: Predicting minority patients at risk for death or complications after surgery based on data from disparities databases and clinical records
- Histology: Automatic annotation of cellular structures and/or gene expression in histological images
- Microbiology: Inferring interactions among bacterial species and with their human hosts
- Chemistry: In silico drug-like molecule creation; retrosynthetic planning for drug-like molecules; physical property prediction; toxicology predictions
- Medical images: Detecting the presence of abnormalities in radiographic images; real-time simultaneous spatial and temporal cellular images
- Clinical data: Predicting patient outcomes from longitudinal electronic health record data
- Sensors: Inferring health attributes from digital health sensors (wearable, contactless, …)

What about my existing data?

(2) Develop and publish criteria for ML-friendly datasets.

Publish criteria for evaluating datasets based on their value for ML-based analysis.

- what makes a dataset most useful for ML-based analysis?
- what attributes are and aren't addressed by existing datasets?
- start as guidelines; within two years recommend a subset as requirements

Examples of potential criteria:

- **clear provenance:** as much metadata as possible, to detect & correct for batch effects
- **well-described data:** what does each variable mean? what's the distribution of values?
- **accessible data:** flexible data access policy, reasonable data access process
- **large sample size:** to allow training (and evaluation) without overfitting
- **multimodal data:** to study complex systems from multiple perspectives
- **perturbation data:** includes outcomes ("outputs") as well as measurements ("inputs")
- **longitudinal data:** to allow modeling and prediction of progression
- **active learning:** data grows over time, incorporates new data-gathering techniques, and uses ML-based analysis of existing data to inform future data generation

(3) Design and apply "datasheets" and "model cards" for biomedical ML.

- Develop and publish best practices for:
  - "datasheets" that describe & evaluate training datasets
  - "model cards" that do the same for generated models
- Test the best practices in the real world:
  - build after-the-fact examples for existing datasets
  - apply to new datasets, and update the best practices
- Once best practices have been updated:
  - require datasheets and model cards for all NIH extramural grant applications and NIH intramural projects that involve ML research
  - encourage journals to do the same for paper submission and publication

Potential datasheet best practices:

- demographics and UBR characteristics
- privacy, consent, and copyright issues
- known blind spots, which could otherwise create hidden biases

Potential model card best practices:

- what training data was used
- how training and validation were done
- known limitations on applicability
- intended use, and potential harms of inappropriate use

(4) Develop and publish consent and data access standards for biomedical ML.

Charge a working group to address the substantial gap between consent standards typically required in biomedical research and consent standards typically applied in ML.

- Standards should ensure appropriate consent for biomedical ML, by reconciling:
  - common ML practices;
  - existing biomedical best practices; and
  - ongoing efforts in the global biomedical community to harmonize consent and data use standards to facilitate the widest responsible use of data, while ensuring protections against potential harms
- Once standards are developed and refined, implement appropriate mechanisms to ensure adherence.

Background: it is common practice for ML developers to create training sets for ML models by scraping the internet for public text, images, and videos without explicit additional consent for such reuse.

(5) Publish ethical principles for the use of ML in biomedicine.

Charge a working group to develop ethical principles for ML in biomedicine.

- address unique ethical challenges in this space, that add to existing challenges for the use of ML in other public and private sector settings
- include expertise in ML, biomedicine, law and public policy, Science and Technology Studies
- include representation of communities that could be negatively impacted by ML in biomedicine
- refine the draft principles by testing them against new publications, and reviewing the gaps
- once refined, recommend appropriate mechanisms to ensure adherence

Principles are likely to cover:

- **fairness and equity:** avoid reinforcing existing biases; don't contribute to future health disparities
- **privacy and consent:** coordinate with Recommendation 4
- **reliability, safety, and security:** extend existing mechanisms as needed to work with ML-powered tools
- **accountability and governance:** extend existing mechanisms as needed to account for the unique attributes of ML-powered tools
- **education:** coordinate with Recommendation 6 to include ethics content in curricula; inform broader audiences

(6) Develop curricula to attract and train ML-BioMed experts.

Curricula goals are to:

- **help experts from all fields successfully collaborate across disciplines**
- entice upcoming and established data experts into biomedicine
- inform upcoming and established biomedical experts about modern ML techniques, strengths, and limitations
- invite social scientists and humanists to collaborate and inform on best practices
- raise the awareness of biomedical policymakers and decision makers about ML opportunities and risks

Potential curriculum elements include:

- emphasis on fundamental concepts, not technical details or fact memorization
- risks of mis-applying ML in biomedicine, such as hidden bias from training sets.
- (in ML courses) motivating examples from biology and clinical medicine
- (in ML courses) hands-on exercises to build models from real-world biomedical datasets
- (in bio courses) hands-on exercises to use real-world ML tools to train and evaluate biomedical models

(7) Expand the pilot for ML-focused trainees and fellows.

- Piloted in summer 2019 with three ML projects in the Civic Digital Fellowship and two in the Graduate Data Science Summer Program.
- Make ML a major ongoing focus of these and similar programs going forward.

(8) Convene cross-disciplinary collaborators.

- Support biomedical tracks and workshops at leading computational conferences.
  - piloted with workshop at NeurIPS in December 2019
- Expand to other types of conferences.
  - including general biomedicine and scientific conferences
- Explore other opportunities for convening experts from different fields.

Conclusion

Recent advances in data generation and data analysis have brought biomedicine to the cusp of a new world of ML-BioMed. The computational and biomedical communities are poised to jointly drive transformative progress in biomedical research -- leading to new insights into how all living systems work -- and in care delivery -- leading to improvements in the health of all humans and all communities.

The NIH is well positioned to accelerate that progress, by supporting the three complementary areas of **data** to fuel the analysis engines, **ethics** to always be steering in accordance with our highest values, and **people** to drive projects forward. The eight recommendations in this report suggest specific ways to propel progress. We look forward to seeing the results unfold.

Thank you!

- Questions?
- Discussion

AI WG kick-off meeting at NIH in February 2019

Recommendations

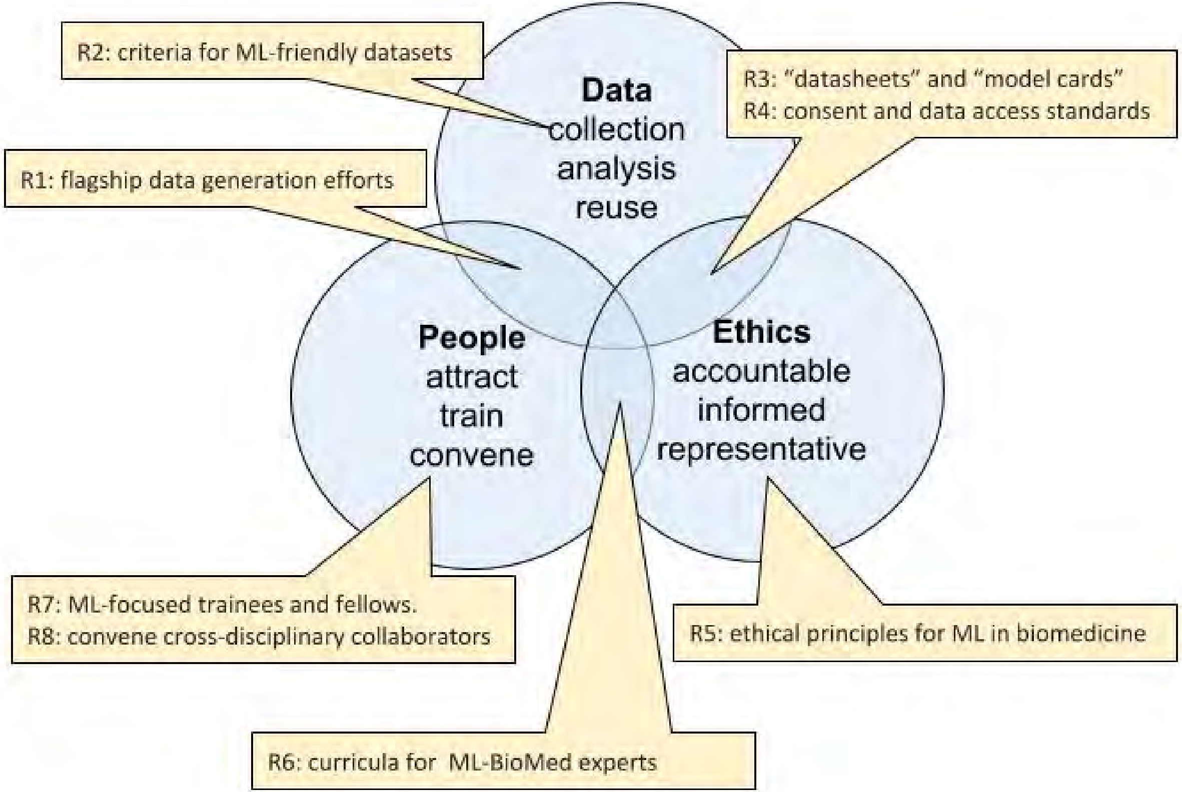

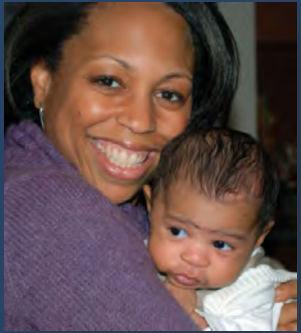


NIH...

Turning Discovery Into Health