Introduction to AI for Public Health

PHAC Speaker Series
2026-09-23

Professor Laura C. Rosella
Tier 1 Canada Research Chair in Population Health Transformation & Analytics
Dalla Lana School of Public Health, University of Toronto
Education Lead, Temerty Centre for Artificial Intelligence Research and Education in Medicine (T-CAIREM)
NPA for the AI for Public Health (AI4PH) Research Training Platform (https://ai4ph-hrtp.ca/)

AI4PH AI FOR PUBLIC HEALTH

UNIVERSITY OF TORONTO
DALLA LANA SCHOOL OF PUBLIC HEALTH

Today I will cover:

Foundations

Concepts and limits

Applications: How is AI being used?

Tasks and added value

Impacts

Risks for health, equity, and the environment

Responsible implementation

Requirements for practice

How will we know if advances in AI are having a positive impact for public health?

⇒Fewer people developing diseases or poor health outcomes that are preventable

⇒Promoting health in our environments and communities

⇒Reduction in health inequities

Unclassified / Non classifié

PURPOSE

To achieve optimal health and well-being for all people living in Canada

AIMS

What the public health system is trying to achieve

- Enhance the health status of populations
- Protect against health emergencies and mitigate the impacts
- Achieve equitable health outcomes

ESSENTIAL PUBLIC HEALTH FUNCTIONS

How the public health system works to achieve the aims

1. Health promotion
2. Health surveillance
3. Health protection
4. Population health assessment
5. Disease and injury prevention
6. Emergency preparedness and response

BUILDING BLOCKS

How the public health system carries out the functions

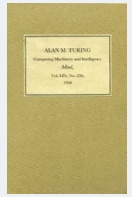
- Policy and program interventions
- Evidence, knowledge, and information
- Medical and digital health technology
- Workforce expertise and human resource capacity
- Financing
- Governance, leadership, and engagement

A brief history

Part 1 · 1950–1999: from big ideas to winning games

1950s

AI is born



1950 Turing: “Can machines think?”
1956 Dartmouth names the field
1957 The perceptron

1960s

Early optimism



Heuristics and symbolic reasoning
1966 ELIZA, the first chatbot
1966 Shakey, a robot that plans

1970s–80s

Winters & experts

1973 first AI winter
MYCIN, XCON expert systems
1987 second AI winter

1990s

Games & statistics



1992 TD-Gammon: early RL success
1997 Deep Blue beats Kasparov
Statistical vision and NLP

A brief history - 2

Part 2 · 2000–today: data, deep learning and generative AI

2000s

Data & compute



2005 DARPA driverless car challenge
2006 Netflix Prize
2009 ImageNet dataset

2010s

Deep learning



2012 AlexNet
2014 GANs, Seq2Seq
2015 ResNet, DQN
2016 AlphaGo
2017 Transformers
2018 BERT

2020s

Generative AI



2020 AlphaFold 2
2022 ChatGPT, Stable Diffusion
2023 Open-weight LLaMA
2024 Nobel Prizes for AI
Now multimodal agents

Select AI Index technical performance benchmarks vs. human performance

Source: AI Index, 2026 | Chart: 2026 AI Index report

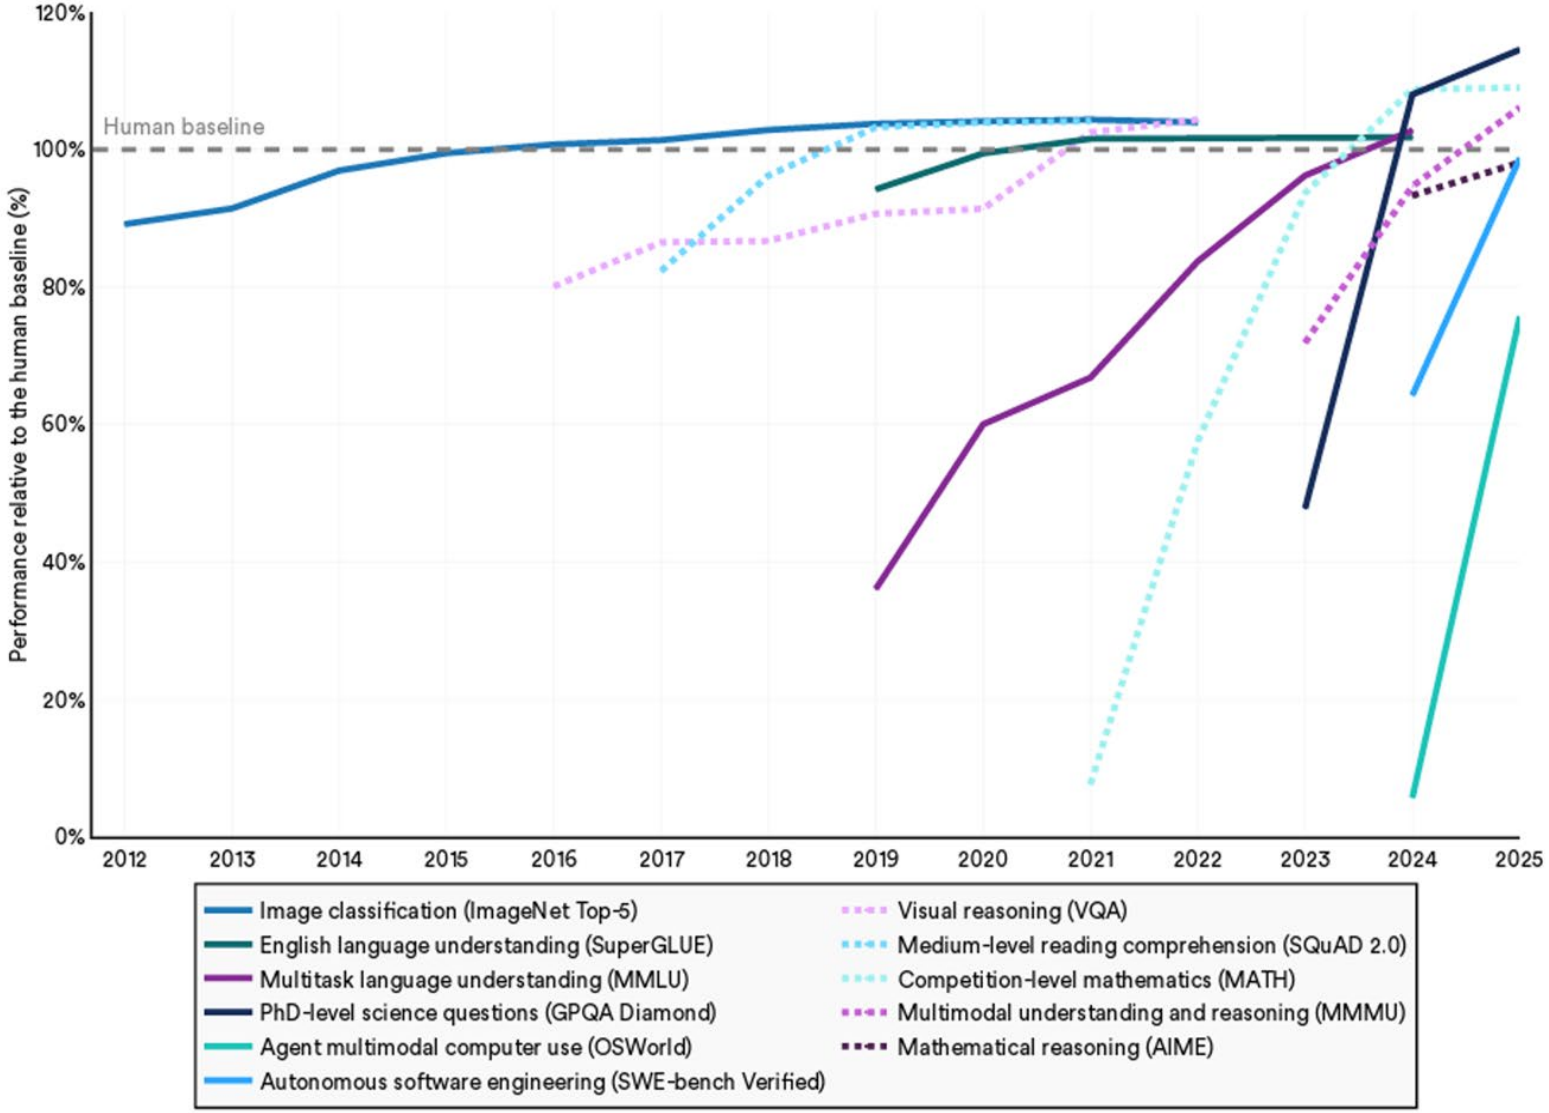
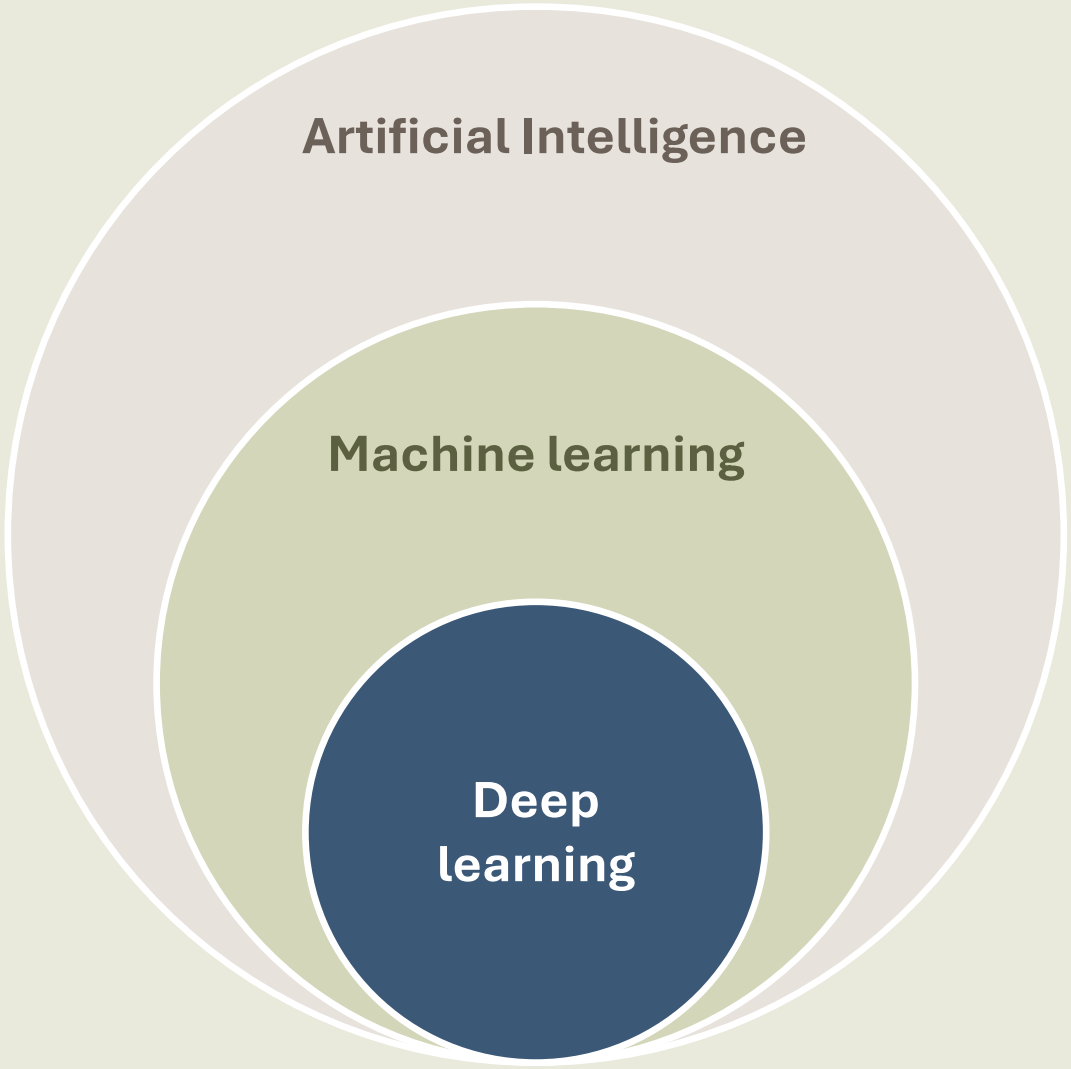


Figure 2.1.1[1]

AI, machine learning and deep learning



Artificial intelligence
Machines doing tasks that need human-like intelligence

Machine learning
Learns patterns from data

Deep learning
Many-layered neural networks for images, text, sound, and other complex data

Generative AI

AI that learns patterns from training data to generate new content in response to instructions.



How it works

- During training, models learn patterns in language, images, or other data.
- Large language models generate text one piece at a time, using the prompt, available context, and preceding text to predict what comes next.
- Instructions and supplied documents shape the response, including its content, tone, and level of detail.

What this enables in public health

- Turn a technical report into a draft briefing or public-facing explanation.
- Adapt existing guidance for different audiences, languages, or reading levels.
- Explore themes in consultation responses or generate code to support analysis.
- Requires checking!

Training, testing, and real-world use

Stage	Question
Training	Can the model learn useful patterns from these data?
Validation and testing	Do the patterns hold in data not used to fit the model?
Real-world evaluation	Does using the model improve decisions and outcomes?
Monitoring	Does it remain useful as the setting changes?

Machine learning in public health

Five approaches, from finding patterns to generating text

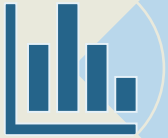
APPROACH	WHAT IT DOES	PUBLIC HEALTH EXAMPLE
Clustering	Finds natural groupings in data	Spotting hotspots and early outbreak signals
Decision trees and ensembles	Predicts an outcome from many factors	Risk prediction, case finding, forecasting
Deep learning	Finds patterns in images, signals and complex data	Imaging, sensors and wearables
Natural language processing	Pulls meaning out of text	Finding cases, symptoms and public sentiment
Large language and multimodal models	Generates and combines text, images and data	Summarizing reports; extracting information

Sources: WHO, 2021; WHO, 2024; Sidey-Gibbons & Sidey-Gibbons, 2019; LeCun et al., 2015; Moor et al., 2023.

Emerging Technology Trends

- Foundation models: general-purpose AI models that can be adapted across many tasks
- Real-time AI: continuously sensing, updating, and informing response
- Multimodal AI: combining data sources from multiple modalities
- Agentic AI: Systems that plan and carry out multiple steps using tools
- Artificial General Intelligence (AGI): Human-level or beyond across most cognitive tasks

Getting more specific on AI applications



Descriptive



Predictive



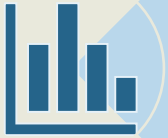
Causal Inference



Prescriptive

Hernan et al. CHANCE 2019

Getting more specific on AI applications – 2



What is happening?



What will happen?



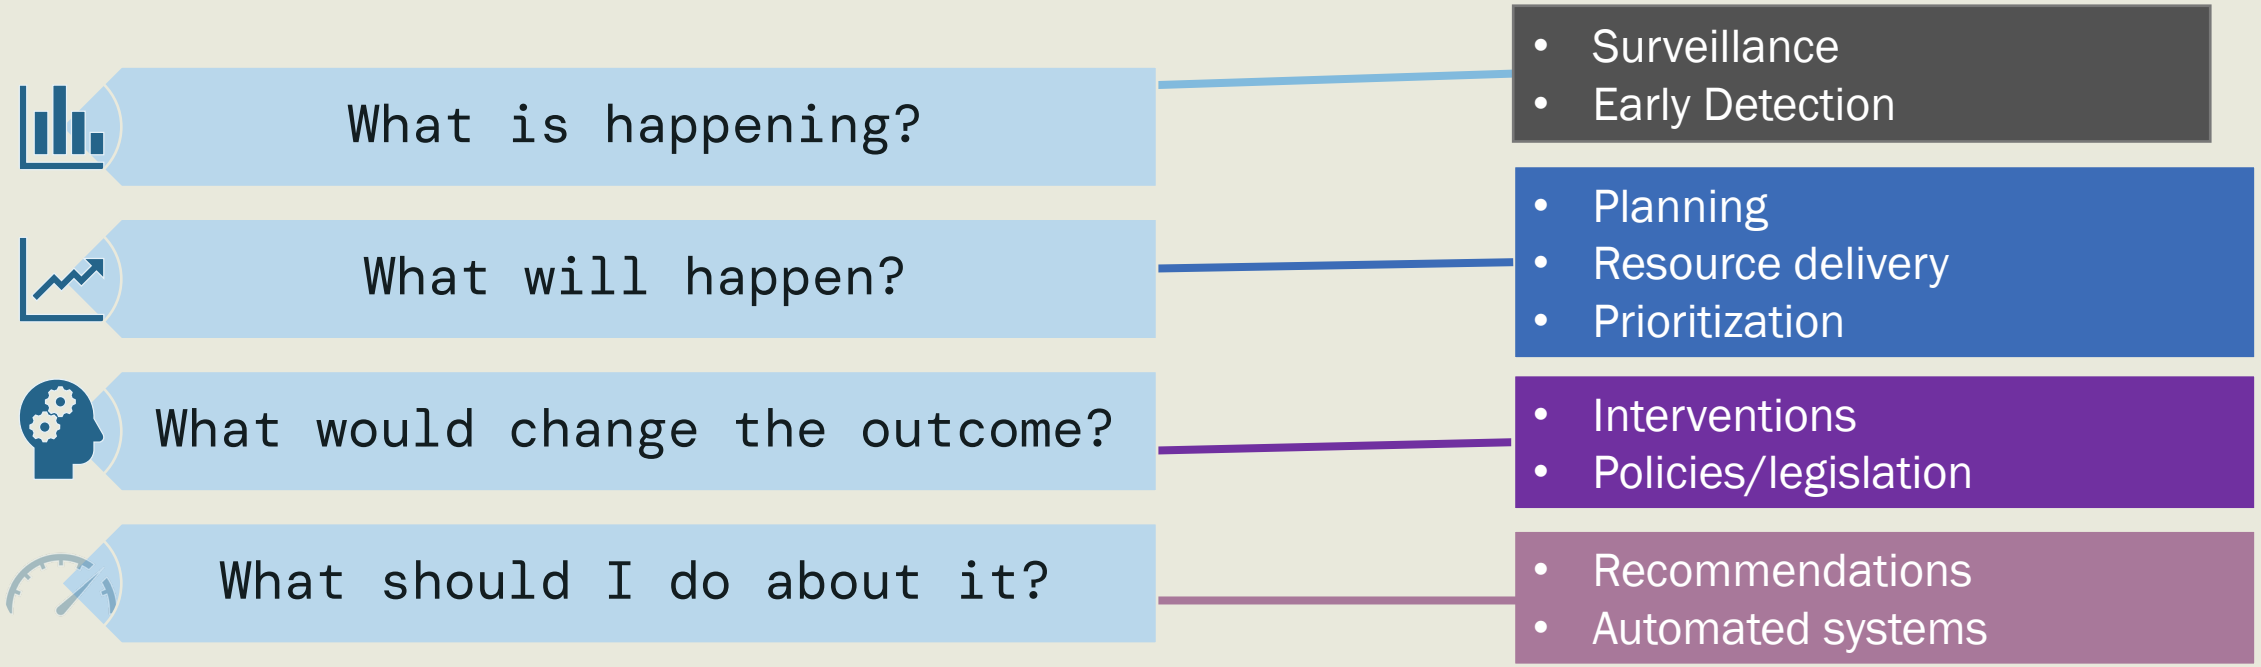
What would change the outcome?



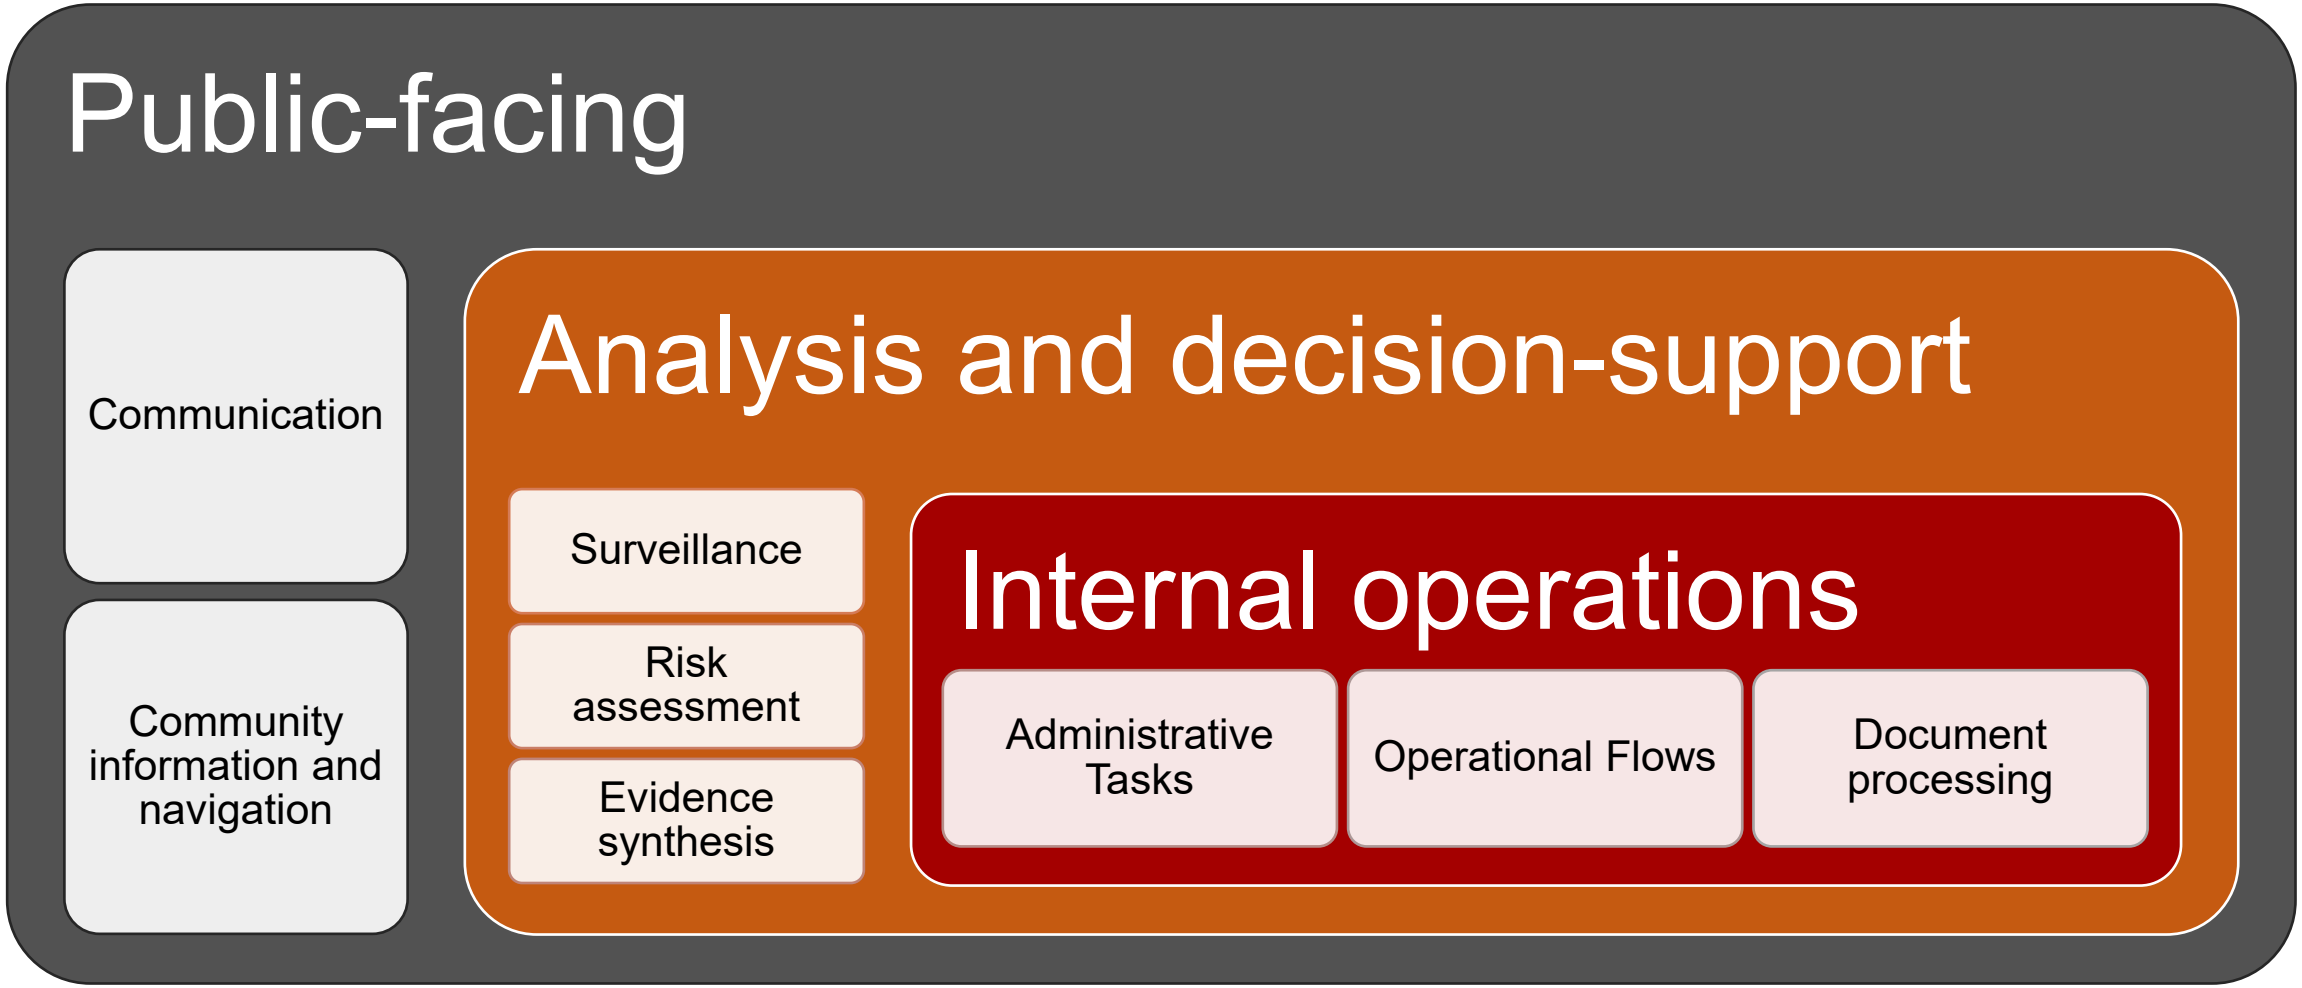
What should I do about it?

Hernan et al. CHANCE 2019

Getting more specific on AI applications for public health



Where AI fits in public health organizations



Text classification for foodborne illness outbreak investigations using deep learning

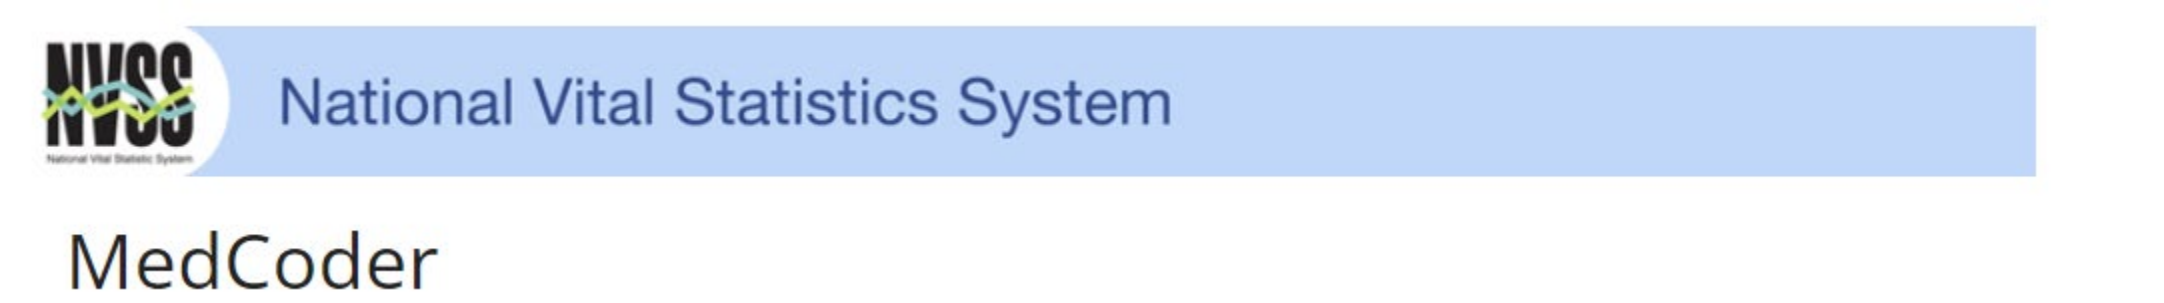
- **Task:** Turn free-text food histories into standardized food categories for outbreak analysis.
- **What AI adds:** A language model assigns food categories to unstructured descriptions, helping investigators analyse exposure information at scale.
  - Using cluster computing, the CDC trained a **deep neural network** language model to predict food categories for free text observations

Descriptive

Surveillance

Khan, MA. et al. Text Classification for Foodborne Illness Outbreak Investigations Using Deep Learning EIS Conference 2023

CDC example: MedCoder



- **Task:** Convert causes of death recorded on death certificates into standardized ICD-10 codes.
- **AI contribution:** language processing and machine learning, combined with coding rules.
  - *According to the CDC, MedCoder can code nearly 90% of records automatically, compared to less than 75% for the previous system*

Descriptive

Surveillance

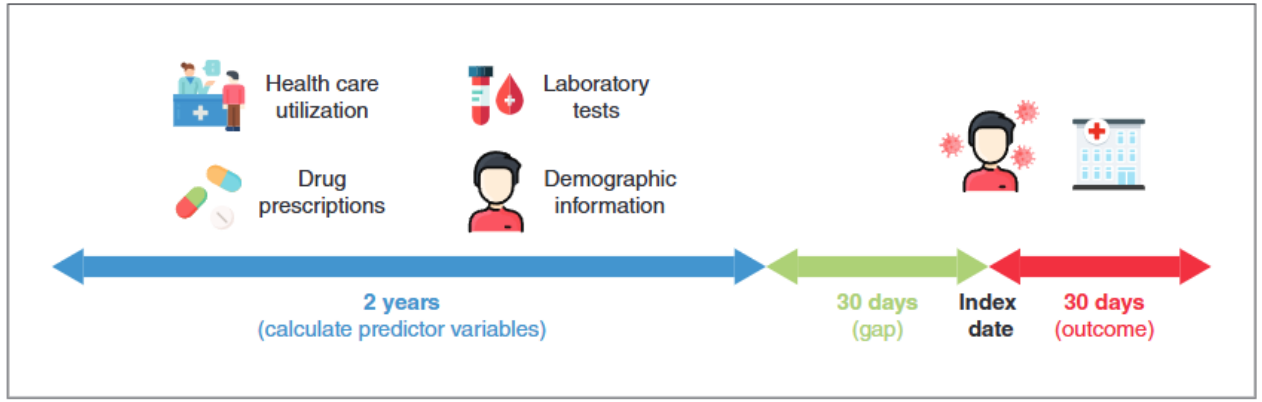
CDC. A Timelier Understanding of Death. 2023
https://stacks.cdc.gov/view/cdc/133554

Risk stratification for COVID-19 hospitalization: a multivariable model based on gradient-boosting decision trees

Jahir M. Gutierrez PhD, Maksims Volkovs PhD, Tomi Poutanen MSc, Tristan Watson MPH, Laura C. Rosella PhD

Predictive

- **Task:** Estimate the risk of hospitalization within 30 days of a positive SARS-CoV-2 test.
- **What AI adds:** A gradient-boosting model combines routinely collected health information to identify high-risk individuals more effectively than the age- and comorbidity-based rules evaluated.



Prioritization

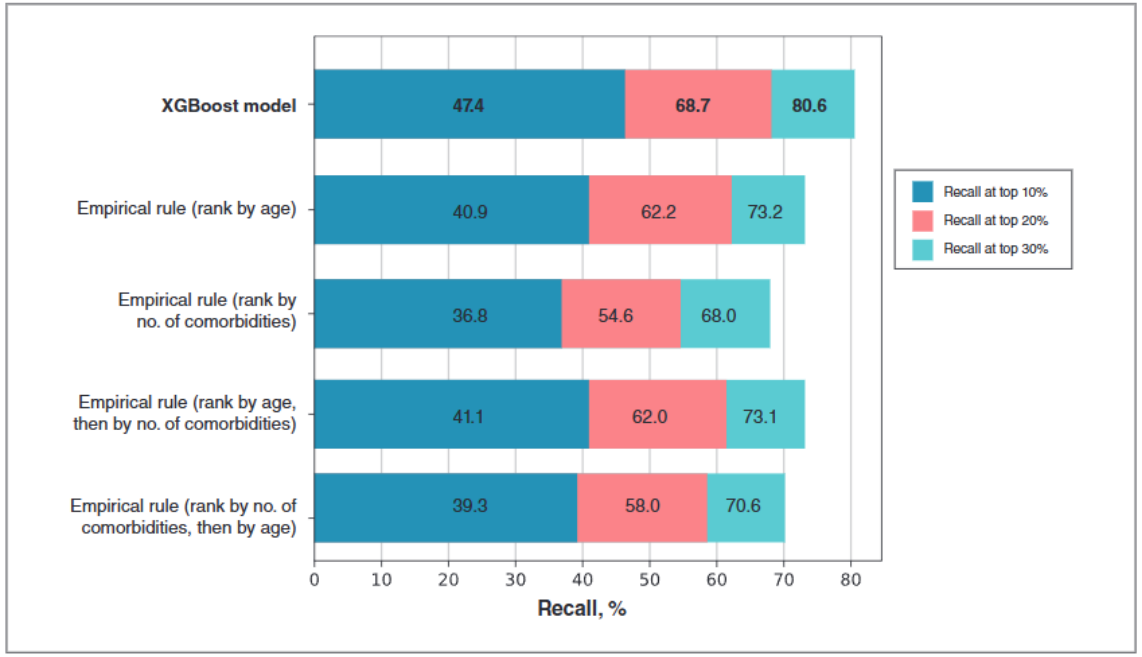


Figure 4: Comparison of recall at top percentiles. The final Extreme Gradient Boosting (XGBoost) model recall (percentage of true hospitalizations recovered in validation data set) was compared against 4 empirical rules.

Surveillance

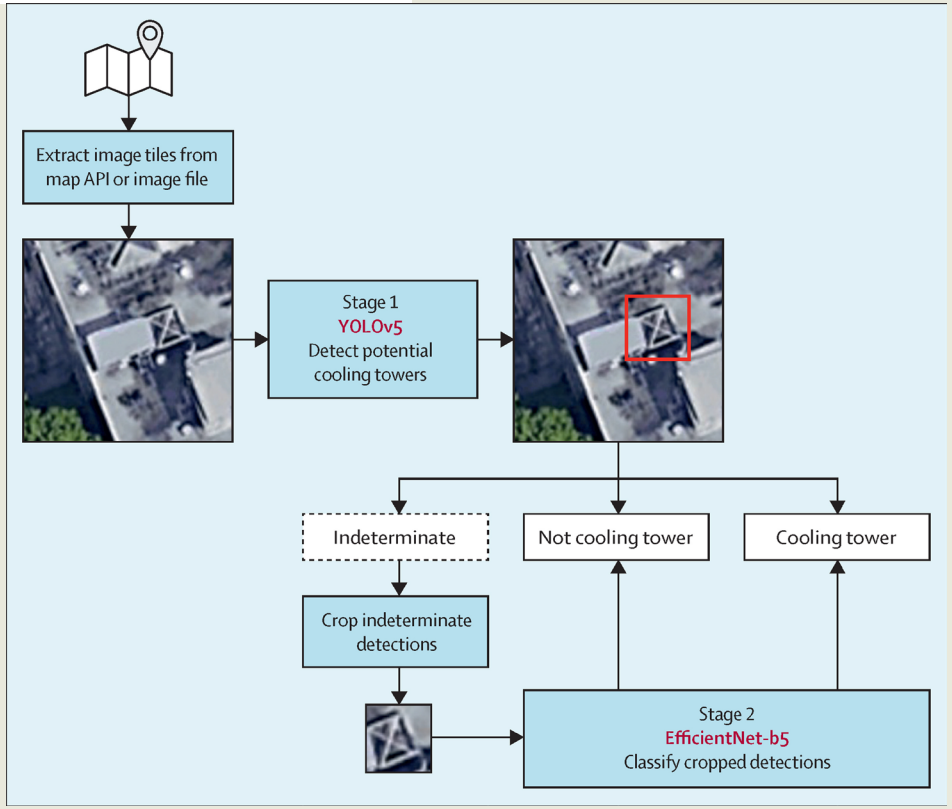
Prioritization

Descriptive

Automated cooling tower detection through deep learning for Legionnaires' disease outbreak investigations: a model development and validation study

Karen K Wong, Thaddeus Segura, Gunnar Mein, Jia Lu, Elizabeth J Hannapel, Jasen M Kunz, Troy Ritter, Jessica C Smith, Alberto Todeschini, Fred Nugen, Chris Edens

- **Task:** Locate cooling towers that may require investigation as potential sources of Legionella exposure.
- **What AI adds:** Computer vision searches aerial imagery for potential cooling towers, narrowing the search for investigators and reducing manual image review.



PLOS DIGITAL HEALTH

RESEARCH ARTICLE

Use of large language models as a scalable approach to understanding public health discourse

Laura Espinosa*, Marcel Salathé

Digital Epidemiology Lab, School of Life Sciences, School of Computer and Communication Sciences, EPFL, Switzerland

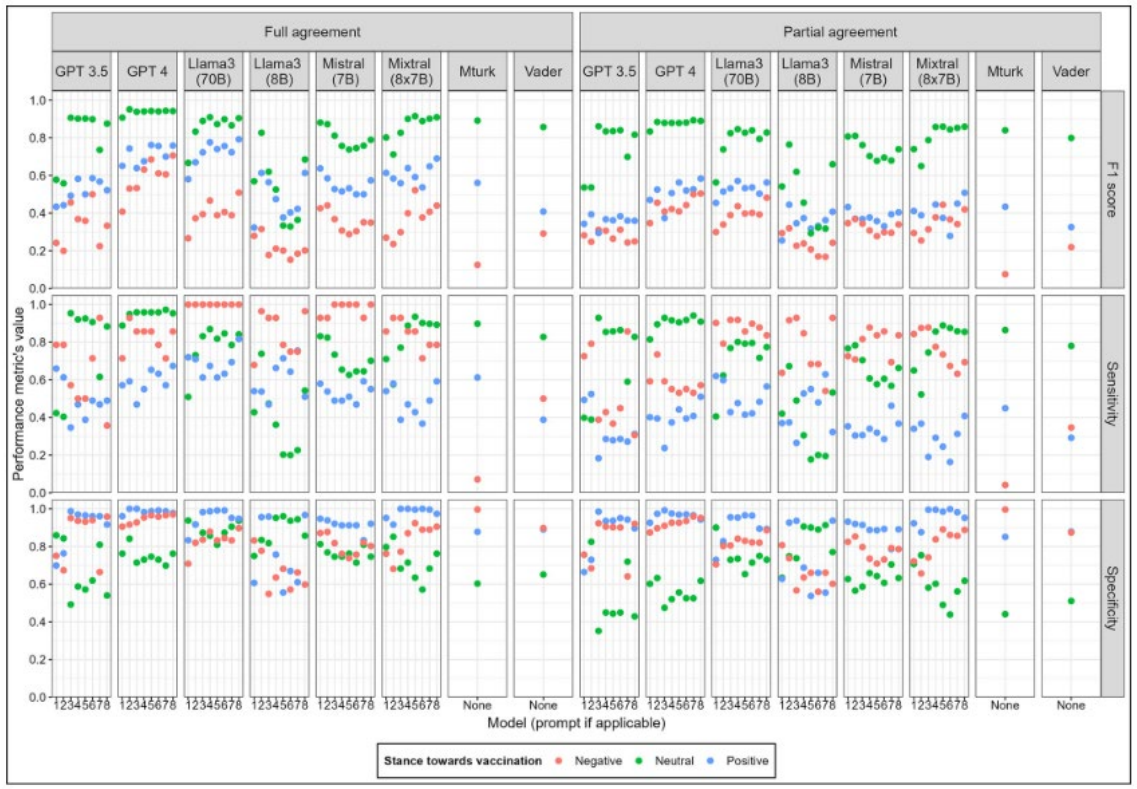


Fig 3. F1 score, sensitivity and specificity per model, prompt (if applicable), stance and experts' agreement.

Descriptive

Communication Planning

- **Task:** Classify whether social media posts express a positive, negative, or neutral stance toward vaccination.
- **What AI adds:** Pretrained language models classify unstructured posts using instructions and examples, offering a scalable alternative to manually labelling every post.

What these examples teach us

- AI can structure information, identify patterns, and support prioritization.
- Each output needs a clear connection to a public health action.
- Evaluation must distinguish model performance, effects on work, and effects on health.
- Benefits and harms must be assessed across populations.

Challenge: Ensuring relevance for action

As a [public health decision-maker],
if I knew [information],
I would do [intervention/action],
to improve [measurable outcome

Data quality, representation, and fit for purpose

How data are collected, measured, and used shapes what AI can tell us

Questions to ask – New, Who, Fast, Last



Will using this data allow for new information?



Who does this data represent?



Will using this data allow for information faster?



Can we maintain the data for ongoing applications?

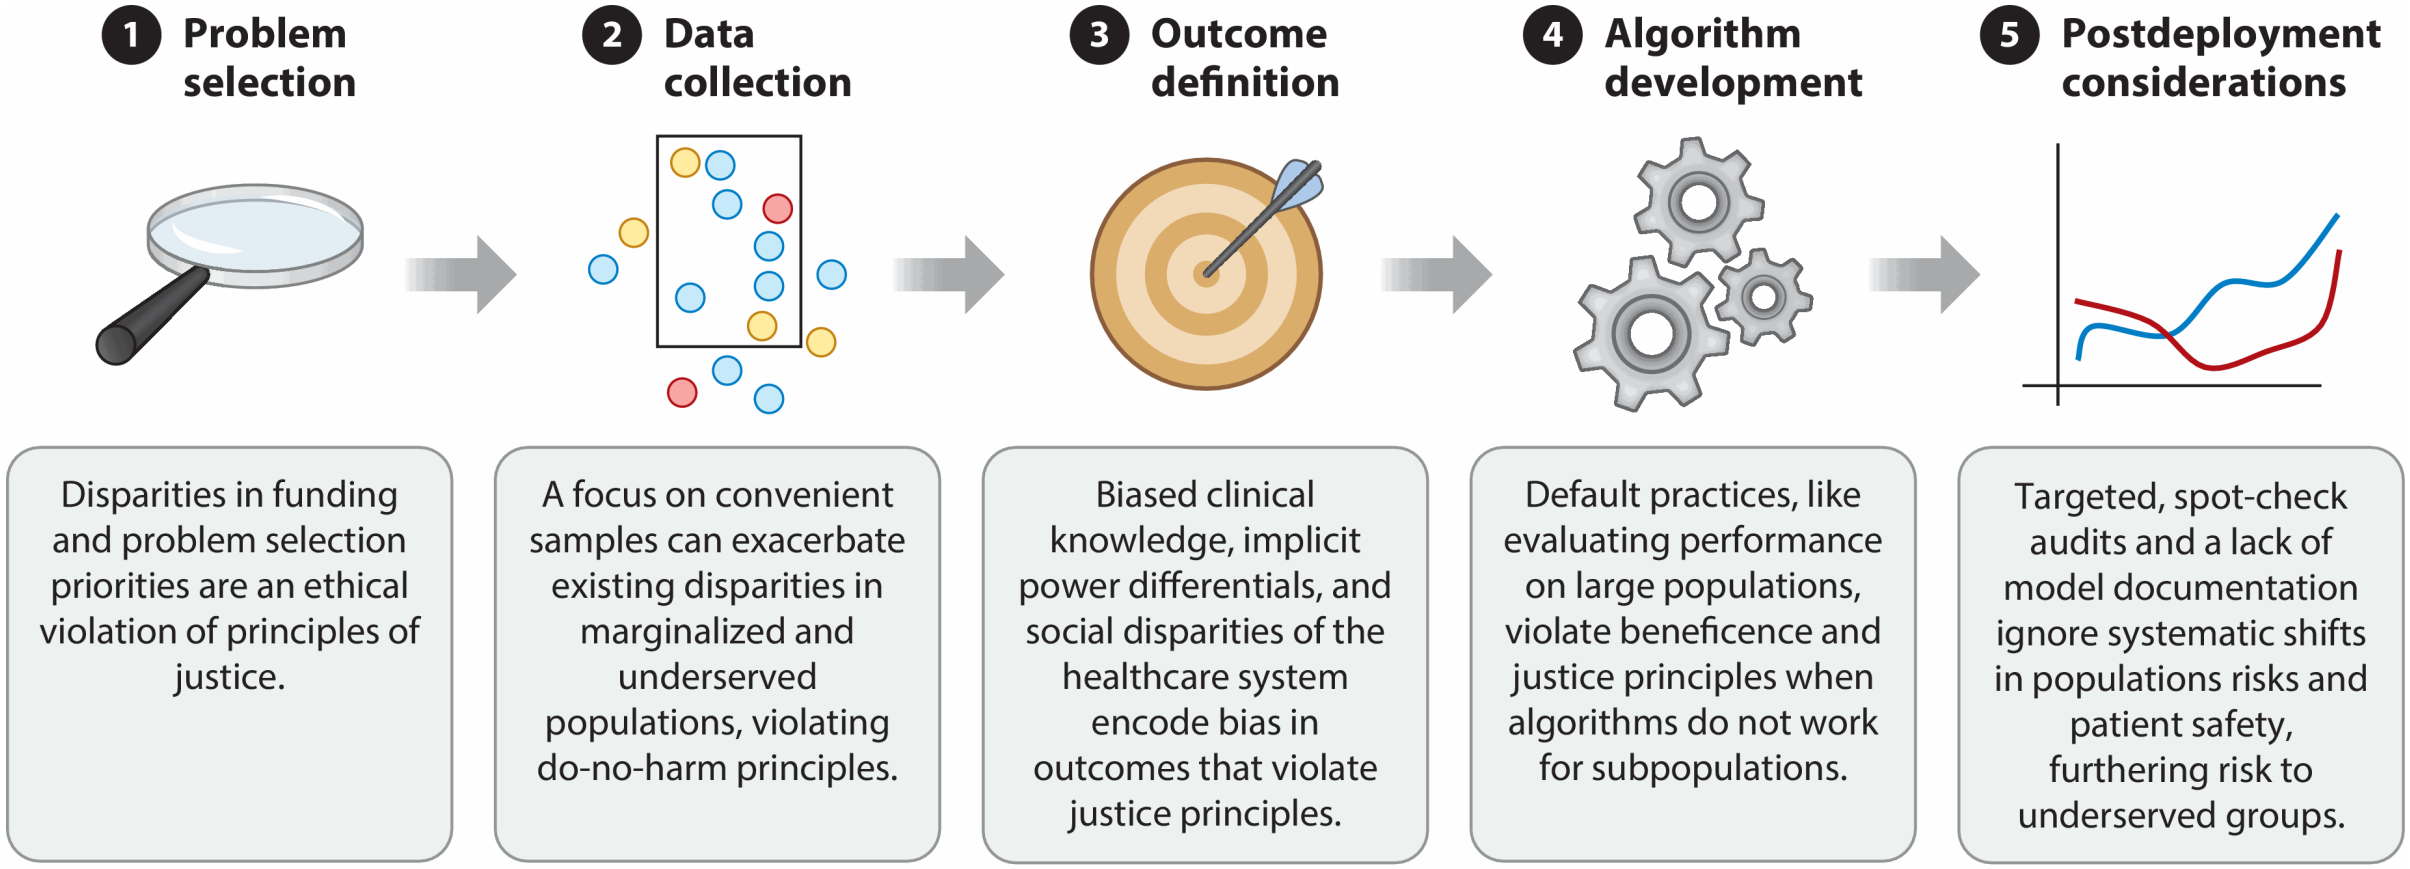
Emerging Data Sources – examples

Wastewater Data

- ✓Timely
- ✓Population Coverage
- ✓Few privacy or legal risks
- ✓Actionable
- ❖Pipeline costly?

Social Media Data

- ✓Timely
- ✓Cost-effective
- ✓New information
- ❖Selection bias
- ❖Privacy questions
- ❖Difficult to action



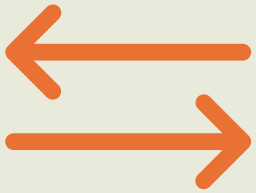
Chen IY, et al. 2021
Annu. Rev. Biomed. Data Sci. 4:123–44

AI has an environmental footprint

Stage	Resources and impacts to consider
Hardware production	Material extraction, manufacturing, and embodied emissions
Training and use	Electricity for computation, including repeated inference
Infrastructure	Cooling, water use, and local electricity supply
End of life	Replacement cycles, reuse, and electronic waste

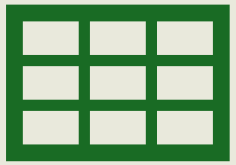
IEA, Energy and AI, 2025

In the way we train our models



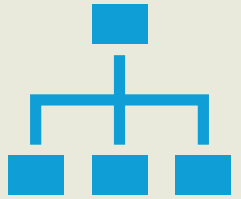
Transfer learning

Using a pre-trained model & fine-tuning it for a new task



Pruning & quantization

Removing lower importance parameters; lower numerical precision



Smaller task-specific models

Designing focused models rather than general purpose

In our workflow

Reduce redundant training

Improve efficiency of data centres

Avoid unnecessary uses

In AI Governance

Policy	Institutions	Technical
• Data center regulation • Water regulation • Responsibility for e-waste	• Procurement standards • Managing compute • Task refinement	• Model size limits • Carbon reporting • Transfer learning default

Unclassified / Non classifié

Requirements for successful deployment in public health

Aligned with Public Health Goals

- Fewer people develop diseases or poor health outcomes that are preventable
- Promoting health in our environments and communities
- Reduction in health inequities
- Sustainable

Effective

- Independently tested and verified
- Consistent and accurate in real-world settings
- Scalable across different populations
- Oversight and responsive to change/model updates

Safe & Ethical

- Data Privacy & Security
- Human oversight on decisions
- Legal and regulatory compliance
- Avoids bias, discrimination; equitable access

Unclassified / Non classifié

Six Priority Areas for Public Health Organizations to Advance AI into Practice*

Governance

Infrastructure

Including secure systems that support integration, testing, maintenance, and monitoring**

Workforce

Partnerships

Reviewers need the expertise, time, and authority to question outputs and act on problems**

Good AI Practices

Equity and Bias

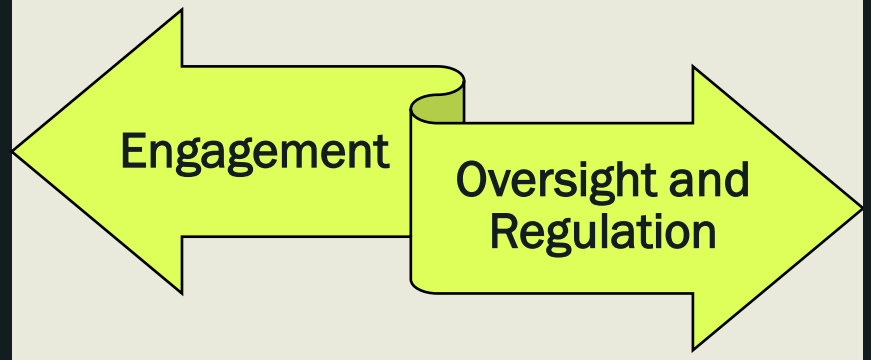
*Fisher S, Rosella LC. BMC Public Health. 2022;22:2146

**Hattab G, et al. AJPH 2025;115(2):123–128.

Building Responsible AI for Health

Centering people, practitioners & operators

- All stages of problem identification, design, implementation, evaluation
- Address real public health and community challenges
- Based on authentic relationships

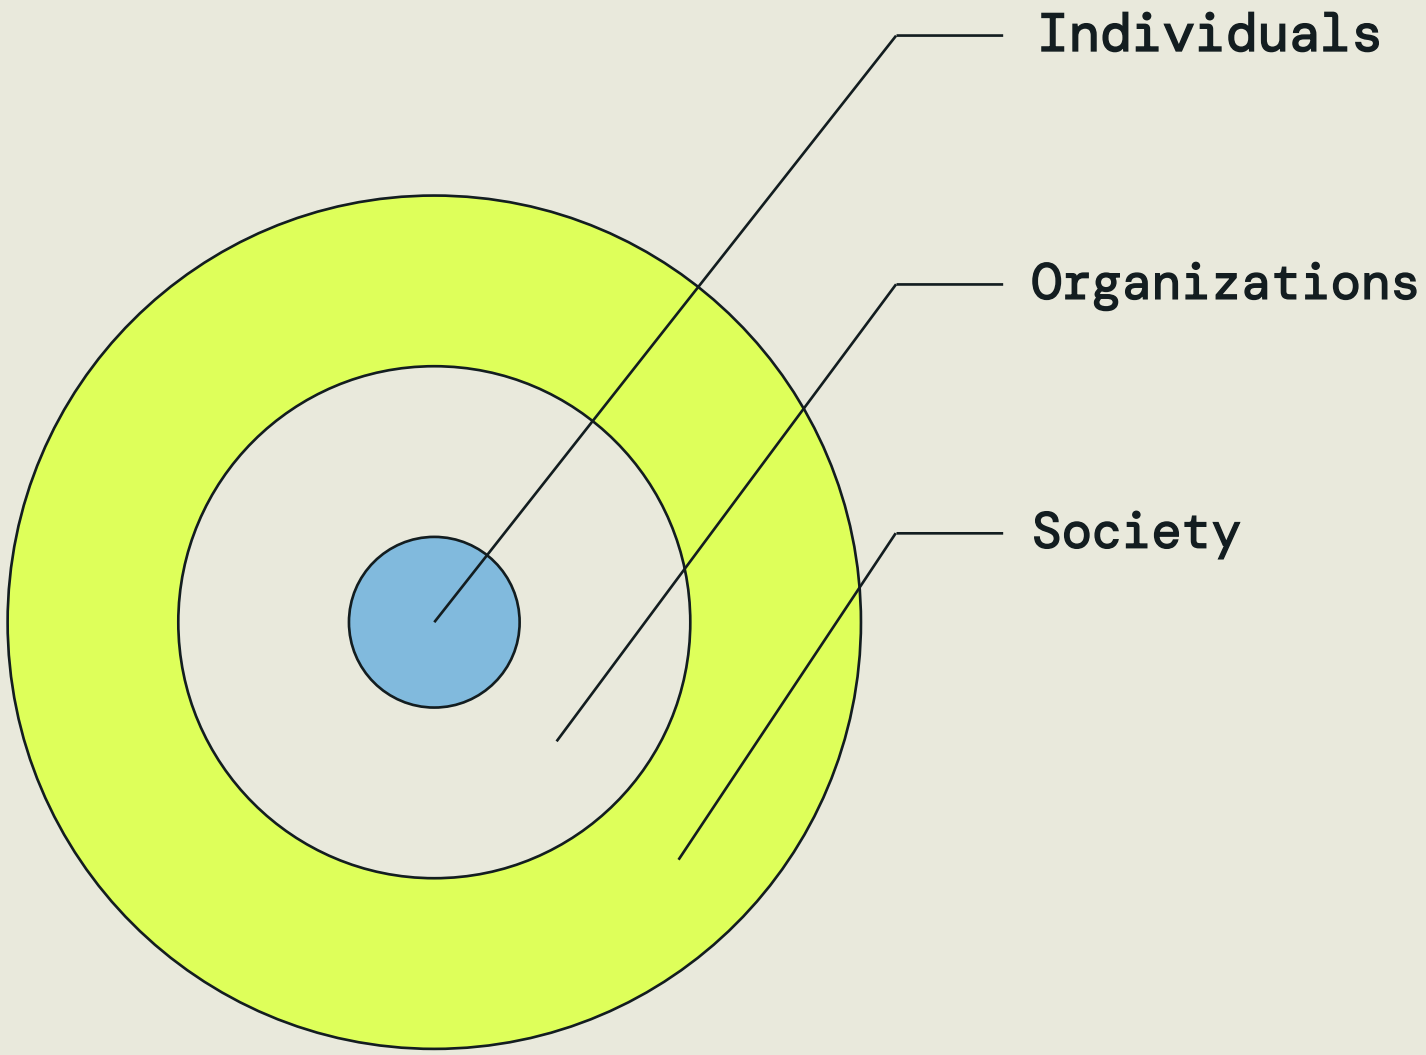


Regulation: Formal rules or laws created by a government or regulatory authority

Guidelines: Guidance or Strong Recommendations without enforcement

Standards: Standards are formal documents that establish specific criteria and requirements; more detailed than guidelines

How do we get there?

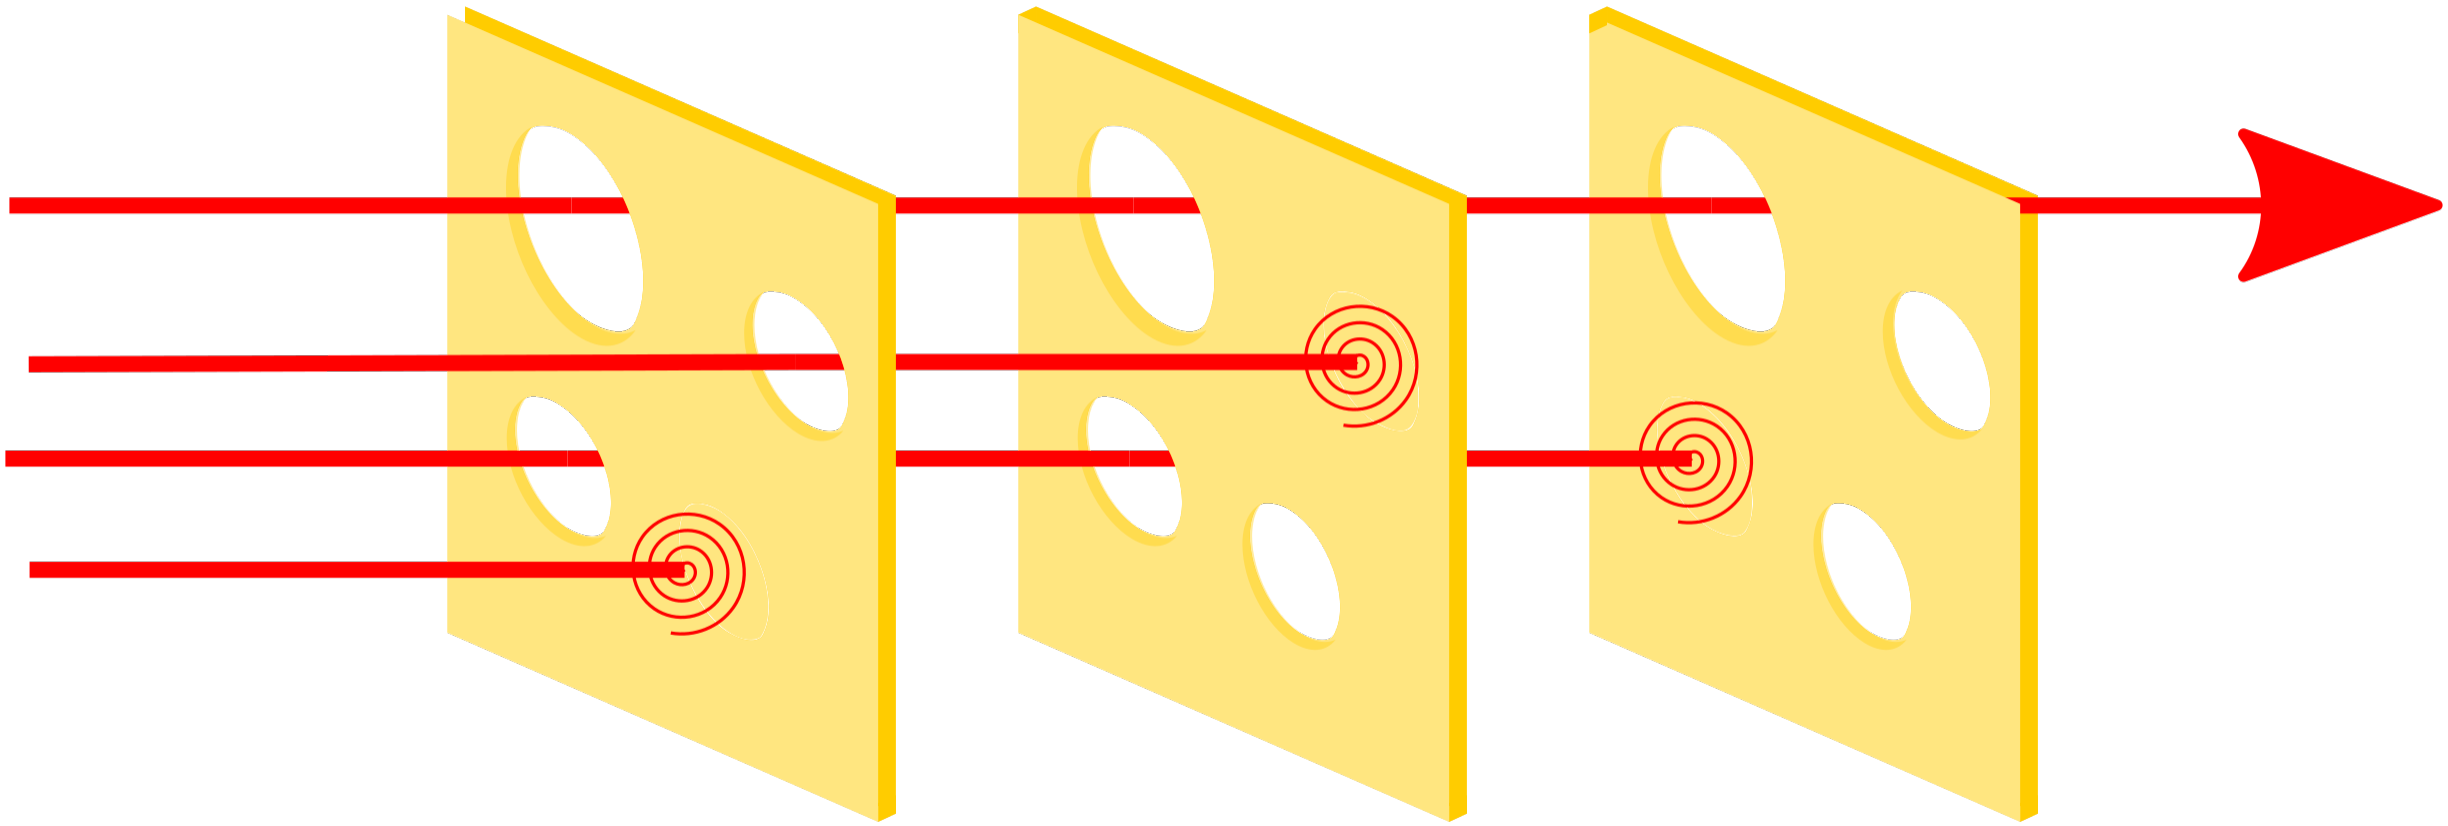


Foundations · Applications · Impacts · **Responsible Implementation**

Regulation

Guidelines/Standards

Monitoring and Oversight



Foundations · Applications · Impacts · **Responsible Implementation**

Image Adapted from BenAveling, Wikimedia Commons, CC BY-SA 4.0a

WHO: six principles for AI in health

1 Protect autonomy

People remain in control of health decisions; privacy and consent are respected

2 Promote well-being, safety and the public interest

Tools meet requirements for safety, accuracy and effectiveness, and prevent harm

3 Ensure transparency and explainability

Enough is disclosed for developers, users and the public to understand and question use

4 Foster responsibility and accountability

Clear responsibility for outcomes, with ways to challenge decisions and seek redress

5 Ensure inclusiveness and equity

Designed for equitable use and access, regardless of age, sex, income, race, ability or other characteristics

6 Promote responsive and sustainable AI

Continuously assessed in real use, with attention to environmental and workforce impacts

World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021.

Coming together to advance artificial intelligence for diabetes prediction and prevention



www.instituteforbetterhealth.com/ai4dpp-meeting/



cifar.ca/ai/cifar-solution-networks/ai-for-diabetes-prediction-and-prevention/

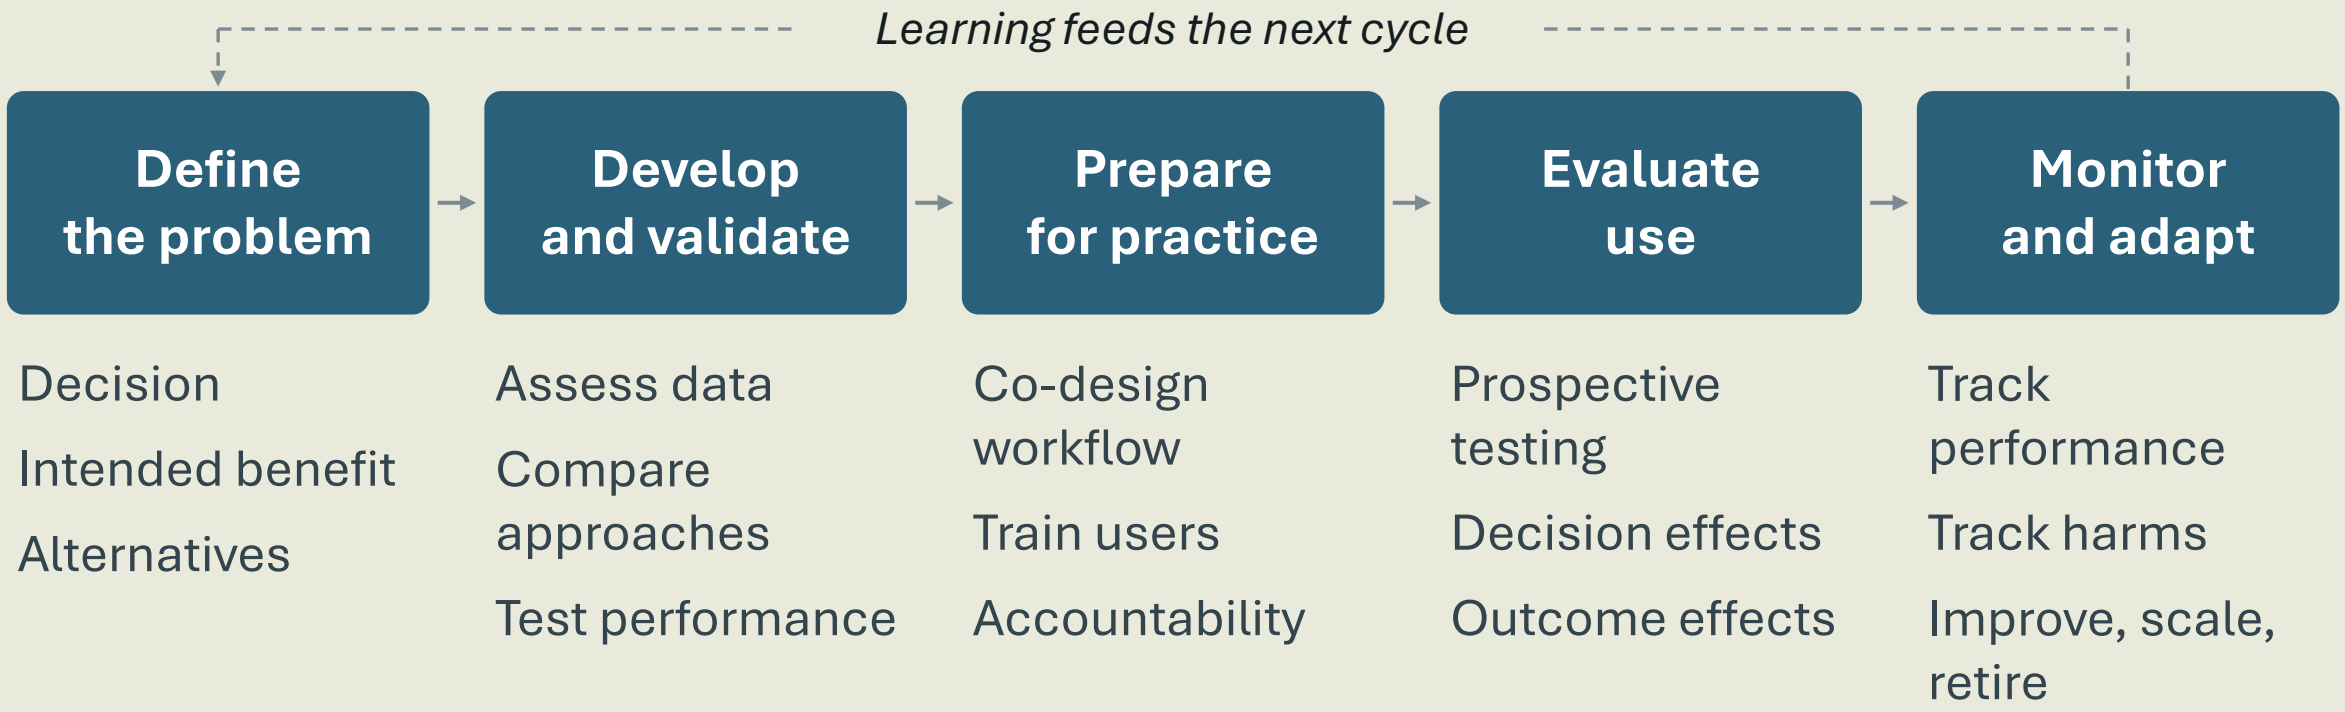
- **Task:** Develop a useful and acceptable approach to diabetes prediction and prevention with practitioners and community partners.
- **What AI adds:** Risk estimates to help prioritize prevention efforts.
- **What participation adds:** Decisions about which outputs are useful, how they should connect to services, how they are communicated and what responsible use requires.

1. Come together, learn, and share ideas around AI and health
2. Co-design AI deployment with practitioners and community
3. Inform AI outputs for diabetes in Peel Region

Abejirinde IO et al J Particip Med. 2025 Oct 2;17.

Rosella LC, et al. *DIGITAL HEALTH*. 2025;11

The lifecycle of responsible deployment



Community participation, equity, and governance continue throughout

What should we evaluate?

Dimension	Key question
Performance	Does the tool work reliably in the intended population and setting?
Practice	Does it improve decisions or reduce workload after accounting for review and correction?
Public health benefit	Does its use improve prevention, timeliness, or health outcomes?
Equity	Who benefits, who is missed, and who experiences harm?
Sustainability	Are the benefits worth the ongoing financial, workforce, and environmental costs?

Rosella LC. Will AI improve public health? Setting the standards for meaningful impact. *CJPH.* 2025;116:815–818.

We have a role in shaping what comes next

Foundations

Understand AI's capabilities… and its limits

Applications

Start with public health goals & functions and ask what AI adds

Impacts

Make protecting people, equity, and the environment part of progress

Responsible implementation

Turn promising tools into lasting public health benefit

Thank you!

Laura C. Rosella, PhD
laura.rosella@utoronto.ca
@LauraCRosella
www.linkedin.com/in/laura-rosella-b68b5b91

UNIVERSITY OF TORONTO
DALLA LANA SCHOOL OF PUBLIC HEALTH

Population Health Analytics Lab
www.pophealthanalytics.com
@PopHlthAnalytic

Population Health Analytics Laboratory

AI4PH AI FOR PUBLIC HEALTH

https://ai4ph-hrtp.ca/
@ai4ph
/ai4ph-health-research-training-platform

Unclassified / Non classifié

AI for Public Health (AI4PH)

Mission & Approach

MISSION

To establish a workforce of public health researchers and practitioners who develop and apply innovations in equitable artificial intelligence (AI) and machine learning (ML) to public health research, policy, and practice in Canada.

APPROACH

Develop a suite of training activities and opportunities dedicated to facilitating the application of equitable AI applied to public health research, policy, and practice.

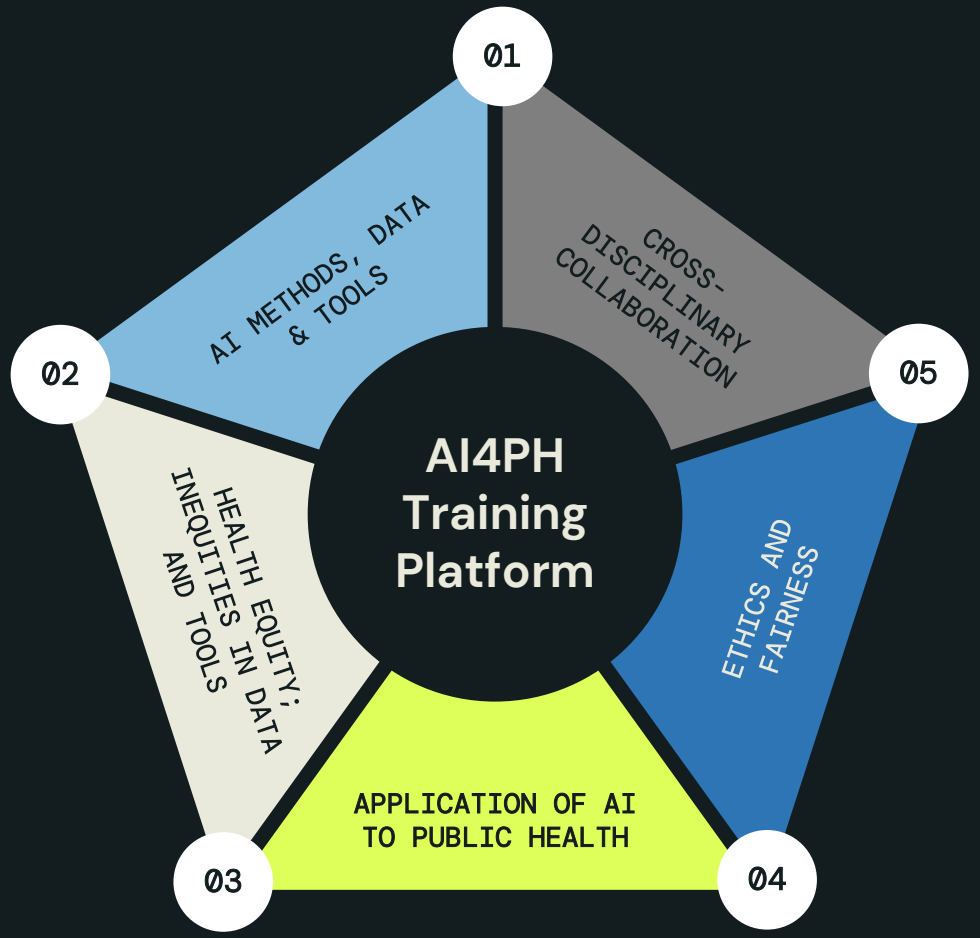
Our platform will focus on skill development and capacity-building to apply AI and ML innovations to Canada's most pressing public health challenges that address the social, cultural, and environmental determinants of health.

CIHR IRSC Canadian Institutes of Health Research / Instituts de recherche en santé du Canada

AI4PH is supported by the Canadian Institutes of Health Research (CIHR)

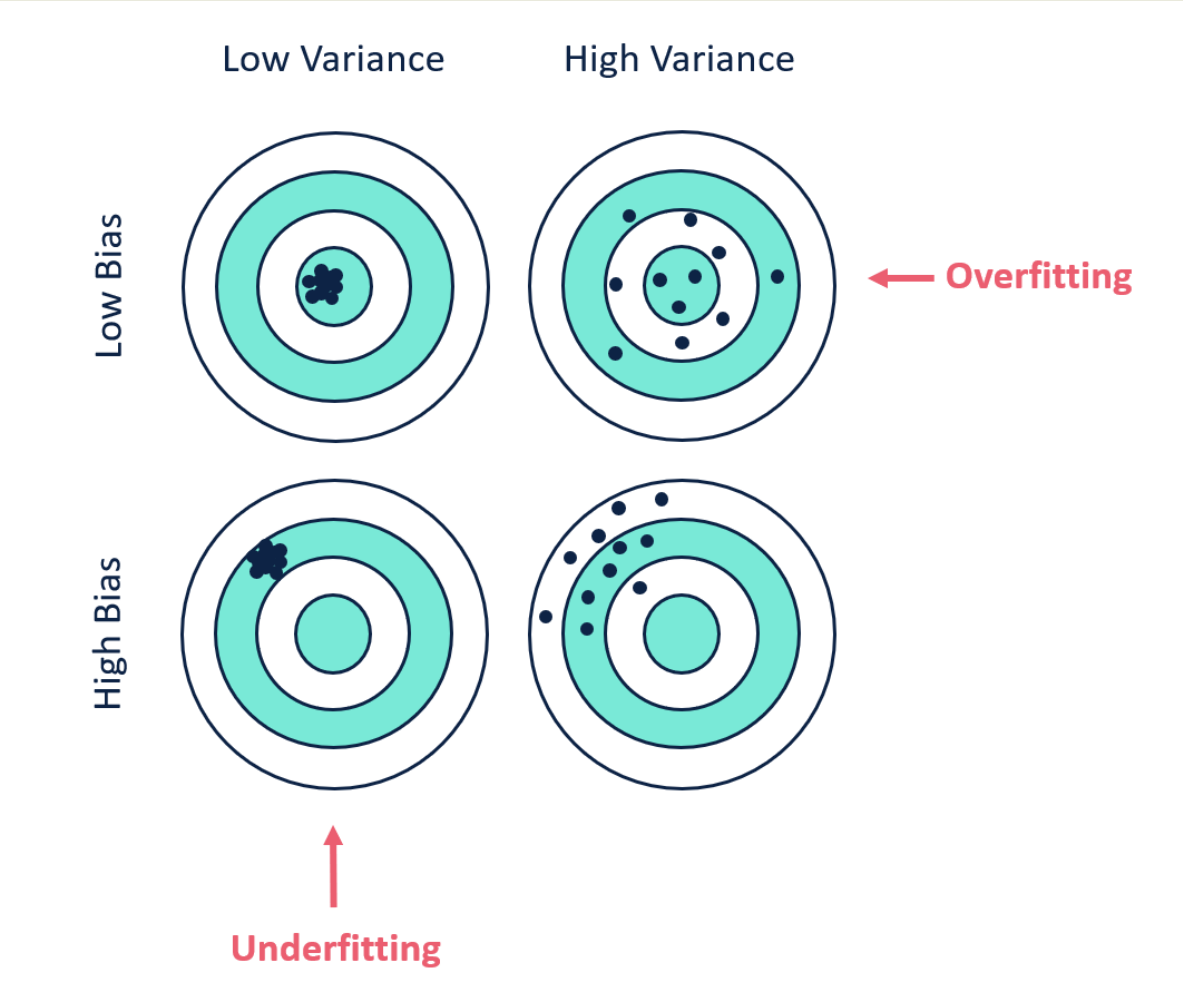
Unclassified / Non classifié

AI FOR PUBLIC HEALTH (AI4PH) TRAINING PLATFORM



- **Short courses**
- **Scholarships**
- **Internships**
- **Community building**
- **Workshops**

Bias-Variance Trade-Off



- Model with high variance means predictions are sensitive to the training sample and may reflect noise
- Bias is when a model is too simple and missing the mark repeatedly
- Known as a trade-off because you need to optimize to get the balance

E.g. simpler model will have lower variance but if you simplify too much you will increase bias

Images:
https://www.javatpoint.com/bias-and-variance-in-machine-learning

Lost in Translation?

Computer Science Term	Epidemiology Term
Features	Independent variables
Labels	Outcomes
Noisy labels	Measurement error
Learning	Fitting
Data mining	Exploratory analyses
Classification algorithm	Algorithm with a categorical outcome
Regression algorithm	Algorithm with a continuous outcome
Precision	Positive predictive values
Recall	Sensitivity
Dimensionality	Number of covariates
Imbalanced data	Unequal outcome distribution
One-hot encoding	Creation of dummy variables

Qifang et al. What is machine learning? A primer for the epidemiologist. AJE, 2019.