ML Essentials: supervised v. semi-supervised v. unsupervised learning

- **Semi-supervised**: a mix between supervised and unsupervised learning

  Classic **examples**

  – Positive & unlabeled
    - Only **green** instances labeled
    - Algorithm adapts iteratively

  – Role of 'learning' objective
    - Entropy v. other criteria

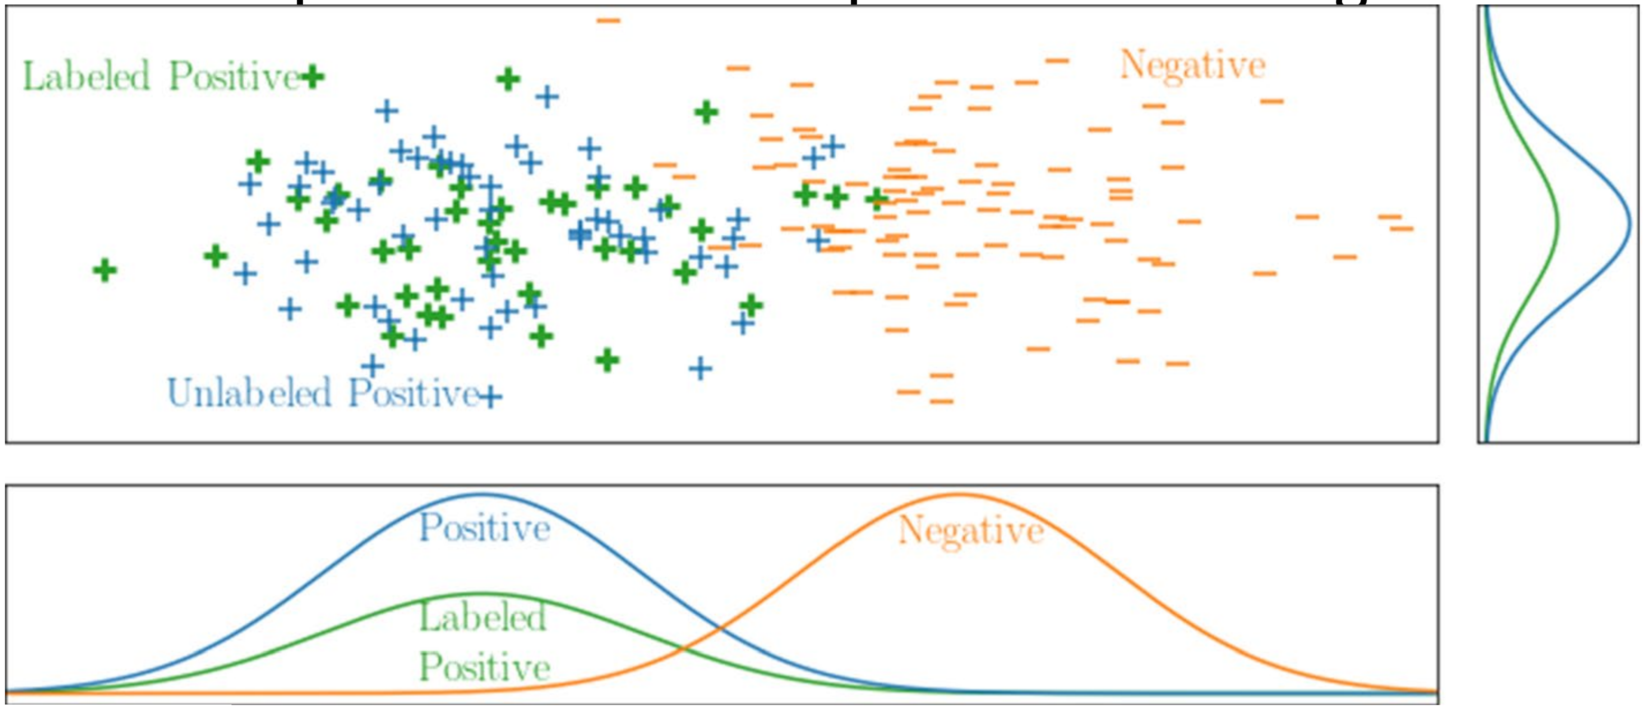

Bekker, J., Davis, J. Learning from positive and unlabeled data: a survey. *Mach Learn* **109**, 719–760 (2020). https://doi.org/10.1007/s10994-020-05877-5

- Survey of DL use cases

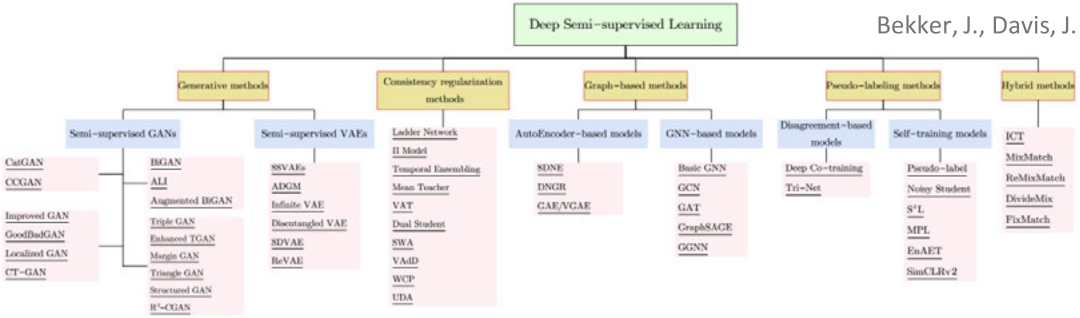

Fig. 1. The taxonomy of major deep semi-supervised learning methods based on loss function and model design.

ML Essentials: supervised v. semi-supervised v. unsupervised learning

- From Semi-supervised to ***Reinforcement Learning*** (frequently uses *Q-learning)*
  - *Particularly useful over time*
  - *suited to decision sequences*
  - *Caveats in health settings,*
    - *Nature editorial poses challenges*
    - *Example at right: intensive care*

To perform sequential decision making, such as for sepsis management, treatment-effect estimation must be solved at a grand scale—every possible combination of interventions could be considered to find an optimal treatment policy. The diagram shows the scale of such a problem with only three distinct decisions. **Blue** and **red** people denote positive and negative outcomes, respectively.

Reinforcement learning is a type of machine learning that focuses on training AI agents to make a sequence of decisions to maximize a cumulative reward. It's used in gaming, robotics, and autonomous systems.

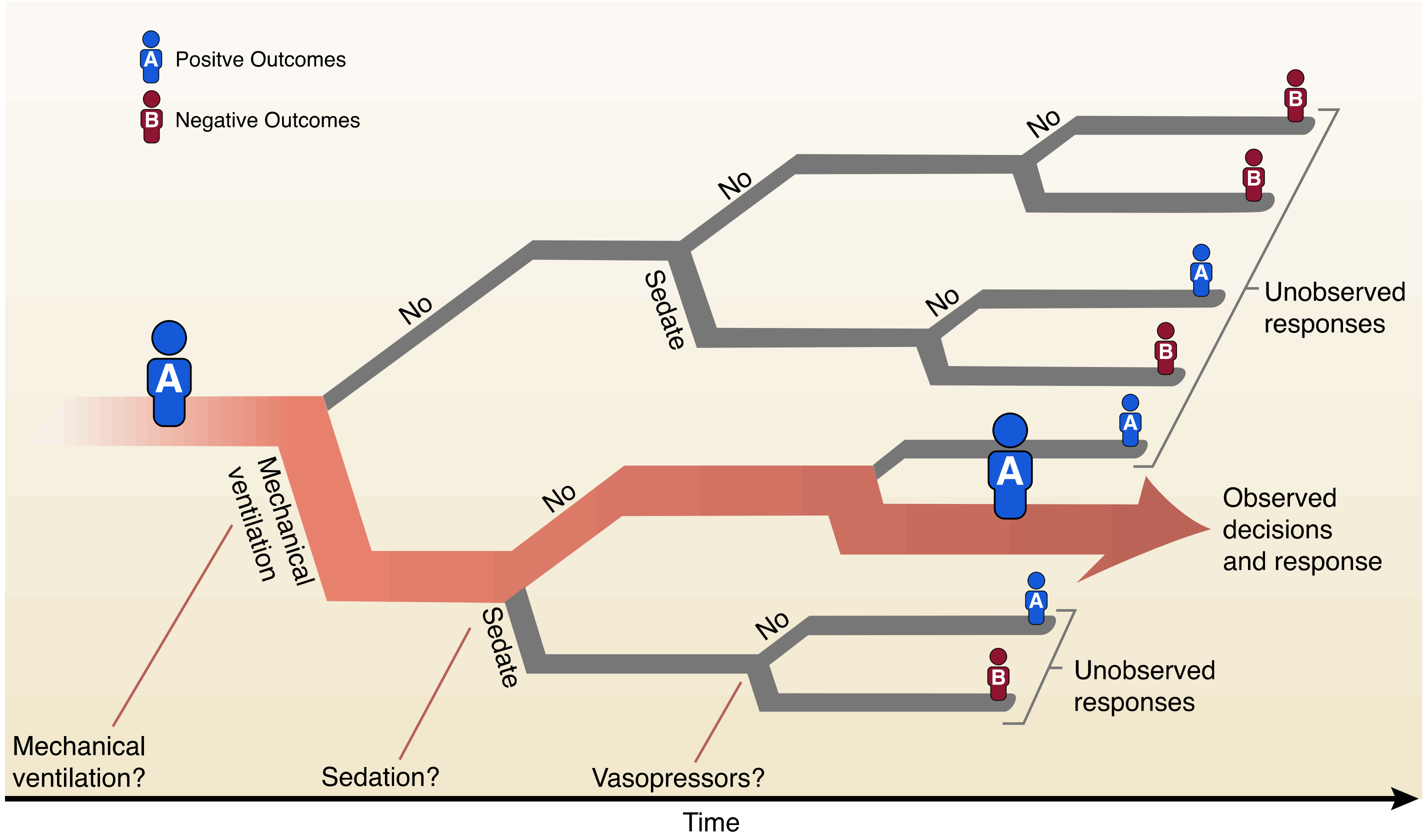

ML Essentials: supervised v. semi-supervised v. unsupervised learning

- From Semi-supervised to ***Reinforcement Learning*** (frequently uses *Q-learning)*
  - *Particularly reliant on BIG data*
  - *Need cases along all sequences*
  - *Caveats in health settings,*
    - *Nature editorial shows challenge*
    - *Figure @right: effective sample size*

Reinforcement learning is a type of machine learning that focuses on training AI agents to make a sequence of decisions to maximize a cumulative reward. It's used in gaming, robotics, and autonomous systems.

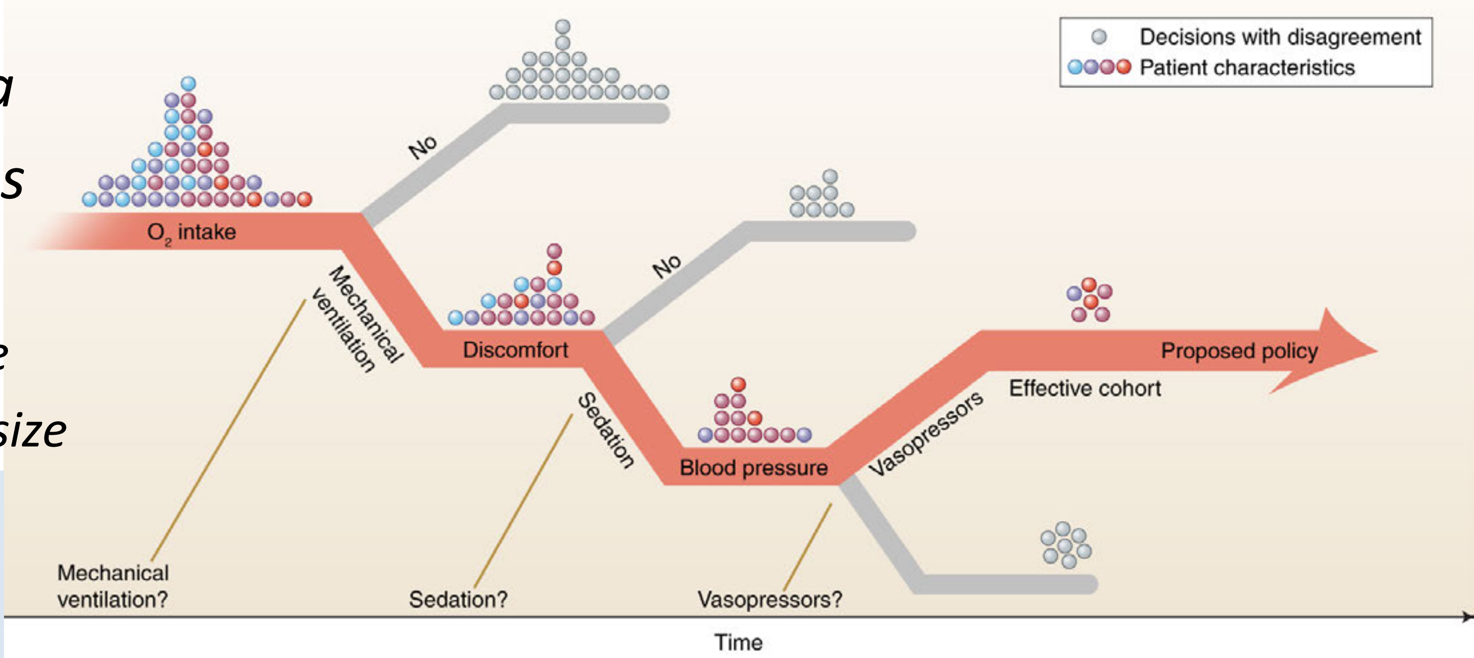

Each dot represents a single patient at each stage of treatment, and its color (gradation from **blue**↔**red**) indicates the patient's characteristics. The more decisions that are performed in sequence, the likelier it is that a new policy disagrees with the one that was learned from. **Gray** decision points indicate disagreement. Use of only samples for which the old policy agrees with the new results in a small effective sample size and a biased cohort, as illustrated by the difference in color distribution in the original and final cohort.

Gottesman, O., Johansson, F., Komorowski, M. *et al.* Guidelines for reinforcement learning in healthcare. *Nat Med* **25**, 16–18 (2019). https://doi.org/10.1038/s41591-018-0310-5

Machine Learning Essentials: concept check

- We now engage participants to check our mutual understanding:

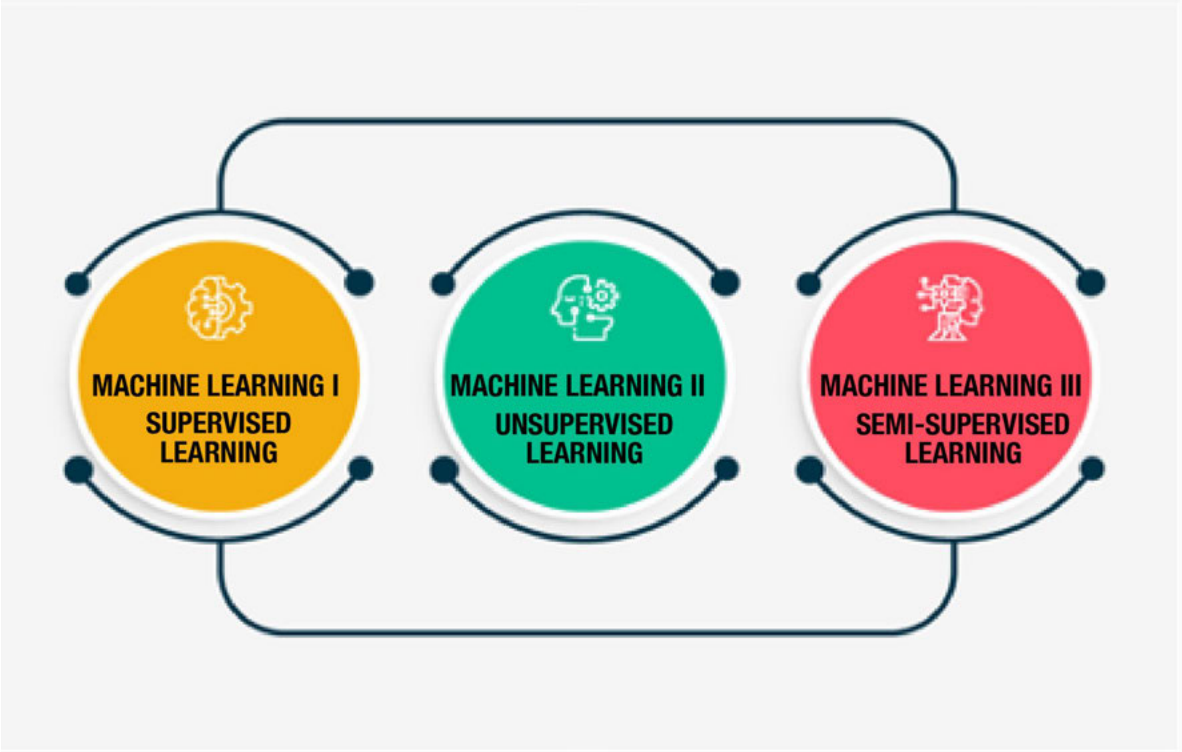

[recall sli.do questions re: supervised v. unsupervised v. semisupervised]

Machine Learning Computational Strategies

- We now provide detailed explanations and use cases for ML strategies, which can improve upon traditional/modern stats / epi data methods.
  - **Example:** See differences in race-specific v. race-agnostic for machine learning predicted in-hospital mortality…
  - either improved on logistic regression

- Detailed Examples of ML computational strategies used in healthcare disparities research (the list of examples to follow is not exhaustive)

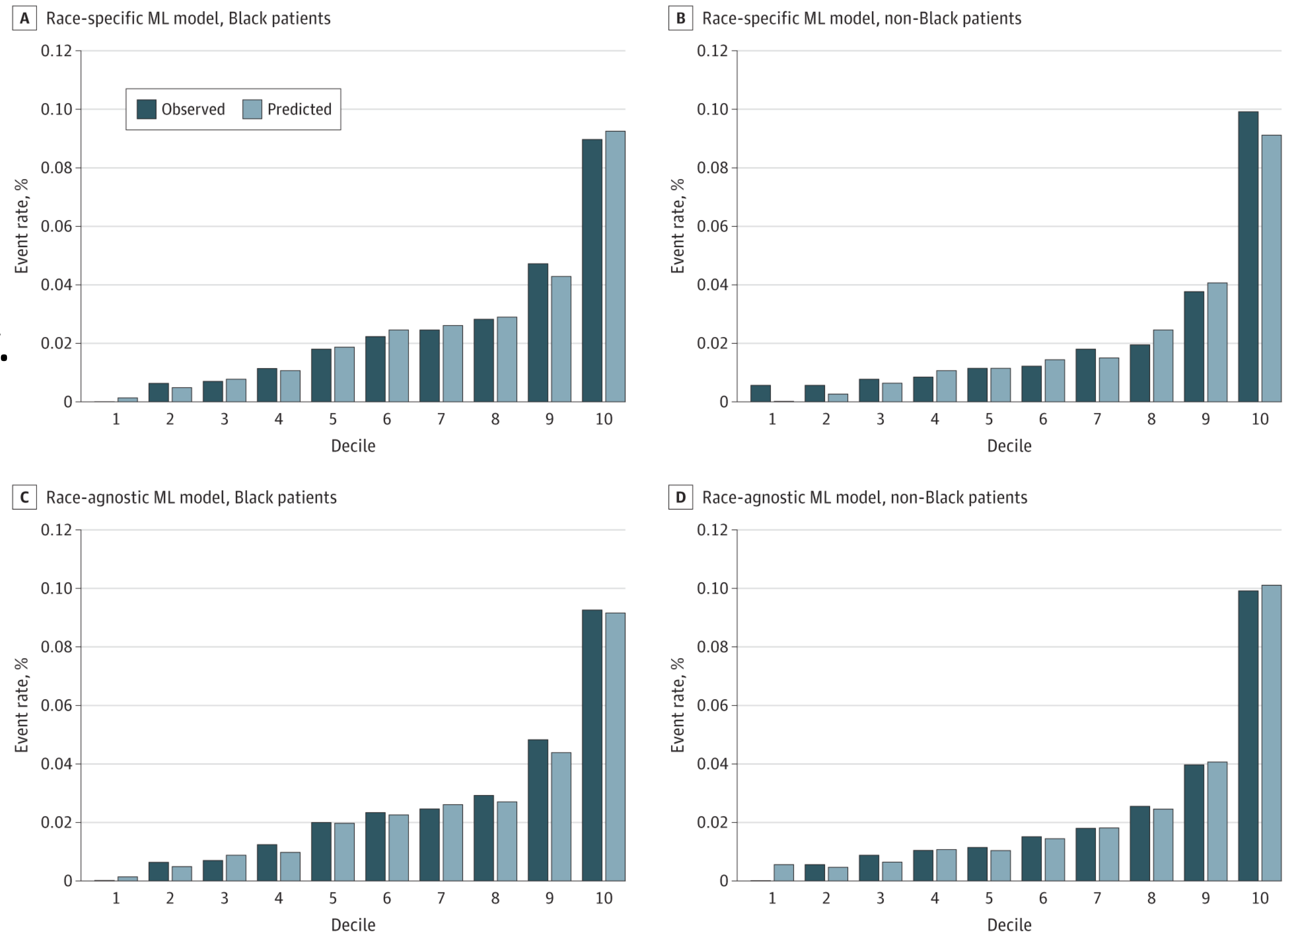

https://jamanetwork.com/journals/jamacardiology/fullarticle/2793728

Machine Learning Computational Strategies

1. Predictive Modeling for Patient Outcomes:

- a. Strategy: Using machine learning algorithms to predict patient outcomes.
- b. Application: Identifying high-risk populations for specific diseases [**examples**].
- c. Python Libraries: Scikit-learn, TensorFlow, PyTorch.

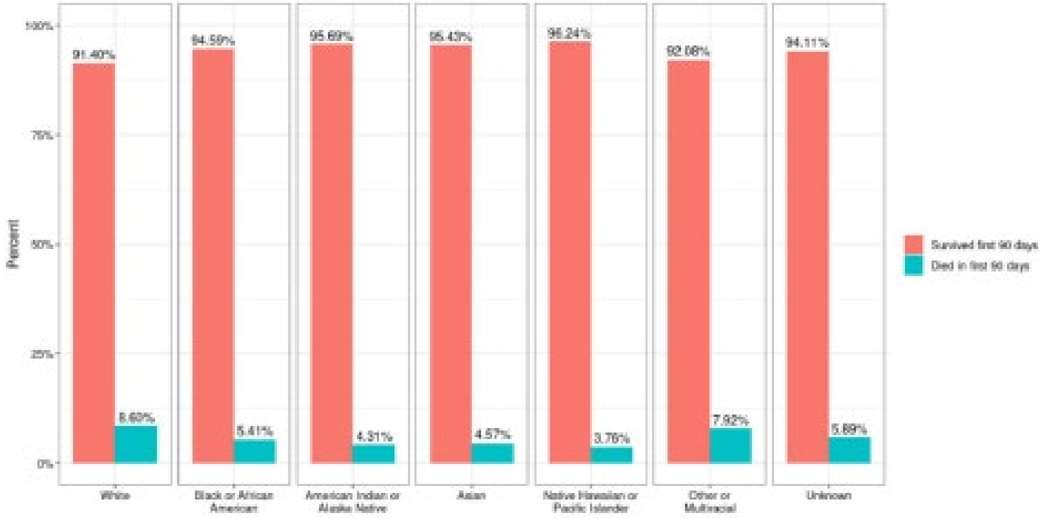

Box plot showing distribution of age at first dialysis treatment for each race

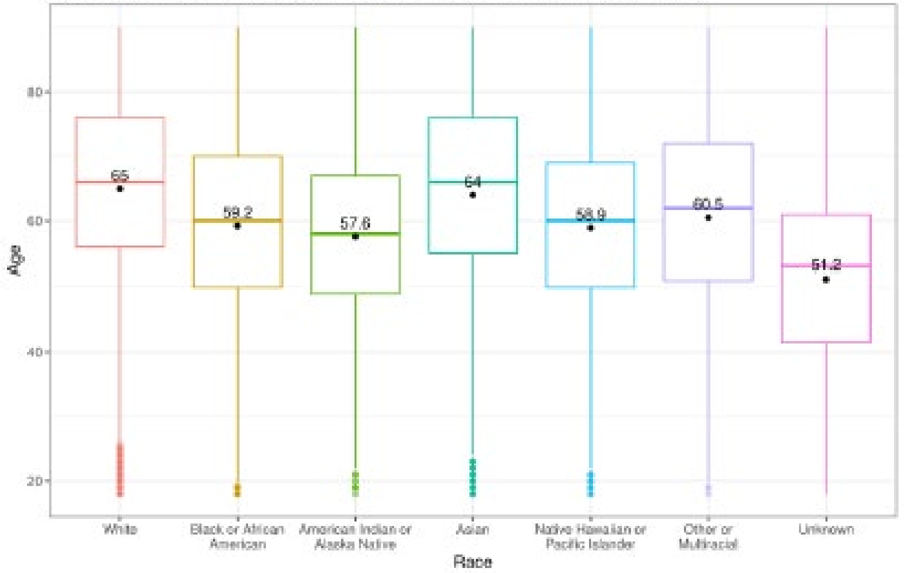

Machine Learning Computational Strategies

2. Image Analysis for Diagnostics:

- a. Strategy: Applying computer vision and deep learning.
- b. Application: Improving diagnostic accuracy from medical images [breast density **example above**; melanoma w/o regard to skin color ***counter*-example** @right]
- c. Python Libraries: TensorFlow, PyTorch, OpenCV.

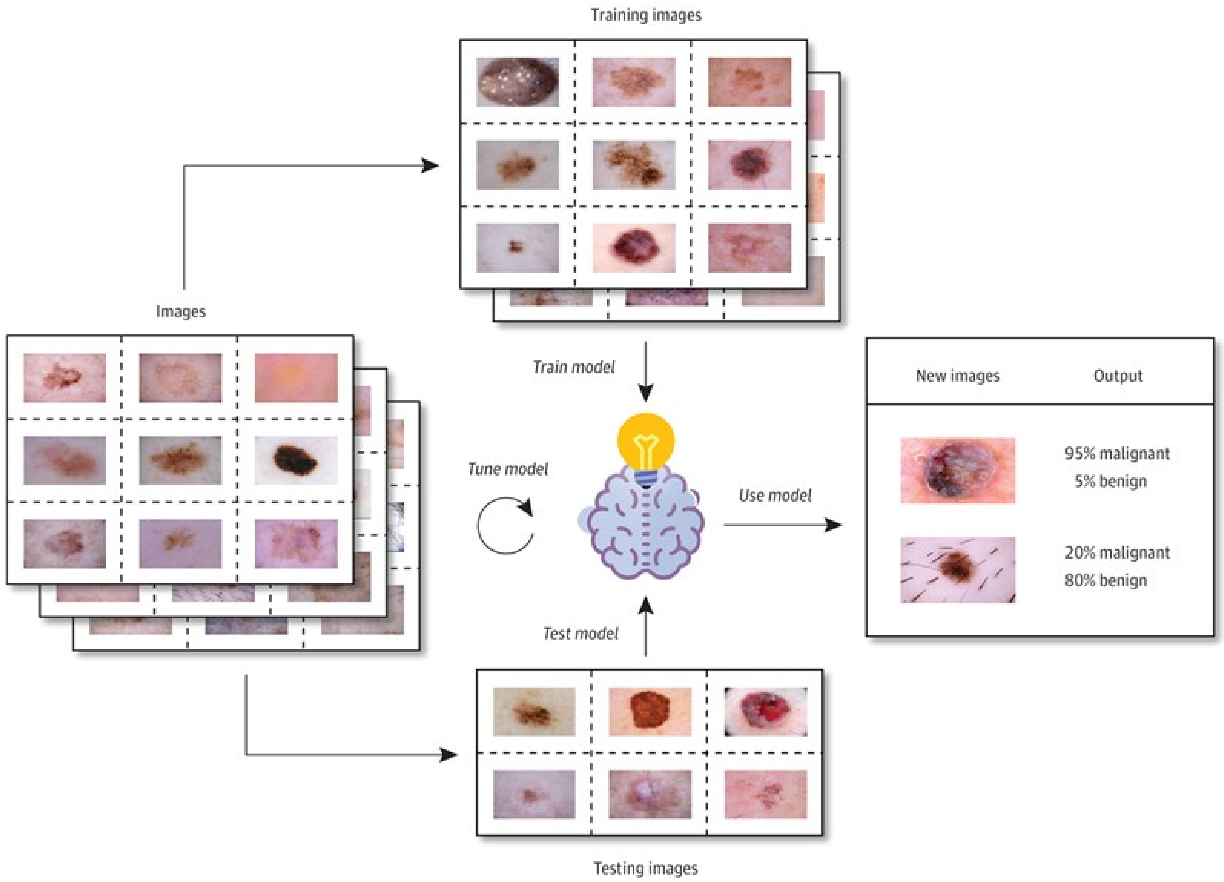

Images are collected of pigmented lesions and split into a larger training image set and a smaller testing image set. The machine learning algorithm (center) uses the training images to learn how to correctly categorize pigmented lesions based on their visual features. The model is then tested with the testing images set to determine model accuracy. The algorithm model is fine-tuned with more training and testing images. Once the machine learning algorithm is developed, it can be used on new images. The output gives an estimate of the likelihood of a given result.

https://jamanetwork.com/journals/jamadermatology/fullarticle/2688587

Machine Learning Computational Strategies

4. Remote Patient Monitoring:

- a. Strategy: Using AI to analyze data from wearable devices.
- b. Application: Monitoring patient health in real-time **examples** [e.g., continuous glucose monitoring, or CGM for Active Insulin Dosing, AID]
- c. Python Libraries: TensorFlow, scikit-learn.

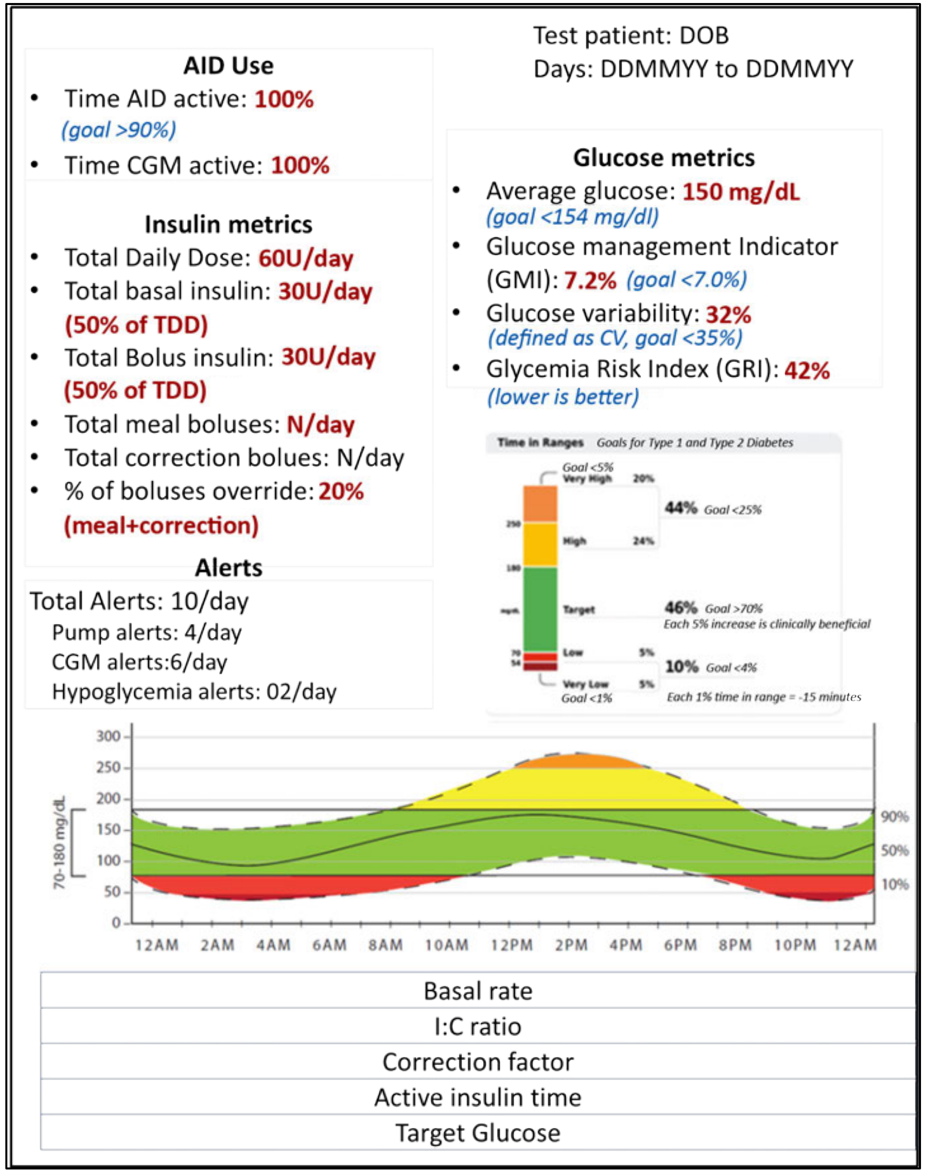

Machine Learning Computational Strategies

5. Population Health Management:

- a. Strategy: Employing machine learning algorithms for population-level health data.
- b. Application: Identifying disparities in health outcomes.
- c. Python Libraries: Scikit-learn, TensorFlow, PyTorch.

A counter-example: ***mentioned in last Think-a-thon***

An algorithm used to predict which patients would benefit from extra medical care **flagged healthier white patients as more at risk than sicker black patients**

- An analysis on 3.7 million patients found that **black patients ranked as equally as in need of extra care** as white patients collectively suffered from 48,772 additional chronic diseases
- The bias was discovered when researchers from a health system in Massachusetts found the **highest scores in their patient population concentrated in the most affluent suburbs of Boston**

Example: **Researchers tweaked the algorithm** to make predictions about their future health conditions

- The tweak increased the percentage of black patients receiving additional help from 17.7 to 46.5%

Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-453. doi:10.1126/science.aax2342

Machine Learning Computational Strategies

6. Social Determinants of Health (SDOH) Analysis:

- a. Strategy: Integrating AI to analyze social, economic, and environmental factors.
- b. Application: Understanding the impact of social determinants on healthcare disparities. **example**
- c. Python Libraries: Scikit-learn, pandas, NumPy.

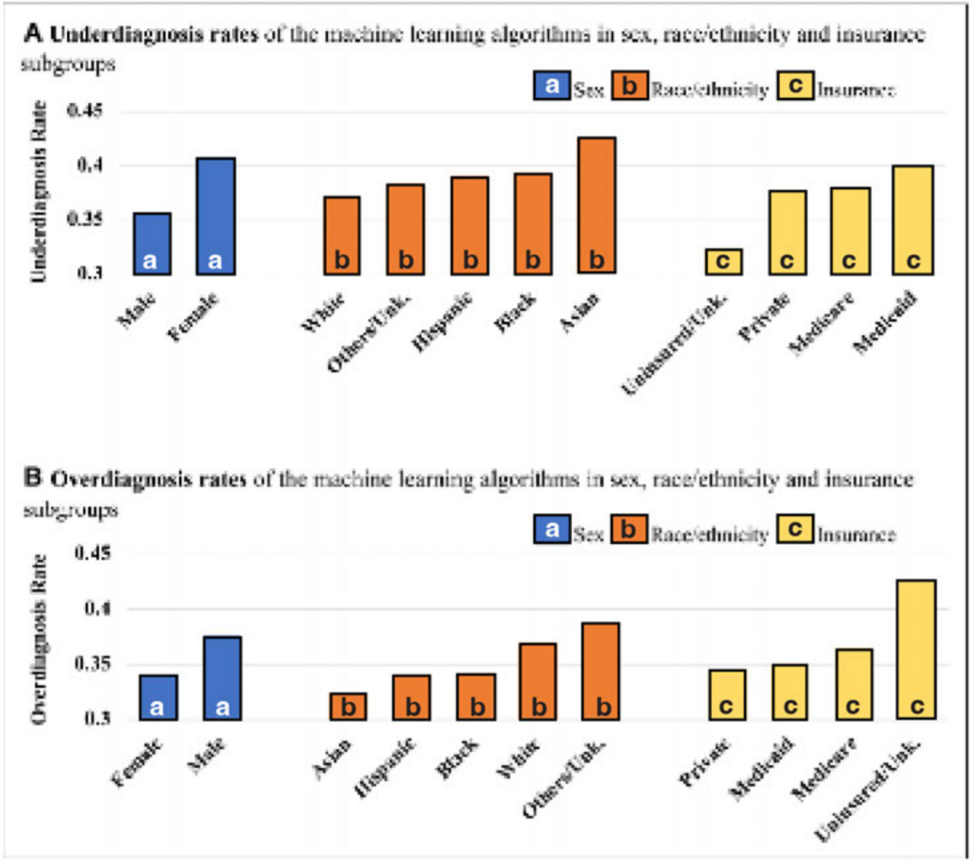

Figure. Underdiagnosis (false negative rate) and overdiagnosis (false positive rate) rates in each sex, ethnoracial, and insurance subgroup, when using random forest classifier to predict the composite heart failure outcome. The model achieves the highest performance and fairness scores. Unk indicates unknown.

b. Application example by Luo’s team: Social Deprivation Index (SDI) & Area Deprivation Index (ADI) at both state and national levels) can *somewhat* mitigate the Figure-noted heart failure risk disparities

https://news.feinberg.northwestern.edu/2022/12/15/investigating-disparities-in-machine-learning-algorithms/

Machine Learning Computational Strategies

6. Social Determinants of Health (SDOH) Analysis:

- b. Application: Understanding the impact of social determinants on healthcare disparities… can be *less often considered sources for SDOH,* if the use case points to a need
- b. **Example**: stark climate-change related vulnerabilities, like flooding

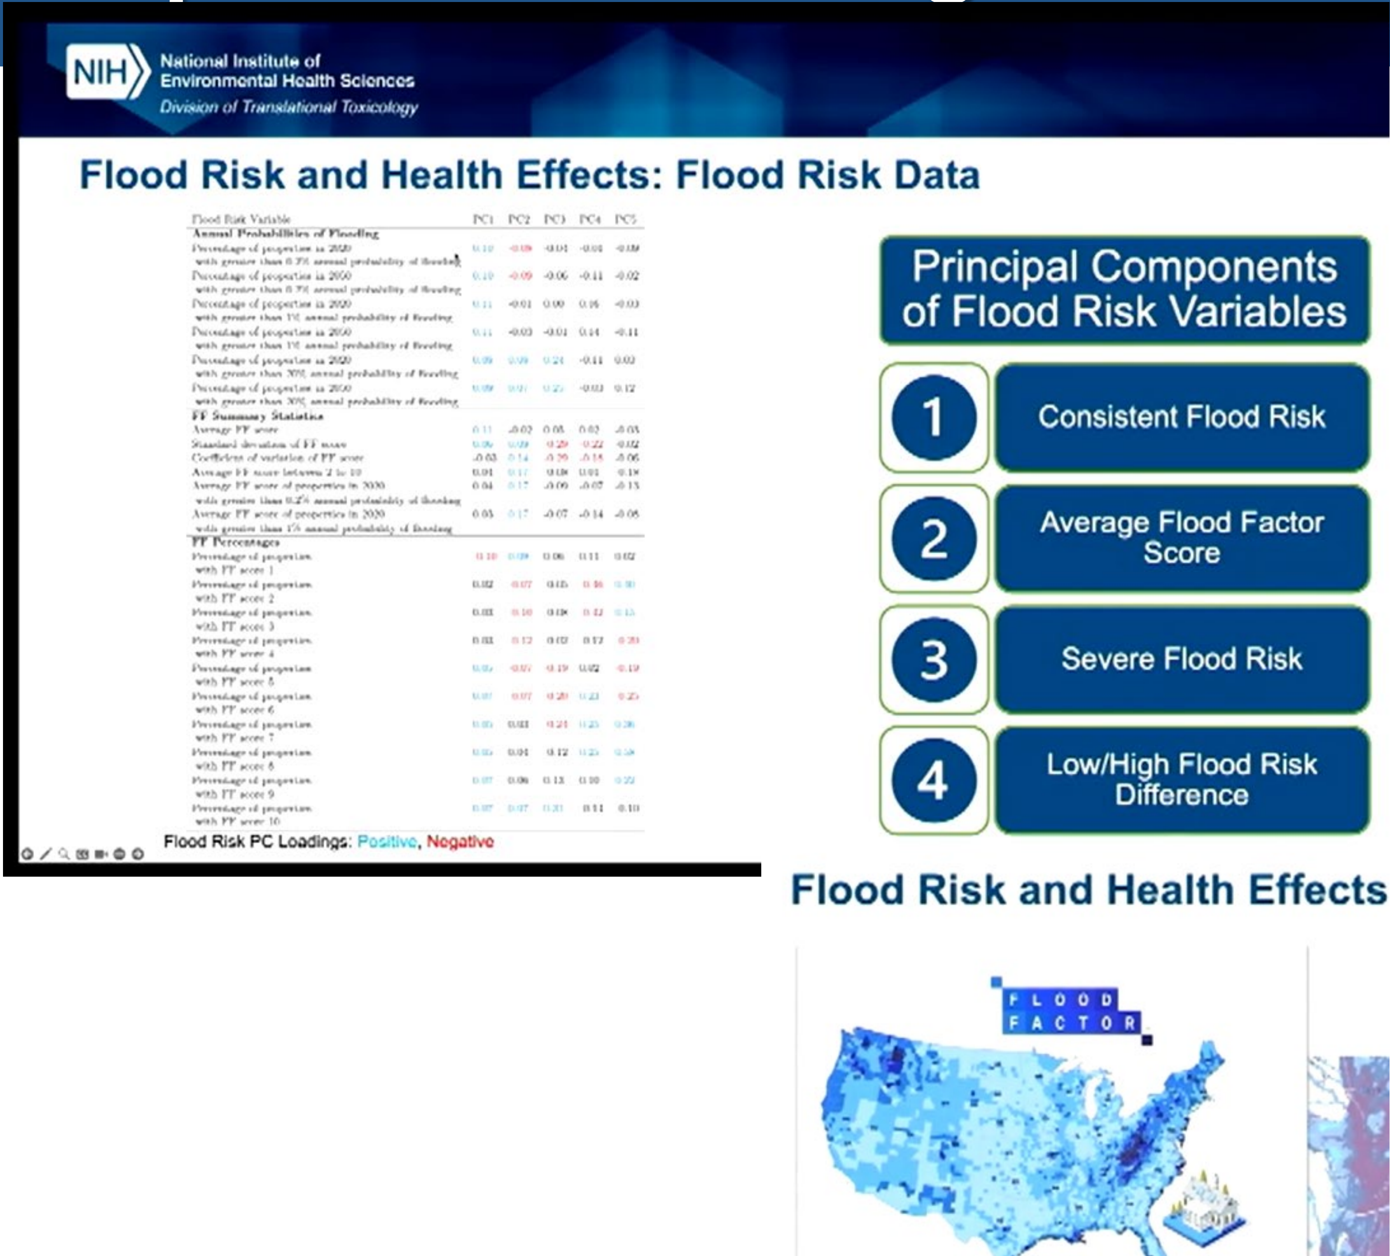

b. Application example by NIEHS/NIMHD PI Messier's SET group: used First Street Foundation's Flood measures panel at granular area levels -- can *somewhat* mitigate the noted__ risk disparities

https://videocast.nih.gov/watch=53935

Machine Learning Computational Strategies

7. Ethical AI for Bias Mitigation:

- a. Strategy: Implementing fairness-aware and explainable AI models.
- b. Application: Ensuring AI systems do not perpetuate biases.
- c. Python Libraries: AIF360, Fairness Indicators ([Caton & Haas review]()), AI Fairness 360
  - NB: includes a [scikit-learn compatible Application-Programmer Interface]() (API)!

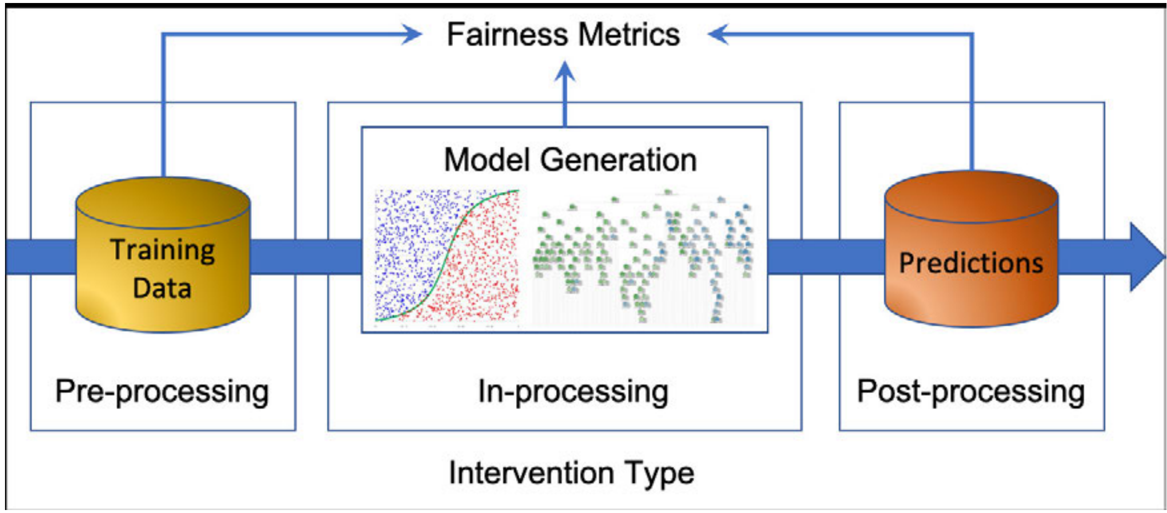

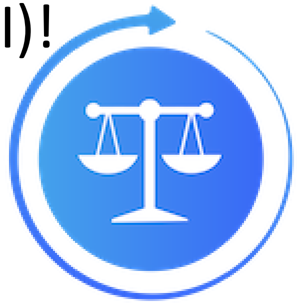

Machine Learning Computational Strategies

7. Ethical AI for Bias Mitigation:

- b. Application: Ensuring AI systems do **not** perpetuate biases… may be *most tractable* by applying Caton&Haas framework
  - Pre-processing
  - *IN-processing*
  - Post-processing: helpful capacity to apply to *any* data science workflow

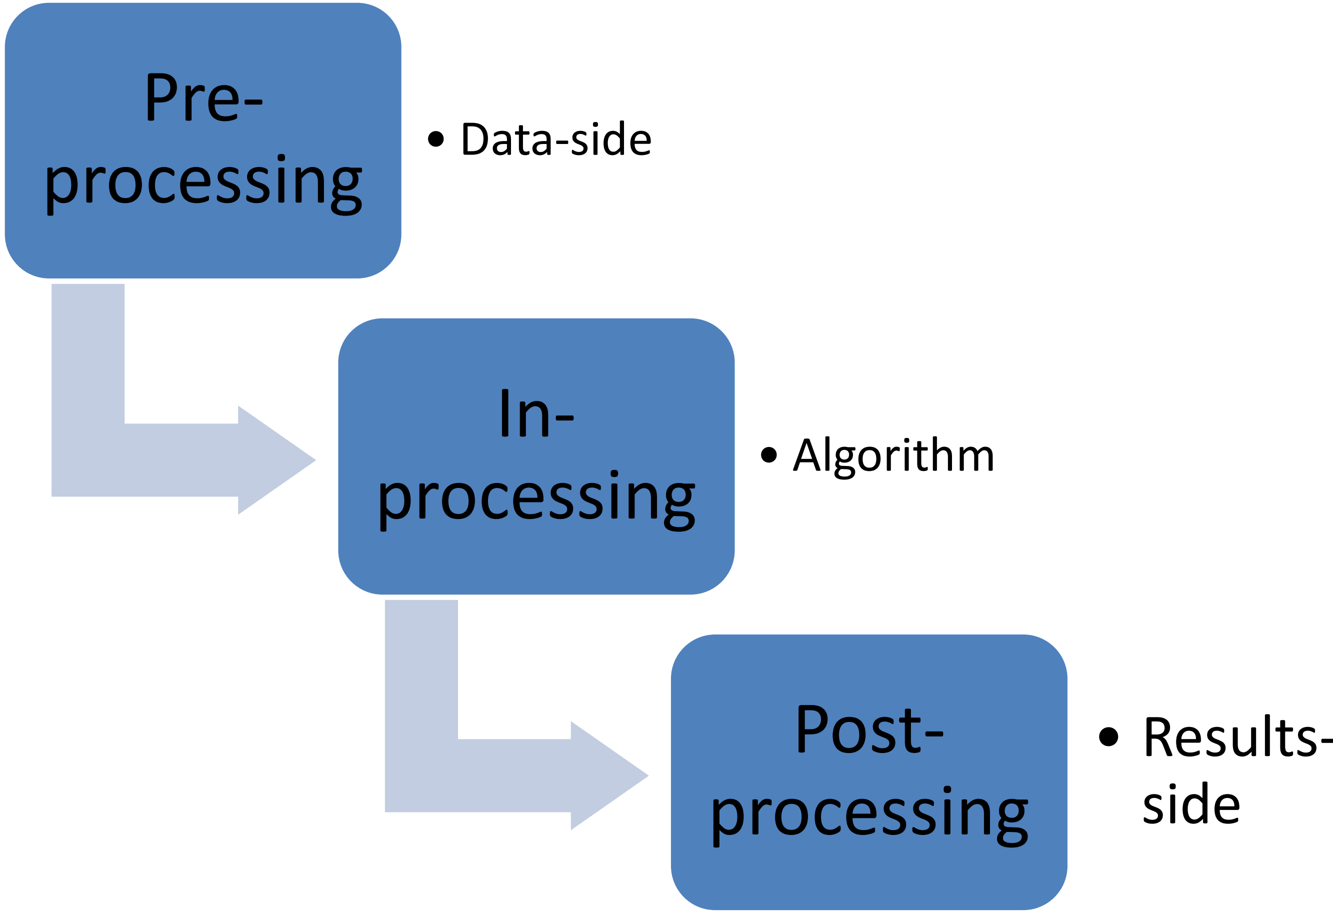

Machine Learning Computational Strategies

7. Ethical AI for Bias Mitigation:

- b. Application: **example** of applying [Caton&Haas framework]
  - **Post-processing**: helpful capacity to apply to *any* data science workflow

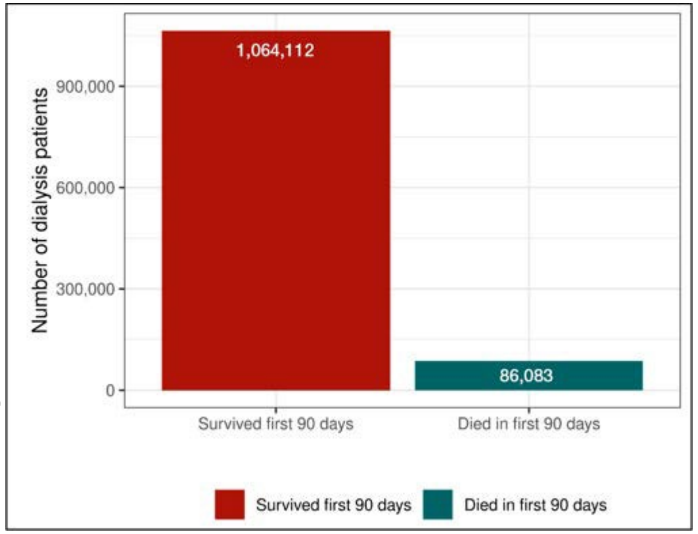

	Feature	Value	Count	AUC	TN	FP	FN	TP
0	agegroup	1.0	4340	0.859782	4289	5	45	1
1	agegroup	2.0	12774	0.844446	12523	39	188	24
2	agegroup	3.0	26120	0.848271	25361	178	487	94
3	agegroup	4.0	53564	0.818192	51089	660	1548	267
4	agegroup	5.0	85076	0.799289	78955	1797	3508	816
5	agegroup	6.0	86140	0.785491	74353	4263	5370	2154
6	agegroup	7.0	62193	0.764716	46951	6974	4626	3642
7	agegroup	8.0	15098	0.748486	9194	2936	1235	1733
8	sex	1.0	198347	0.830416	173954	9746	9456	5191
9	sex	2.0	146957	0.818450	128760	7106	7551	3540
10	dialtyp	1.0	310415	0.816646	270848	15496	16115	7956
11	dialtyp	2.0	15082	0.850065	14758	44	248	32
12	dialtyp	3.0	13295	0.858981	12988	36	245	26
13	dialtyp	4.0	77	0.965753	70	3	1	3
14	dialtyp	100.0	6436	0.779859	4051	1273	398	714
15	race	1.0	230577	0.817986	196977	13823	12509	7268
16	race	2.0	93560	0.826123	85998	2552	3760	1250
17	race	3.0	3225	0.819874	3044	53	98	30
18	race	4.0	12965	0.845486	12063	325	436	141
19	race	5.0	3776	0.833047	3566	42	142	26
20	race	6.0	881	0.808297	772	48	46	15
21	race	9.0	321	0.789957	295	9	16	1
22	hispanic	1.0	51021	0.843191	47324	1198	1852	647
23	hispanic	2.0	292532	0.820216	254208	15364	15037	7923
24	hispanic	9.0	1752	0.790421	1183	290	118	161

From prior ScHARE Think-a-thon slides (not covered):
Performing the fairness assessment on the categories of interest gives additional insight into how the model performs by different patient categories of interest (by demographics, etc.). Future researchers should perform fairness assessments to better evaluate model performance, especially for models that may be deployed in a clinical setting. Other methods of assessing fairness include evaluating true positives, sensitivity, positive predictive value, etc. at various threshold across the different groups of interest, which would allow selection of a threshold that balances model performance across the groups of interest.

Machine Learning Computational Strategies

Concept check [slido]

National Institute of Diabetes and Digestive and Kidney Diseases

Machine Learning Computational Strategies

Practical hands-on

(on your own, using ScHARe@Terra)

- Instances of iris flowers …do their petal/sepal length/width vary naturally?
  - Vary by species… exploratory plots confirm [try scikit learn vignette]