ClosedLoop.ai
HEALTHCARE'S DATA SCIENCE PLATFORM

Building a System of AI Drive Intelligence for Healthcare Leadership Institute 11/3/2021

- Introduction
- CMS AI Challenge
- Real world Healthcare AI
- Bias & Fairness
- Emerging Technologies

Dave DeCaprio
CTO
dave.decaprio@closedloop.ai

CMS.gov
Centers for Medicare & Medicaid Services

Search

Medicare | Medicaid/CHIP | Medicare-Medicaid Coordination | Private Insurance | Innovation Center | Regulations & Guidance | Research, Statistics, Data & Systems | Outreach & Education

Innovation Center Home > Innovation Models > Artificial Intelligence (AI) Health Outcomes Challenge

Artificial Intelligence (AI) Health Outcomes Challenge

AI HEALTH OUTCOMES CHALLENGE

The CMS Artificial Intelligence (AI) Health Outcomes Challenge was an opportunity for innovators to demonstrate how AI tools – such as deep learning and neural networks – can be used to accelerate development of AI solutions for predicting patient health outcomes for Medicare beneficiaries for potential use in CMS Innovation Center innovative payment and service delivery models.

The Challenge was operated by CMS in partnership with the American Academy of Family Physicians and Arnold Ventures.

Challenge Objectives

1. For Stage 1, use AI, including but not limited to deep learning methodologies, to predict unplanned hospital and SNF admissions, and adverse events within 30 days for Medicare beneficiaries, based on a data set of Medicare administrative claims data, including Medicare Part A (hospital) and Medicare Part B (professional services).
2. For Stage 2, use AI, including but not limited to deep learning methodologies, to predict unplanned hospital and SNF admissions, and adverse events, within 30 days for Medicare beneficiaries, as well as 12-month mortality for all Medicare beneficiaries, based on a Part A and Part B data set.
3. For both Stage 1 and Stage 2, develop innovative strategies and methodologies to: explain the AI-derives predictions to front-line clinicians and patients to aid in providing appropriate clinical resources to model participants; and increase use of AI-enhanced data feedback for quality improvement activities among model participants.

Participants also were required to address implicit algorithmic biases that impact health disparities in their submissions.

Model Summary

Stage: No Longer Active
Number of Participants: 7
Category: Initiatives to Accelerate the Development and Testing of New Payment and Service Delivery Models
Authority: America COMPETES Reauthorization Act

Milestones & Updates

April 30 2021
Announced: Artificial Intelligence (AI) Health Outcomes Challenge Winner & Runner-Up Announced

October 29 2020
Announced: Seven finalists to advance to Stage 2

October 31 2019
Announced: Top 25 participants selected to advance to Stage 1

April 10 2019
Announced: Informational overview webinar April 18

Challenge Overview

- Use AI to predict unplanned hospital and SNF admissions and adverse events
- Explain the AI-derived prediction to front-line clinician and patients
- Discuss how you addressed algorithmic bias

CMS AI Challenge Participants

300+ entrants nationwide

Top 25

MAYO CLINIC
IBM
Jefferson
MERCK
NORTHROP GRUMMAN
Northwestern Medicine
Penn Medicine
Deloitte.
Booz | Allen | Hamilton
strategy and technology consultants

Top 7

Ann Arbor Algorithms
ClosedLoop.ai
Deloitte Consulting
Geisinger/Medial
Jefferson Health
Mathematica
University of Virginia

1st place
ClosedLoop.ai

2nd place
Geisinger | Medial EarlySign

March 2019

October 2019

October 2020

April 2021

Keys to Success

We didn't build models. We built a machine that builds models

- ClosedLoop does this every day - the challenge was a way to push forward our core product - justified more investment
- Realize accuracy isn't the only goal - extensive user testing on the UI
- Rapid experimentation - Created a list of 257 ideas and tried them all
- Deep understanding of the data - RUPP codes used for home health
- Best experiment
  - Sample each patient at many time points - produced 13 million data points from 2.5 million people
- Underlying algorithm was XGBoost
  - Bayesian hyperparameter optimization
  - Process tuned on a 25 separate healthcare data sets
  - Final result was stacked with a transformer network
    - XGBoost ROC: 0.867
    - NN ROC: 0.864
    - Stacked: 0.869

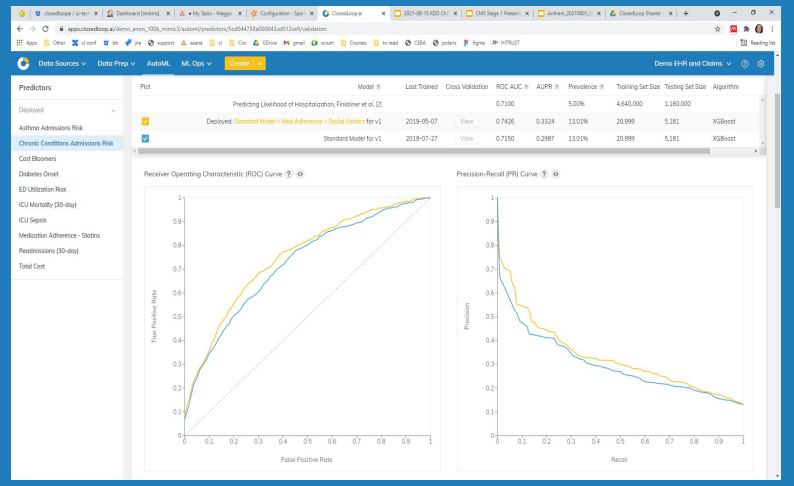

- **End to end experiment management**
  - Hundreds of models, most with different data sets
  - Experiment encompasses all decisions from initial ETL through final trained model
- **Standardized evaluation and visualization**
- **Team-based approach**

Healthcare's Data Science Platform

ClosedLoop makes it easy and affordable for healthcare organizations to use data science to improve quality and reduce costs.

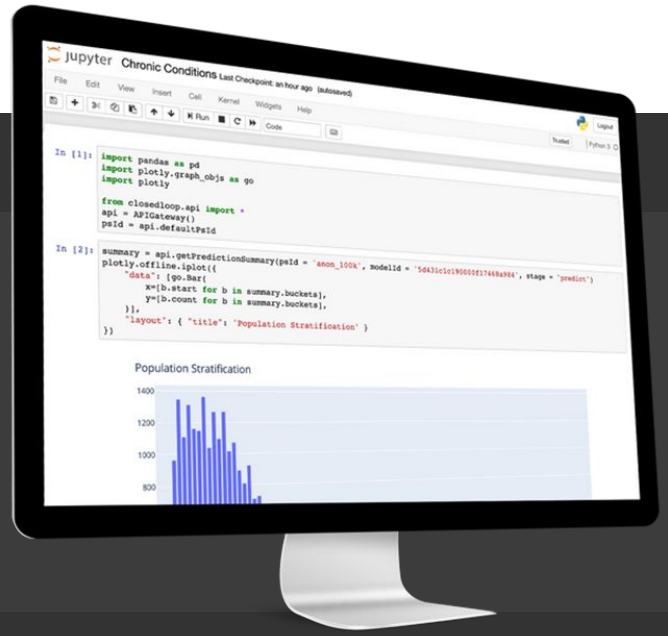

ML / AI Automation

- Build & Deploy New Models In 24 Hours
- Easily Handle Messy Healthcare Data
- Automate Clinical Feature Engineering
- Direct Support for ICD9/10, SNOMED, GPI, RXNorm
- Enhance Explainability
- Enable Seamless Model Deployment
- API integration to Jupyter, Python, etc.

Model Template Library

- Admissions / Readmissions
- Total Utilization / Total Risk
- ED Utilization
- Out of Network Leakage
- Appointment No-Shows

- Chronic Disease Onset
- Risk Adjustment
- Suspect Diagnosis
- …and dozens more

Or build custom models using our clinical expression language:

CL-SCRIPT

Raw Data To Actionable Insights

Data Sources	Healthcare's Data Science Platform	Comprehensive Risk Profiles	Workflow System

Raw Patient Data ➡ Actionable Insights

Comprehensive Risk Profiles

"Baseball Card" previews are the primary day to day point of interaction.

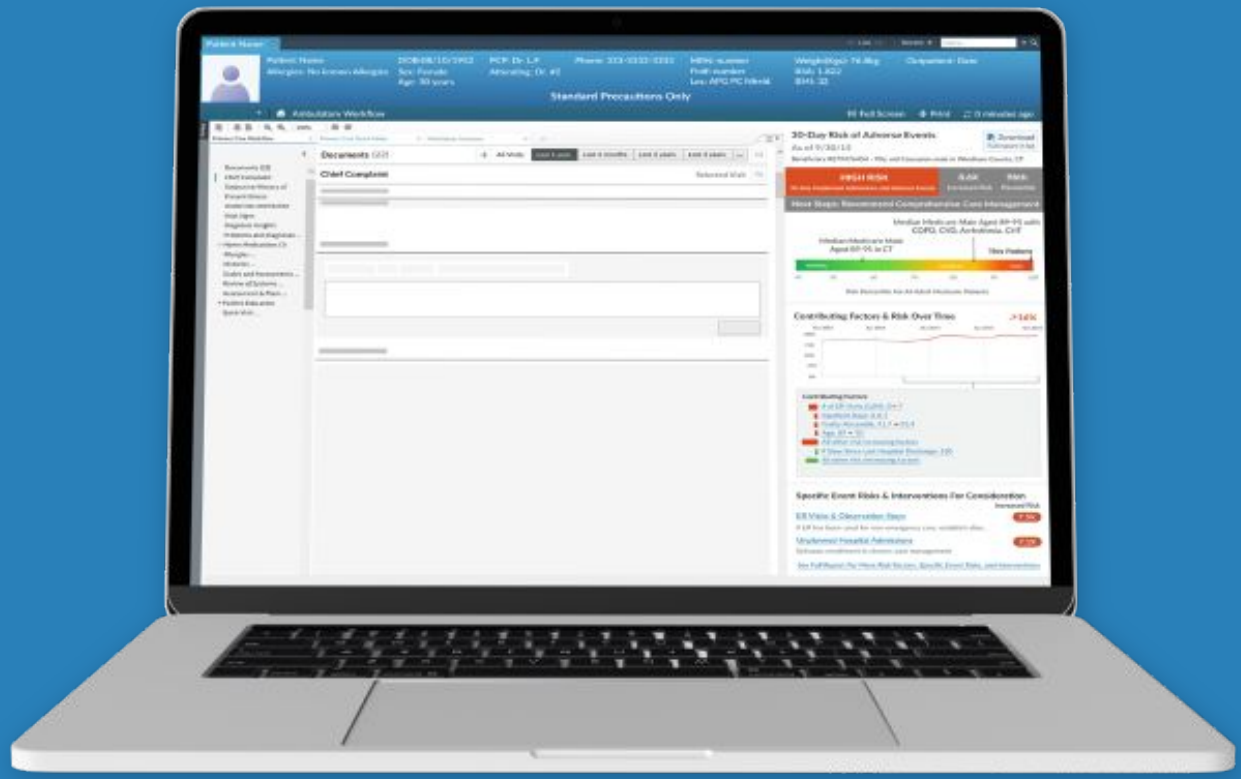

Cerner Epic

Patient Centric Digital Risk Forecasts

Full 10-15 page risk report allows clinicians to view full detail of risk over time, contributing factors, data used, potential bias, etc.

Risk Trend Line

Risk trend line shows how risk has changed over time.

Contributing Factors

how which variables contributed to increased/ decreased risk at various points in time. *Patent Pending

Specific Event Risks

Specific event risks highlight additional predictive models for which the patients was flagged as high risk and associated interventions for those specific events (e.g. avoidable ED, FRIs, etc.)

Composite Model with Component Predictions

The following events constitute a broad portfolio of models designed to predict multiple, distinct unplanned and adverse events within the next 30 days. Each event definition has been validated by published research and is associated with evidence-based interventions that can potentially avoid or mitigate their impact.

Unplanned Hospital & SNF Admissions

Potentially Avoidable Admissions (8 specific PQIs)

A narrower definition of avoidable admissions for patients with ambulatory care sensitive conditions (ACSCs).
Support: AHRQ, StatBrief #195, Rollow et al.
Beneficiaries Impacted: **1.8 million**

Unplanned Hospital Admissions

A broader definition of UHAs including admissions and readmissions for acute illness or complications of care.
Research Support: QualityNet.org, IHI.org, Wells et al
Beneficiaries Impacted: **6.3 million**

Admission to a Skilled Nursing Facility

Any SNF admission (Medicare requires prior hospital admission for 3+ days making SNF admissions "planned" in some sense).
Research Support: MedPAC, Ouslander et al., Frick et al.
Beneficiaries Impacted: **480,000**

Adverse Events

Hospital-Acquired Infections (CLABSI, CAUTI, MRSA, C. diff)

Inpatient stays getting a secondary diagnosis of HAI.
Research Support: CDC, AHRQ, Nat'l HAI Plan, HAI disclaimer
Beneficiaries Impacted: **600,000**

ER Visits & Observation Stays

Any emergency room (ER) visit or Observation Stay.
Research Support: KHN.org, Nagurney et al., Fleischman, W.
Beneficiaries Impacted: **11.7 million**

Severe Fall-Related Injuries (FRIs)

Any hospital visit with a principal diagnosis of a serious FRI.
Research Support: CDC, Hoffman et al. , Frick et al.
Beneficiaries Impacted: **3.2 million**

30-day Mortality

Death due to any cause within 30 days.
Research Support: Meier et al., Ellershaw, J., Teno et al.
Beneficiaries Impacted: **1.7 million**

Patient Centric ML - Required Tech

In order to deliver on the vision of “Patient Centric ML”, systems that are capable of ingesting HC data, creating clinically relevant features, generating explainable models, and storing predictions and insights over time are required.

Risk Trend Line

The risk trend line requires that predictions for a model are stored over time, even as that model may be retrained and or improved with subsequent versions, features, etc.

Contributing Factors

Contributing factors require that every prediction include and persist the factors and values used in every prediction every day.

Specific Event Risks

Specific event risks require a common framework and platform for defining, training, explaining, deploying and monitoring a portfolio of dozens of specific predictive models.

Factor Evidence™

ClosedLoop patent pending Factor Evidence™ technology allows Contributing Factors to be traced back to the specific source data and associated timestamps used in calculating those factors.

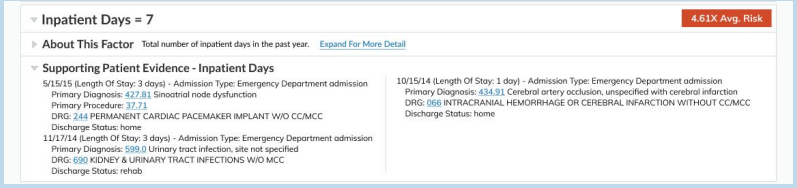

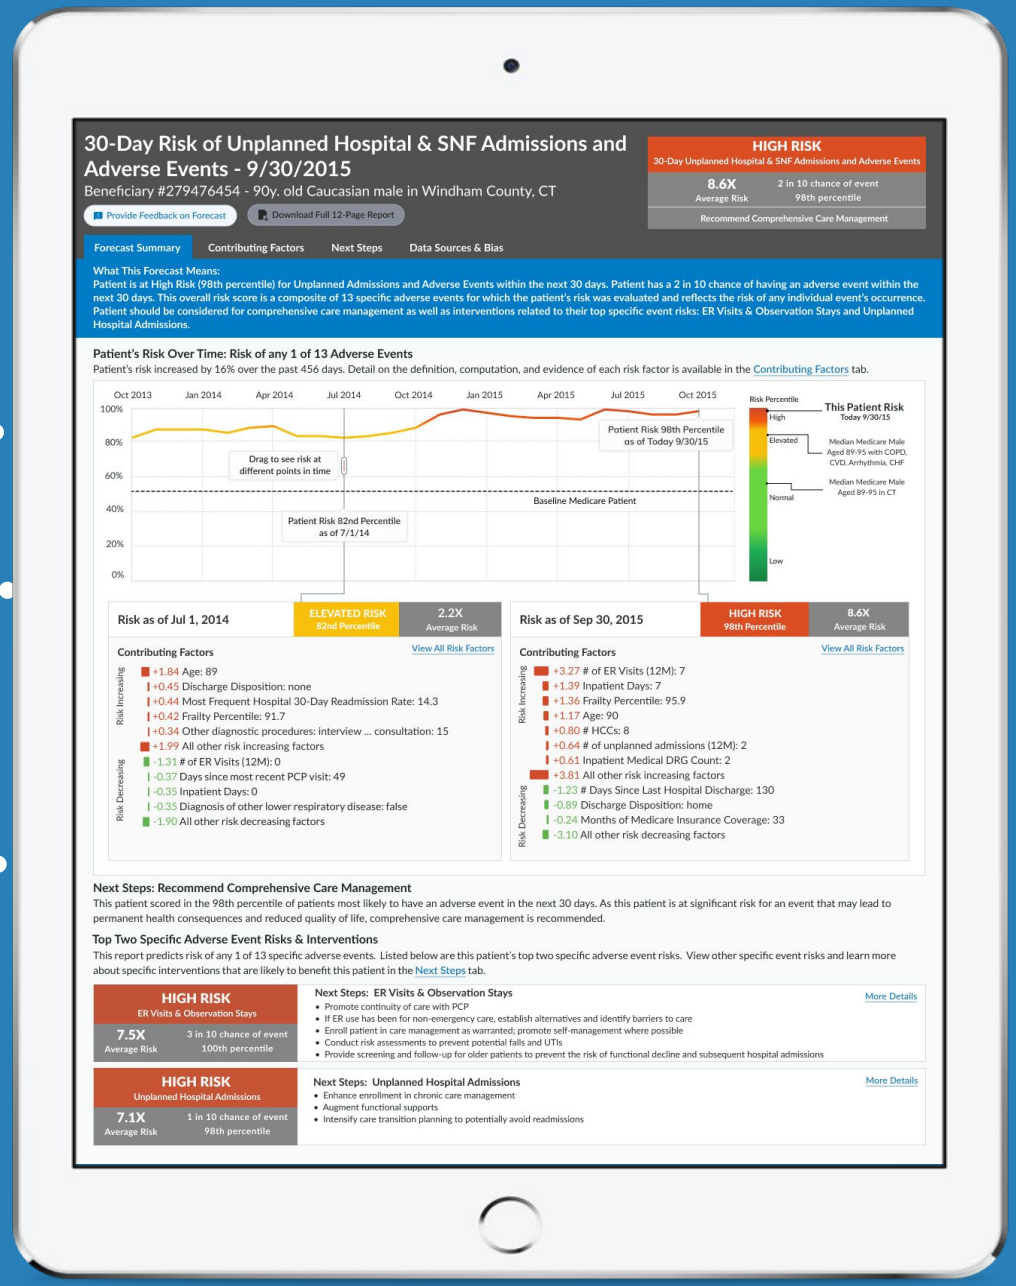

Enterprise Healthcare Feature Store

ClosedLoop's EHFS enables sharing and collaboration of analytic features and eliminates duplicate work

The EHFS makes it easy to:

- **Productionize** new features without extensive engineering support
- **Automate** feature computation, backfills, and logging
- **Share** and **reuse** feature pipelines across teams
- **Track** feature versions, lineage, and metadata
- **Use prediction outputs** as features to other models & reports
- **Achieve consistency** between training and serving data
- **Monitor the health** of feature pipelines in production

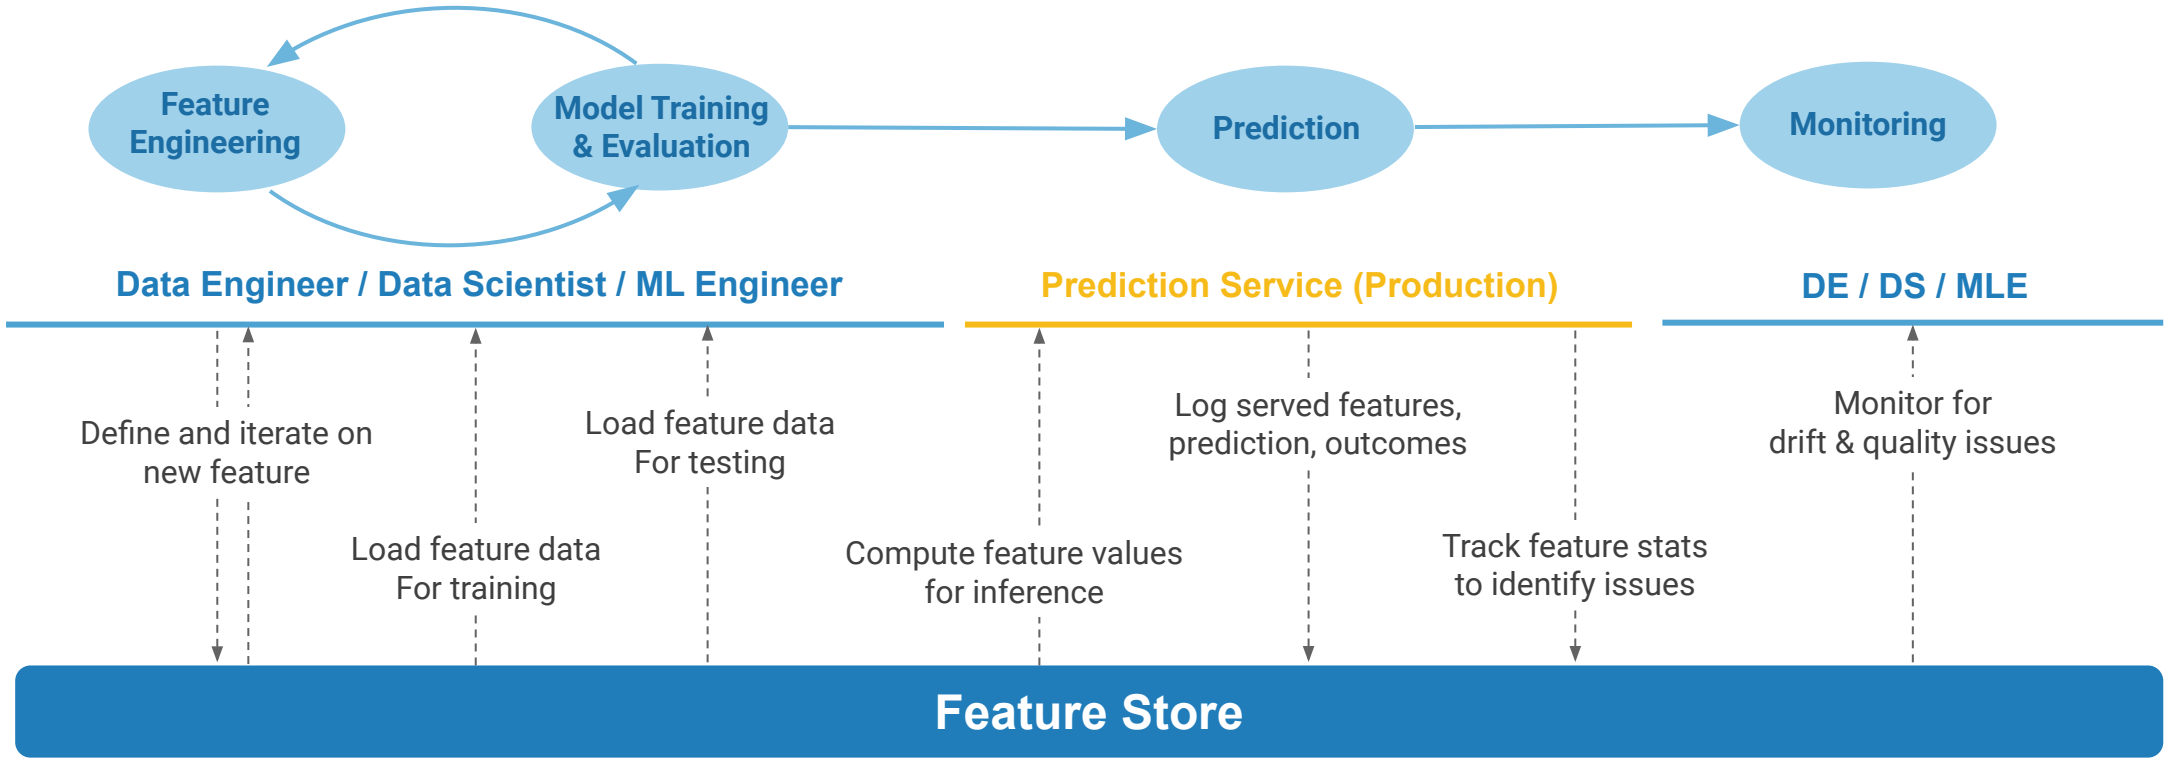

Auto-ML

"Train & Explain" - traditional Auto-ML tools begin with an existing data frame and focus on automating the training and (to some extent) explaining or predicting models. Only when integrated with Enterprise Feature Store and ML-OPS capabilities, a fully automated ML pipeline can be realized.

Preprocessing

- Time Series Expansion
- NLP / NLU Pipelines
- Imputation / Reweighting

Model Training

- Hyperparameter Optimization
- Model Training
- Validation & Testing
- Version Control
- Change Logging
- Bias / Fairness Assessment

Model Explanations

- Human Readable Explanations
- Significance Testing
- Population Level Visualization
- Patient Level Clinical Review
- Factor Evidence™

Accuracy Across Feature Sets

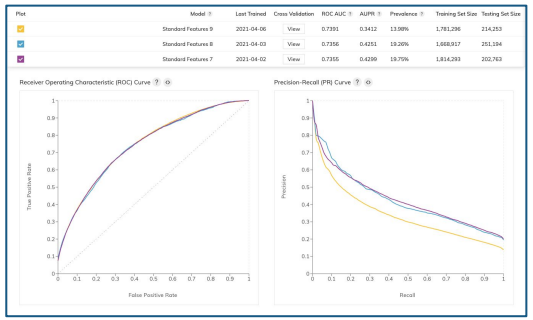

Model / Feature Version Tracking

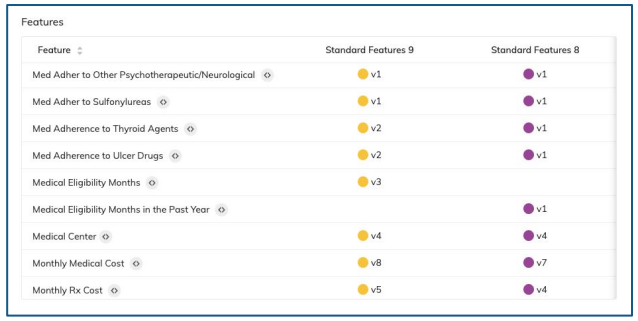

Single Factor Detail

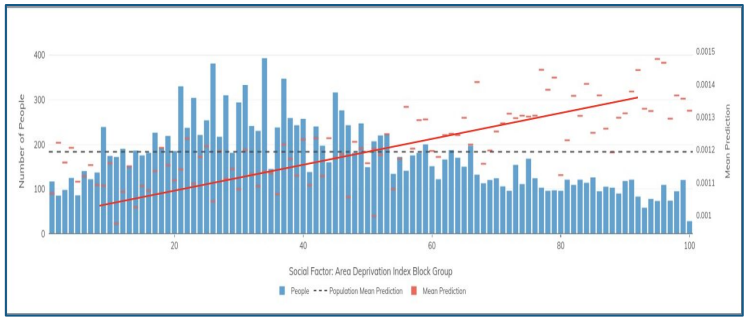

Population Level Feature Importance

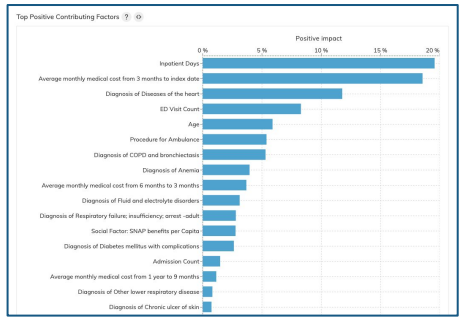

ML OPS

a set of practices that combines Machine Learning, DevOps & Data Engineering, which aims to deploy and maintain ML systems in production reliably and efficiently. - *Towards Data Science, 1/2020*

ML Deployment

- Define ML & Data Processing Pipelines
- Diagram data ingest, model run, report creation workflow
- Manage staged versus active deployments
- Manage model retraining, release, deprecation

ML Health Monitoring

- Monitor Individual model performance
- Track predicted versus actual performance
- Track feature drift
- Monitor service quality (speed, errors, etc)

ML Governance

- Access event log
- Manage individual permissions
- Configurable warnings and permissions
- Mange versioning of features, models
- Evaluate and report on potential bias / fairness concerns

Sample Deployment

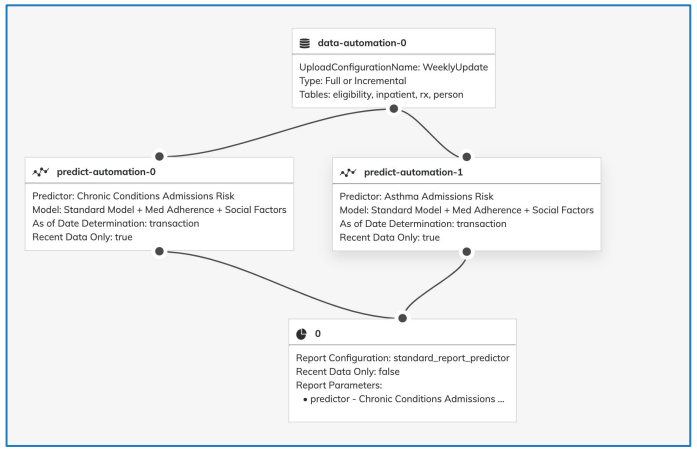

Drift Versus Importance

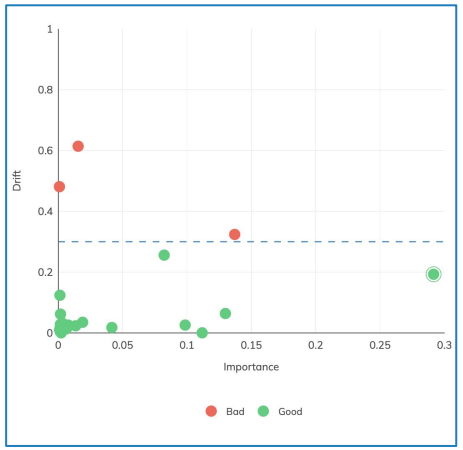

Accuracy over time

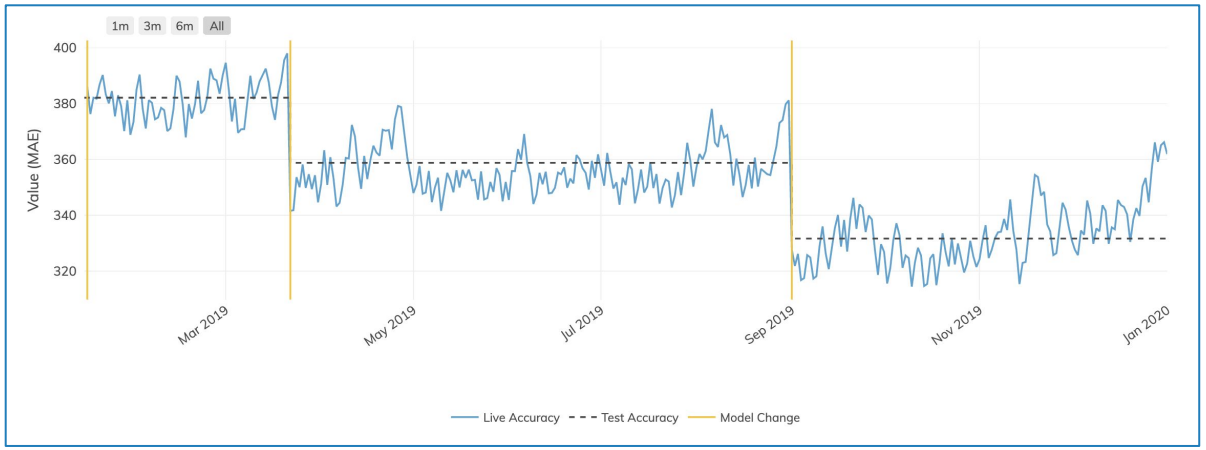

Clinical Feedback & Continuous Learning

Clinical & Business User Feedback

- Validation of contributing factors learned by AI models
- Validation of new surfaced cases vs. legacy system results
- Adjustments to predicted event definitions focused on impactability
- Adjustments to intervention strategies based on surfaced results

Content Creation

- Re-usable features (diagnosis, procedures, drugs, SDoH, labs, vitals, cost…)
- Data adapters (Claims, EHR, Labs, ADT, Device…)
- Model templates (Outcome + Population)
- Linkable data (SDoH, environmental, consumer…)

Increased Value

- Increased accuracy
- Increased explainability
- Increased impactability
- Decreased integration time

PLATFORM ENHANCEMENTS

Scalability – data prep plans to handle multi thousand feature models
Accuracy – new feature classes – generators, time-series, relative time, geo-coding, NLP
Explainability – factor drill downs, prediction history, bias
Monitoring and operations – input and output quality controls, prediction acceptance tracking

Ex. Medicaid CBO Focused System Of Intelligence

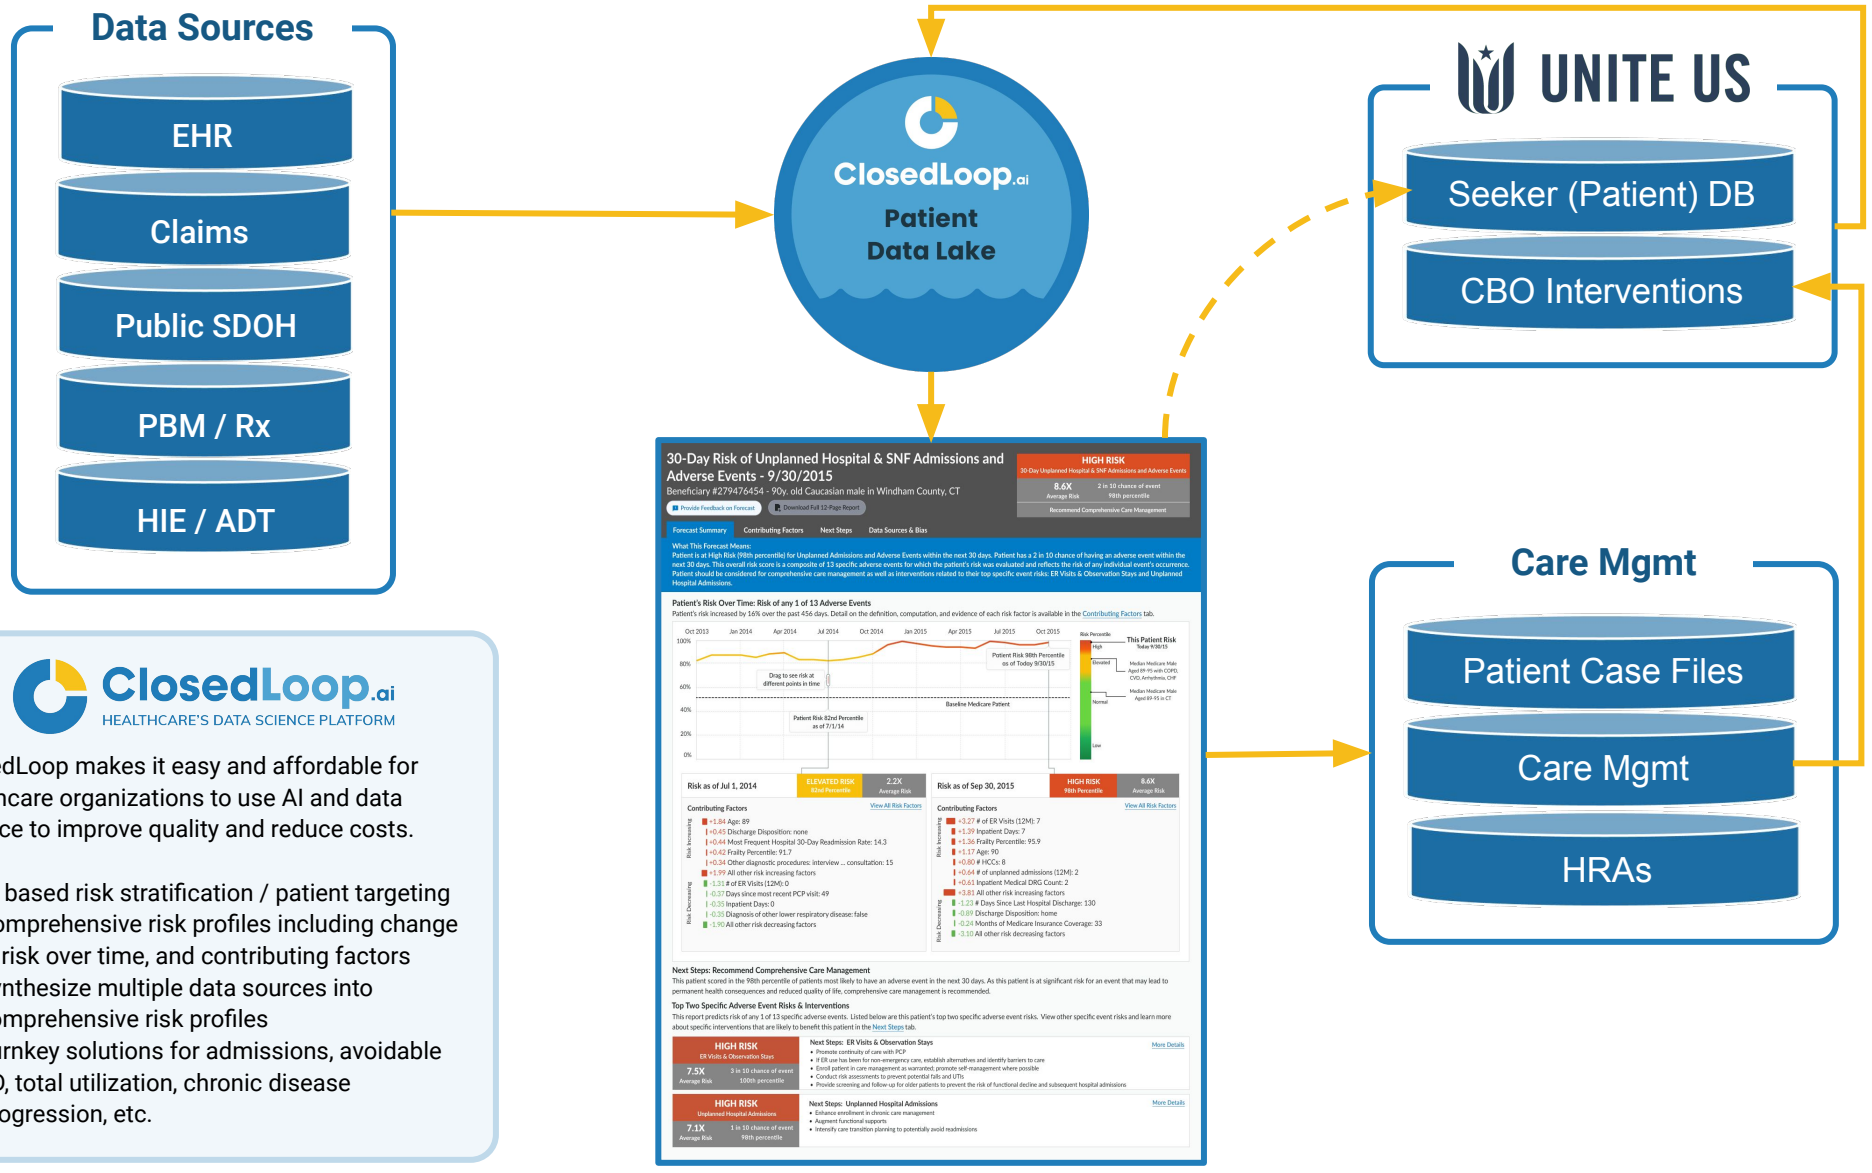

UniteUs enables individuals to find and apply for government and charitable social service programs.

- National footprint of CBO programs
- Proactive network monitoring
- Free intake and referral tools for CBOs

Case management systems are used by care coordinators to track patient progress and make referrals to CBOs. ClosedLoop will prioritize patients for referral. ClosedLoop and UniteUs connect via APIs to various care management platforms.

ClosedLoop.ai
HEALTHCARE'S DATA SCIENCE PLATFORM

ClosedLoop makes it easy and affordable for healthcare organizations to use AI and data science to improve quality and reduce costs.

- AI based risk stratification / patient targeting
- Comprehensive risk profiles including change in risk over time, and contributing factors
- Synthesize multiple data sources into comprehensive risk profiles
- Turnkey solutions for admissions, avoidable ED, total utilization, chronic disease progression, etc.

Benefits of a Machine Learning Approach

- **Identifies relationships that a “rules based” approach will miss**
  - Machine Learning will uncover relationships and deliver more granular predictions of risk
- **Leverage vastly more data (when it is available)**
  - Multiple types and sources, thousands of columns, millions of rows
- **Provide explainability at the population and individual level**
  - Transparency and explainability are capabilities that you should pursue
- **Rapidly built, iterated upon, and deployed**
  - Data-to-prediction in days or weeks, not months or never
- **Continually evolve in response to the situation**
  - Machine Learning models can be rapidly updated in response to changing data

Bias & Fairness

nature

Explore content · About the journal · Publish with us · Subscribe

nature > news > article

NEWS | 24 October 2019 | Update 26 October 2019

Millions of black people affected by racial bias in health-care algorithms

Study reveals rampant racism in decision-making software used by US hospitals — and highlights ways to correct it.

Heidi Ledford

Sections

The Washington Post
Democracy Dies in Darkness

Health

Racial bias in a medical algorithm favors white patients over sicker black patients

By Carolyn Y. Johnson

October 24, 2019

The algorithm wasn't intentionally racist — in fact, it specifically excluded race. Instead, to identify patients who would benefit from more medical support, the algorithm used a seemingly race-blind metric: how much patients would cost the health-care system in the future. But cost isn't a race-neutral measure of health-care need. Black patients incurred about $1,800 less in medical costs per year than white patients with the same number of chronic conditions; thus the algorithm scored white patients as equally at risk of future health problems as black patients who had many more diseases.

Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019 Oct 25;366(6464):447-453

ClosedLoop Framework for Algorithmic Accountability

Reducing Algorithmic Bias

For models used to aid decisions for balancing harms-benefits when decision maker and subject are aligned

Use Representative samples for model development

Diagnose subgroup validity: Evaluate performance of the model (e.g. discrimination and calibration) overall and in each class (e.g. race, gender)

Consider the potential for **labeling bias**

Evaluate the potential for **feature bias:** Consider interaction terms with protected attribute or stratified models

Consider potential for **sampling bias**

Reducing Algorithmic Unfairness

For models used for rationing

Approach 1: Restrict model inputs

Models should include only well-established, causal risk factors.

- Models should therefore be ***race-unaware***, and also exclude non-causal variables that may be race proxies (e.g. zip code)

Approach 2: Restrict model outputs

Examine models using fairness criteria. Ensure "fair" distributions of services by either:

- Using different decision thresholds
- Applying fairness constraints to strategically reclassify based on the protected attribute

Collaboration Partners

Cory Booker | Booz | Allen | Hamilton | PACE The Predictive Analytics & Comparative Effectiveness Center | Berkeley Public Health | CHICAGO BOOTH The Center for Applied Artificial Intelligence

Emerging Technologies

Important Emerging Technologies

- **Build Models Without Sharing Data**
  - Several technologies can potentially address this
    - Synthetic Data
    - Homomorphic encryption
    - Differential Privacy
    - Federated Learning
  - These are all early stage technologies that have yet to be demonstrated to work in healthcare
- **AI-driven Interoperability**
  - Computers write the code instead of people
- **AI-driven Evidence Generation**
  - Smarter causal modeling driven by observational data to produce more focused studies

ClosedLoop.ai
HEALTHCARE'S DATA SCIENCE PLATFORM

Thank you!

Dave DeCaprio
CTO
dave.decaprio@closedloop.ai