Examining Model Drift

PUTTING ML MODELS INTO PRODUCTION AND KEEPING THEM THERE

NOVEMBER 2025

Megan Robertson

AMLC

MEGAN ROBERTSON

- Chief Data Scientist - MegRob DSA
- 8+ years experience in data science and machine learning
    - Building products to support stakeholder and business strategy
    - Intersection between technical and non-technical teams

MODEL DRIFT: DEGRADATION IN ML MODEL DUE TO CHANGES IN DATA

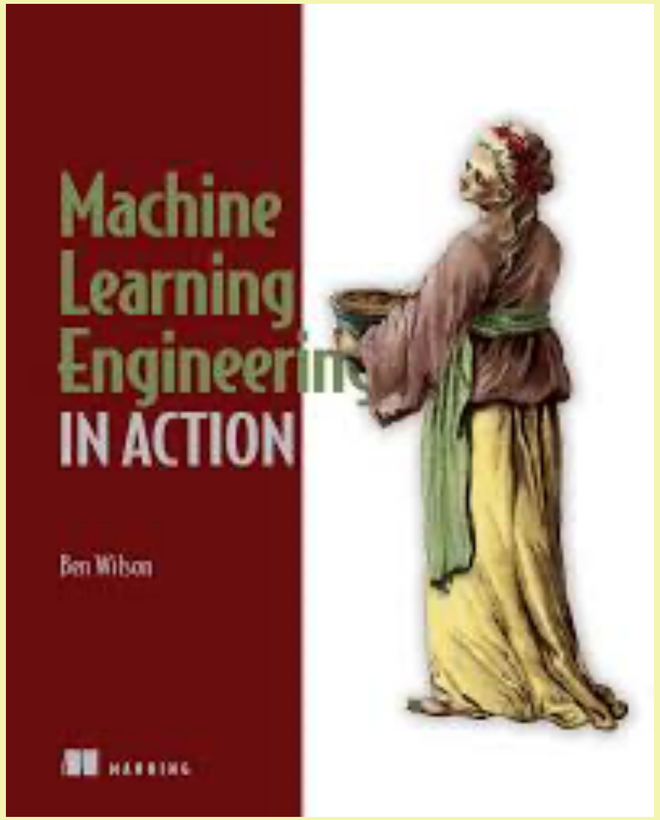

“MODELS WILL DRIFT. THERE IS NO SUCH THING AS STATIC IMPLEMENTATION.”

CONSEQUENCES OF MODEL DRIFT

Incorrect Predictions

Poor Performance

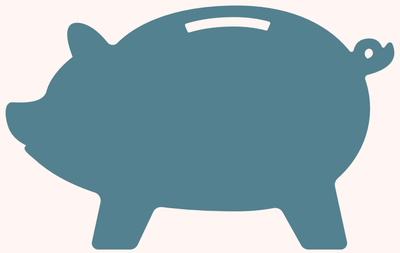

Financial Losses

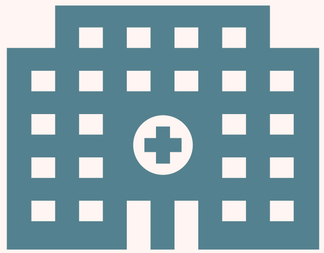

Physical Harm

Loss of Trust

CONCEPT DRIFT

Divergence between input variables and target variable

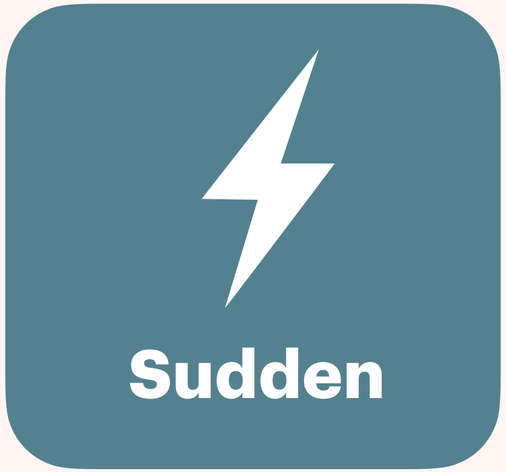

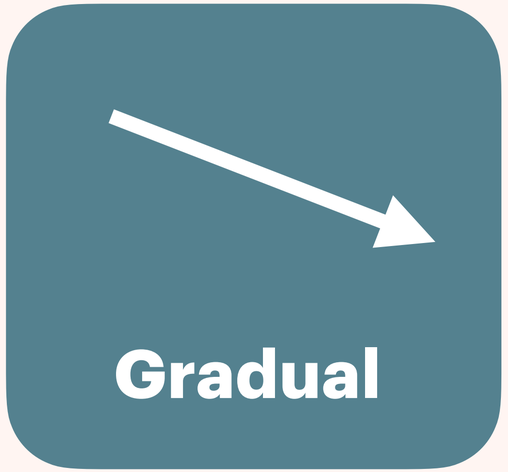

DATA (COVARIATE) DRIFT

Change in underlying data distribution of the input features, relationship between features and target variable unchanged

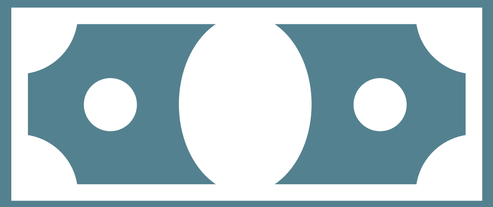

Increase in younger population that has a fundamentally different spending habits than older generations

LABEL DRIFT

Prior probability shift - distribution of target variable changes over time with no change between input features and target variable

Change in prevalence rates of a disease due to a public health campaign or improved testing

REALITY DRIFT

HOW CAN YOU IDENTIFY MODEL DRIFT?

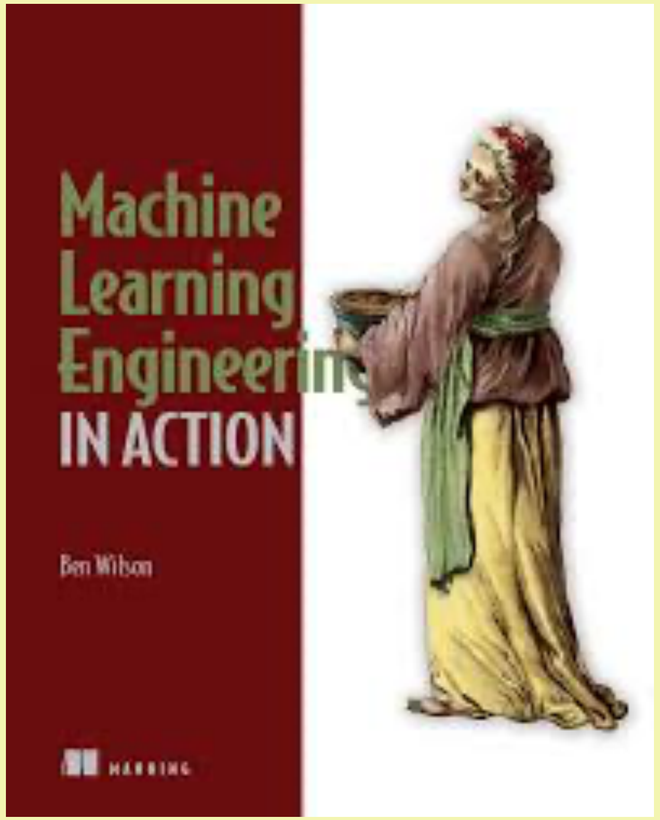

“CONTINUOUS MONITORING AND A SUSPICIOUS AMOUNT OF DISTRUST OF EVERYTHING GOING INTO AND COMING OUT OF A MODEL.”

DRIFT MONITORING SYSTEM
DATA

Quality Control

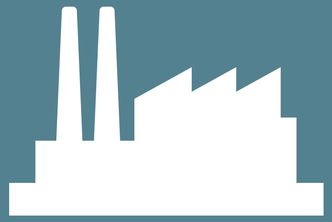

Pipeline Breaks,
Missing Values

Distribution Drift

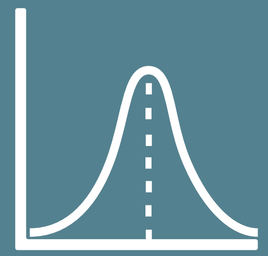

Chi-Square,
Kolmogorov-Smirnoff

DRIFT MONITORING SYSTEM
MODELING

Prediction and Output

Classification Ratios,
Shifting Predictions

Performance

Precision, Recall, F1,
MSE, MAE, Accuracy

INDUSTRY USE CASE

What is the probability that a customer will make a purchase in the next 90 days?

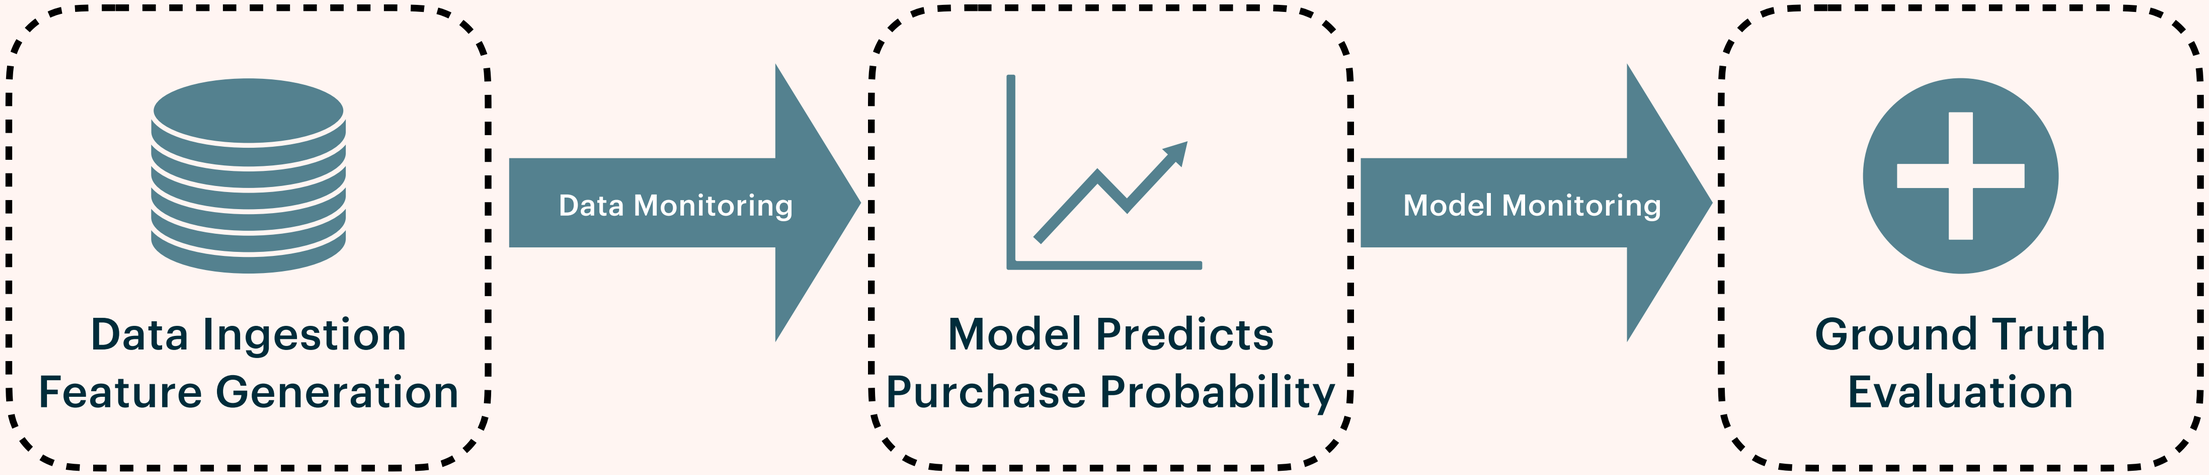

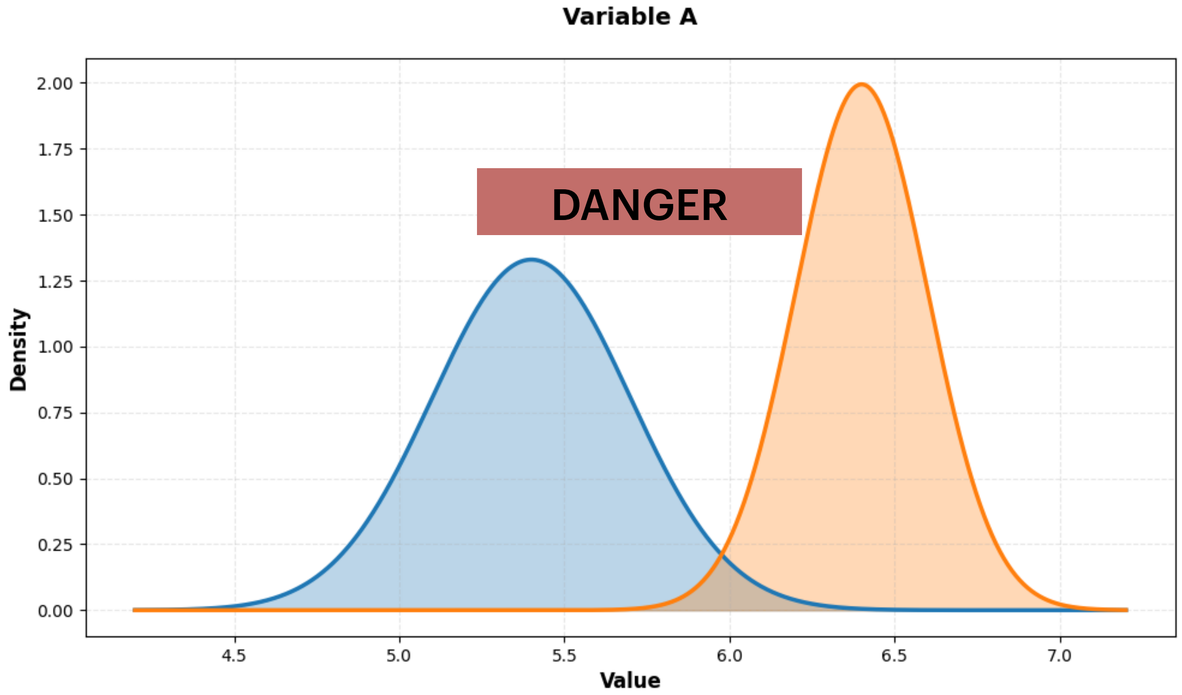
Variable A
DANGER
Density
Value

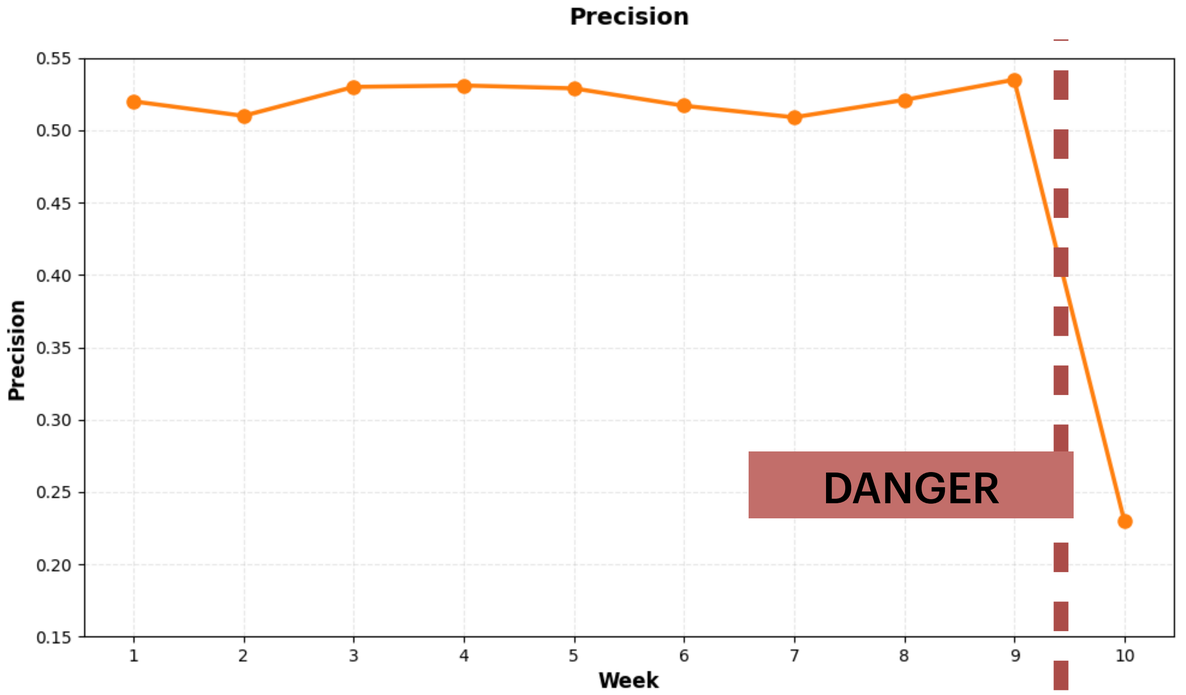
Precision
DANGER
Precision
Week

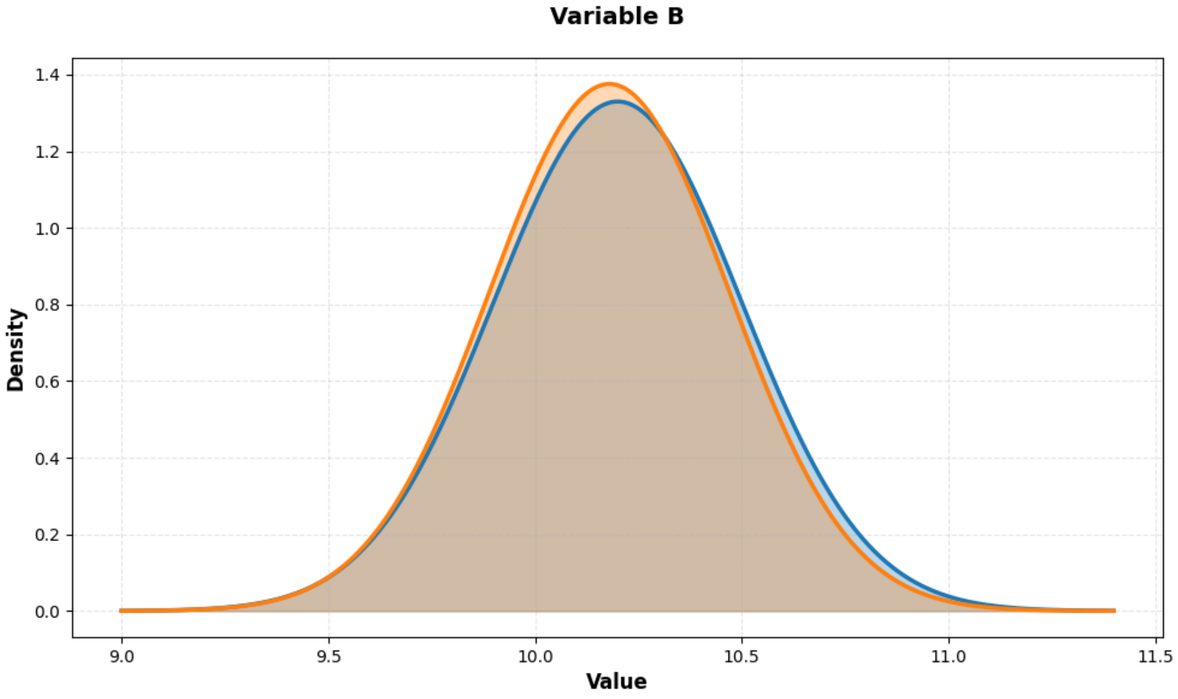
Variable B
Density
Value

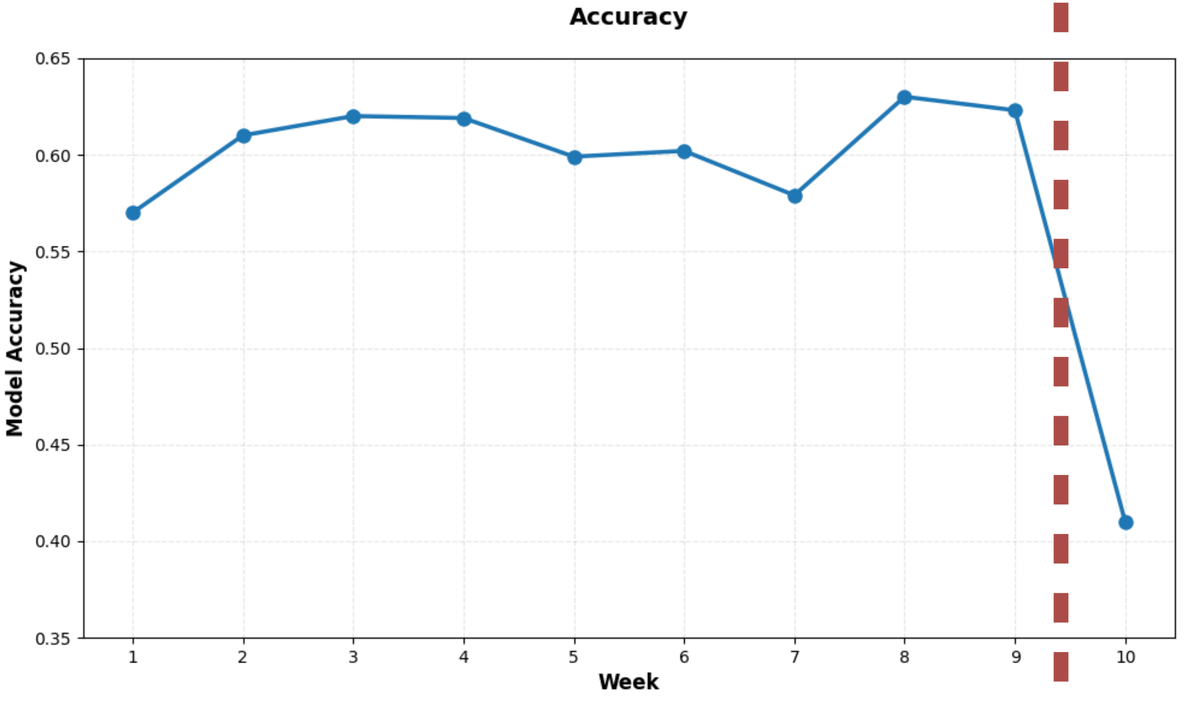
Accuracy
Model Accuracy
Week

HOW CAN WE CORRECT MODEL DRIFT?

ADDRESSING MODEL DRIFT

Infrastructure Check

Feature engineering

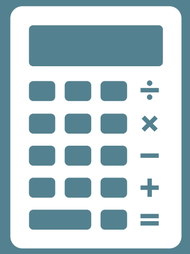

New algorithm

Retrain model

MODEL RETRAINING

Passive: Scheduled Job
Active: Attribution Monitoring

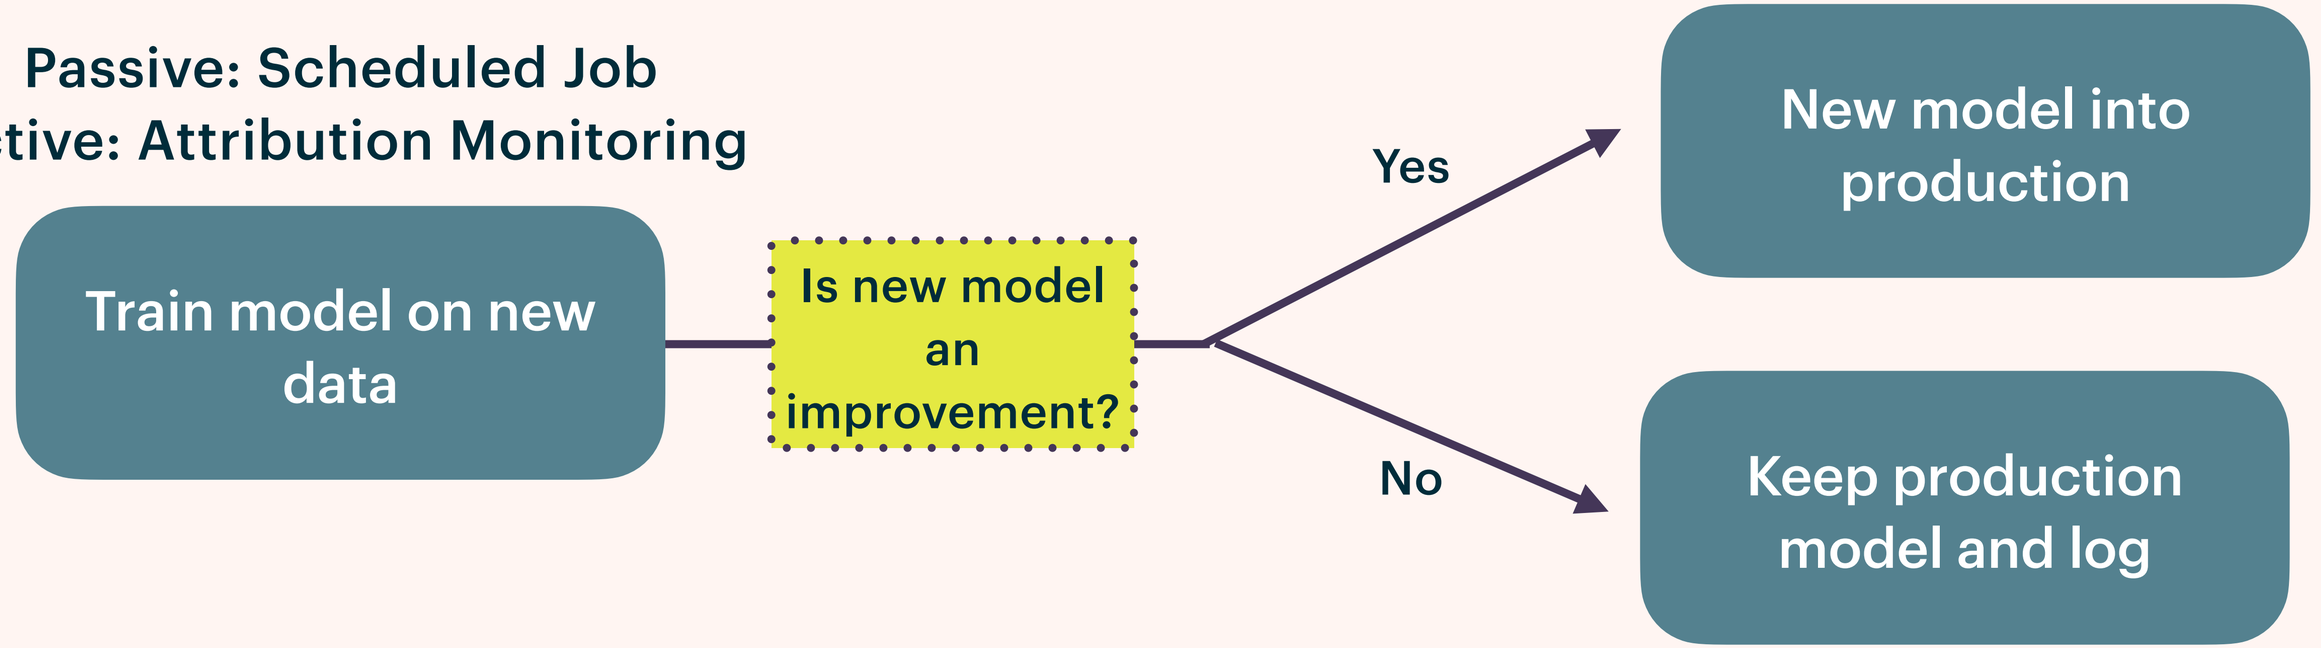

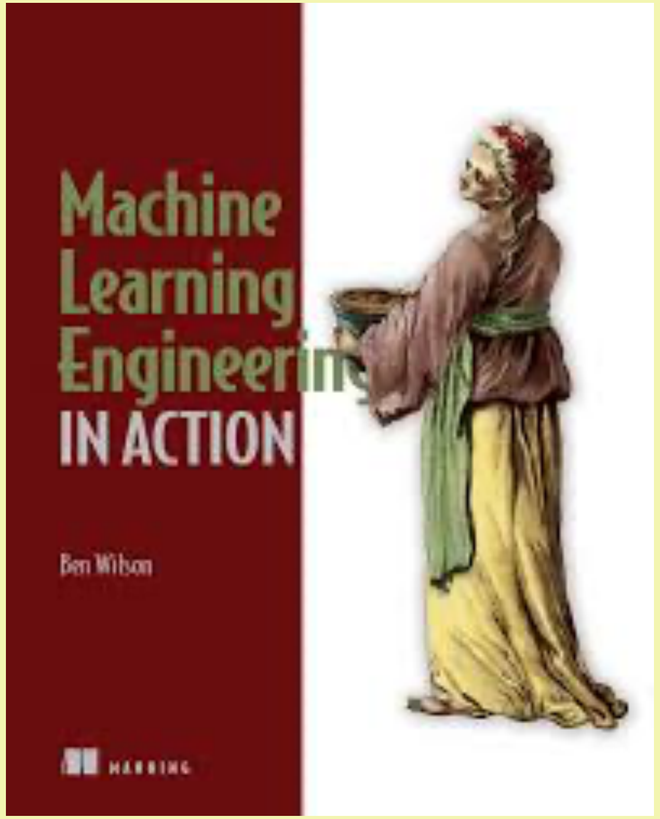

“IF THE BENEFIT OF THE MODEL IS NO LONGER PRESENT, SHUT IT DOWN PERMANENTLY.”

RESOURCES

- Machine Learning Engineering in Action - Ben Wilson (Chapter 12)
- IBM (https://www.ibm.com/think/topics/model-drift)
- Lumenova (https://www.lumenova.ai/blog/model-drift-concept-drift-introduction/)
- Datadog (https://www.datadoghq.com/blog/ml-model-monitoring-in-production-best-practices/)

megrobdatascience.com

megan@megrobdatascience.com

linkedin.com/megan-robertson-ds