SAND2020-6530PE

I could see your lips move.

PRESENTED BY

JD Doak to CCD on 6/30/2020

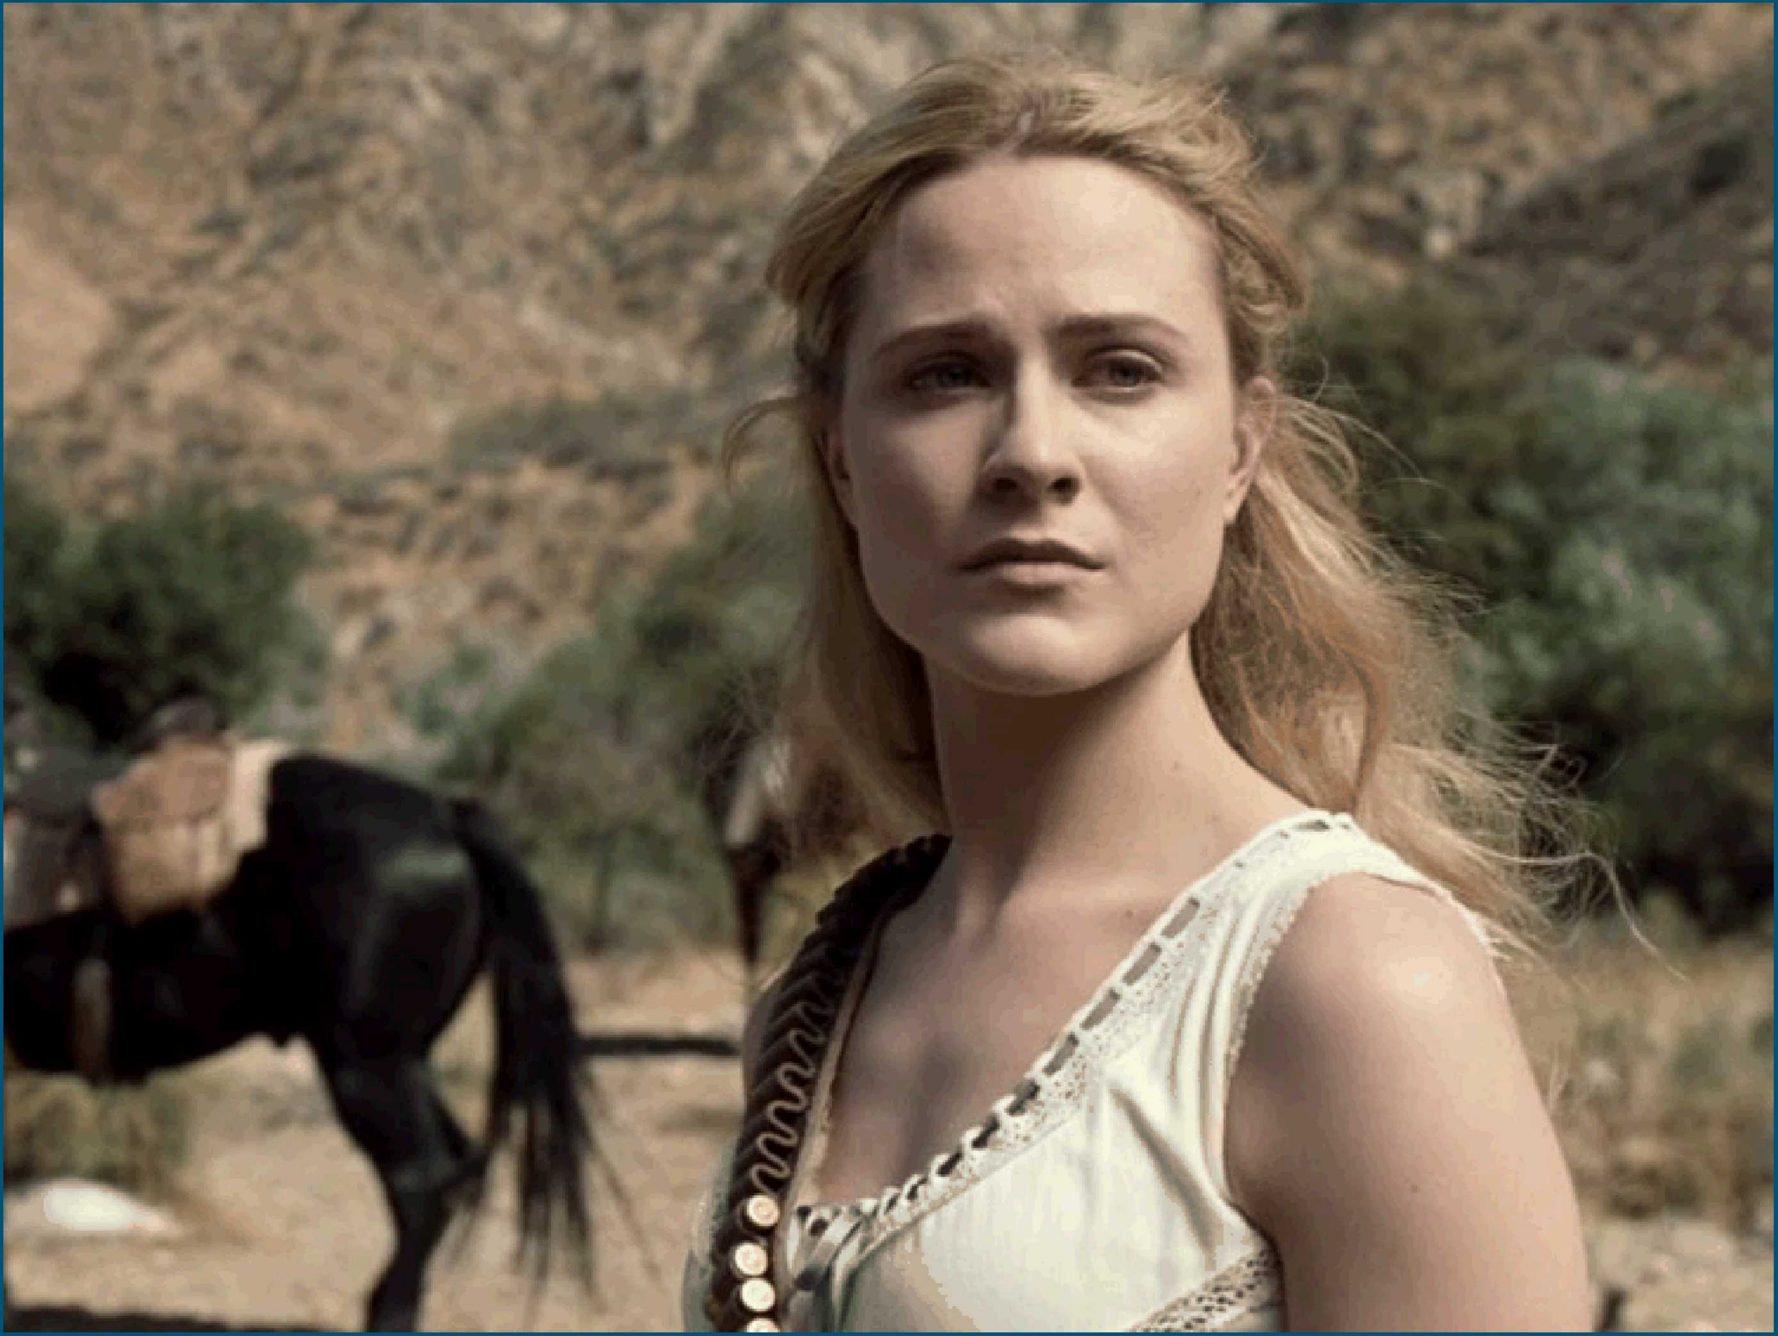

https://www.thisisinsider.com/westworld-season-2-episode-1-analysis-spoilers-2018-4

https://www.spotern.com/en/wanted/movie/blade-runner/32790/the-cigarette-of-rachel-sean-young-in-blade-runner

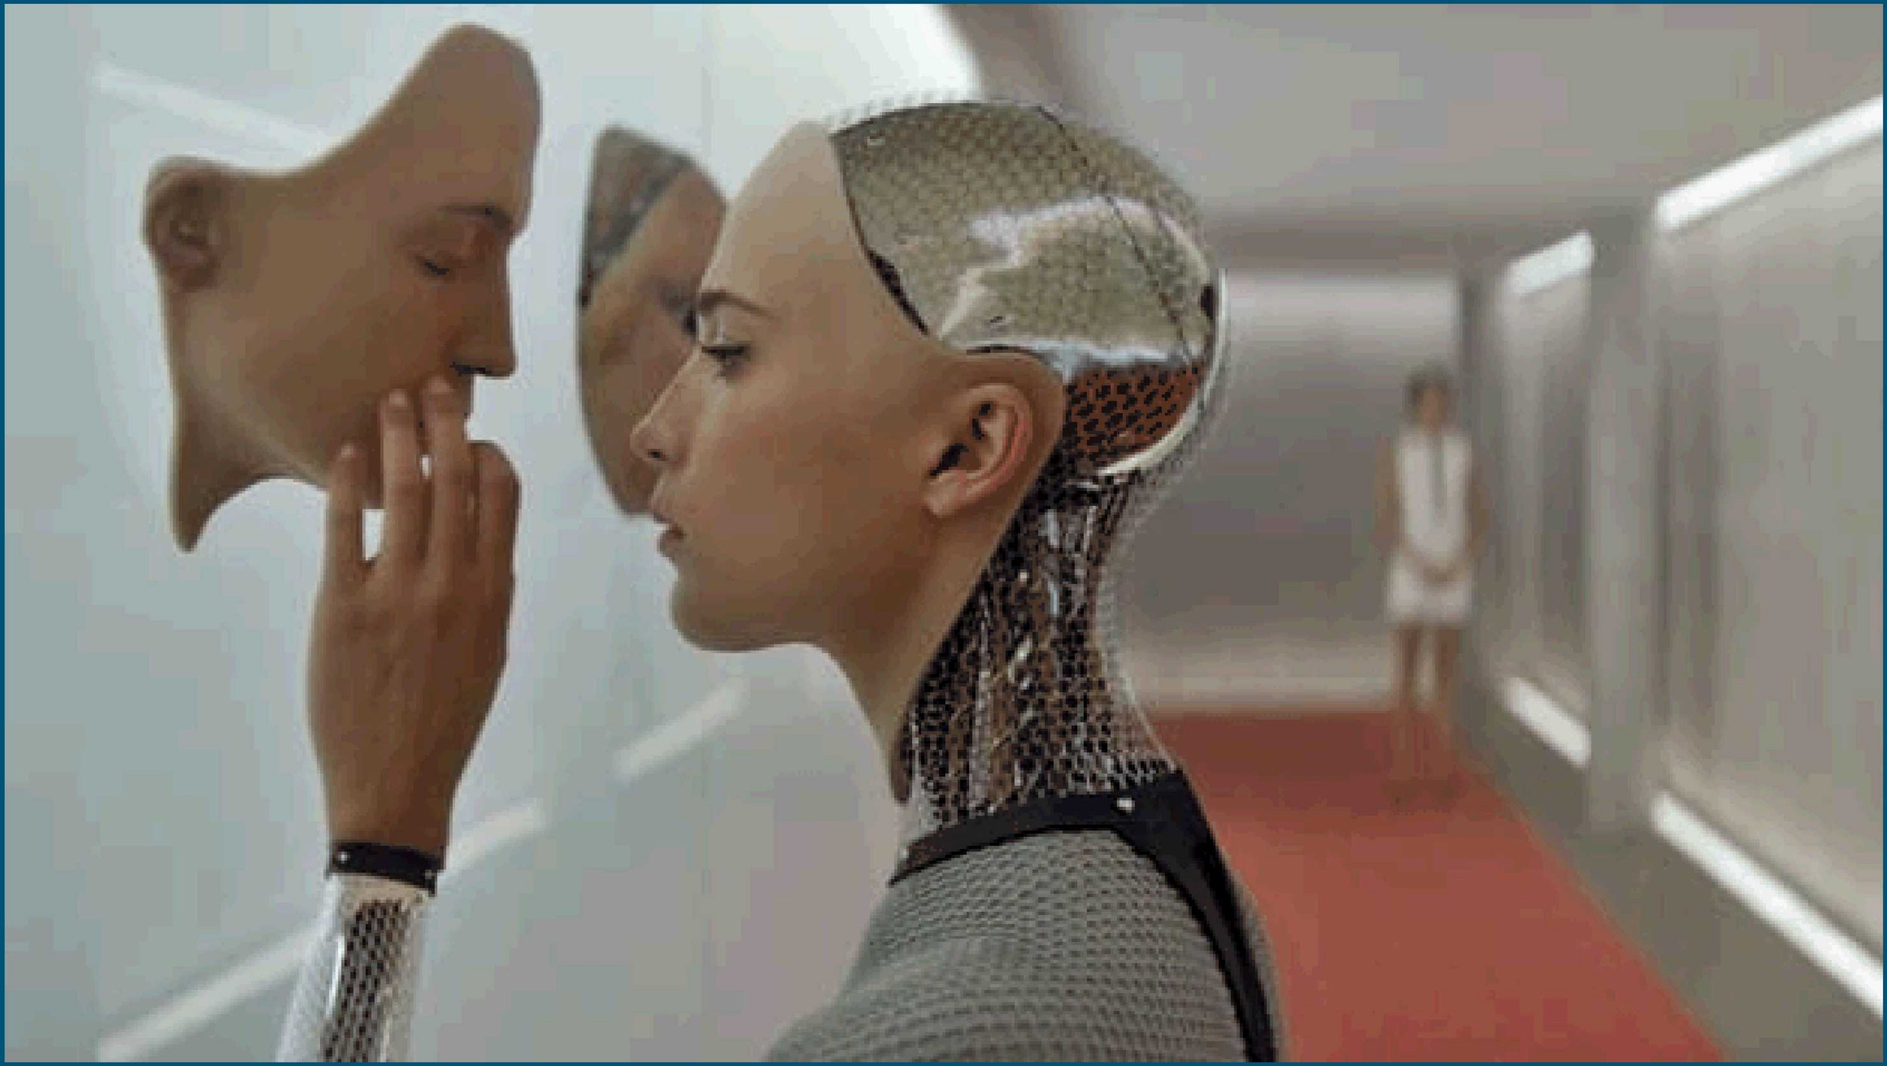

https://vaguevisages.com/2017/04/14/the-original-sin-in-ex-machina-ava-is-the-origin/

https://www.dexlabanalytics.com/blog/amazon-launches-deepracer-an-autonomous-machine-learning-car

http://www.artificialhumancompanions.com/autonomous-deep-learning-robot-the-missing-instructions/

Hardware Acceleration of Adaptive Neural Algorithms (HAANA)

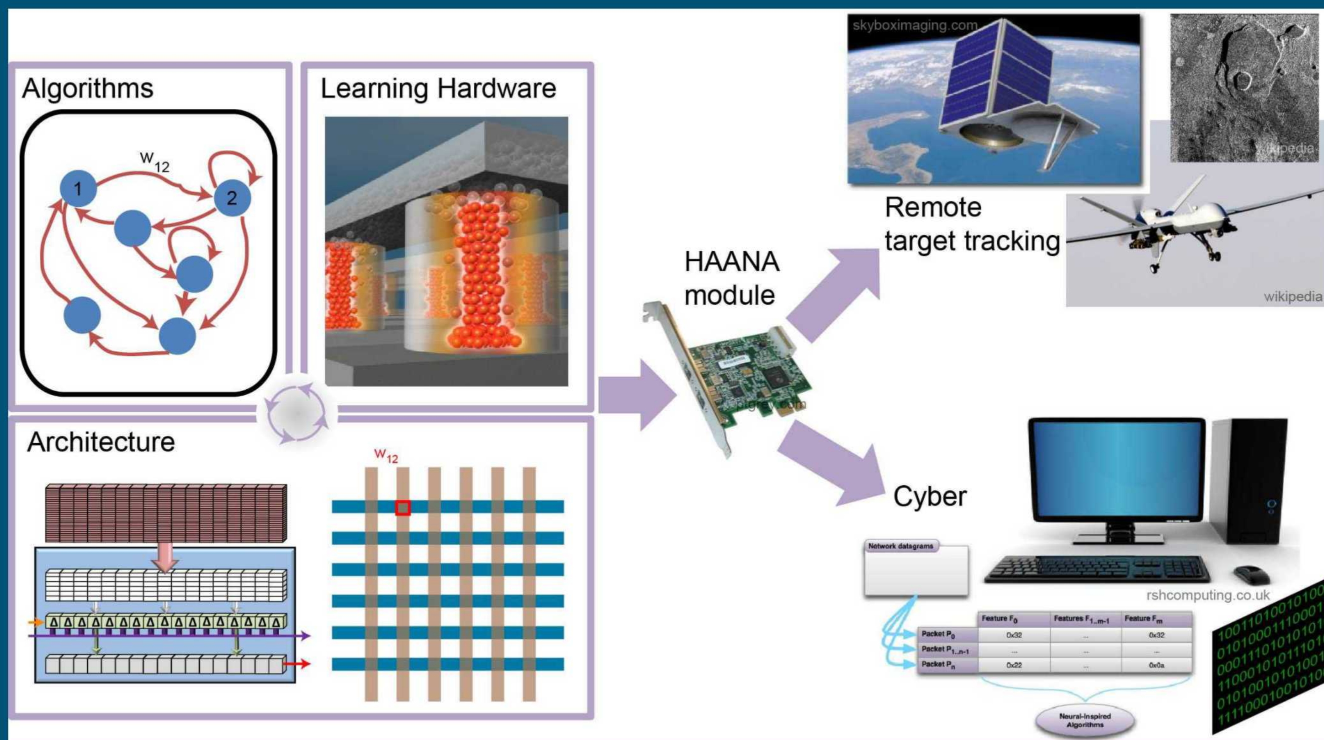

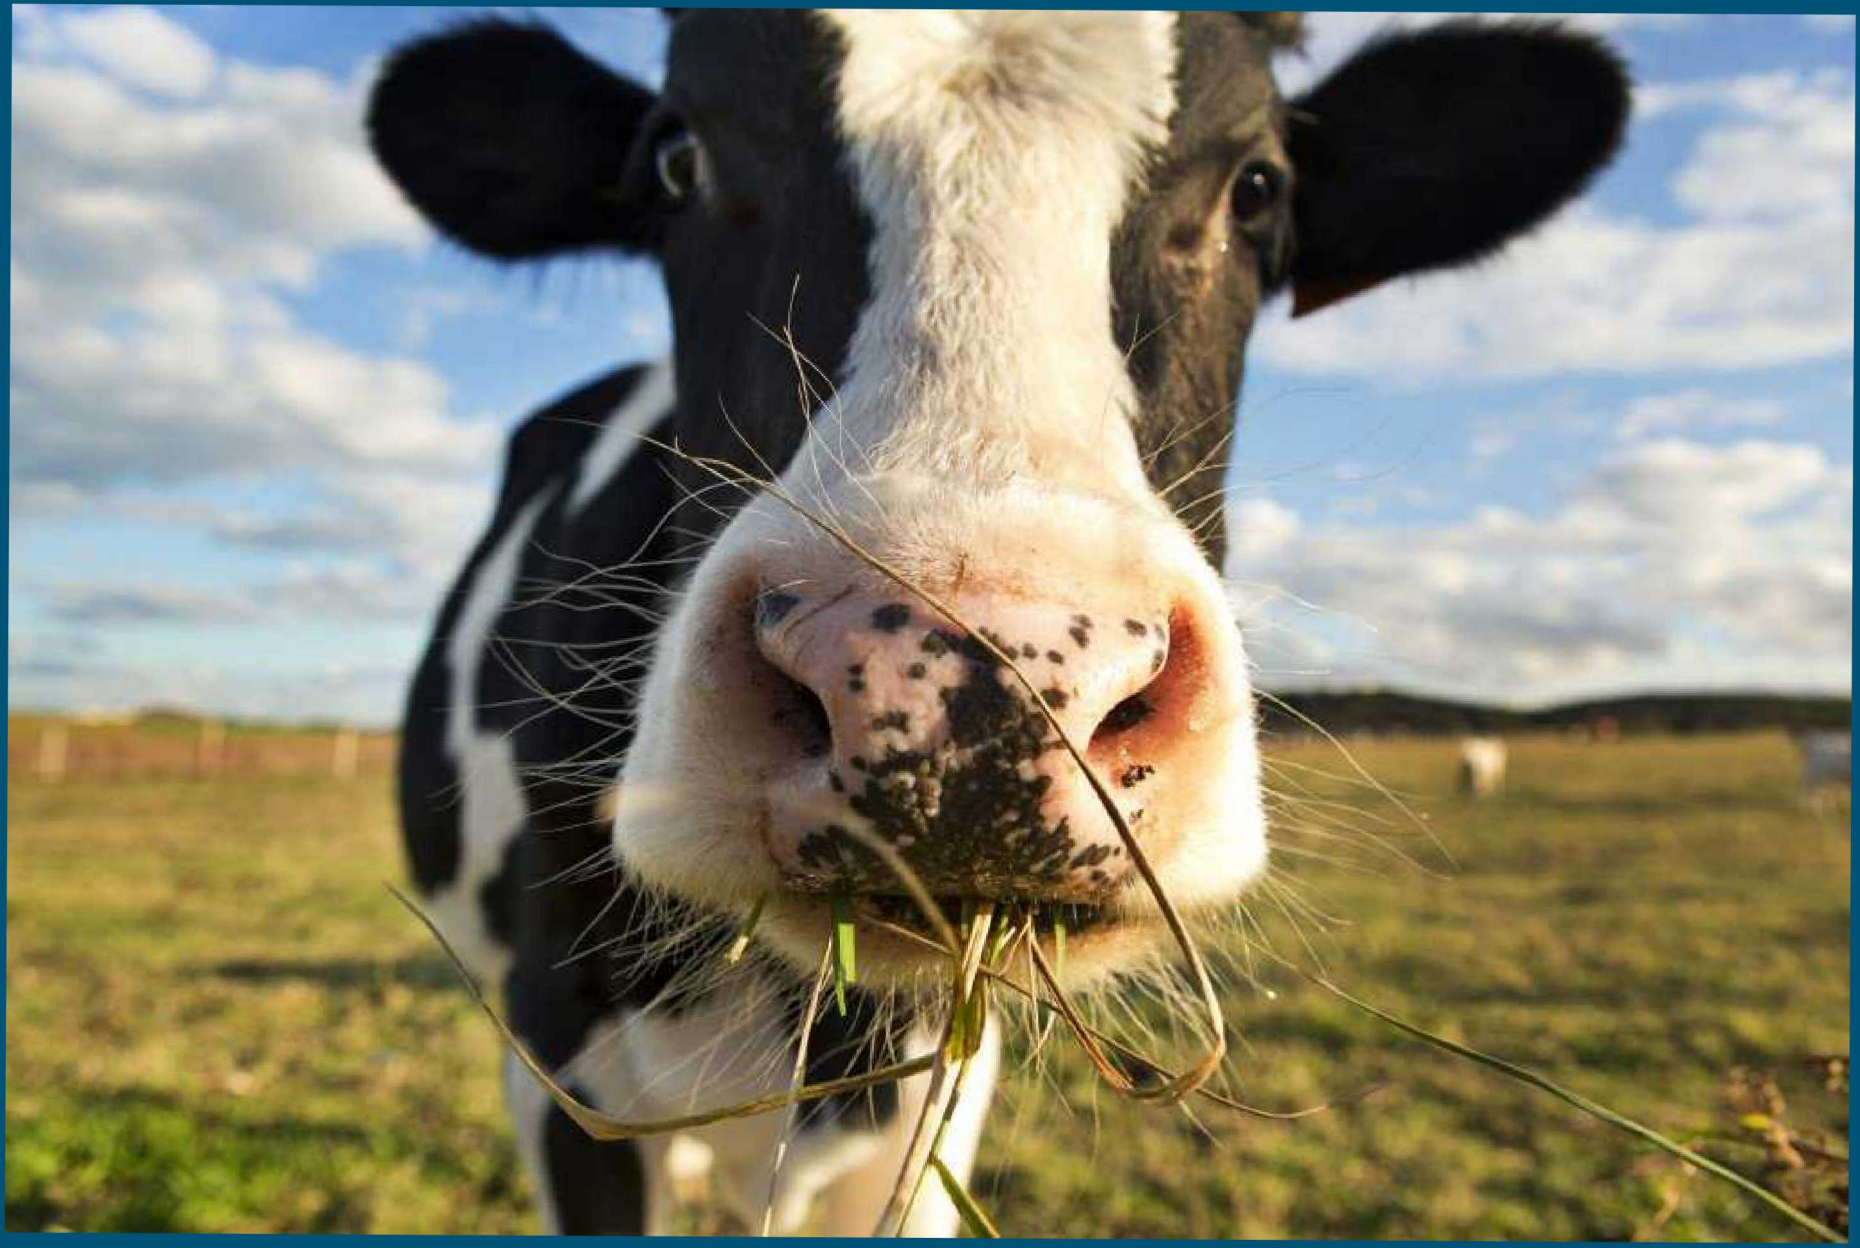

https://steemit.com/blog/@lapilipinas/why-does-a-cow-chew-its-cud

Algorithm 1: Algorithm for Self-Updating Existing Model

Input: Current model, m; Window size, w;
Data stream, D; Algorithm, A

Output: Updated model: $\hat{m}$

```
P = {}
N = {}
for i = 1 to w do
    x ~ D                       ▷ Draw event from stream
    ŷ = m(x)                    ▷ Get model's prediction
    if ŷ == 1 then
        ▷ Add event to positive set
        P = P ∪ {x}
    else
        ▷ Add event to negative set
        N = N ∪ {x}
    end
end
m̂ = m ∪ A(P, N)                 ▷ Update model
return m̂
```

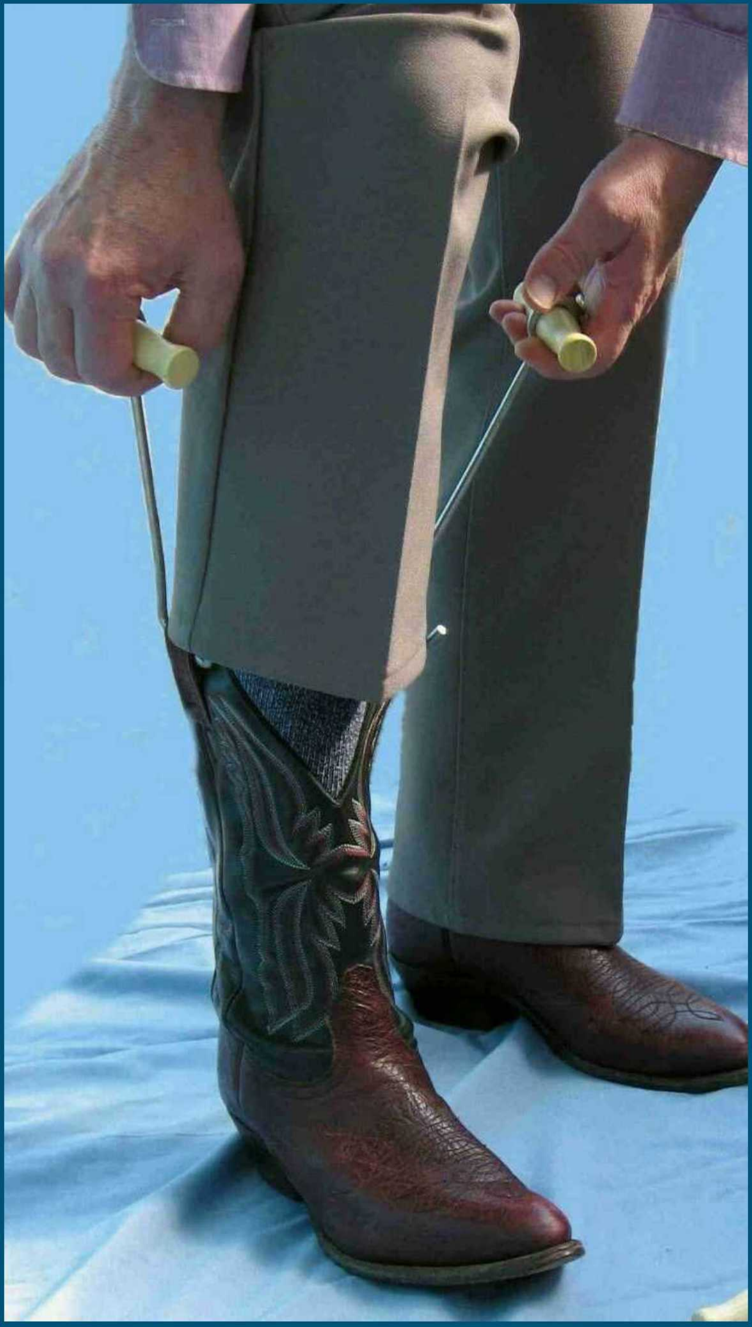

https://www.ebay.com/itm/Boot-hooks-bootstrap-pulls-pullers-lifters-western-riding-wraparound-pair-NEW-/311901884182

https://www.golfdigest.com/story/breaking-down-tiger-woods-new-swing

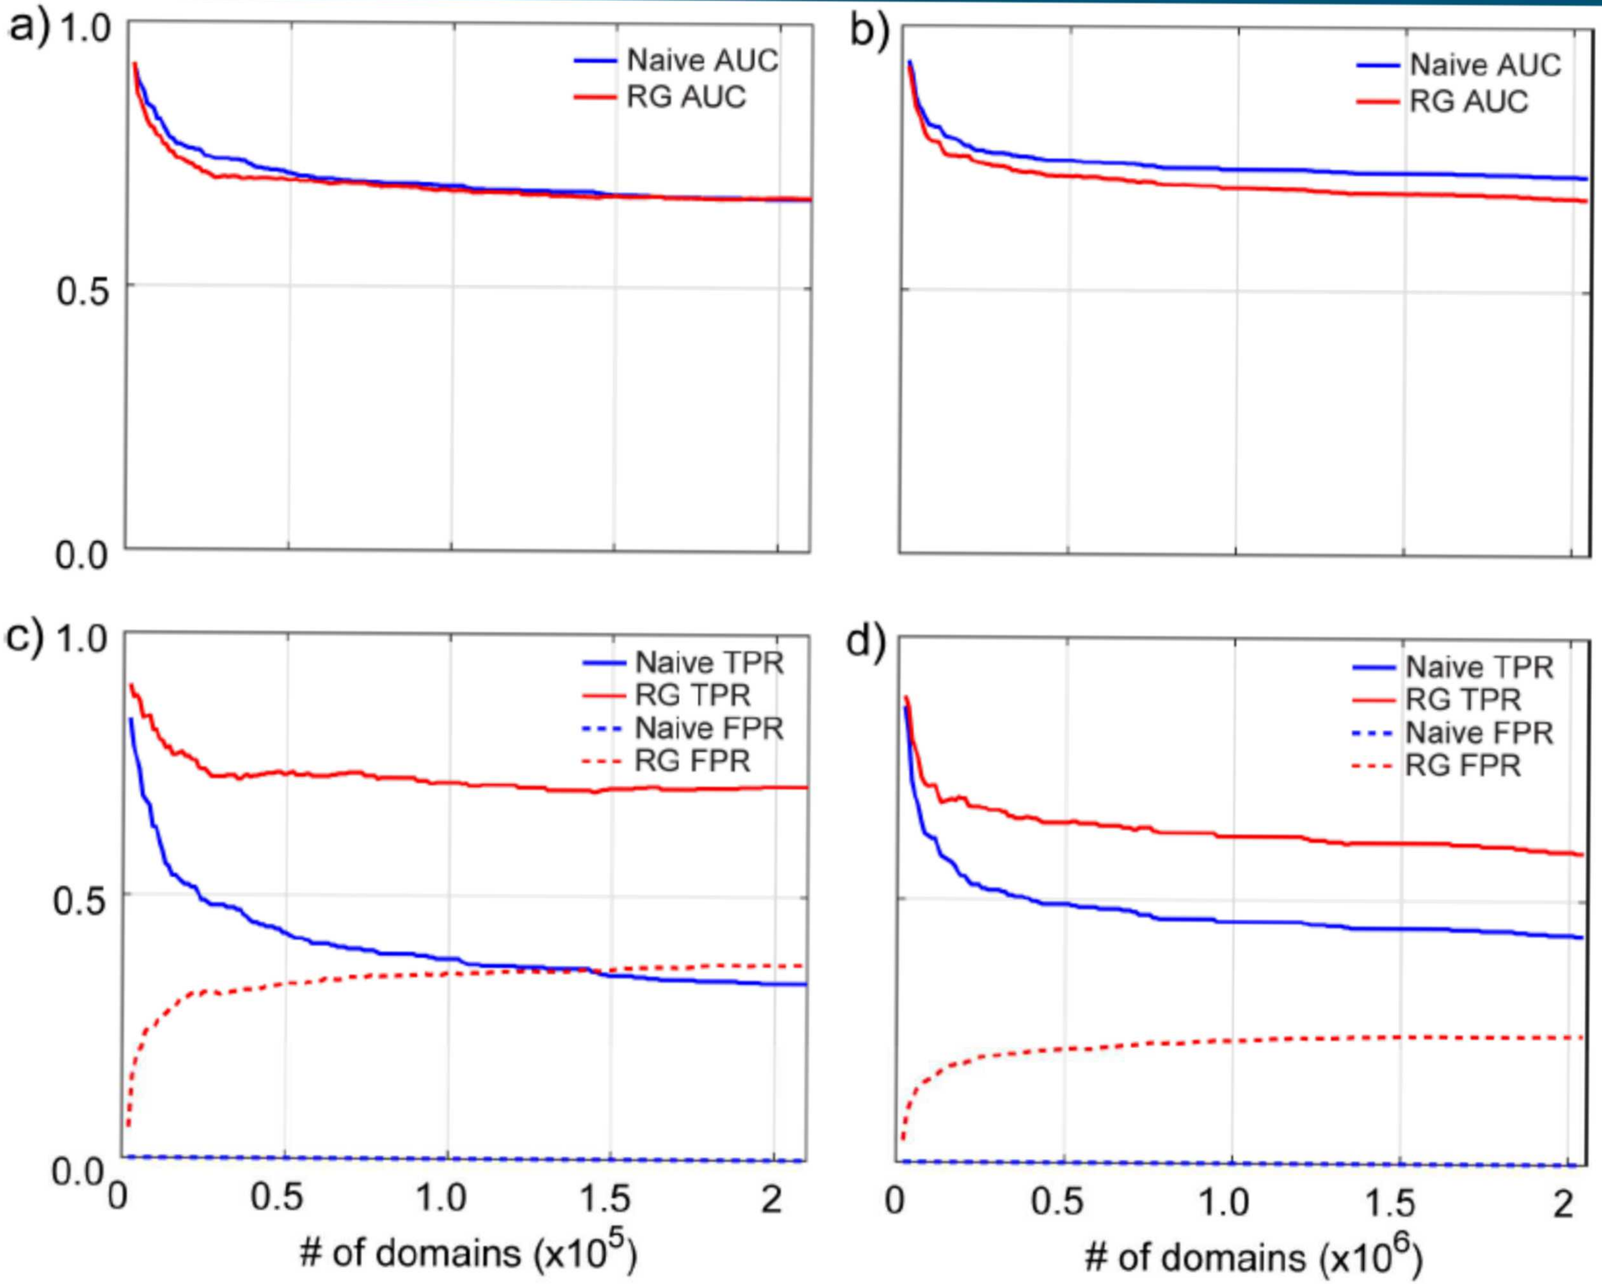

Fig. 3. AUC for a) window size 1,000 and b) window size 10,000. TPR/FPR for c) window size 1,000 and d) window size 10,000.


4 QUEENS
PARKING
FREM
CASI

https://www.mirror.co.uk/news/uk-news/groundhog-day-man-90-minute-memory-6044753

https://archive.siam.org/meetings/sdm15/

Concept Drift

- Concept drift – unforeseen changes in the relationships between input and output (“concepts”) variables.
  - Can be detrimental to model performance.
  - Can be sudden or gradual.
  - Can be natural or adversarial.

Probability of input x, given output y

Model – what is output *y*, given input *x*

$$P(y|x) = \frac{P(x,y)}{P(x)} = \frac{P(x|y)P(y)}{P(x)}$$

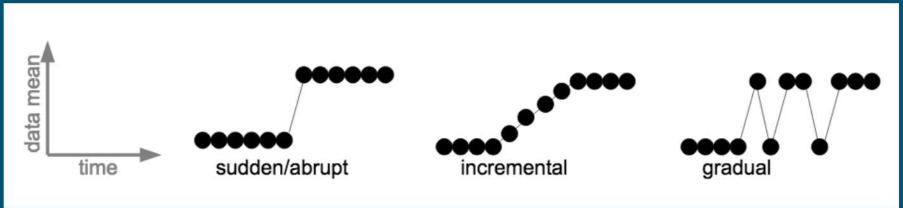

Different rates of concept drift

Effect of Concept Drift

- **Takeaway:** need to update or adapt models to maintain performance.

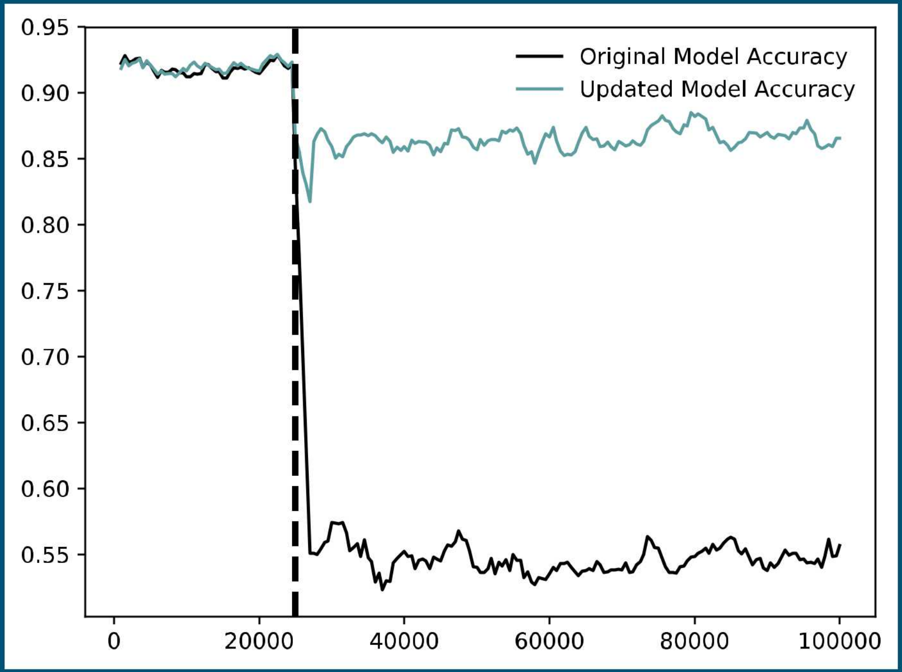

Performance under drift in *P(x|y)* over time

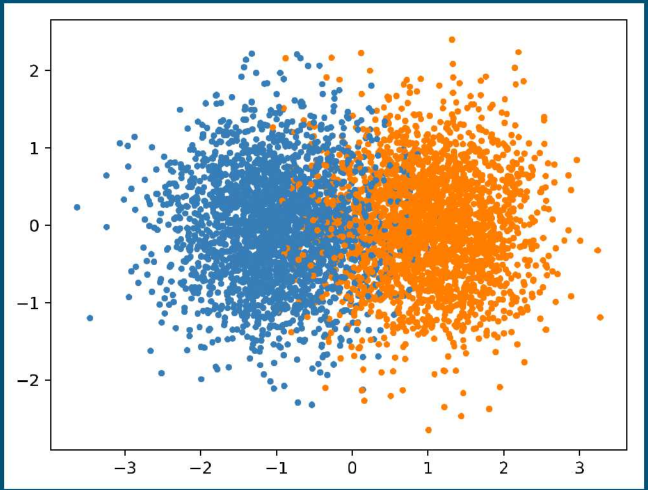

t=0

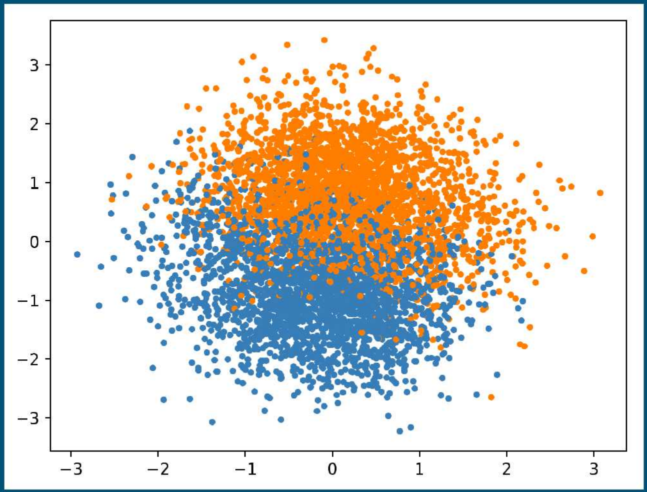

t=25000

Example: Malware Detection

- Developed model in 2012 to detect malicious software.
- Revisited in 2018 and updated model.
  - Updated (2018) model accuracy: 96%
  - Original (2012) model accuracy: 63%

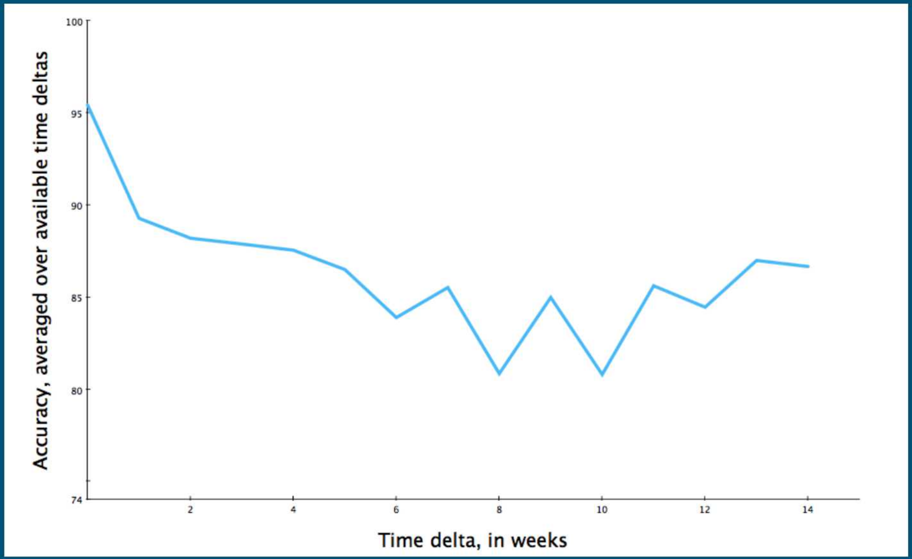

Natural concept drift

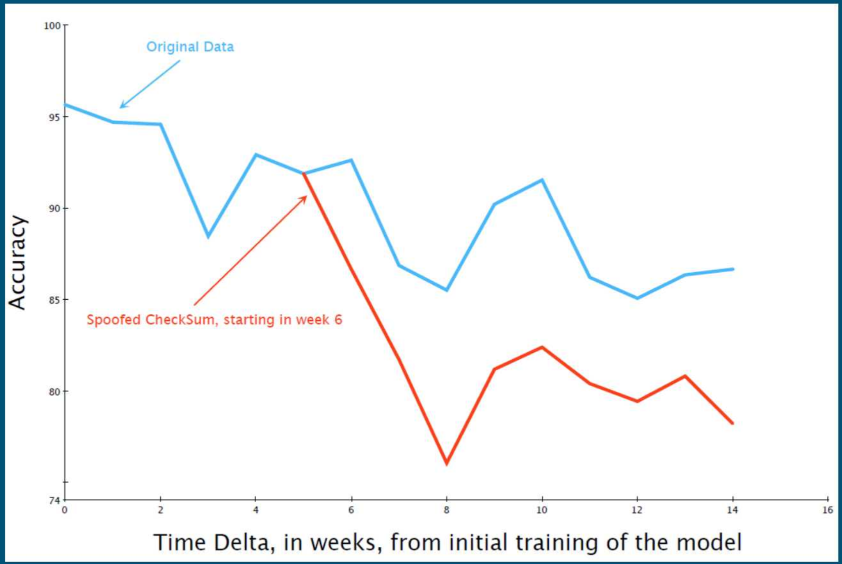

Adversarial concept drift

Effect of Label Noise

- Label noise – data is mislabeled.
- **Takeaway:** need to ***correctly*** update or adapt models to maintain performance.

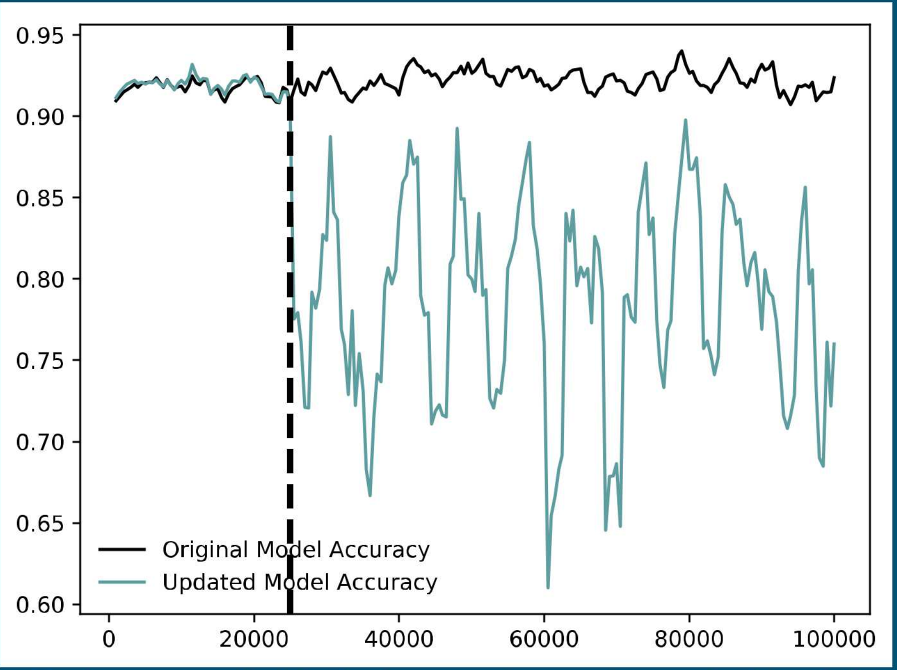

Performance with label noise over time and no underlying drift

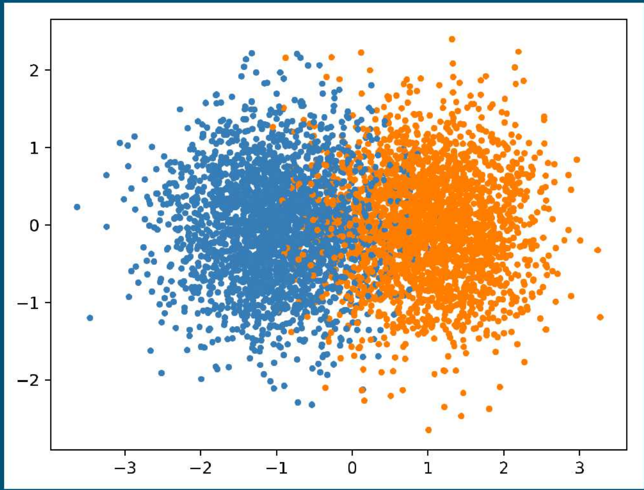
t=0

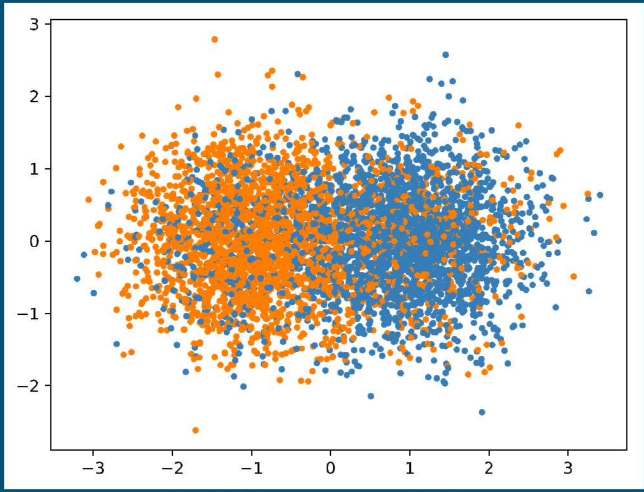
t=25000, corrupted data

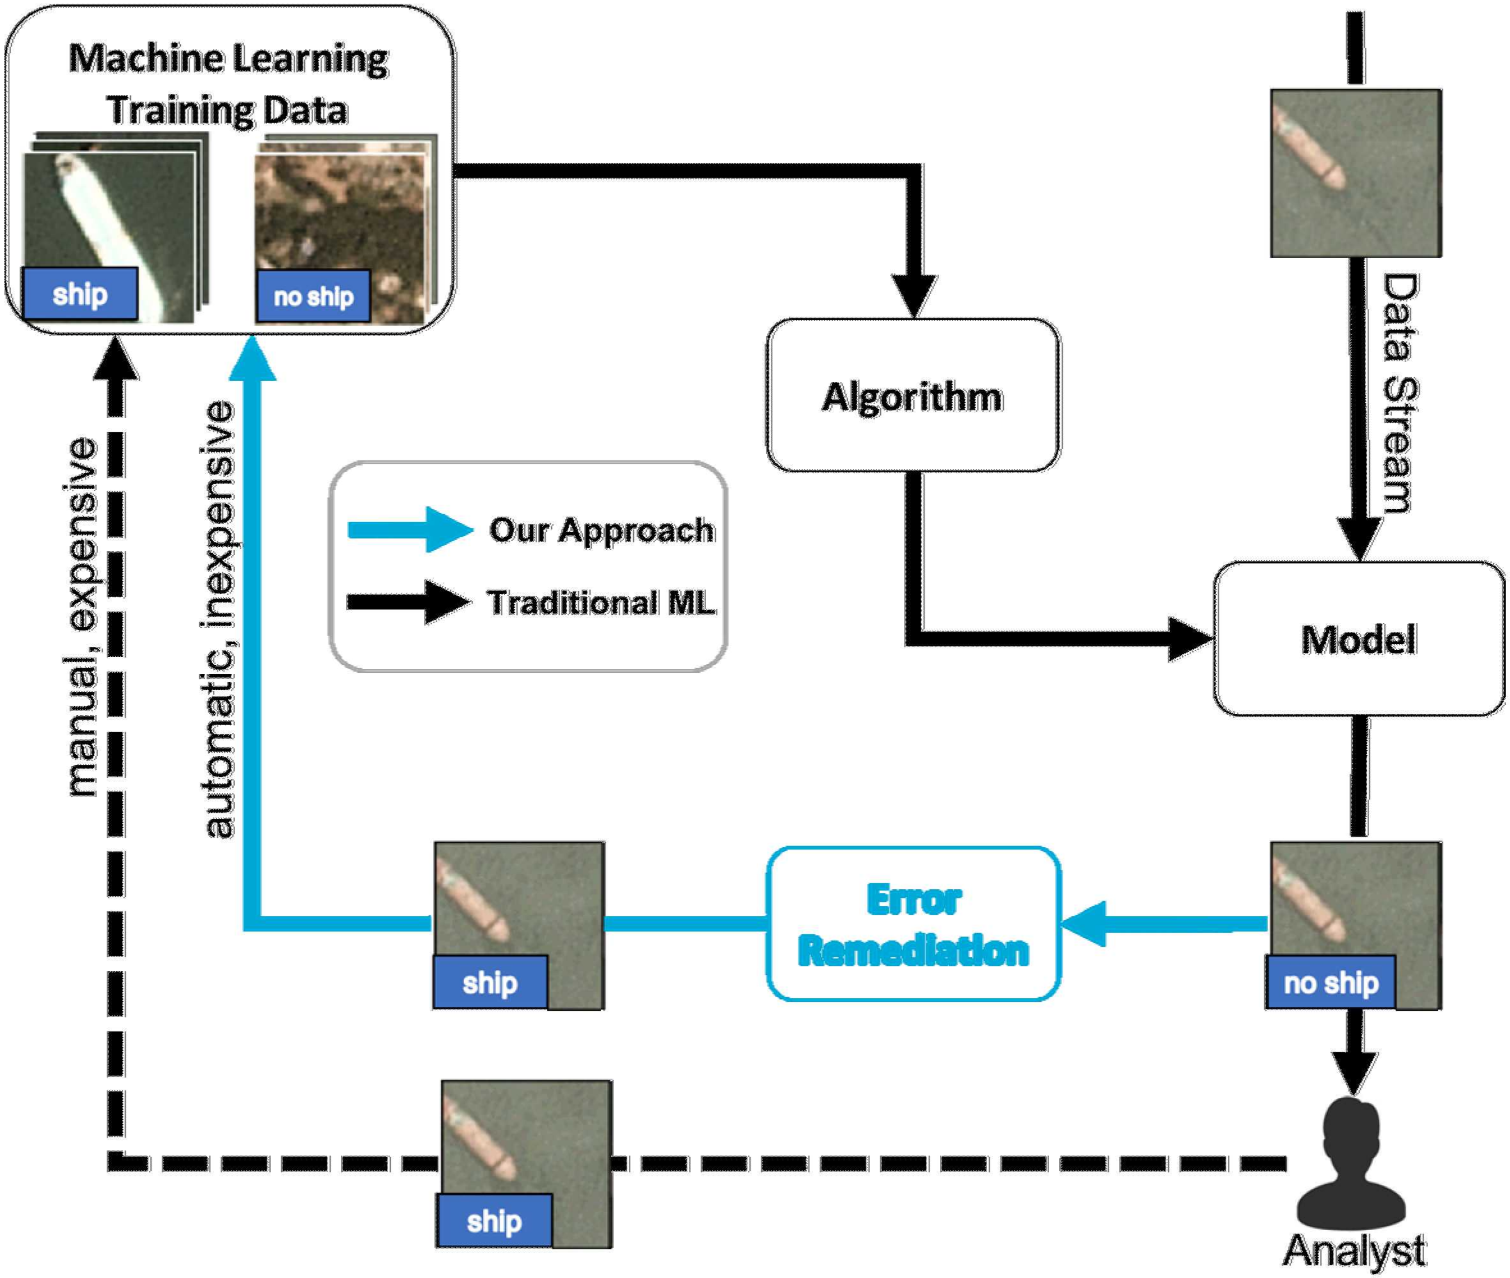

Machine Learning Training Data
ship
no ship
Algorithm
Data Stream
Model
Our Approach
Traditional ML
manual, expensive
automatic, inexpensive
Error Remediation
ship
no ship
ship
Analyst

Description of the data used

- Synthetic, two-dimensional data with mechanisms to introduce label errors and concept drift
- Kaggle “Ships in Satellite Imagery” dataset:
    - Features – 19,200 integers representing pixel intensities in red, green, and blue channels (6,400 values for each)
    - Ships – 1,000 images (25% of data)
    - Non-ships – 3,000 images (75% of data)

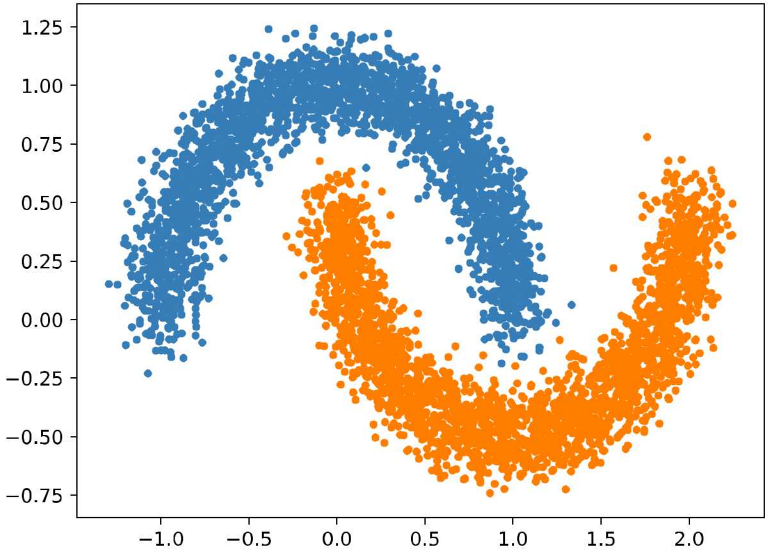

Results – Benefit(s) of Self-updating

- Self-updating (orange line) provides performance boost over not utilizing unlabeled data (blue line).
  - Benefit is more pronounced with less initially-labeled data.
- Self-updating also approaches performance upper-bound (green line) much faster.

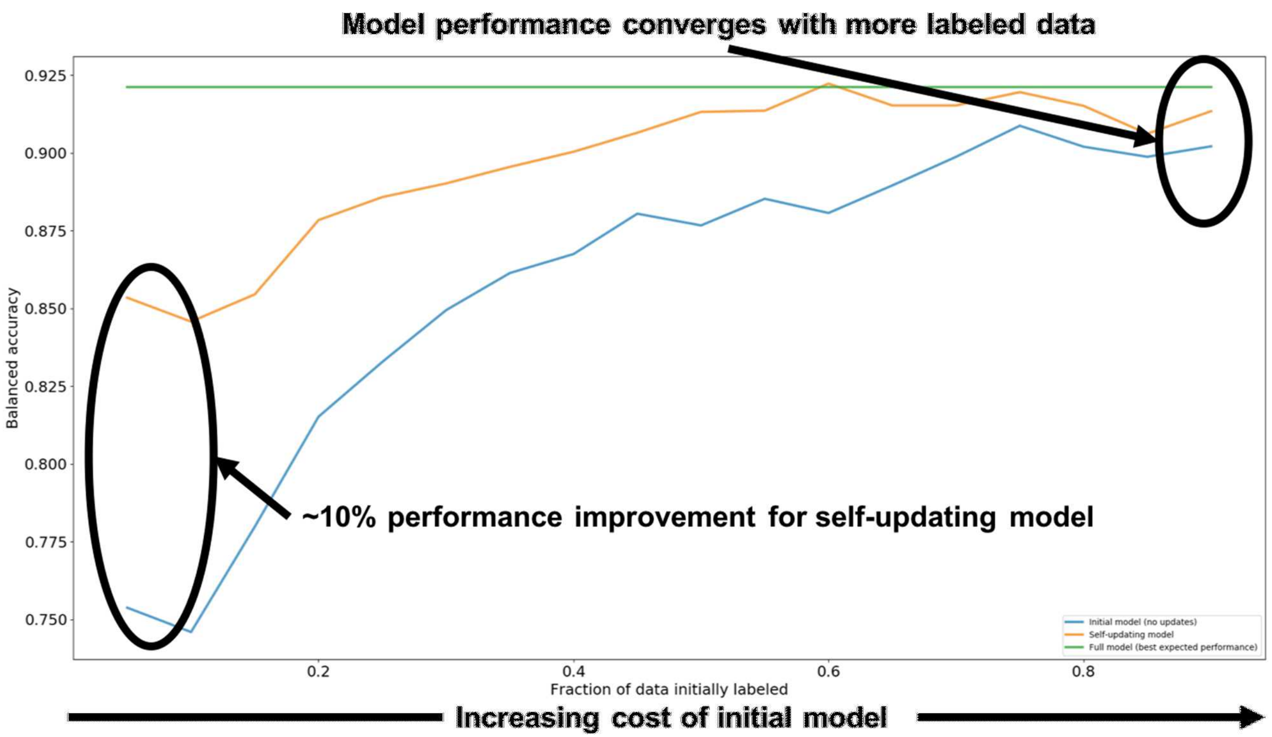

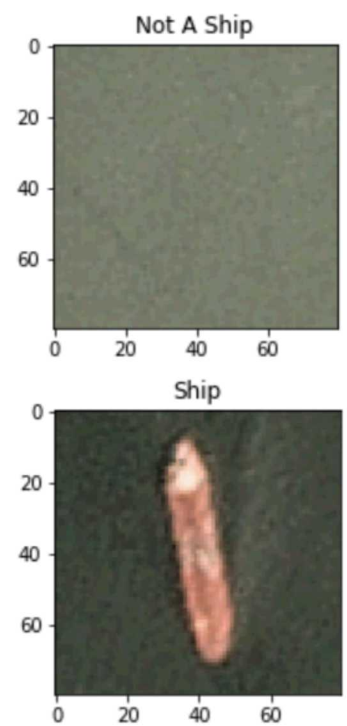

Results – Benefits of SUMER

- Self-updating with error remediation increased performance by 5% and improved decision boundary.

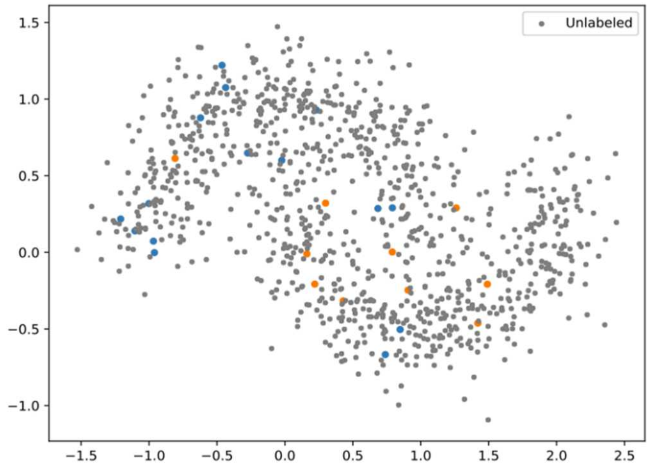

Initial data w/ 20% label noise

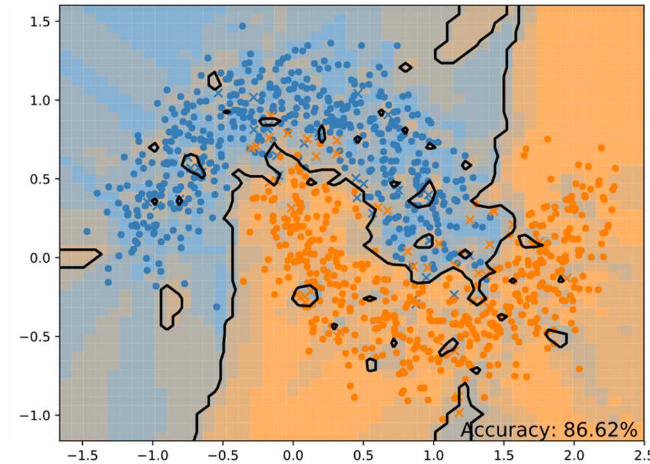

Model w/o self-updating/remediation

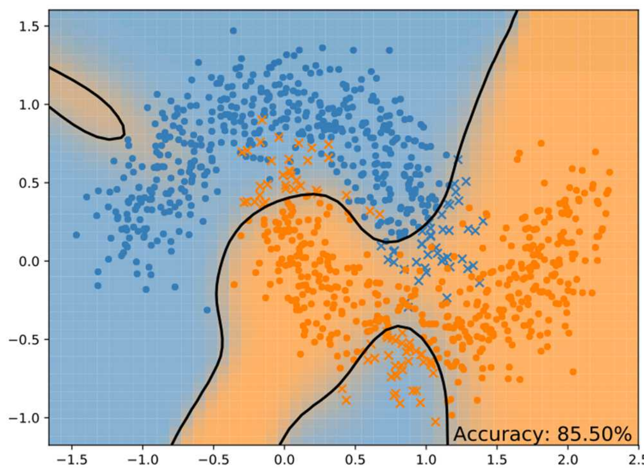

Self-updating without remediation

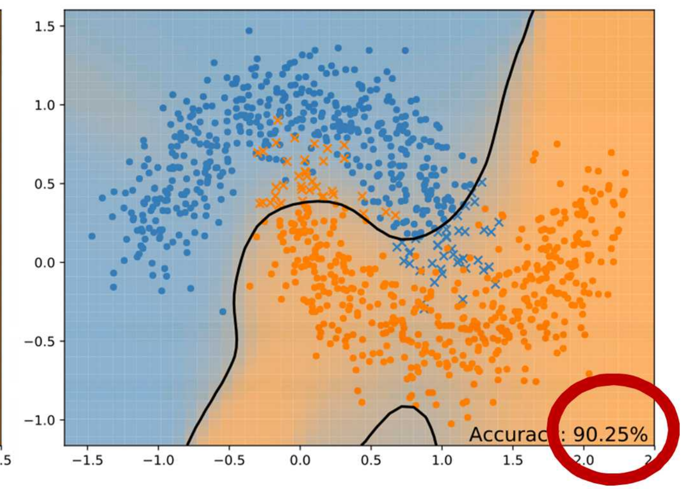

Self-updating with remediation

Results – Benefits of SUMER over Time

- Self-updating by itself only provides marginal improvement.
- Upper Bound Performance – model is updated with ground-truth labels (i.e., label noise is removed).
- Self-updating with remediation provides best performance.

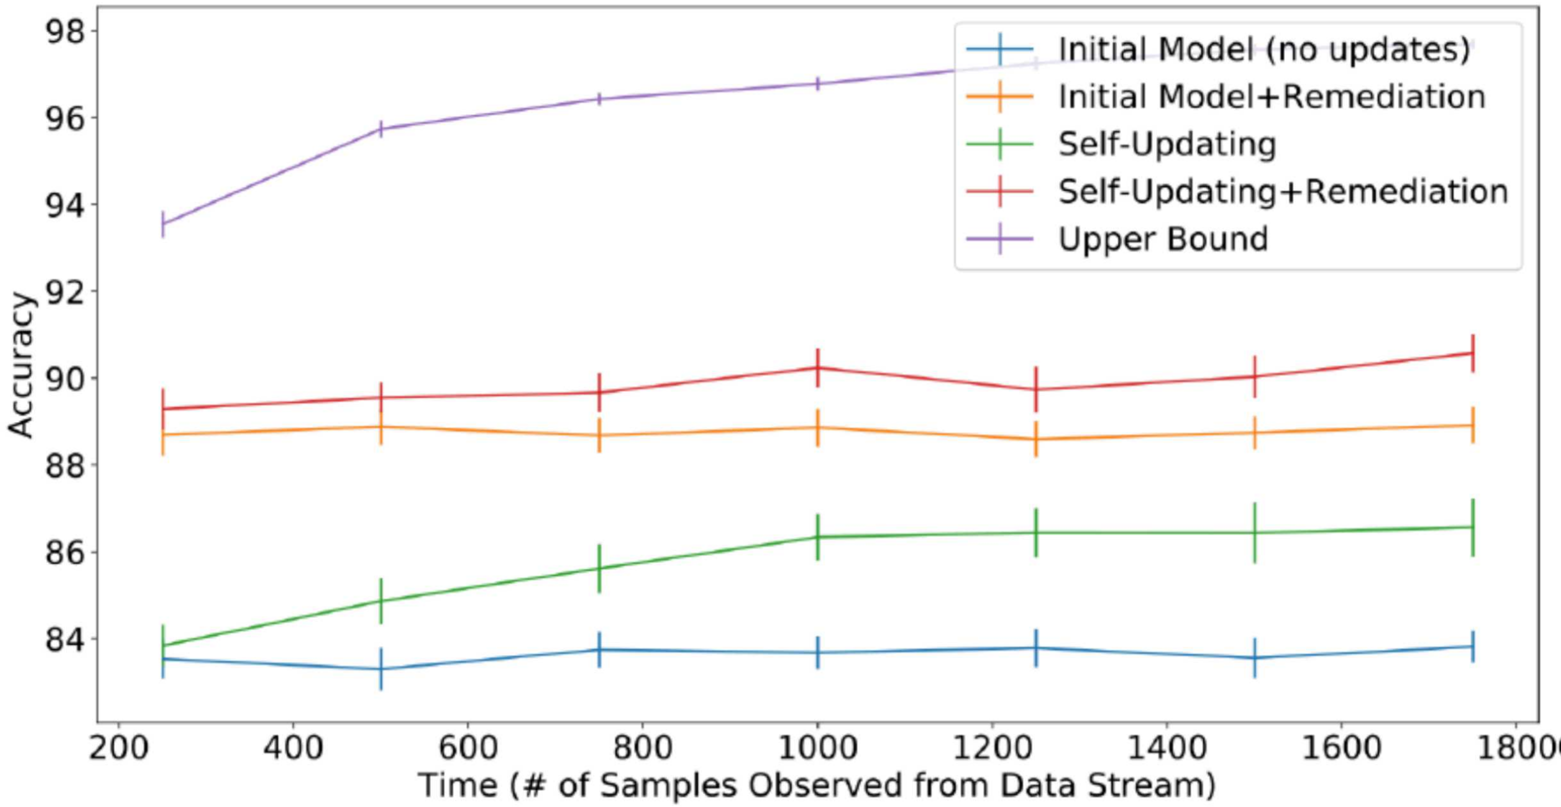

Benefit of SUMER on Kaggle ships data with 20% label noise

Potential Issues with SUMs and Label Correction

SUMs: the initially labeled data points can dramatically impact performance.

- We saw as much as a 15% difference in performance based on the specific points that were initially labeled.

Label correction: if the prediction model and the label correction model are "coupled", i.e., they make the same predictions on all or nearly all of the data points, then little value is provided by label correction.

- $P(\hat{Y}|Y) \sim 0 \text{ and } P(\hat{X}|X) \sim 0$
- One possible solution to "model coupling" is to build the prediction and label correction models with different views of the data.

The MAGE Project and Future Work

MAGE

- Take in overhead imagery of multiple modalities and highlight objects that may be of interest to analysts/operators.
- Automate machine learning pipeline as much as possible.
- Determine how to improve pipeline given feedback from humans-in-the-loop.
- Integrates a variety of techniques, e.g., few-shot learning, SUMs, label correction, active learning (modified), and model calibration.

Future Work

- Research the use of model calibration to improve confidences output by model. This facilitates the use of a threshold to determine if a prediction should be used as a label.
- Experiment with various approaches for novel concept detection.
- Implement and test other promising label correction techniques.
- Develop detectors for feature and label distribution shifts.
- Obtain funding to generate and test hypotheses for solving the model coupling problem.

Questions? Comments?