SAND2019-8352C

Self-updating Models with Error Remediation for Bandwidth-constrained Environments

JD Doak, Joe Ingram, Michael R. Smith, Craig Vineyard

2019 IEEE Space Computing Conference (SCC)

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Need to Update Deployed Models

Training data is taken from a specific location at a specific day and time.

Current data may differ from the training data.

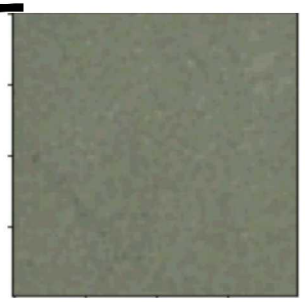

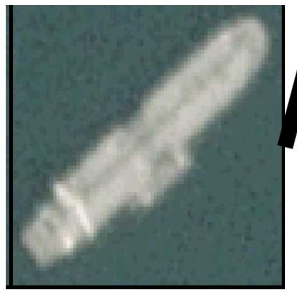

Algorithm

Ship/No Ship

Labelling all new data is infeasible in many cases:

- Lack of analyst time
- Inability to transfer all new data back
- May not have targets of interest

Traditional Method for Deployed Machine Learning

1. Induce model from an existing, labeled dataset.
2. Deploy model operationally to categorize data stream for analysts.
3. Human analyst verifies model and labels new data points.
   - Time intensive
   - Prone to human error
4. Determine if the model's performance has dropped below an acceptable level.
5. Repeat steps 1 - 4.

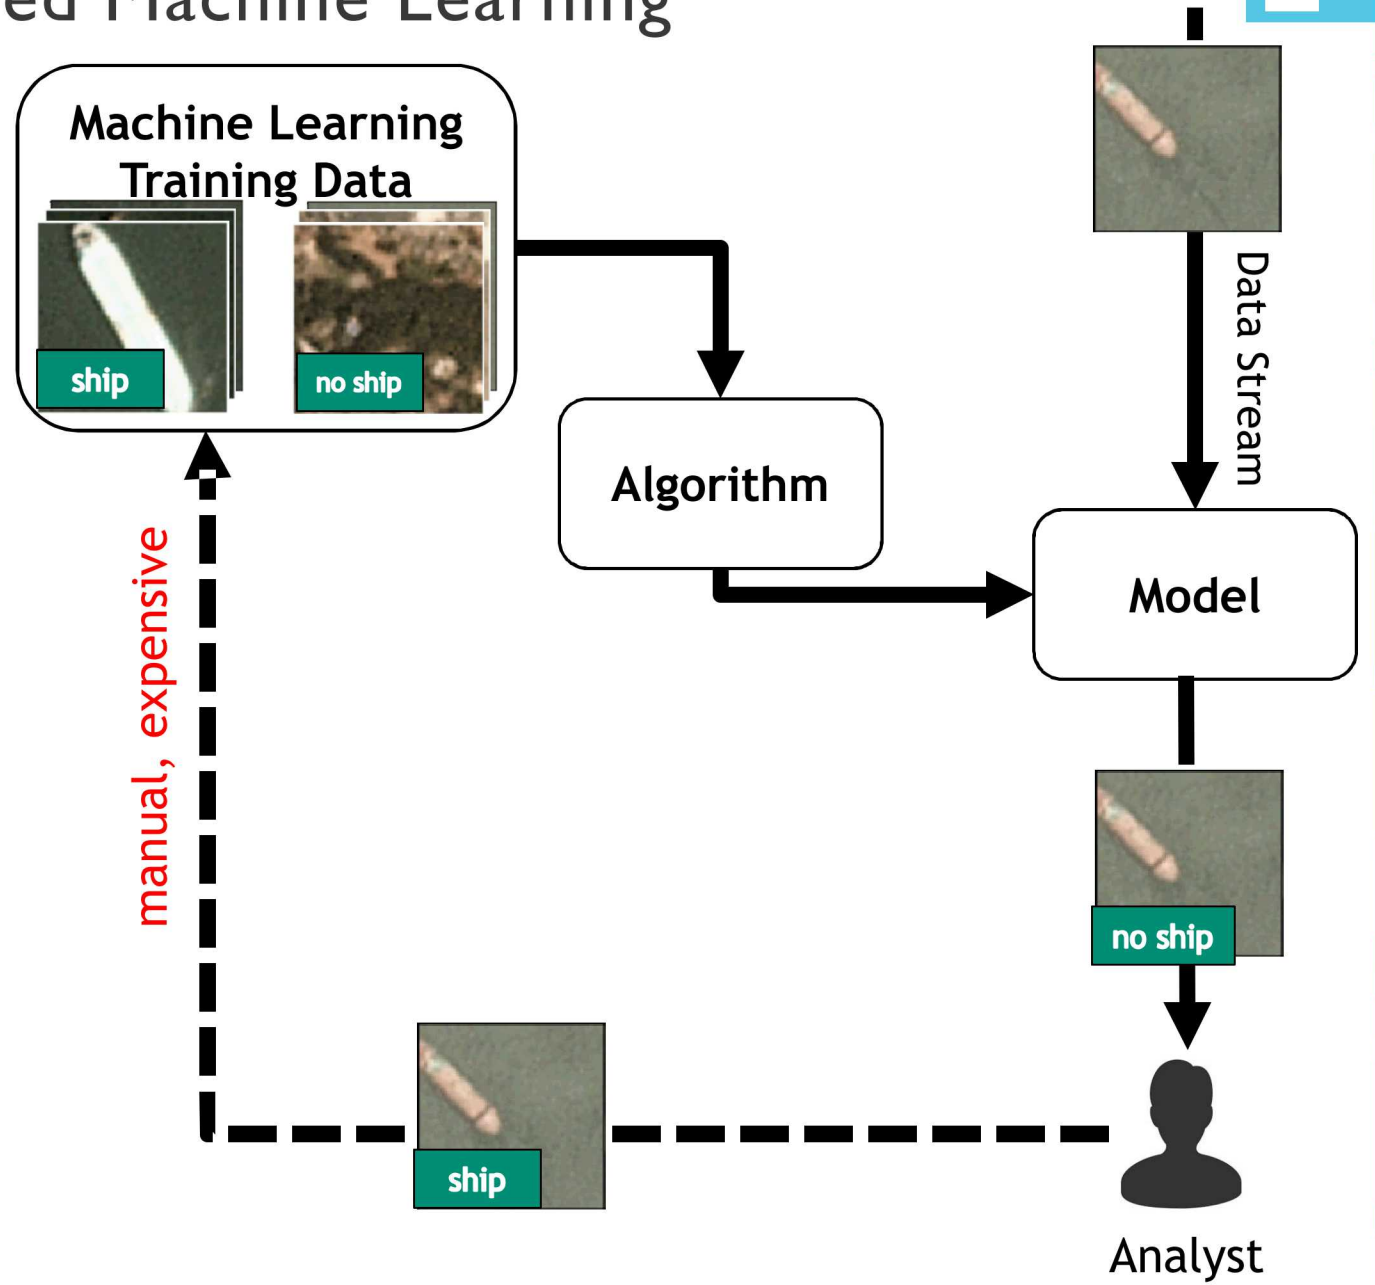

Self-Updating Models with Error Remediation (SUMER)

1. Induce model from an existing, labeled dataset.
2. Deploy model operationally to categorize data stream for analysts.
3. Perform error remediation as necessary using one or more methods:
   - Augment feature space.
   - Use auxiliary model(s) to detect/correct extrapolation.
   - Use algorithm that is robust to label noise.
4. Add samples to retraining dataset that error remediation algorithms predict will most improve the performance of the learned model.
5. Determine if the model's performance has dropped below an acceptable level.
6. Repeat steps 1 - 5.

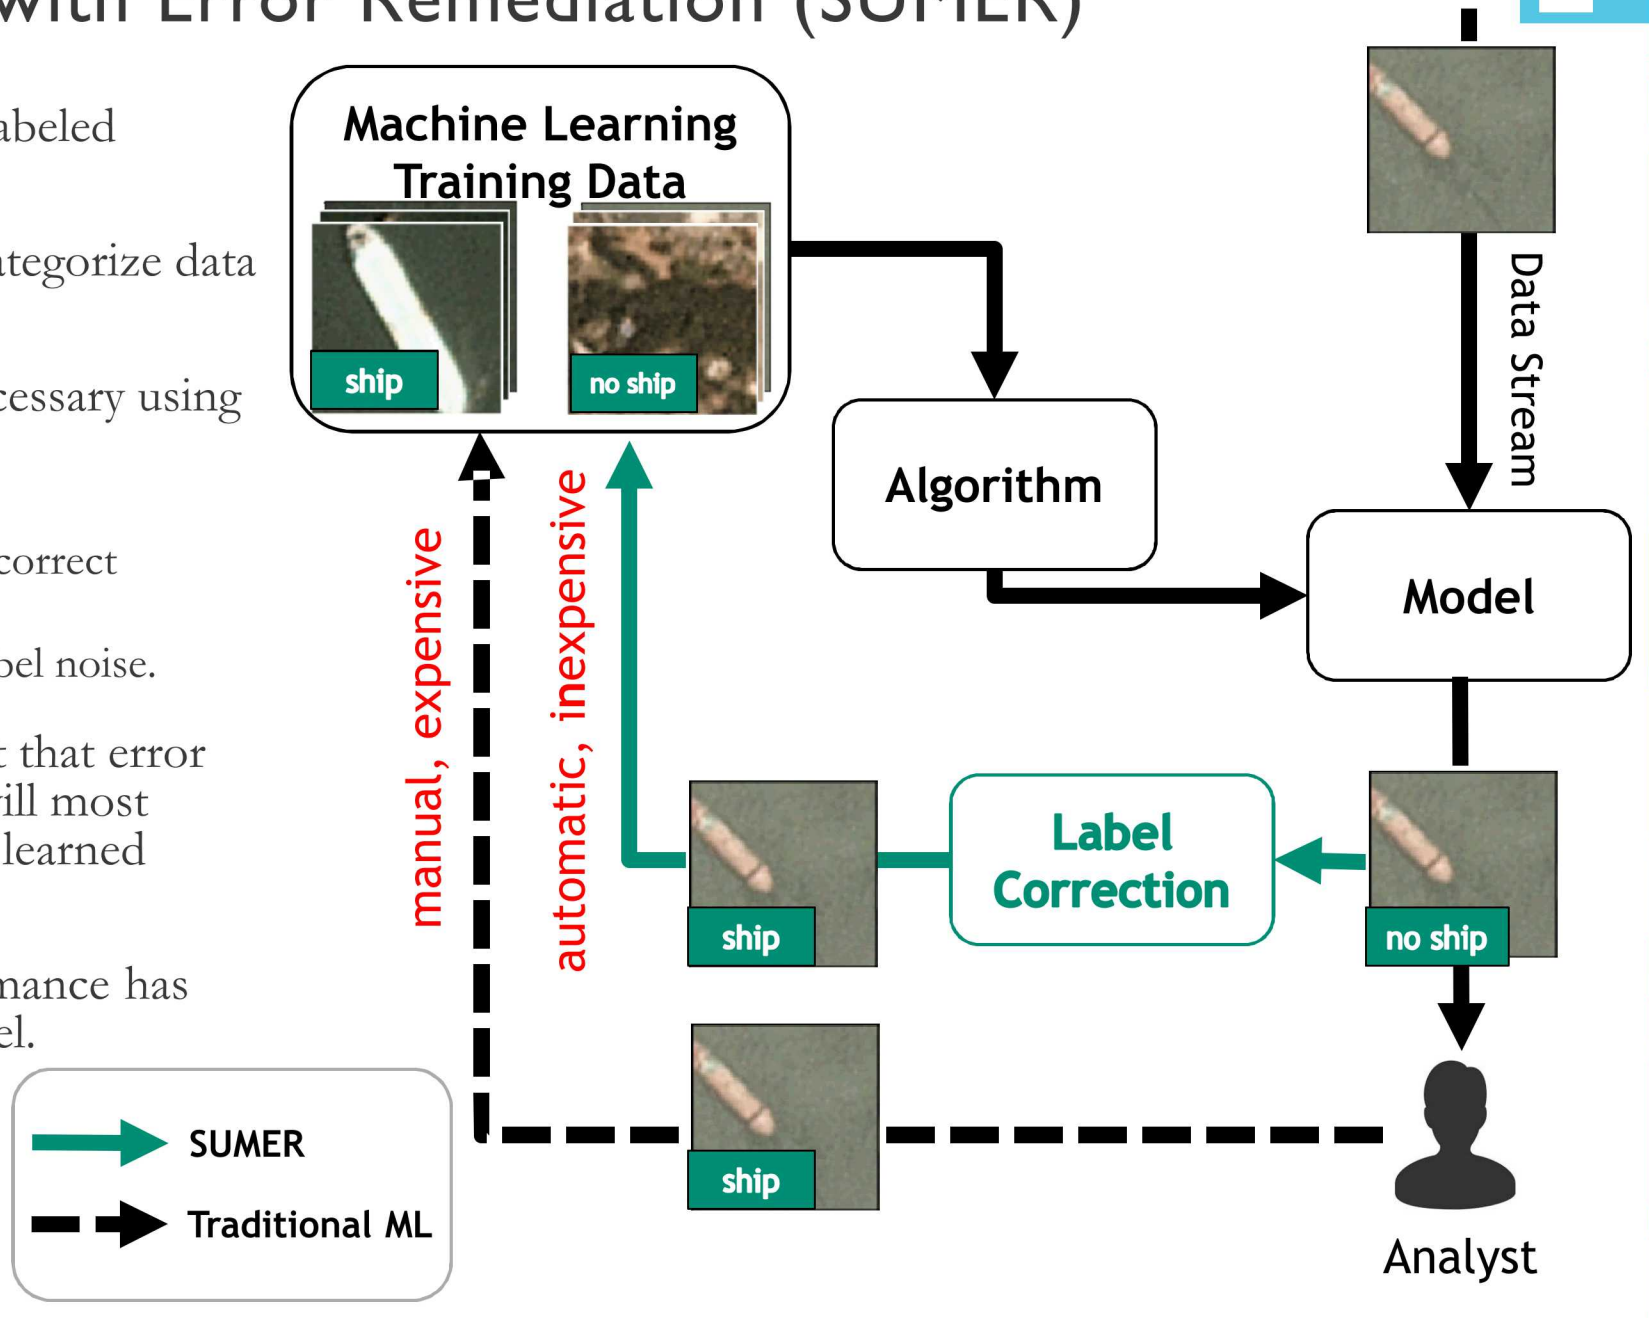

Remediation Methods

Semi-supervised learning

Model output calibration

Uncertainty quantification

Noise filtering

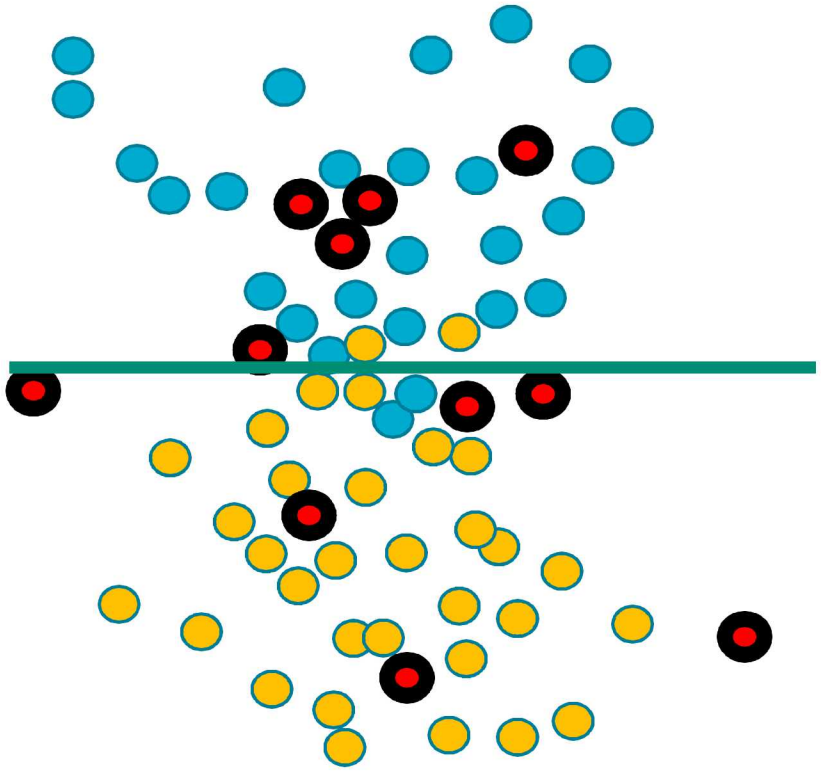

Discriminator between yellow and blue

Red are unlabeled data points.

Remediation Methods

Semi-supervised learning

Use the label from the model for subsequent training.

- Pseudo-labelling
- Self-training
- Co-training
- Active learning

Model output calibration

Uncertainty quantification

Noise filtering

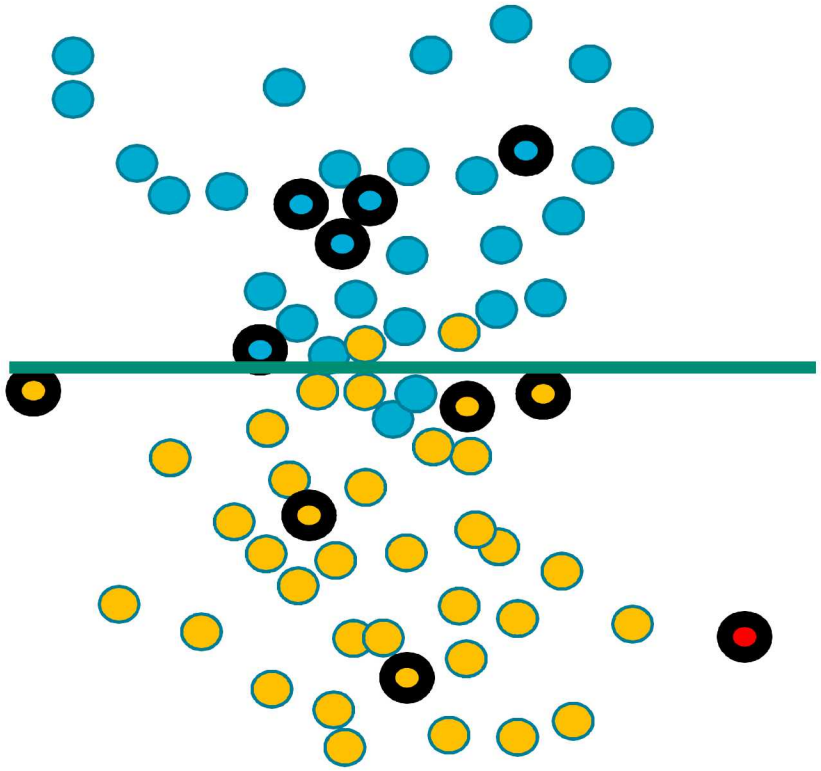

Discriminator between yellow and blue

Red are unlabeled data points.

Remediation Methods

Semi-supervised learning

Model output calibration

The confidence from a model can be highly uninformative (e.g., distance from the classification boundary).

- Trust

Uncertainty quantification

- For example, evaluate predictions from various models in an ensemble.

Noise filtering

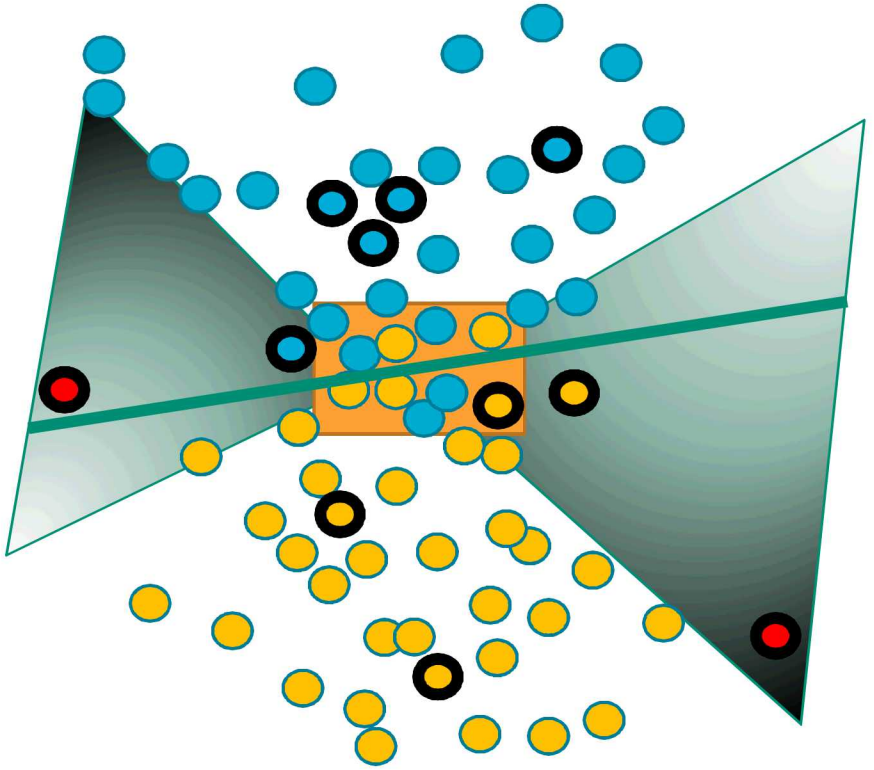

Discriminator between yellow and blue

Red are unlabeled data points.

Remediation Methods

Semi-supervised learning

Model output calibration

Uncertainty quantification

Noise filtering

Remove confusing data points and clean up the classification boundary.

- Rank Pruning

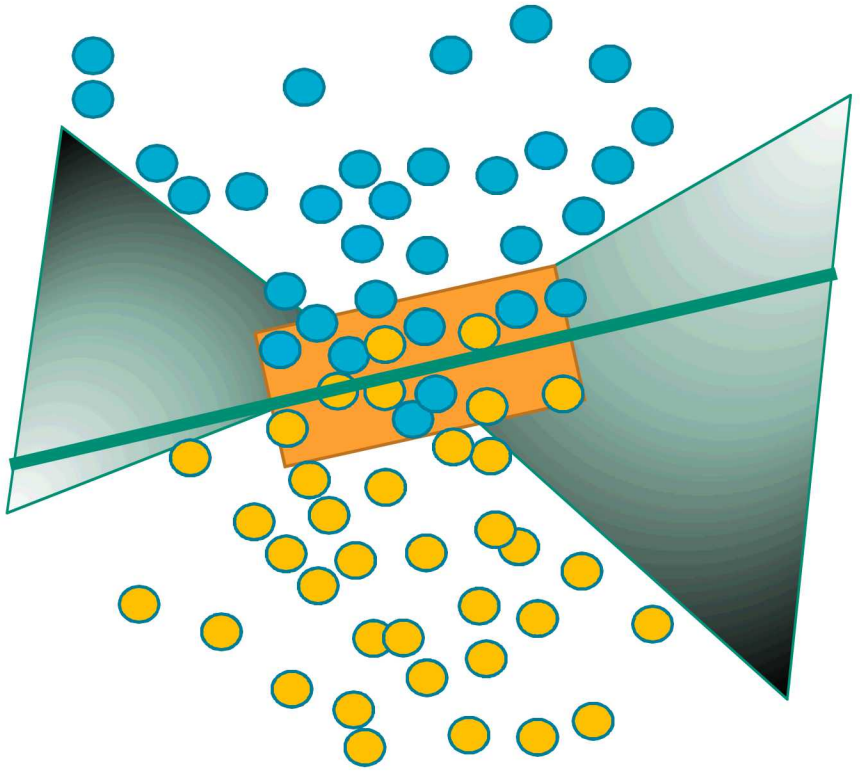

Discriminator between yellow and blue

Red are unlabeled data points.

Data

Synthetic data set

- Allows for systematic experimentation and has mechanisms to introduce label errors and concept drift.

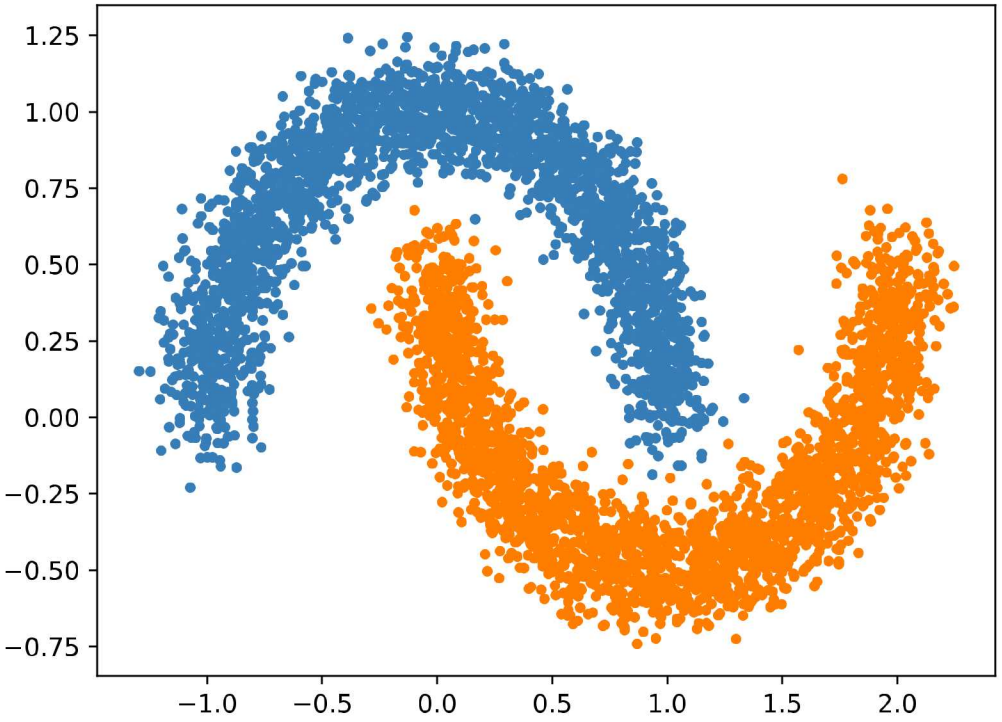

Kaggle "Ships in Satellite Imagery" dataset:

- 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: Synthetic Data with 20% Label Noise

Accuracy results:

- Baseline model: 86.62%
- Self-updating (no error remediation): 85.50%
- SUMER: 90.25%

Error remediation improves the decision boundary.

Baseline classifier and the base algorithm in SUMER is a Random Forest.

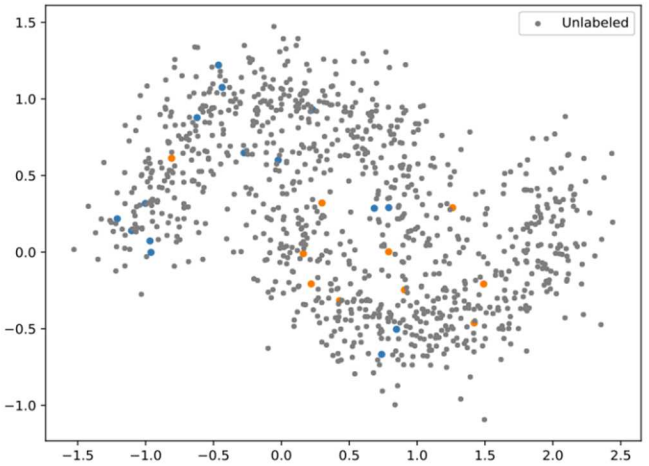

Initial data w/ 20% label noise

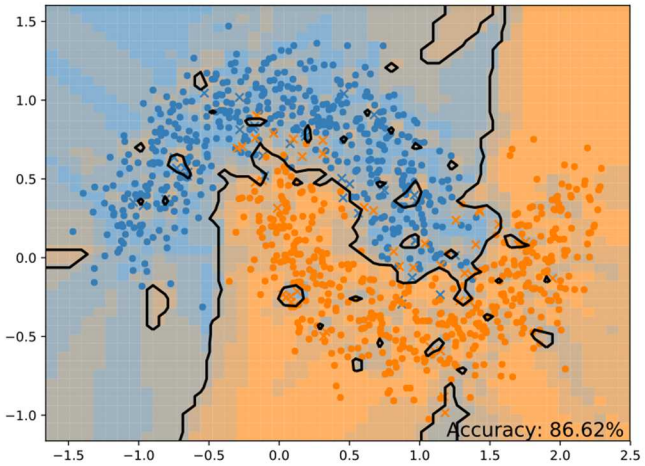

Model w/o self-updating/remediation

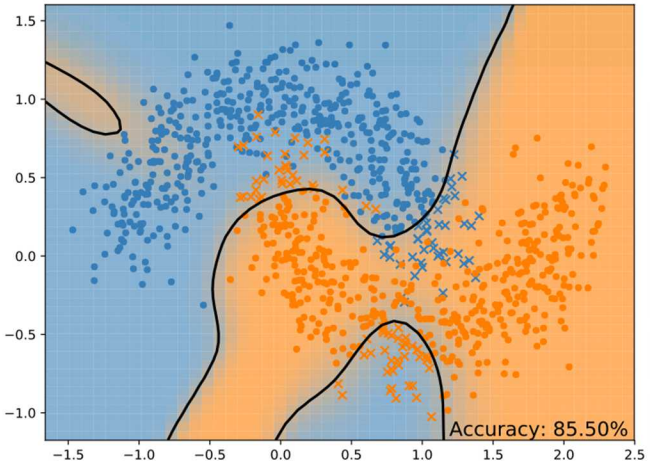

Self-updating without remediation

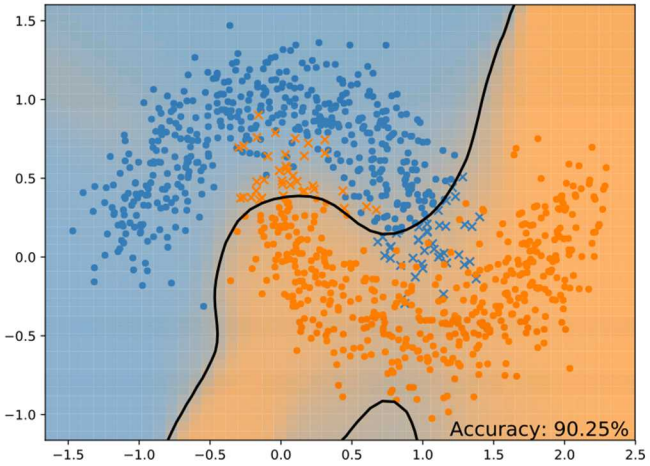

Self-updating with remediation

Results: Synthetic Data over Time

All models initially-trained on 200 data points with 25% label noise.

- Upper Bound Performance assumes no label noise.
- Each subsequent point represents an additional 200 unlabeled data points.
- Original model predicts on the new data points with no update to the model.

SUMER provides the highest classification accuracy.

The self-updating model without error remediation shows modest performance increases.

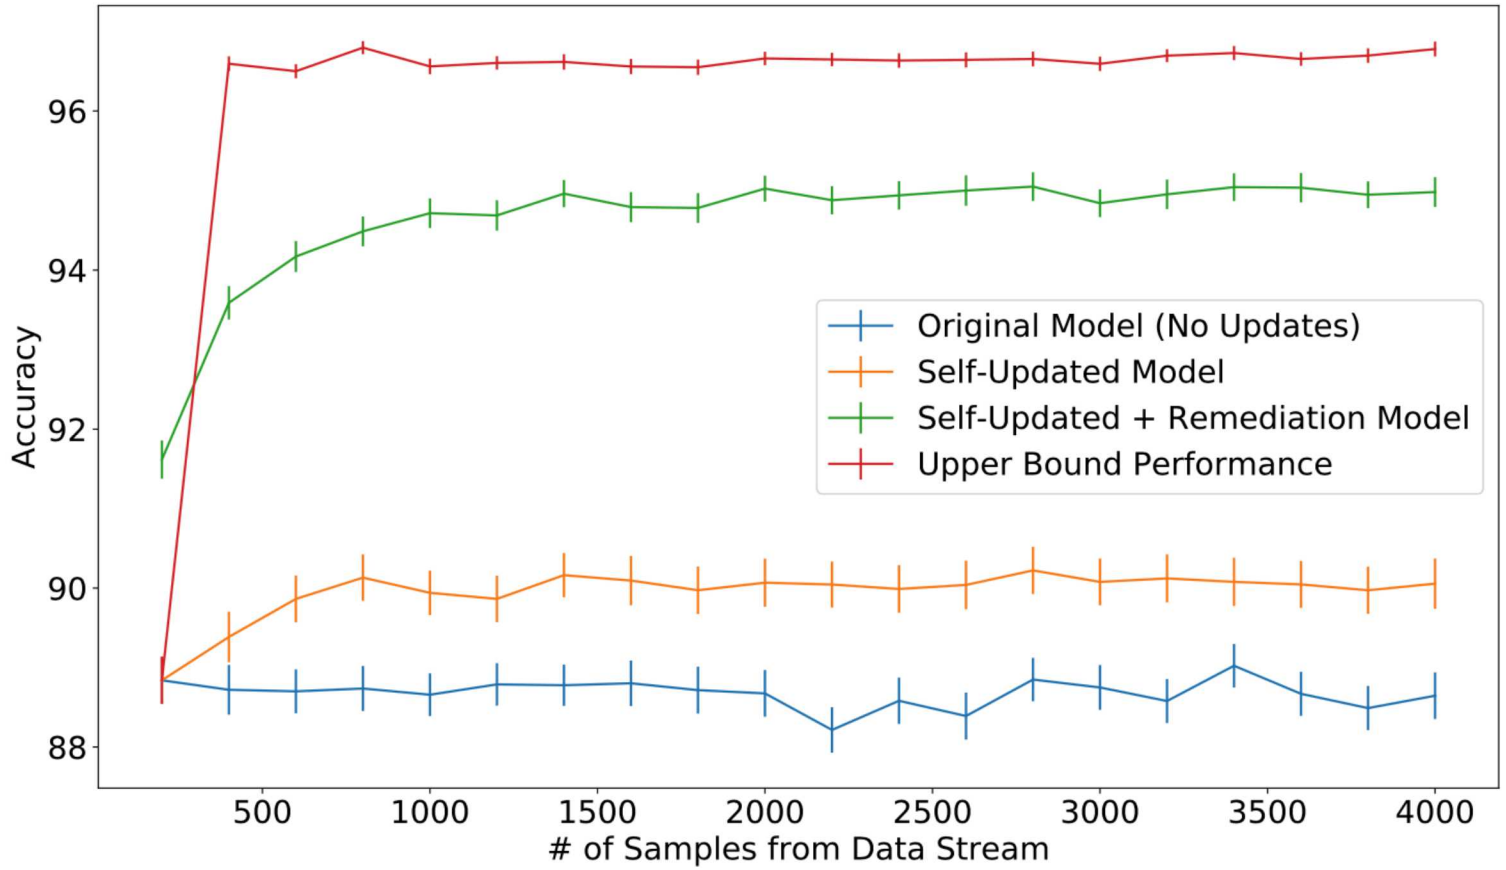

Results: Kaggle Ships Dataset (SUM, not SUMER)

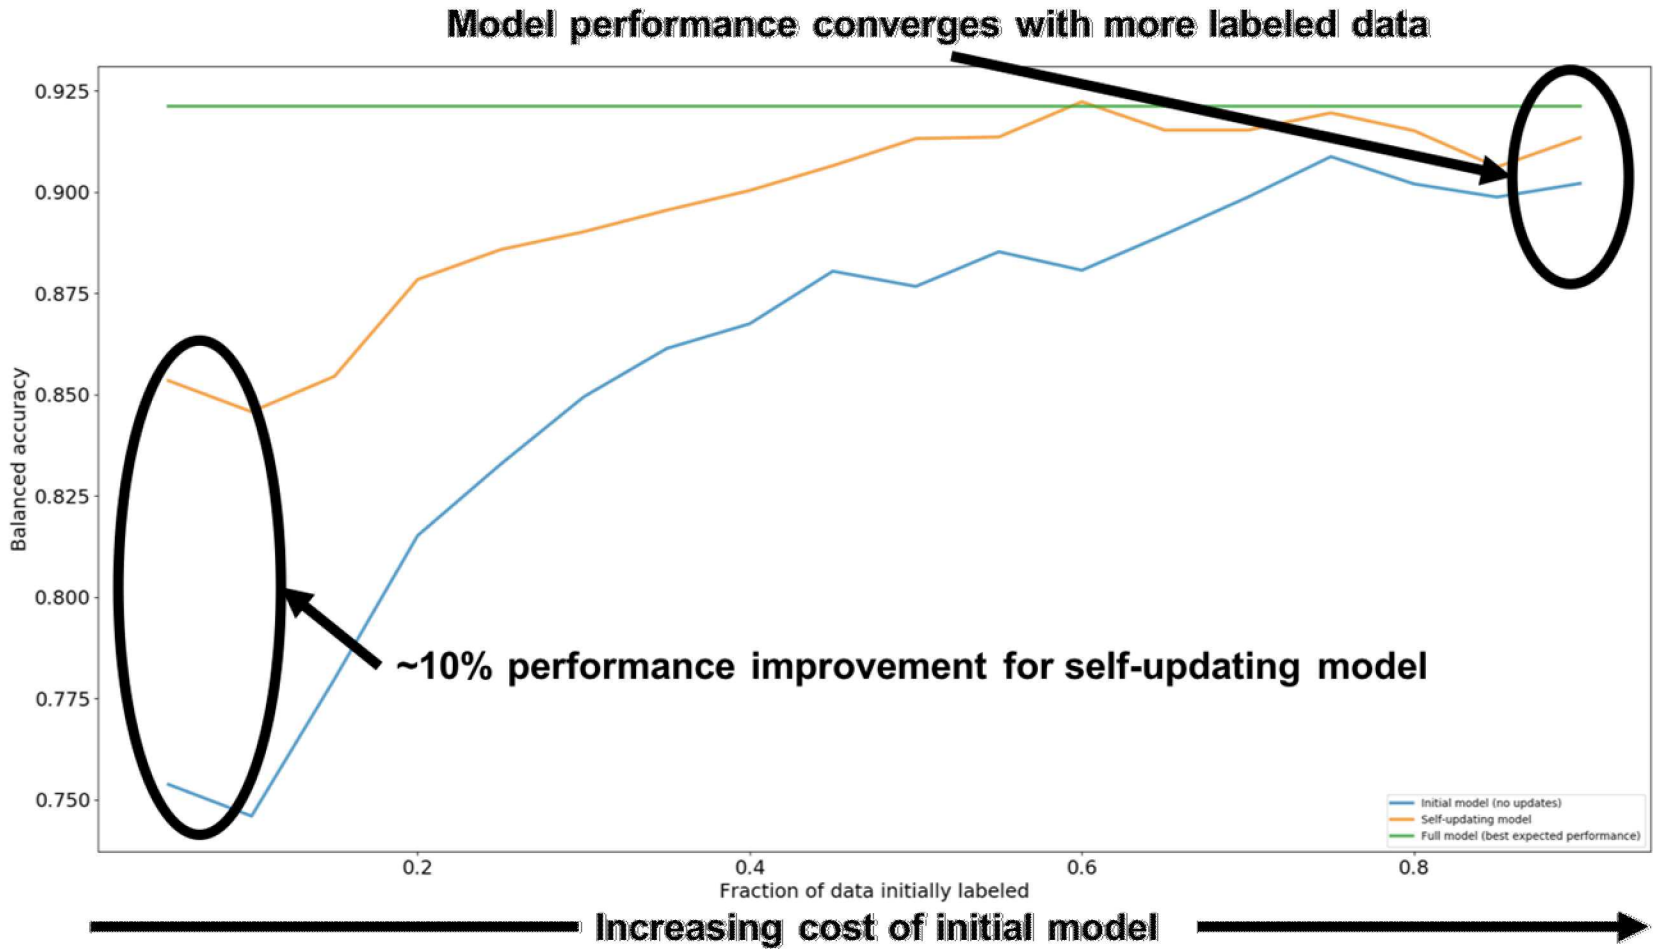

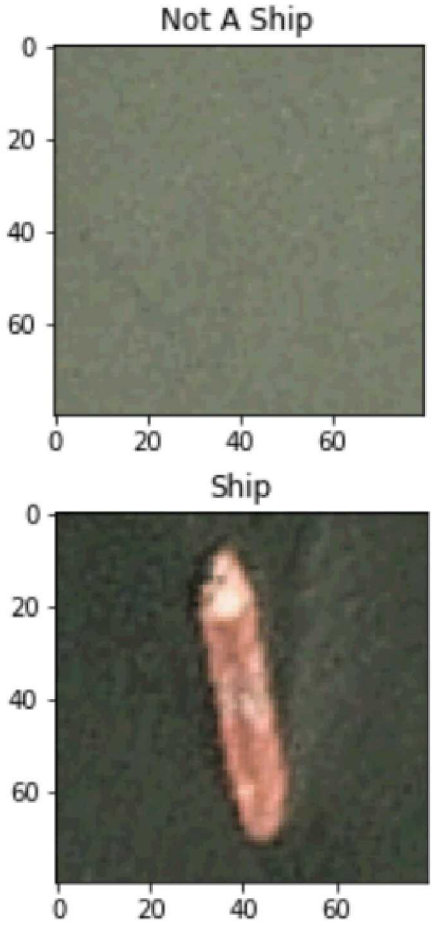

Conclusions and Future Work

SUMER is a framework for updating models utilizing techniques from, e.g., semi-supervised learning, model output calibration, uncertainty quantification, and noise filtering.

Self-updating models help to alleviate the human burden of obtaining labelled data and addressing concept drift.

Self-updating models improve performance over static models when unlabeled data is available.

Self-updating models with error remediation further improve the performance.

FUTURE WORK

Apply SUMER to a relevant, real dataset.

Add concepts from detecting and handling concept drift and novel class detection.

Add concepts from life-long and on-line learning.

Develop strategy to integrate analysts into the SUMER framework, similar to active learning, to leverage what humans are good at while still leveraging the benefits of SUMER.