

Novelty Detection in Streaming Learning using Neural Networks

David Munechika

2019 Summer Intern

Rochester Institute of Technology

Advisors: Ryne Roady, Dr. Kanan

What's all the buzz about AI?



What just happened? The rise of interest in Artificial Intelligence | TheHill

Artificial Intelligence, or AI for short, has Companies and investors are pouring in 1 day ago



Why an AI arms race with China would be bad for

technology" we need to keep from China. In experts, China ramped up its ...

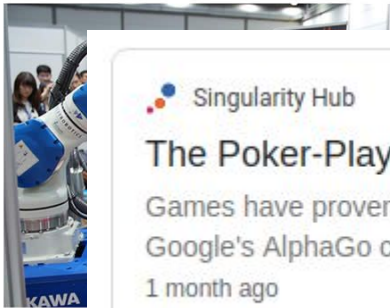
Artificial Intelligence May Doom The Human Race Within A Century, Oxford Professor Says

Posted: 08/22/2014 12:04 pm EDT Updated: 08/25/2014 7:59 pm EDT



Amazon's latest robot champion uses to stock shelves

The Verge • July 5, 2016 • 5 min read



Singularity Hub The Poker-Play

Games have proven a popular test-bed for AI in recent years, and when Google's AlphaGo cracked the ancient Chinese board game Go it was ... 1 month ago

Both methods are dead-end. ...
all investors have been chasing and about their strategies with demanding Washington ...
The financial industry's leaders should enjoy their moment of triumph. They need to know machine ...
The big companies are beginning to recognize the competitive challenges posed by advances in machine learning. Their internal algorithms give them strong market positions and data. However, not their algorithms and slow reaction times. Wall Street should be spending more on R&D, but might as well get it done on one track and regulatory disputes over who should pick the bid.
Portfolio managers and market strategists are only starting to consider which groups of machine learn-

OPINION

poses challenges for Wall Street

Are the machines coming for IT traders' jobs?

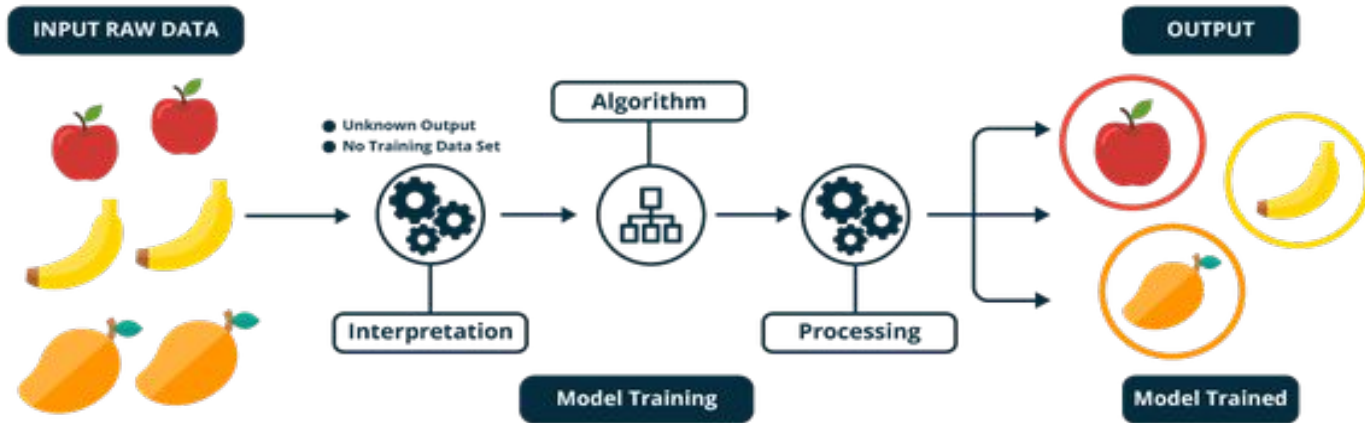
The researchers measured the edge in each day and monthly algorithm performance from 2010 to 2013. They found that the alpha would have performed a couple of dollars per contract per month, or about 10 percent per year. The researchers also checked together to see if machine learning could be used to predict the market. They found that machine learning programs did not outperform a simple buy-and-hold strategy in any of the markets. It would be implausible for such a research program to create a bar that would outperform the world, and it is not what the researchers

AI Hype vs. Reality



Machine Learning

An application of **artificial intelligence** (AI) that involves programming computers so they can learn from data



The Goal of Machine Learning



**Learned
ML
System
(Function)**




Output

“Dog”

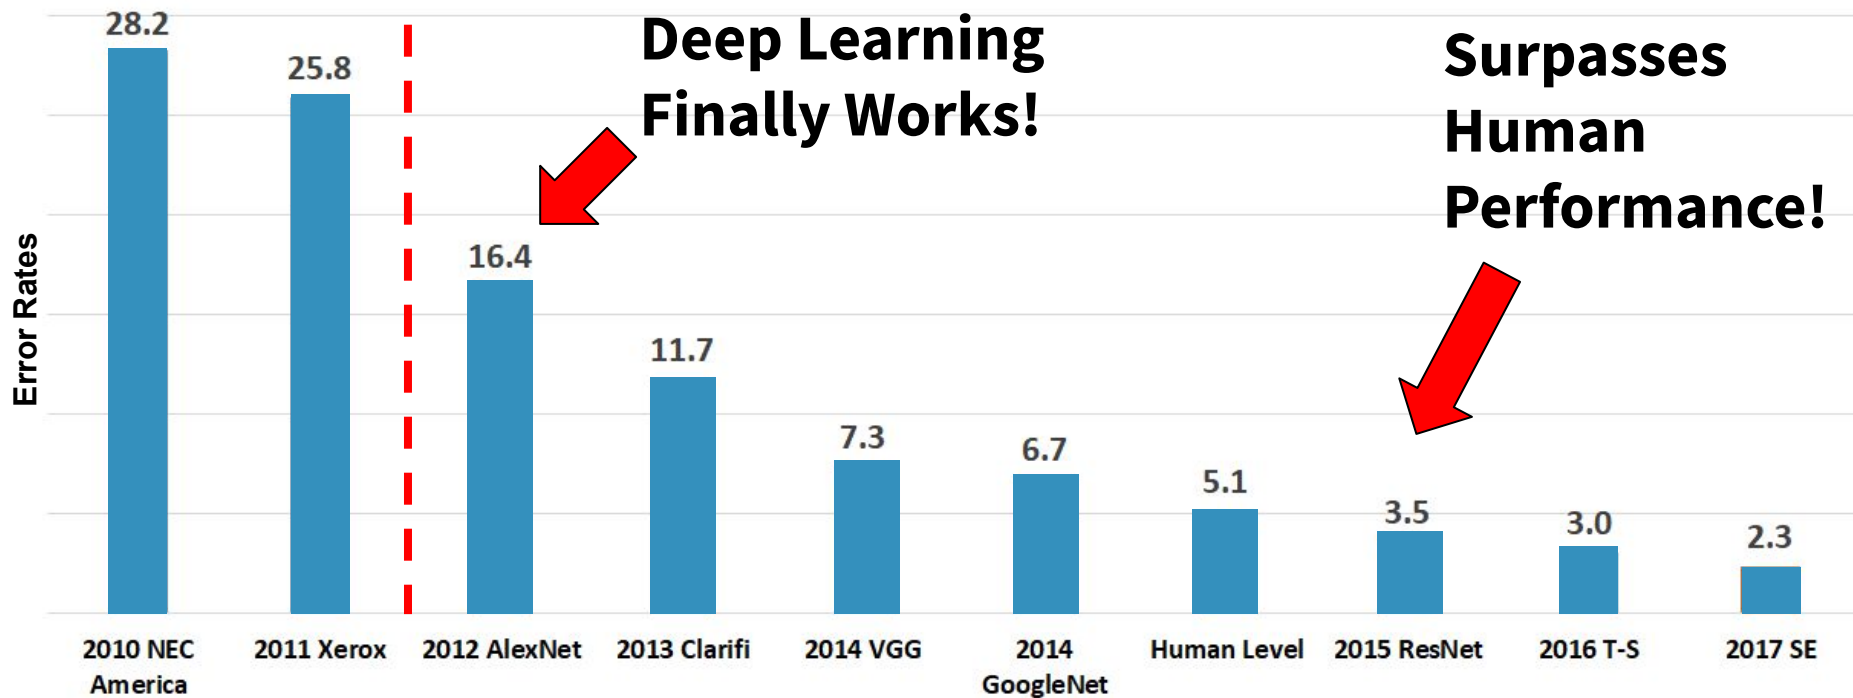


Deep Learning

A subfield of **machine learning** concerned with algorithms inspired by the structure and function of the brain called artificial neural networks



Deep Learning Gold Rush



DEMO: Human Learning vs. Supervised Machine Learning

Let's learn about some animals...



Emperor
Tamarin

DEMO: Human Learning vs. Supervised Machine Learning

Let's learn about some animals...



Tarsier



DEMO: Human Learning vs. Supervised Machine Learning

Let's see what you've learned...



Question 1

Which animal is this?



Emperor
Tamarin

?

Tarsier

Question 2

Which animal is this?



Emperor
Tamarin

?

Tarsier

Question 3

Which animal is this?



Emperor
Tamarin

?

Tarsier

Question 4

Which animal is this?



Emperor
Tamarin

?

Tarsier

Current Problems With Supervised Machine Learning



- Agents are trained on a static dataset and then deployed
- If we want an agent to learn novel information, it will **catastrophically forget** previous knowledge



Current Problems With Supervised Machine Learning

- If given novel information it lacks the ability to identify as it as novel or say “I don’t know”
- We want to develop agents that can recognize and learn from novel information
 - Similar to human learning

GOAL= Lifelong Learning






How do we solve these problems?

Continual Learning and Novelty
Detection Solutions

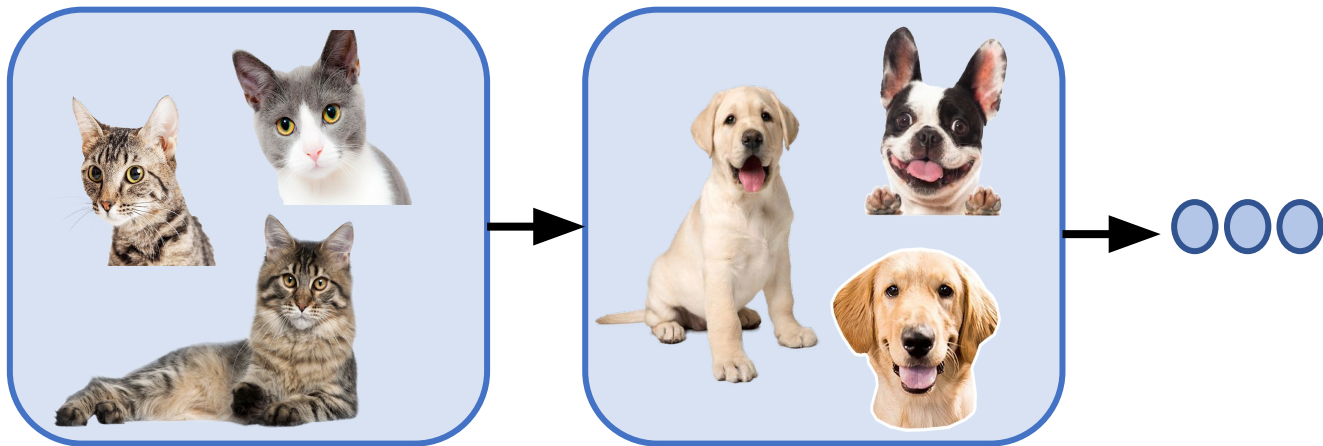
What is Continual Learning?

- Networks are able to learn and adapt continuously over time like humans
- They use previous knowledge to make informed predictions in the future



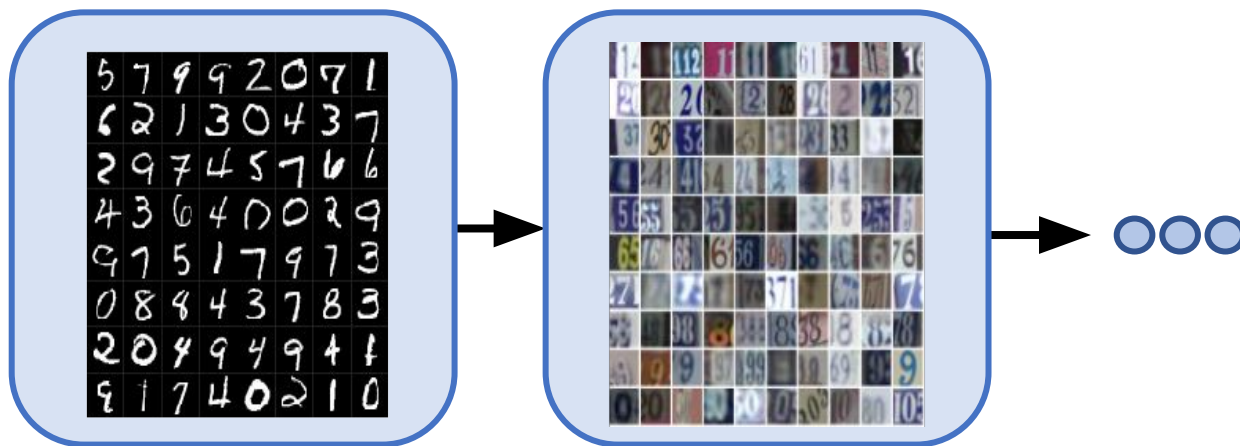
Continual Learning Tasks

- Incrementally learning classes one at a time



Continual Learning Tasks

- Incrementally learning datasets one at a time



Streaming Learning

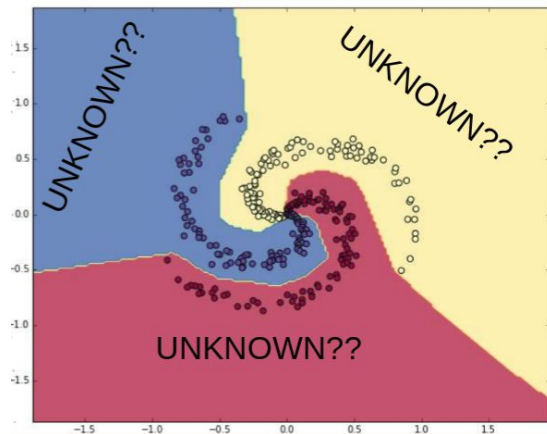
- Instances must be learned immediately
- Batch size of 1 (sample-by-sample)
- 1 epoch through dataset
- Inference can be performed at any time during training



What is Novelty Detection?

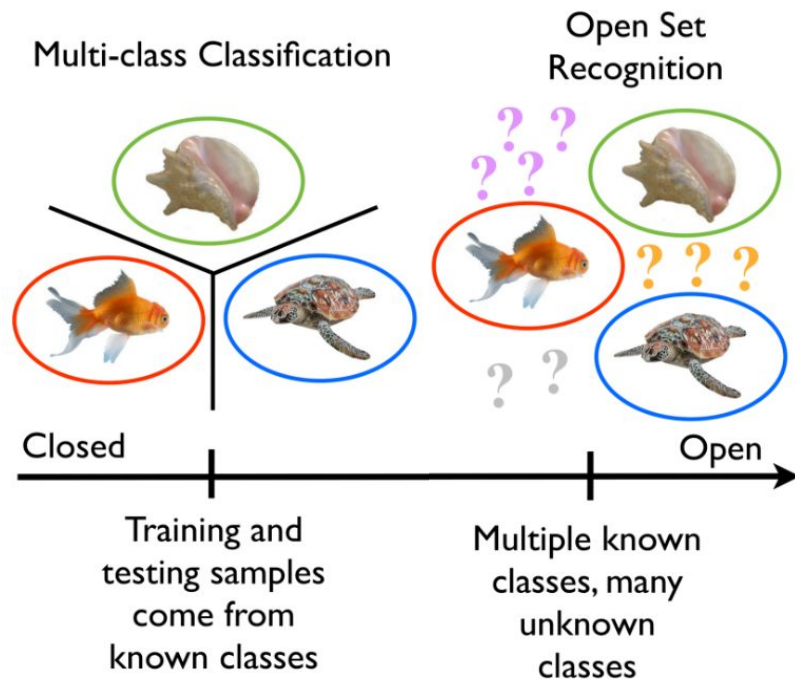


- In addressing general object recognition, there is a finite set of known objects in myriad unknown objects
- Labeling something new, novel, or unknown should always be a valid outcome



Bounded Classification

- Give classifier the ability to say “I don’t know” for novel objects that it has not learned
- Novel information can then be grouped together for continual learning





My Project

Combine current methods for continual learning and novelty detection into single model

Dataset

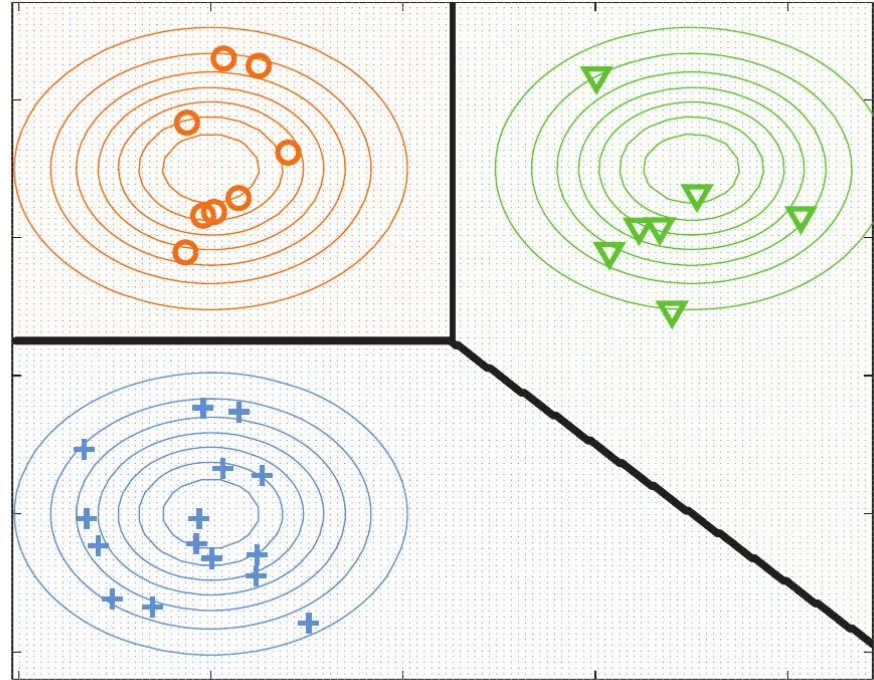
Caltech-UCSD Birds 200 (CUB-200)

- Number of classes: 200
- Number of images: 6,033



Methods for Continual Learning Evaluated

- Offline Model
- No Rehearsal
- Full Rehearsal
- Streaming Linear Discriminant Analysis (SLDA)¹



1. Hayes, Tyler L. and Christopher Kanan. "Lifelong Machine Learning with Deep Streaming Linear Discriminant Analysis." *3rd Conference on Robot Learning (CoRL 2019)*, In Review.

Experimental Setup

1. Trained the model to classify CUB-200 dataset in incremental batches of 20 classes
2. After each batch, tested on all classes learned up to that point
3. Computed classification accuracy and novelty detection accuracy after each batch
4. Repeat until all 200 classes have been learned (10 batches)

Evaluating Novelty Detection

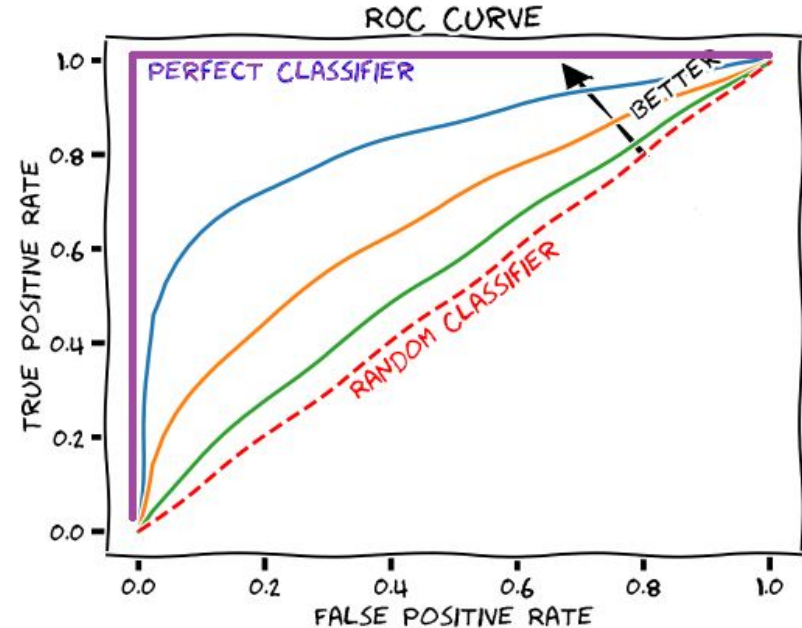
- Outputs of model thresholded to detect novel inputs¹
- Out-of-distribution (OOD) detection was evaluated against two different novel inputs:
 1. Noise (Gaussian distributed)
 2. Inter dataset (samples from the Oxford Flowers dataset)



1. Hendrycks, Dan, and Kevin Gimpel. "A baseline for detecting misclassified and out-of-distribution examples in neural networks." *arXiv preprint arXiv:1610.02136* (2016).

Performance Metrics- OOD Detection

- The Area Under the Receiver Operating Characteristic (AUROC) curve indicates how effectively a model can separate two classes
- The higher the AUROC, the more effective the model is at novelty detection

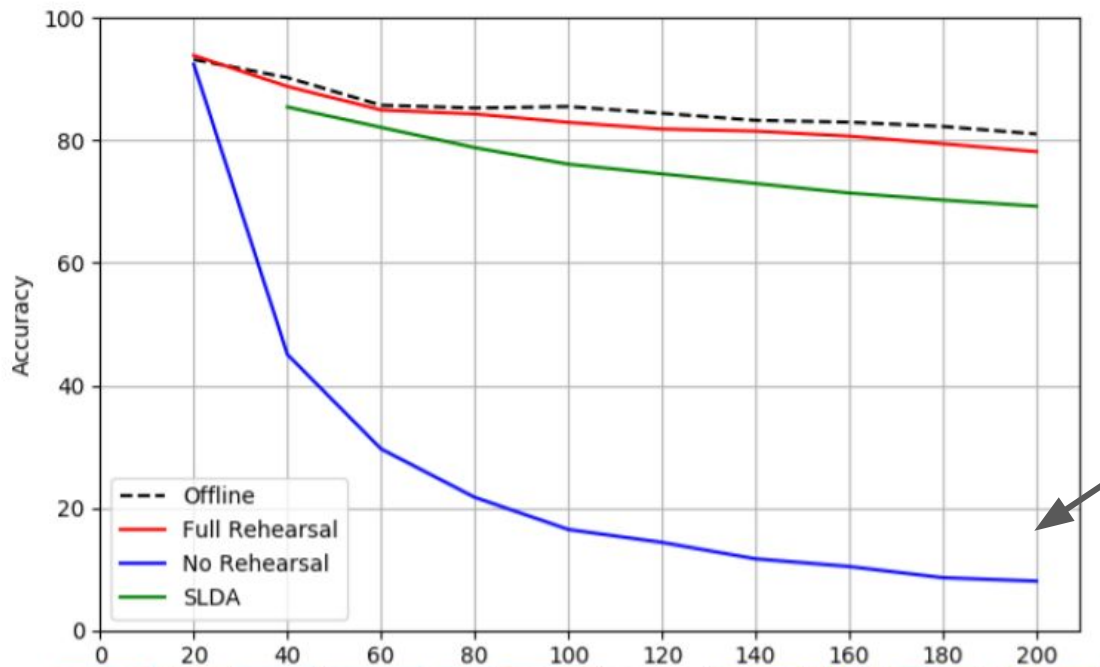


A decorative network diagram in the top-left corner, consisting of various sized nodes (some solid grey, some hollow white) connected by thin grey lines, forming a complex web structure.


The Results

Conclusions of comparing different models

Closed Set Accuracies



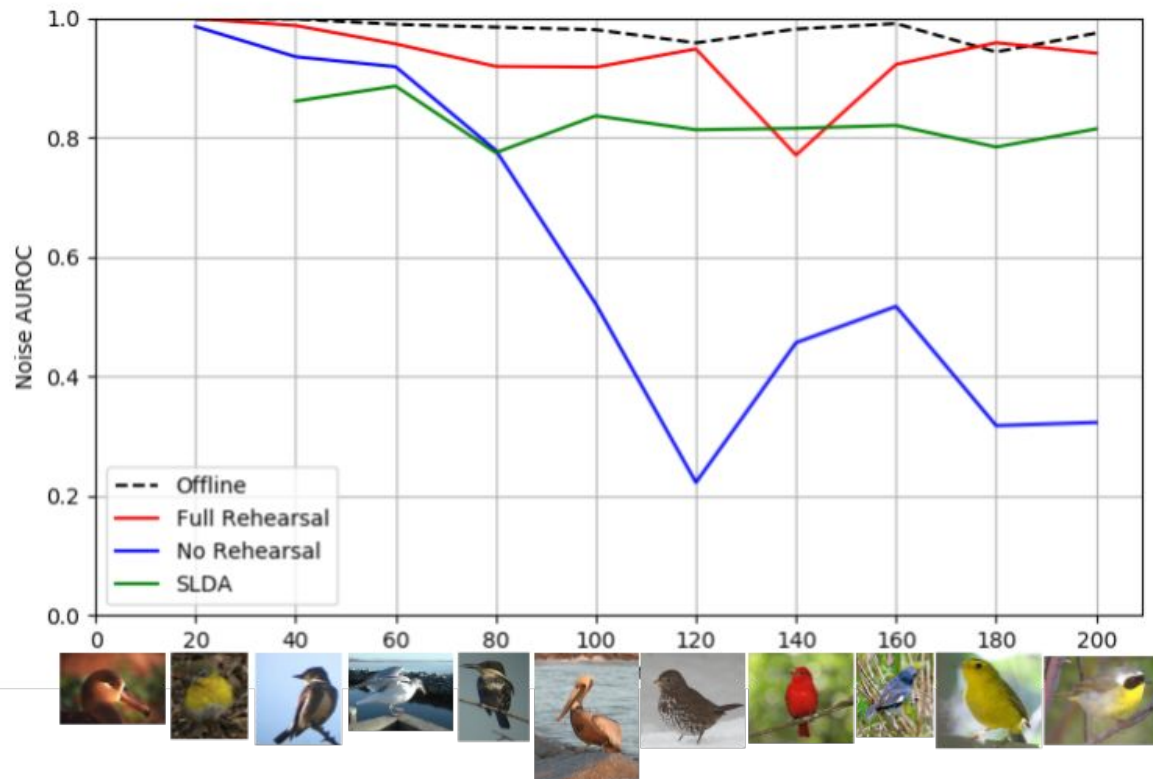

Catastrophic Forgetting
without rehearsal on
previous samples



Novelty Detection Accuracies- Noise



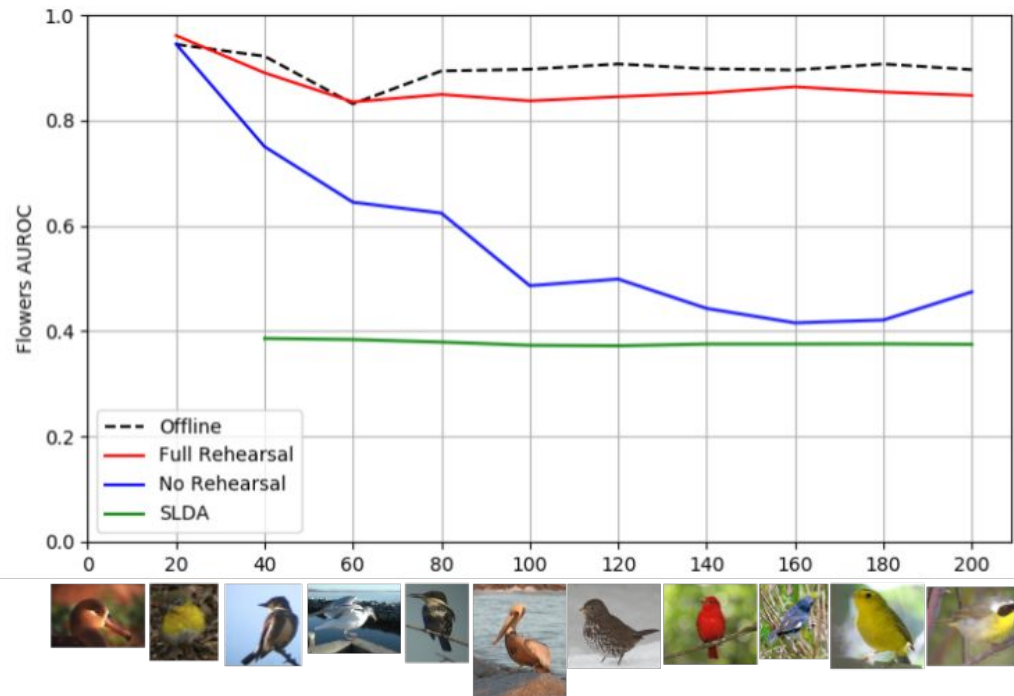
vs.



Novelty Detection Accuracies- Inter




VS.



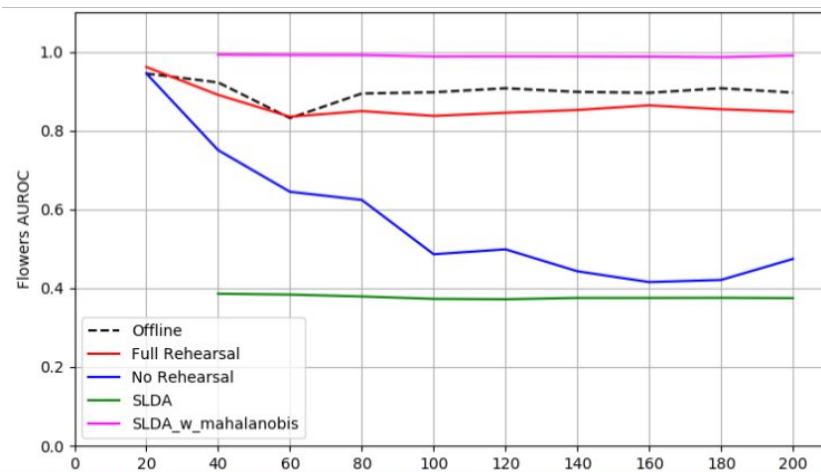
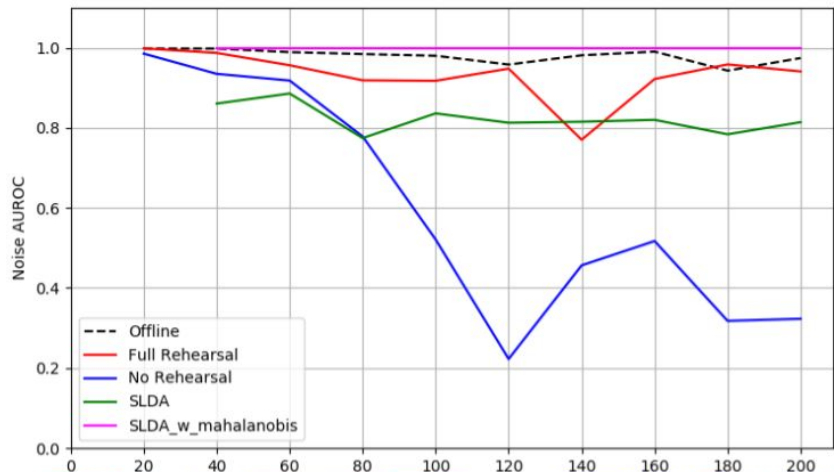
Improving SLDA with Better Novelty Detection

- The Mahalanobis distance measures the likelihood that a sample belongs to a specific class (or none at all) and was shown to be a reliable OOD detector for Deep Neural Networks (DNNs)¹

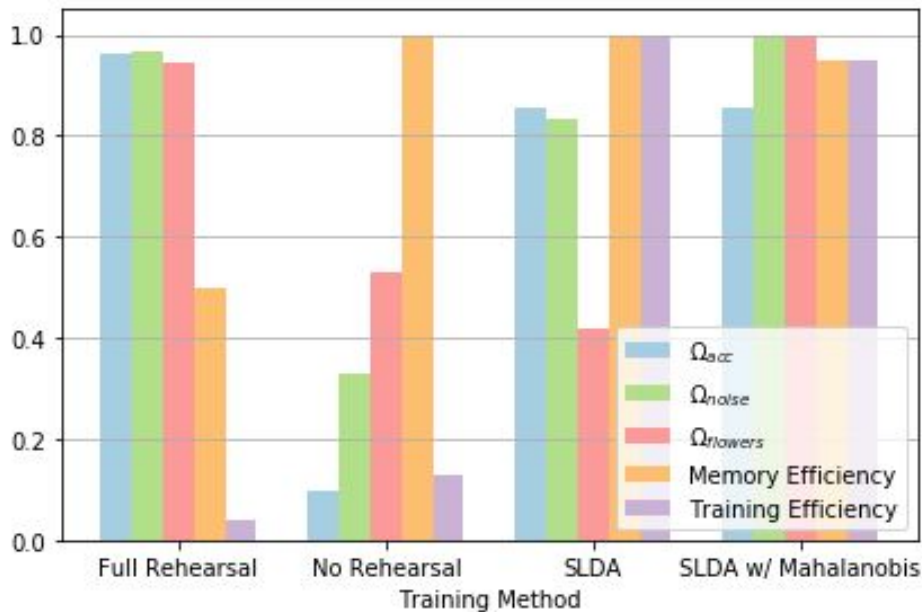


1. Lee, Kimin, et al. "A simple unified framework for detecting out-of-distribution samples and adversarial attacks." *Advances in Neural Information Processing Systems*. 2018.


Performance of Improved SLDA w/ Mahalanobis



Overall Results



SLDA with Mahalanobis Outlier Detection can efficiently learn incrementally and detect novel examples




$$\Omega_{all} = \frac{1}{T} \sum_{t=1}^T \frac{\alpha_{all,t}}{\alpha_{offline,t}}$$

$$\text{Efficiency} = 1 - \frac{\text{MethodTested}}{\text{MaxTested}}$$

Acknowledgements



Special thanks to everyone who helped me with this amazing experience!

- © Ryne Roady
 - © Joe Pow
 - © Dr. Kanan
 - © Tyler Hayes
 - © The other interns- Akul, Amy, Brian, Hannah, Jocelyn, and Varun
- 

Summary + Questions?

- ◎ Deep Neural Networks can be used to solve computer vision problems such as image classification
- ◎ Current machine learning models lack the capability of continual learning and novelty detection
- ◎ SLDA with Mahalanobis has proven it can both efficiently learn incrementally and effectively detect novel examples



Thanks!

Any questions?

You can contact me at:

dampci@rit.edu



Backup slides

Results

Method	Final Accuracy	Omega Accuracy	Final AUROC Noise	Omega AUROC Noise	Final AUROC Inter	Omega AUROC Inter	Final AUROC Intra	Omega AUROC Intra
Offline	0.81	NA	Baseline: 0.97 Mahalanobis: 1.00	NA	Baseline: 0.90 Mahalanobis: 0.99	NA	Baseline: 0.78 Mahalanobis: 0.73	NA
Full Rehearsal	0.78	0.96	Baseline: 0.94 Mahalanobis: 1.00	Baseline: 0.97 Mahalanobis: 1.00	Baseline: 0.85 Mahalanobis: 0.99	Baseline: 0.95 Mahalanobis: 0.99	Baseline: 0.79 Mahalanobis: 0.76	Baseline: 1.01 Mahalanobis: 1.04
MLP	0.08	0.10	Baseline: 0.32 Mahalanobis: 1.00	Baseline: 0.33 Mahalanobis: 1.00	Baseline: 0.47 Mahalanobis: 0.90	Baseline: 0.53 Mahalanobis: 0.91	Baseline: 0.53 Mahalanobis: 0.55	Baseline: 0.68 Mahalanobis: 0.75
SLDA w/ Baseline OOD	0.69	0.85	0.77	0.79	0.38	0.42	0.52	0.67
SLDA w/ Mahalanobis OOD	0.69	0.85	1.00	1.00	0.99	1.00	0.52	0.72

Goal of Computer Vision

To extract meaning from pixels...

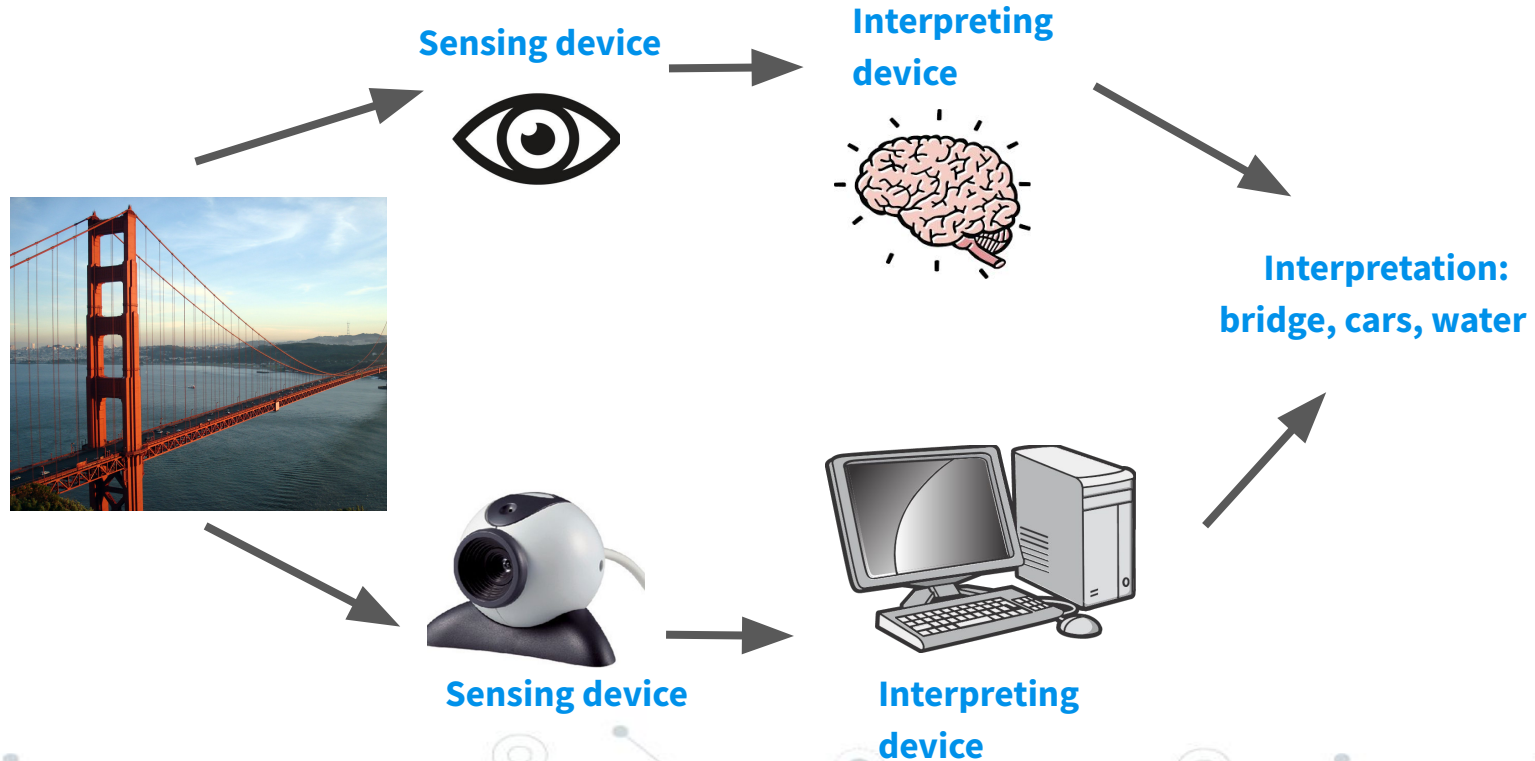


What we see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0


What a computer sees

Computer Vision vs. Human Vision



Limitations of Computer Algorithms

Computers are exceedingly efficient for well-defined tasks:



They perform fairly well on bounded tasks:



Open ended tasks, however, are very difficult for a computer to solve:

Q: What is the red object used for?



How Neural Networks Learn

Takes an input




Makes a prediction based on current knowledge

Dog?

Adjusts the function to make prediction more correct



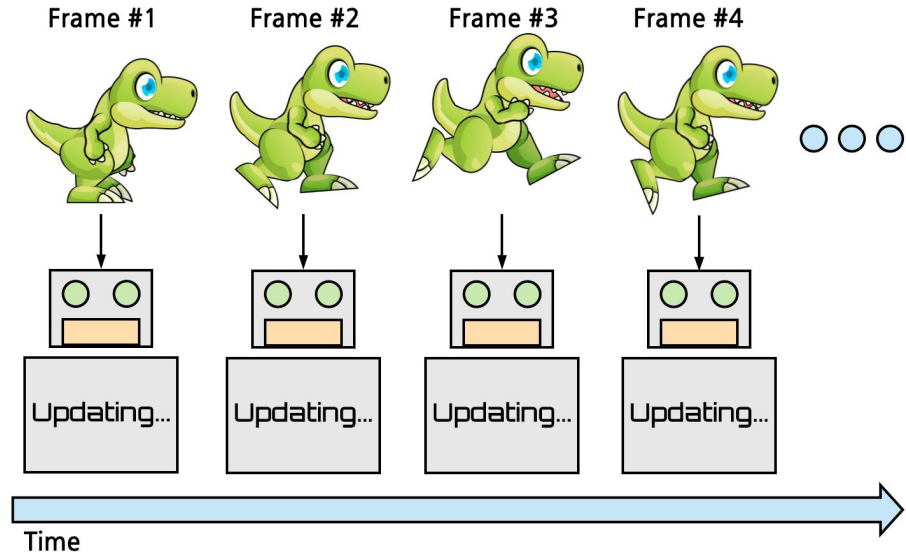
Repeats a **large** number of times



How Streaming Learning Works

1. Data is input into model
2. Model updates using input
3. Repeat

Because the model updates after every example, you are able to **test at any time**



Performance Metrics- Accuracy

- Compared the classification accuracy to the offline performance over time using the metric Ω_{all}

$$\Omega_{\text{all}} = \frac{1}{T} \sum_{t=1}^T \frac{\alpha_{\text{all},t}}{\alpha_{\text{offline},t}}$$

OOD Inference Methods

- Baseline Thresholding
- Mahalanobis

Overall Results

