Tuning online neural networks with reinforcement learning

THE UNIVERSITY OF SYDNEY

Alex Fang *(Dr. Fabio Ramos)*
School of Information Technologies
FACULTY OF ENGINEERING & INFORMATION TECHNOLOGIES

Background

Artificial neural network (ANN)

- A machine learning technique developed with inspiration from biology
- A neural network is an arrangement of individual 'neurons' which respond to input by outputting a value weighted by it's internal state
- By updating the internal states of the individual neurons, the neural network can be trained to respond certain inputs in certain ways – this is known as training and is usually done using stochastic gradient descent

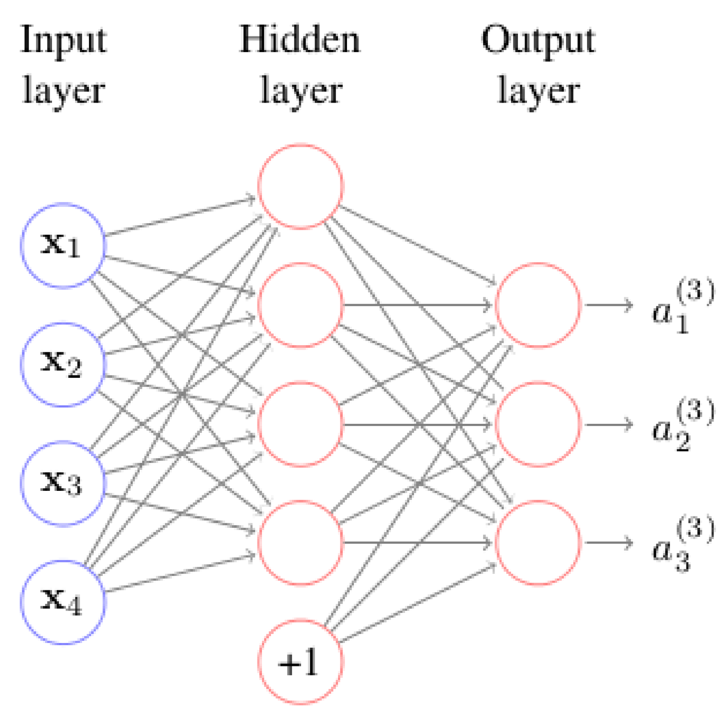

- Deep neural networks (neural networks with more than one layer) have achieved state of the art results many machine learning and computer vision tasks [1]

Online learning

- Online learning is a discipline of machine learning where the algorithm is trained and makes predictions using a continuous stream of incoming data
- It requires the algorithm to quickly adapt with respond with respect to the contents of the data stream

Reinforcement learning

- Reinforcement learning is an approach inspired by behavioural psychology [2]
- It models the learning algorithm as an 'actor' in an 'environment' that attempts to maximize some 'reward' thru it's actions
- In a typical scenario, the algorithm makes a decision based on it's current state and the is given some reward based on it's action
- The learning algorithm uses it's past actions , states and rewards to formulate it's next move

Incremental autoencoders

- The incremental autoencoder is an online learning algorithm that adjusts the architecture of the neural network in order to adapt to changes in the input data [3]
- As its name suggests, the algorithm increases the number of neurons in the network based on the validation error of the current batch of input
- Although simple in concept, it shows that adapting neural network architecture at runtime is a feasible strategy and more sophisticated algorithms can be used to control the adjustments

Motivation

- In normal supervised learning scenarios, Baeysian methods or even simply cross validation is able to find a neural network architecture that works well for a given set of input data [4]
- Online learning's unique in that it is difficult to select the parameters of the neural network because amount of data that is passed thru the neural network [3]
- Previous work on incremental autoencoders addresses some of these concerns, however it still requires deep neural networks to be trained one layer at a time until an optimal configuration for that layer is found, as opposed to training all layers at the same time.

Aim

- To achieve better online learning performance by the improving upon the incremental autoencoder model
- Allow the model to be directly applicable to deep neural networks

Contribution

- A generalised algorithm for incrementing and reducing neurons in the neural network that is directly applicable to deep neural networks
- Using reinforcement learning to automatically learn a strategy for the neural network to adapt it's architecture with respect to variations in the incoming data

Reinforcement learning set up

1. With every batch of incoming data, the algorithm checks the validation error and determines if it is rising or falling
2. Using the history of changes as the state for q-learning, the algorithm chooses from a set of actions: 'increment neurons by 10%', 'decrement neurons 10%', 'merge neurons with similar weights', or 'do nothing'
3. A 'reward' is given to the q-learning algorithm in the form of the negative of validation error
4. An epsilon-greedy exploration strategy is used

Experiments and results

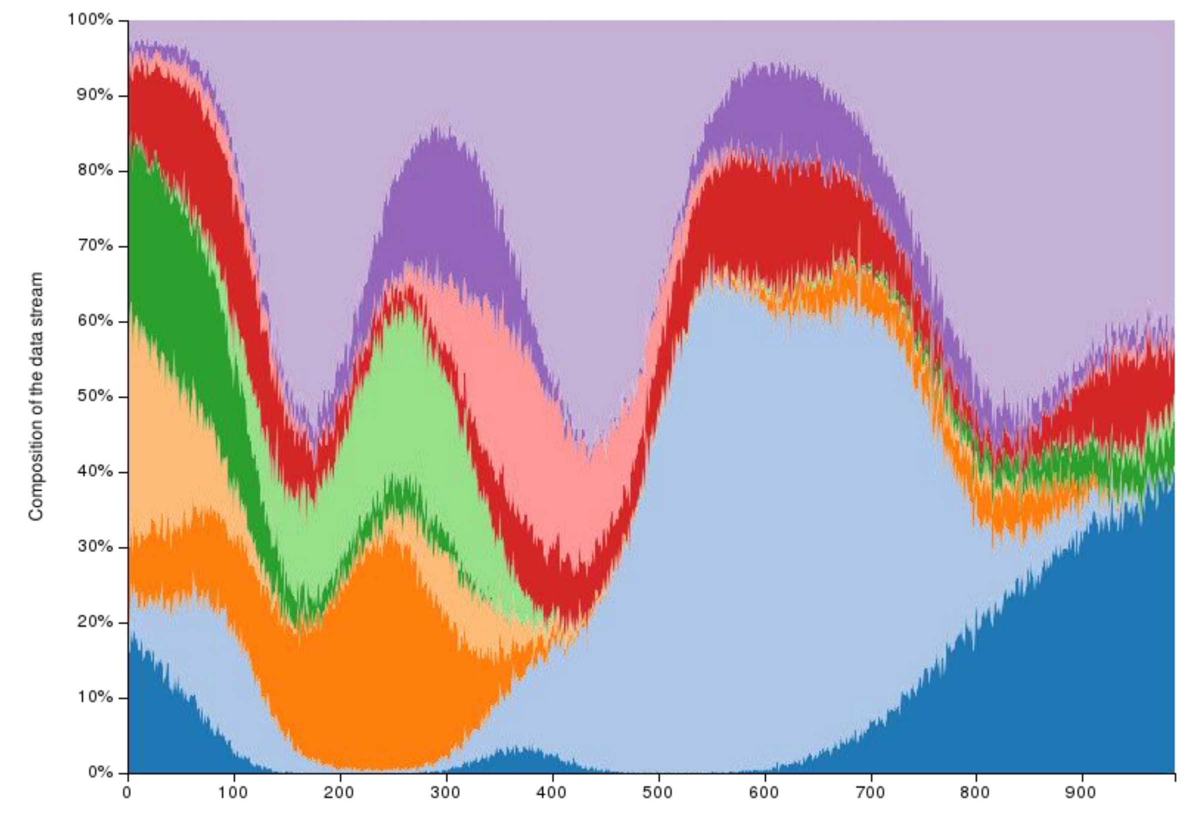

Gaussian processes are used to generate a test data with varying distribution of labels, this is a sample of the data used for the MNIST-noise test

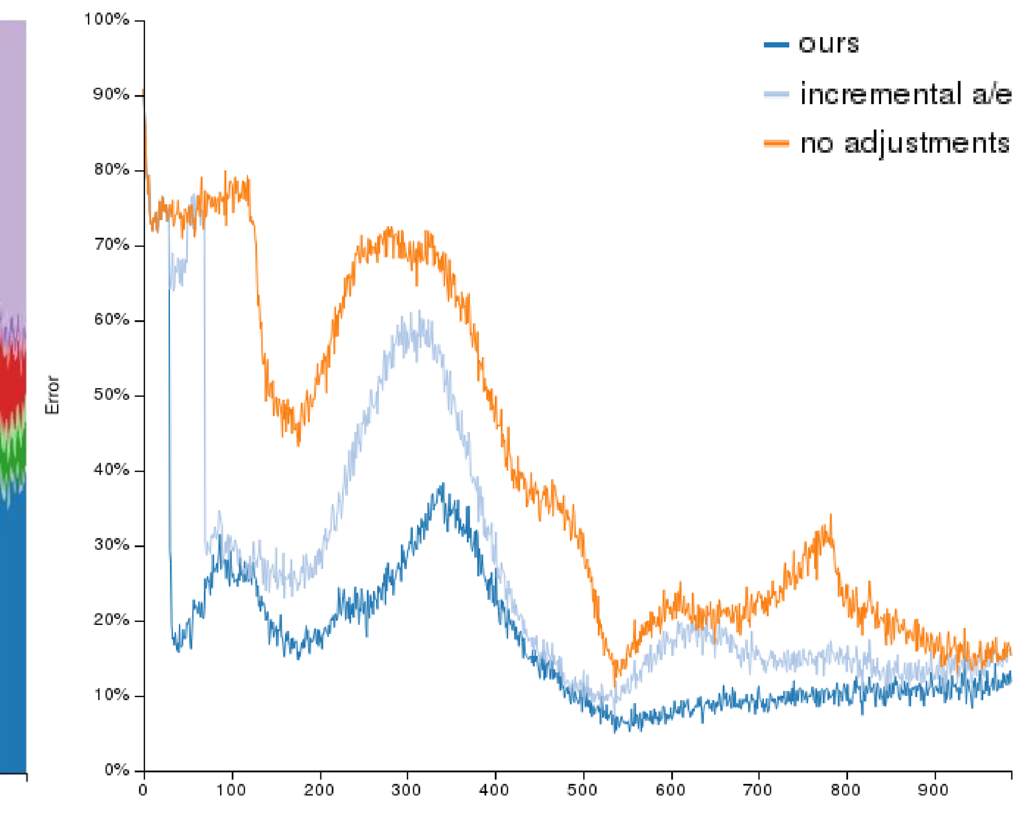

Experimental results with the MNIST-noise (MNIST with background noise overlayed) data set, shallow neural network 1000 neurons in hidden layer.

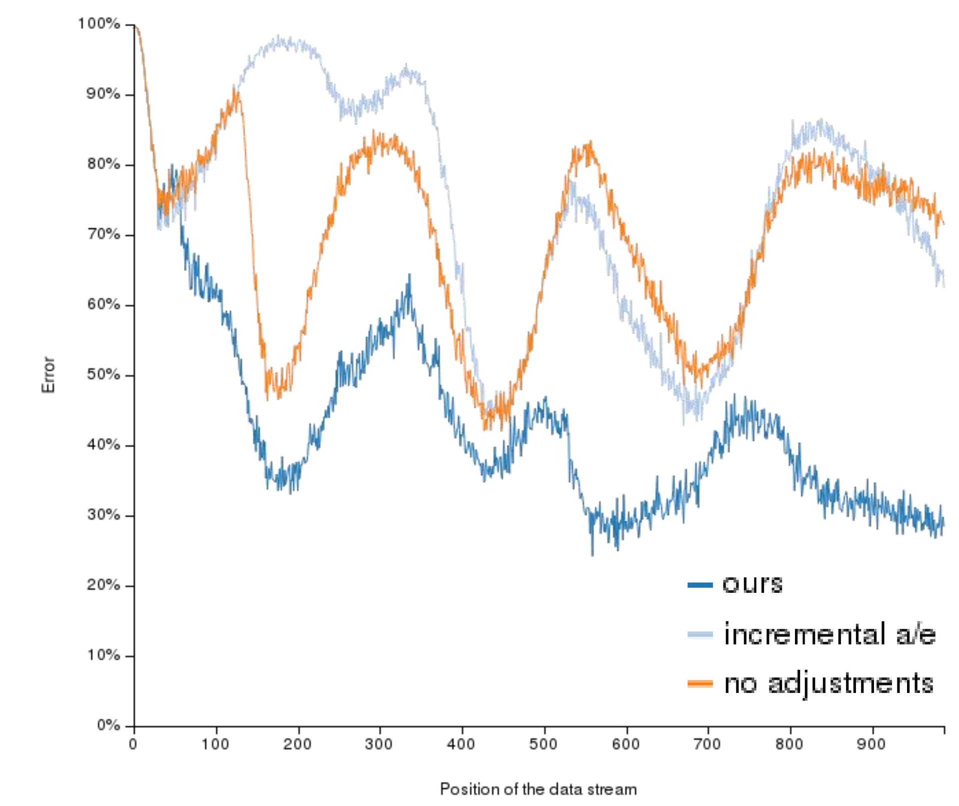

A representative result with the CIFAR-10 data set, shallow neural network, 1000 neurons in hidden layer

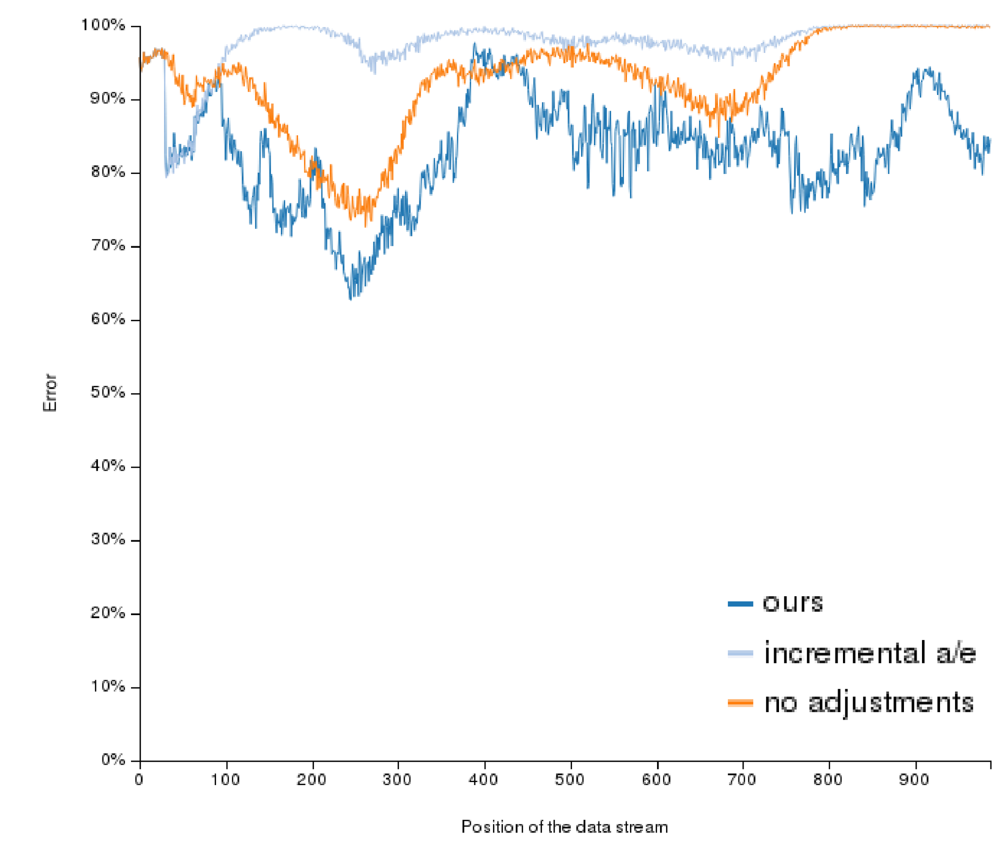

A representative result with the CIFAR-100 data set, shallow neural network, 1000 neurons in hidden layer

Conclusion

- We've shown our method outperforms the incremental autoencoder in all test scenarios, however it comes at the cost of longer computation time per iteration

Future work

- Currently, the action and state values for are discretised, we could explore using a continuous action and state space backed by Gaussian processes
- This would allow the reinforcement learning algorithm to make more precise actions and decisions

References

1. Alex Krizhevsky, Ilya Sutskever, and Georey E Hinton. *Imagenet classication with deep convolutional neural networks*. NIPS
2. Barto, Andrew G. *Reinforcement learning: An introduction*. MIT press, 1998.
3. Guanyu Zhou, Kihyuk Sohn, and Honglak Lee. *Online incremental feature learning with denoising autoencoders*. In AISTATS, 2012.
4. Bengio, Yoshua, Aaron Courville, and Pascal Vincent. "Representation learning: A review and new perspectives." Pattern Analysis and Machine Intelligence, IEEE Transactions on 35.8 (2013): 1798-1828.