BMS Related Work

Design and Implementation of a Continuously Improving Deep Learning Approach

Eiram Mahera Sheikh

About myself

- Master’s in Data Science and AI
- Currently in 5th semester
- Part-time internship at Zalando
- Interests include: Machine Learning, Deep Learning, Computer Vision
- I come from India and love travelling around Europe 😊

Agenda

- Motivation
- Related Work
- Road Map

Retail Store

Stocker Responsibilities:

- Organizing product displays
- Counting Inventory
- Loading product on shelves
- Storing products
- Maintaining cleanliness

Product Shelf

A continuously changing environment

- New types of products stocked
- Product organization might change
- Products may be discontinued
- Special limited season products

Humans Learn Continuously

AI Agent in a dynamic environment

Deep Learning

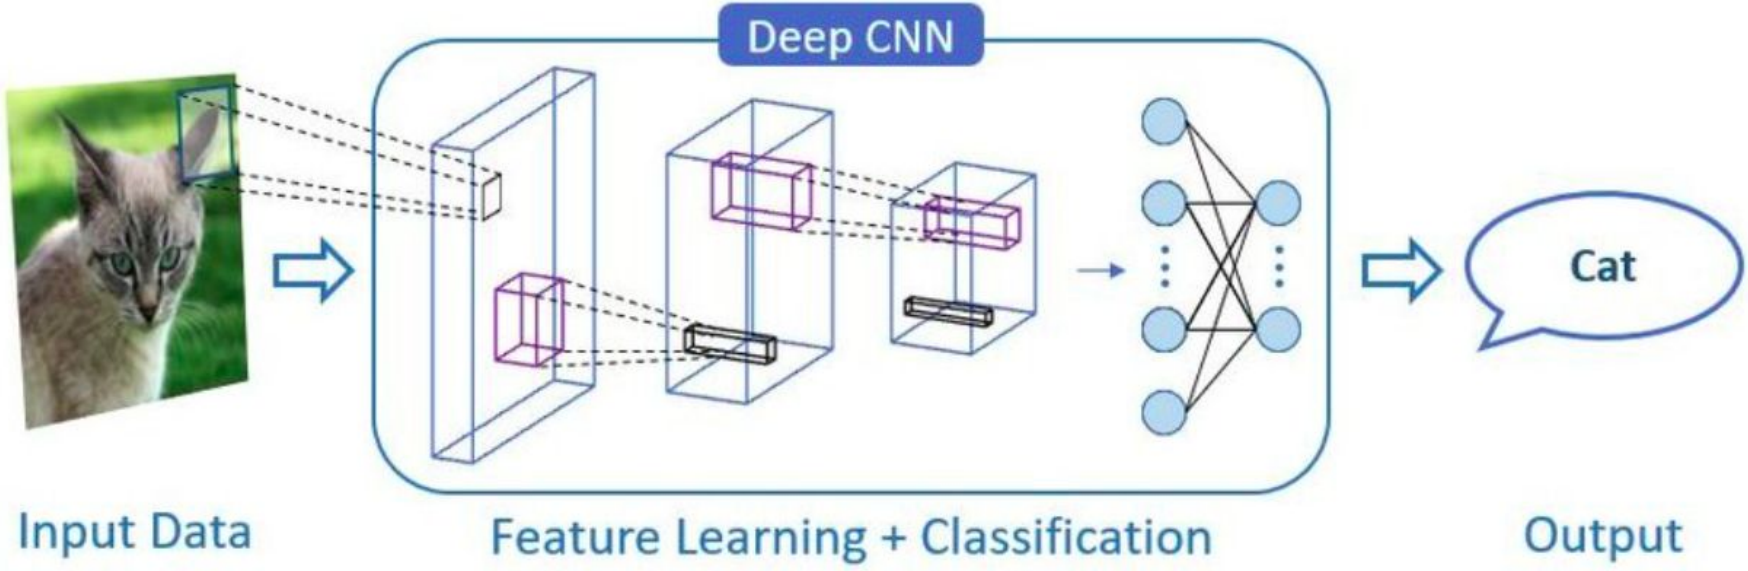

Static Learning

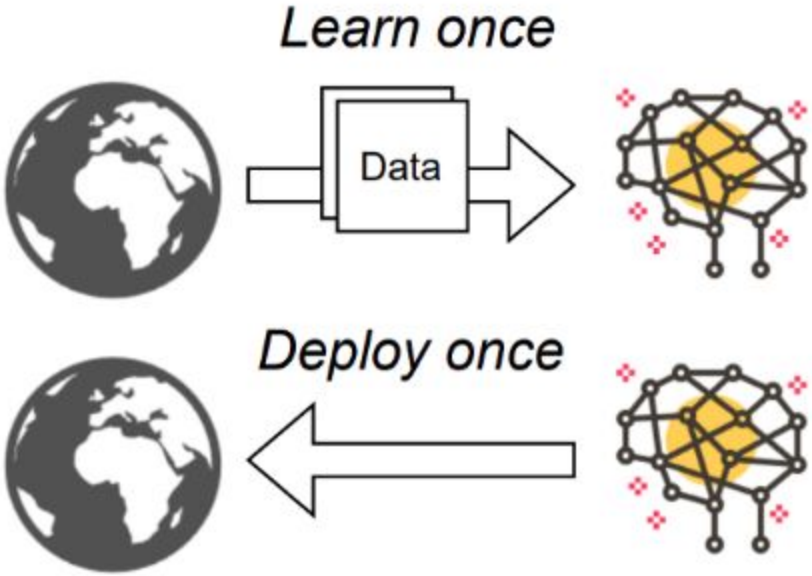

Continual Learning

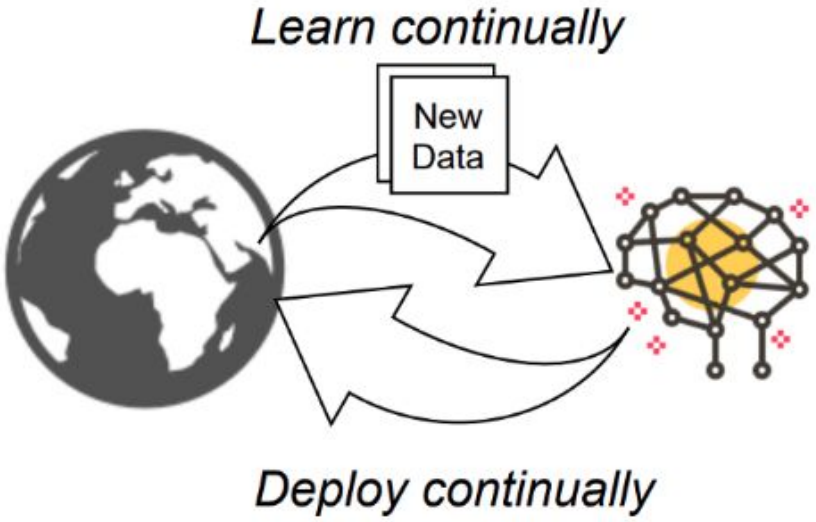

Naive Approach: Finetuning

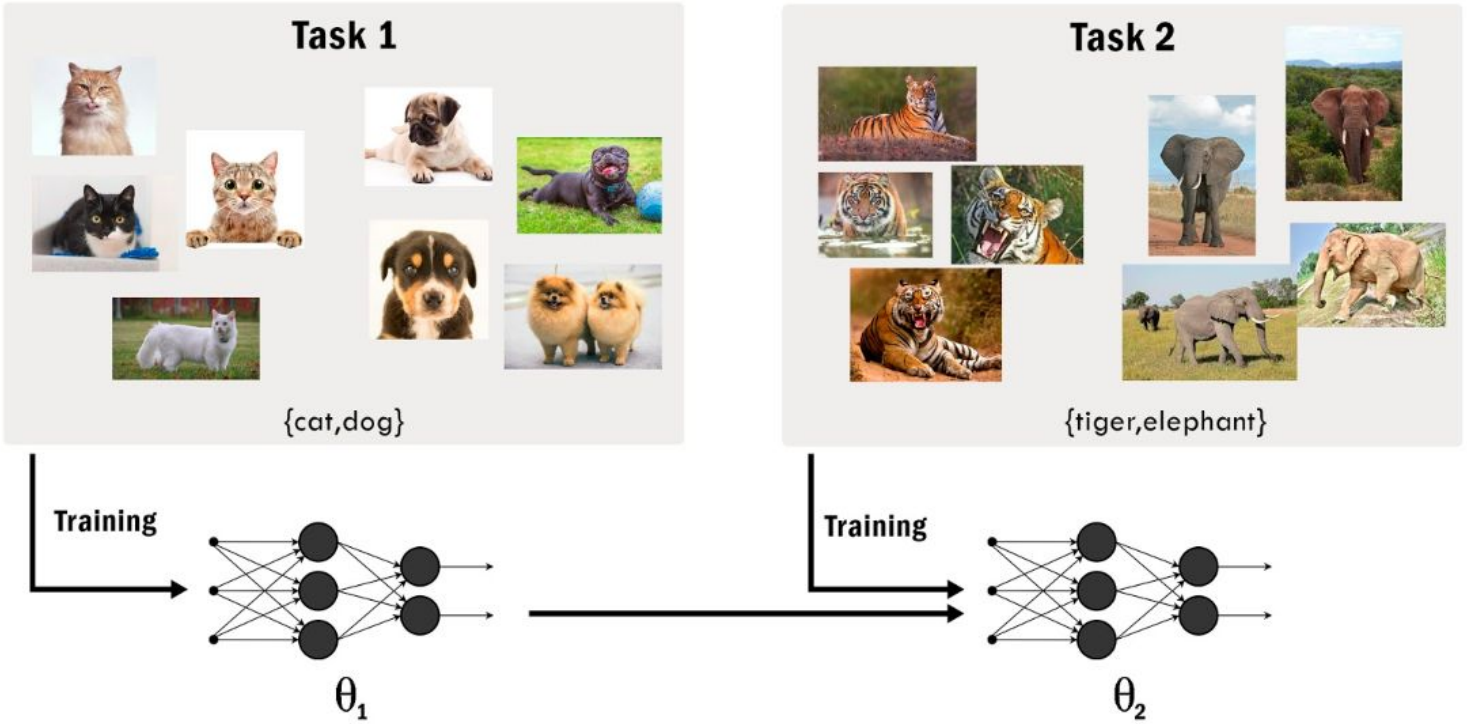

Catastrophic Forgetting

The phenomena where a model tends to forget past knowledge on account of learning new knowledge. (McCloskey & Cohen 1989)

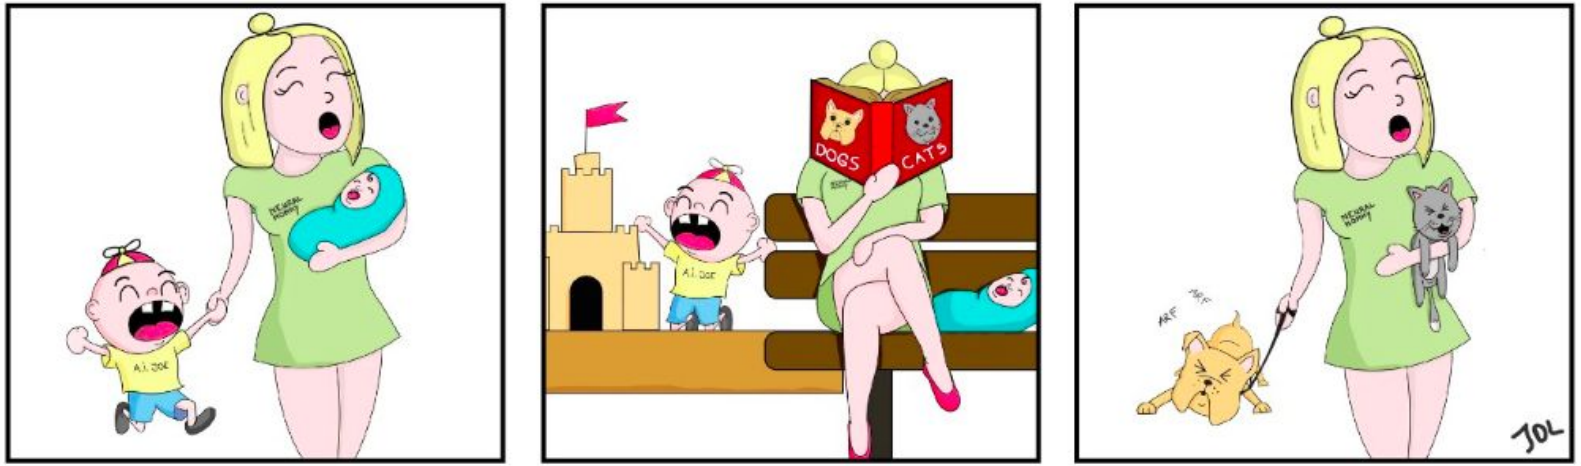

An illustration of catastrophic forgetting in neural networks. Cartoon credits @Jasper De Lange.

Class Incremental Learning

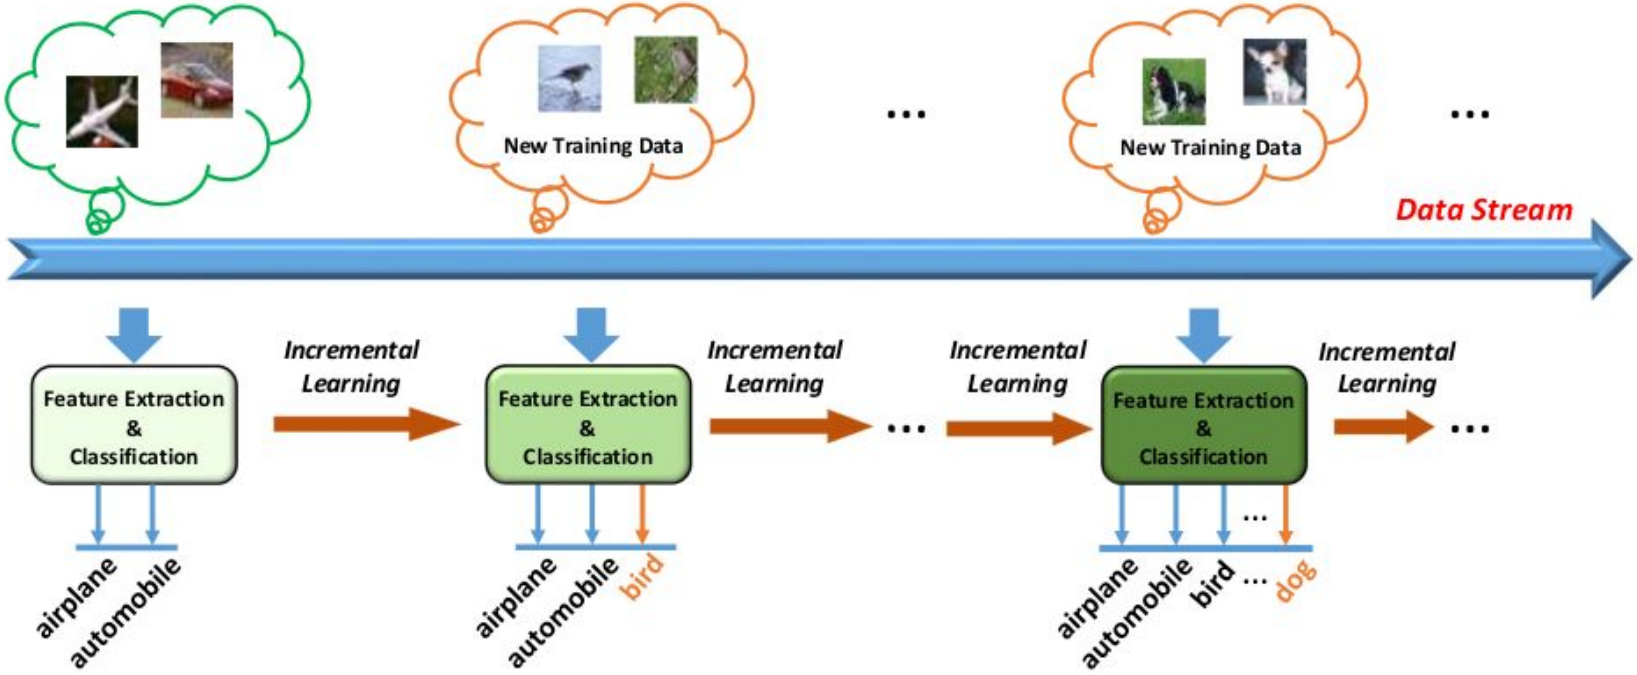

Liu et. al. (2018)

Thesis Goal

Design and implement an Deep Learning model that can continuously learn new knowledge over time.

Related Work

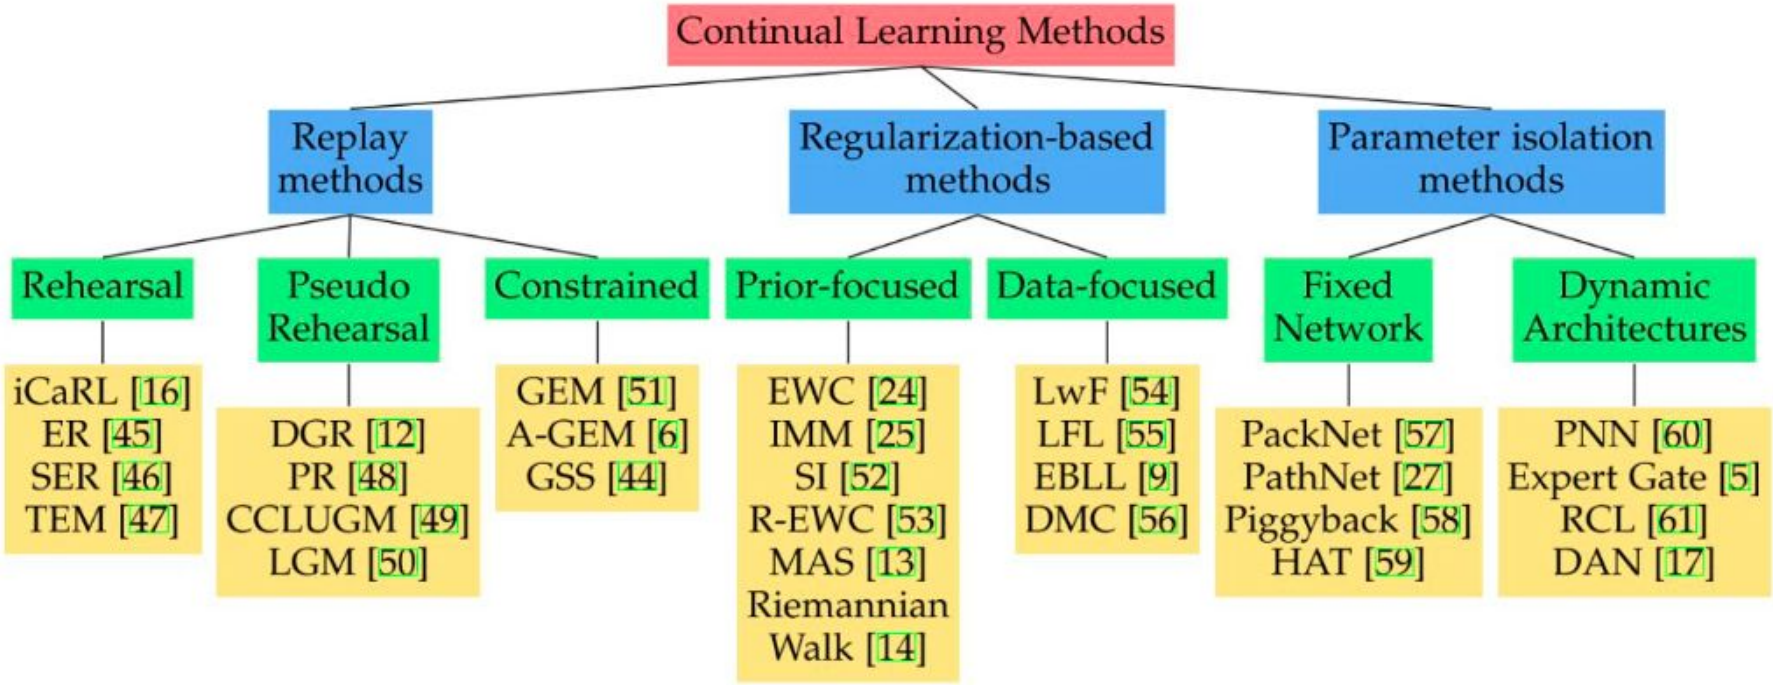

Lange et. al. (2020)

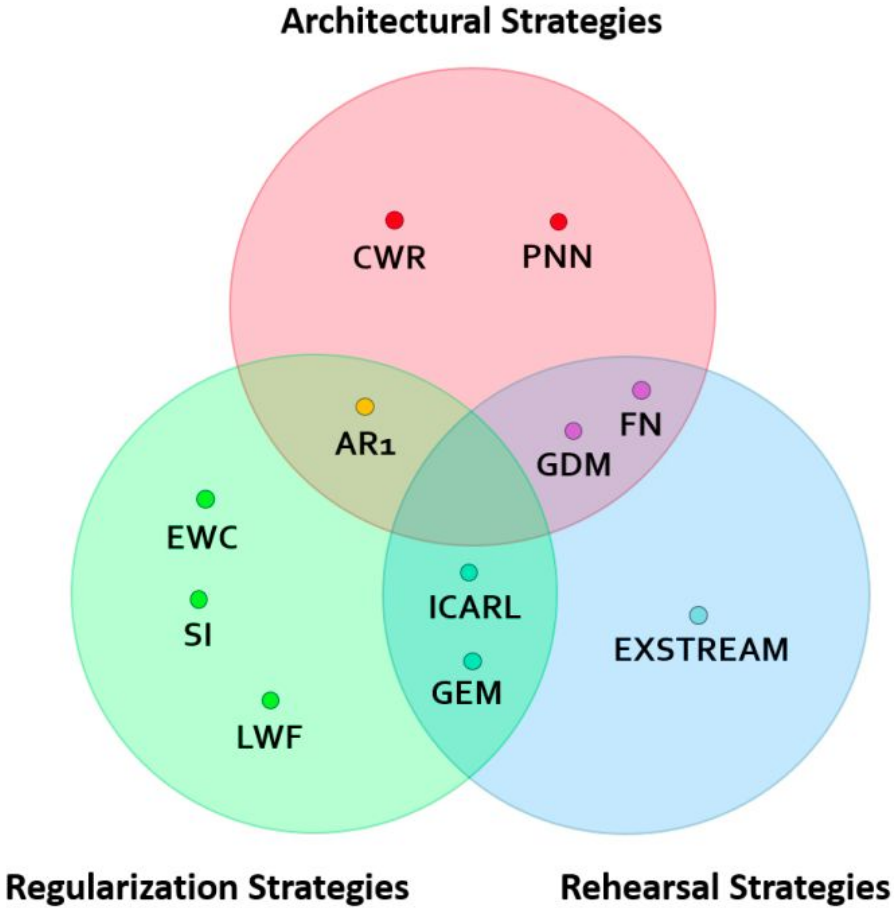

Maltoni & Lomonaco (2019)

Neural Network

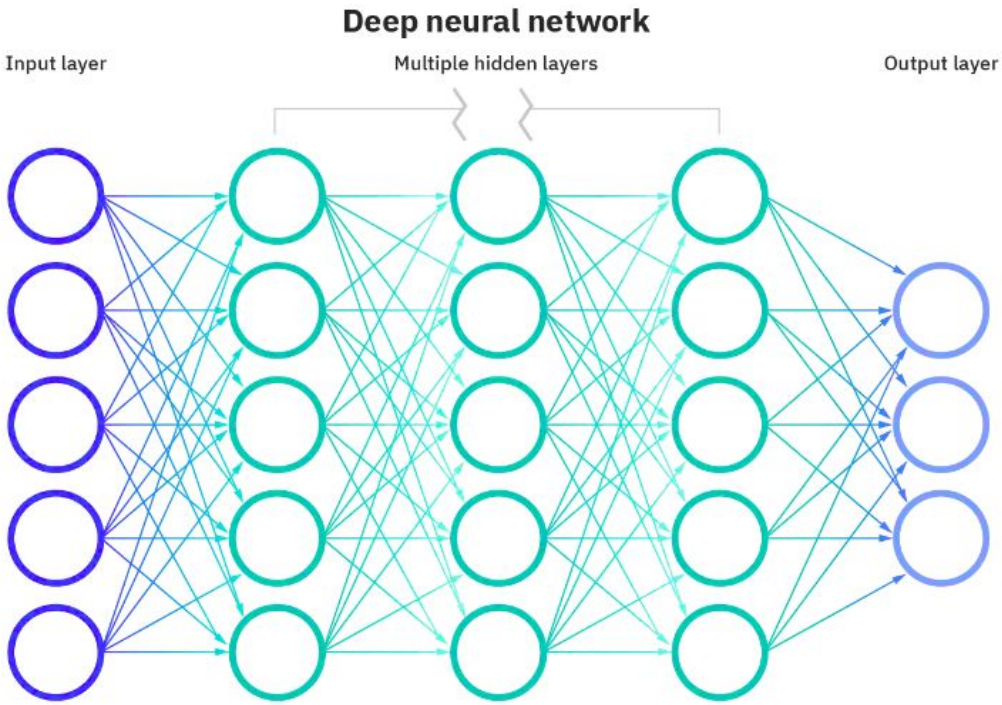

Convolution Neural Network (CNN)

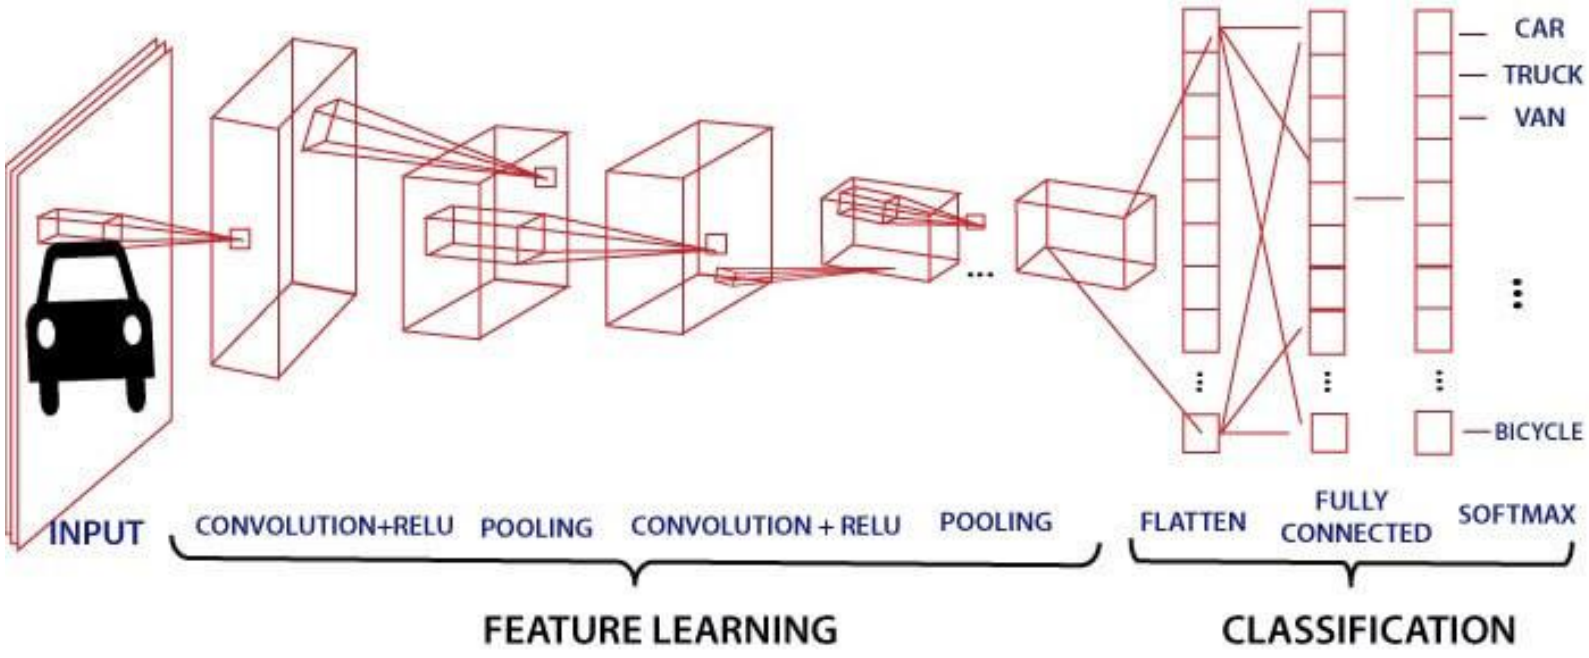

Model-growth Based Approach

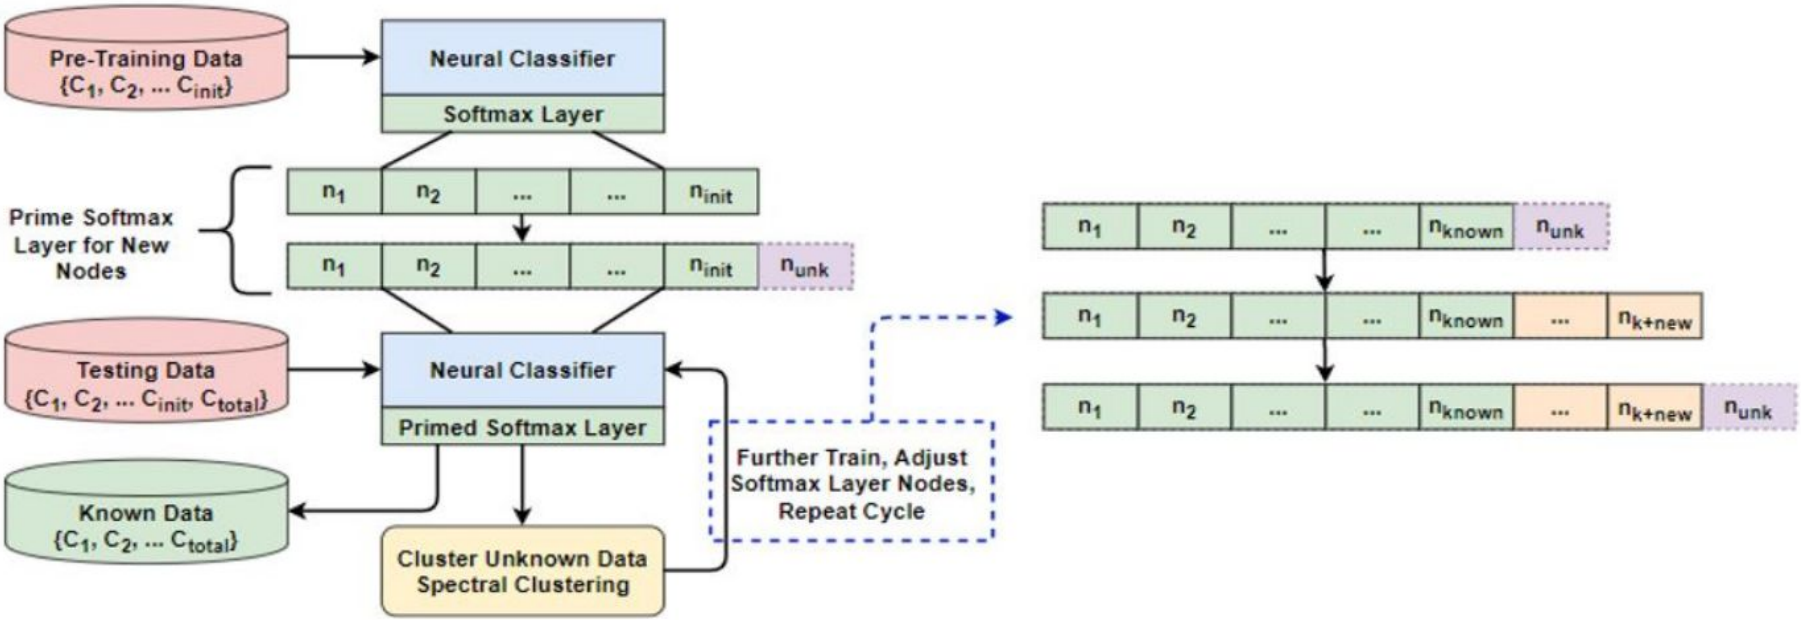

Leo & Kalita (2021)

Rehearsal/Replay Based Approach

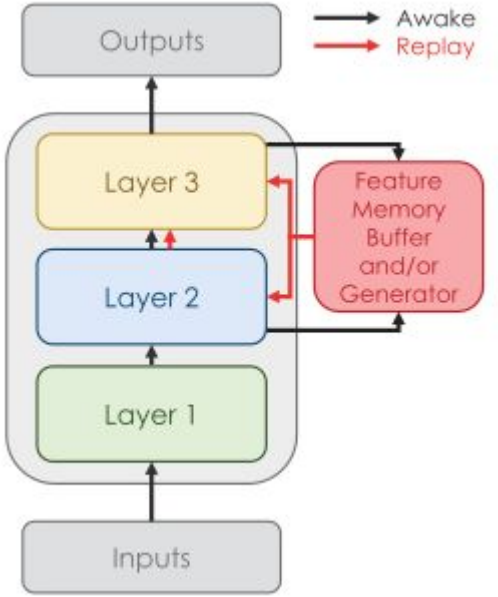

Store old knowledge in memory

- Store some samples of each class in additional memory.
- Store some intermediate representation of each class to save memory.
- Synthesize samples using a generative model.

Hayes et. al (2021)

Regularization Based Approach

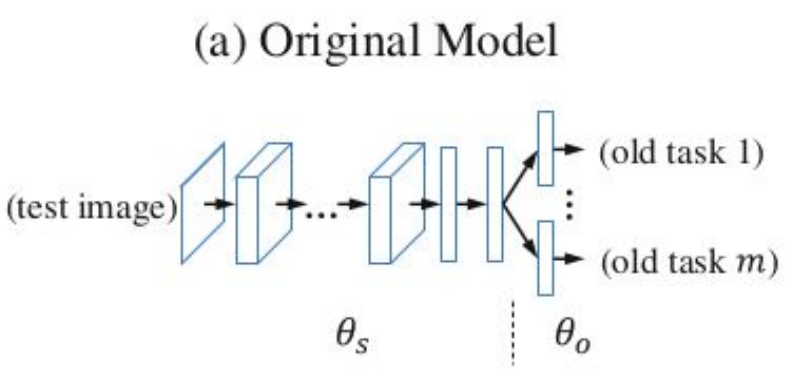

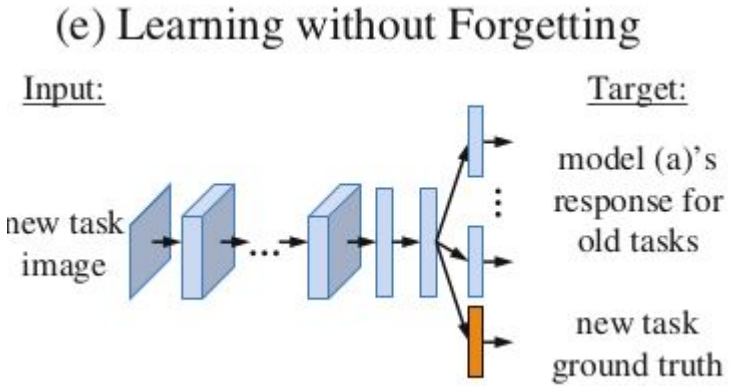

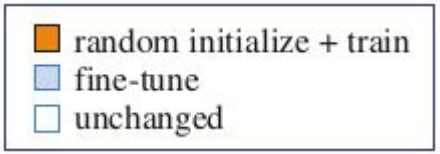

Li & Hoiem (2017)

Batch vs Streaming

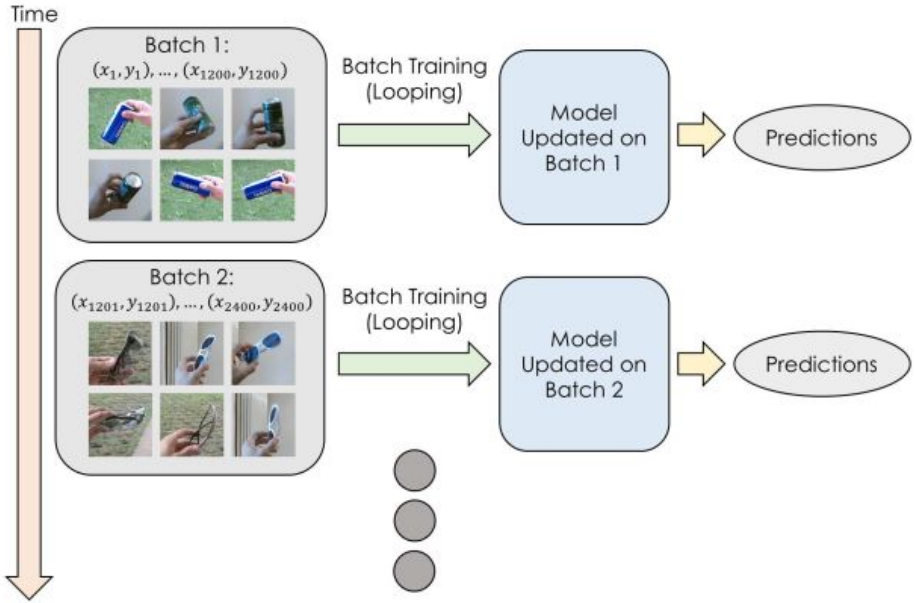

(a) Incremental Batch Learning

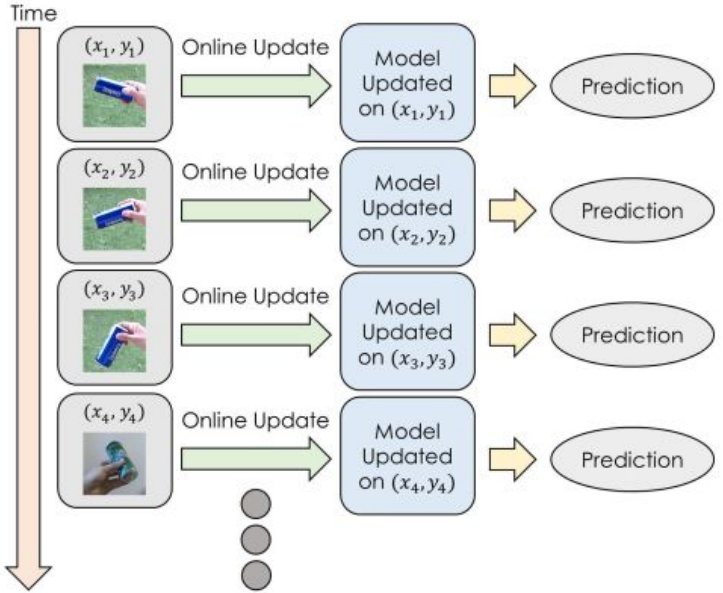

(b) Streaming Learning

Hayes & Kanan (2020)

An Ideal Deep Learning Classification Model

- Continuously learn new classes over time
- Not forget old classes
- Learn a new class from a single sample
- Time to learn a new class must be reasonable
- Limit memory and compute

In literature, this type of learning is termed as “**Online Streaming Class Incremental Learning**”.

Deep-SLDA (Hayes & Kanan 2020)

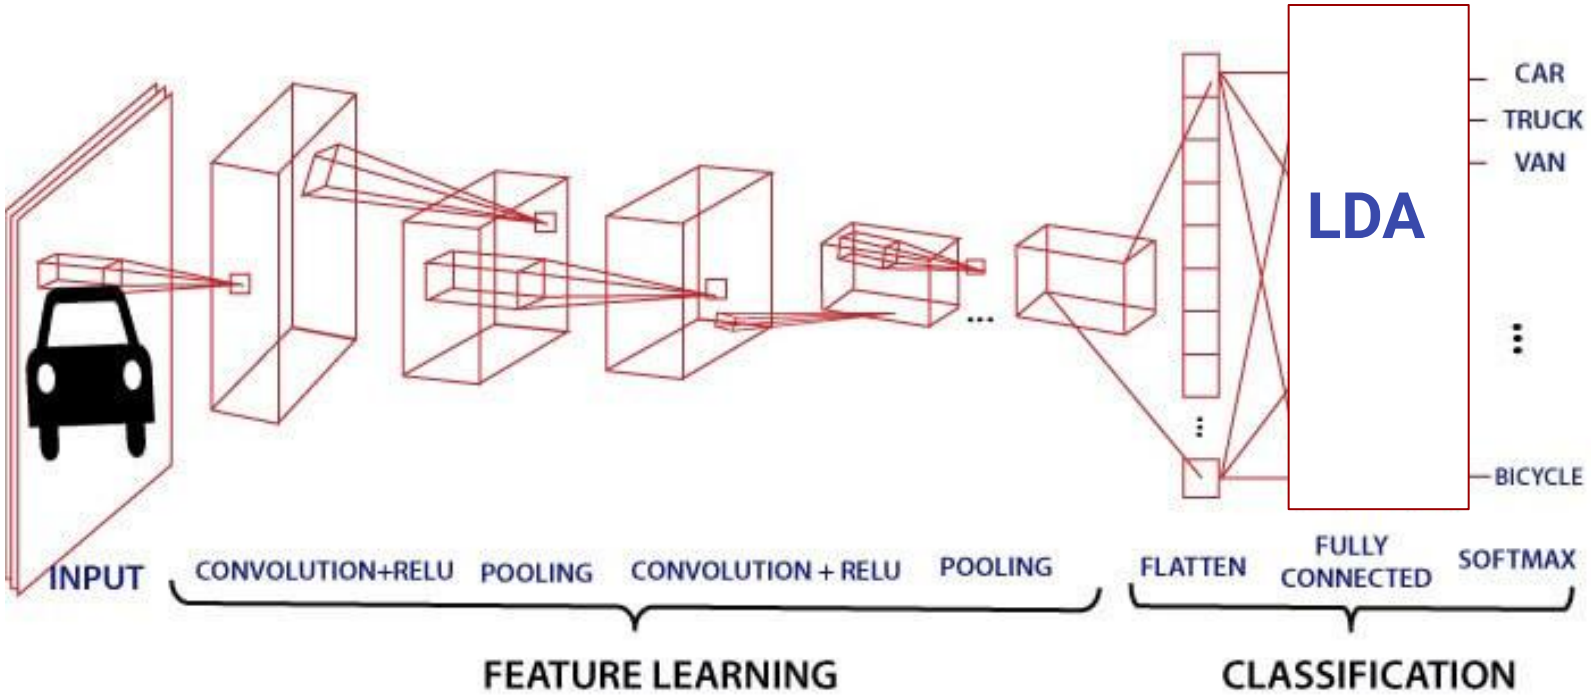

Linear Discriminant Analysis (LDA)

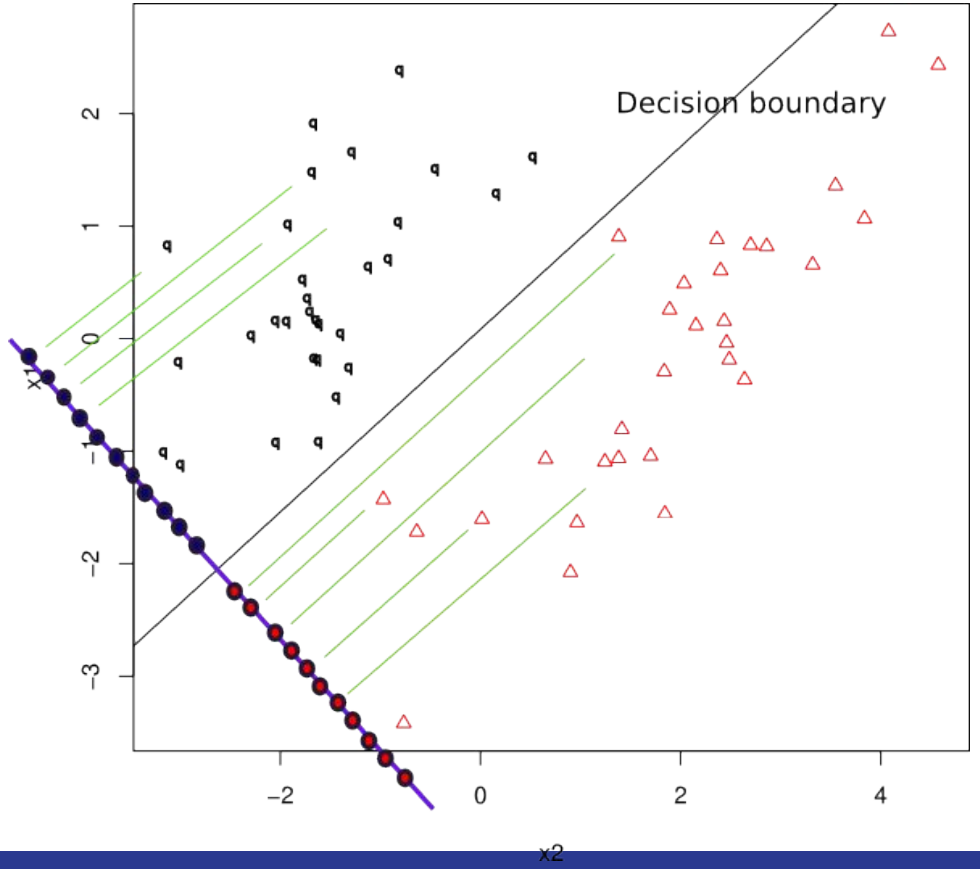

Proposed Design

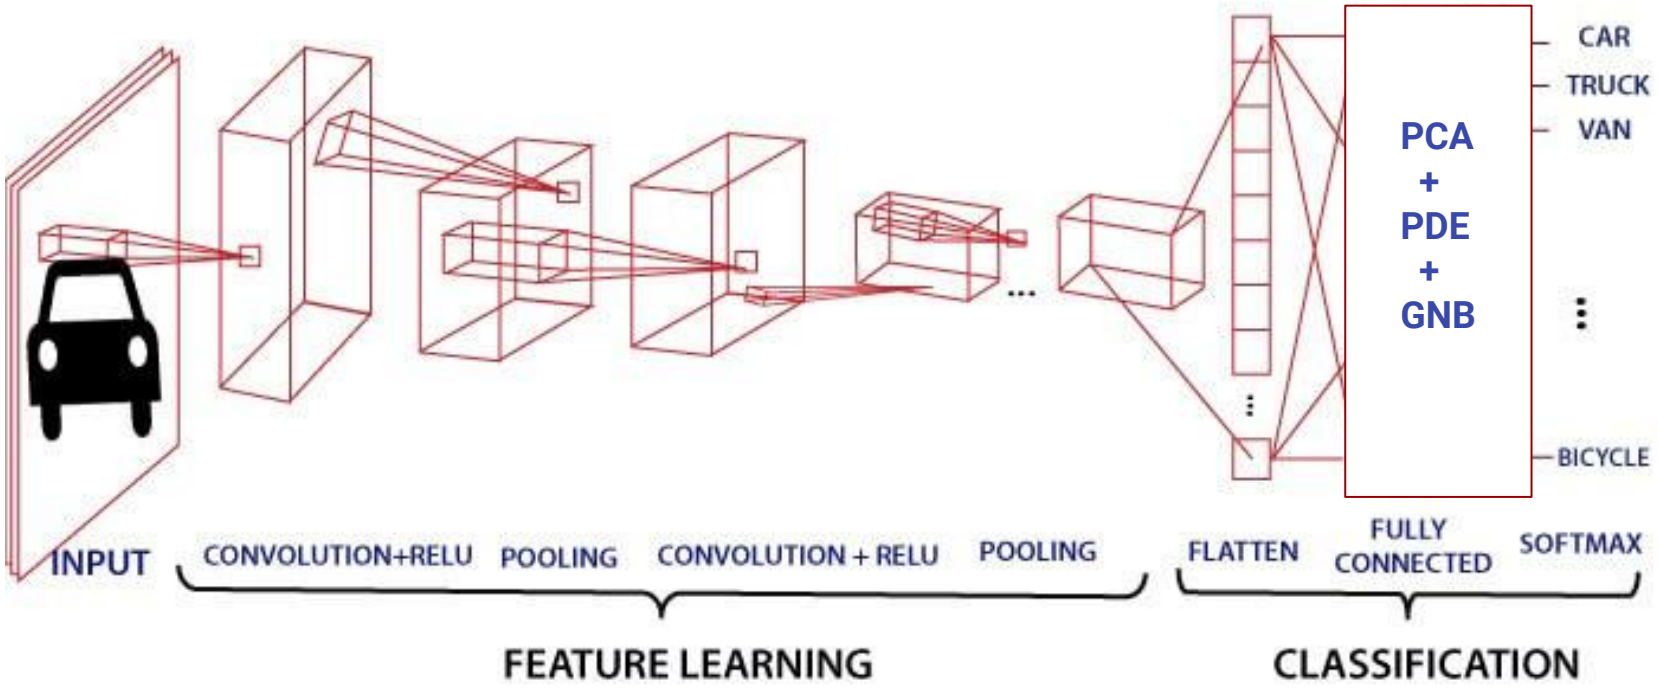

Principal Component Analysis (PCA)

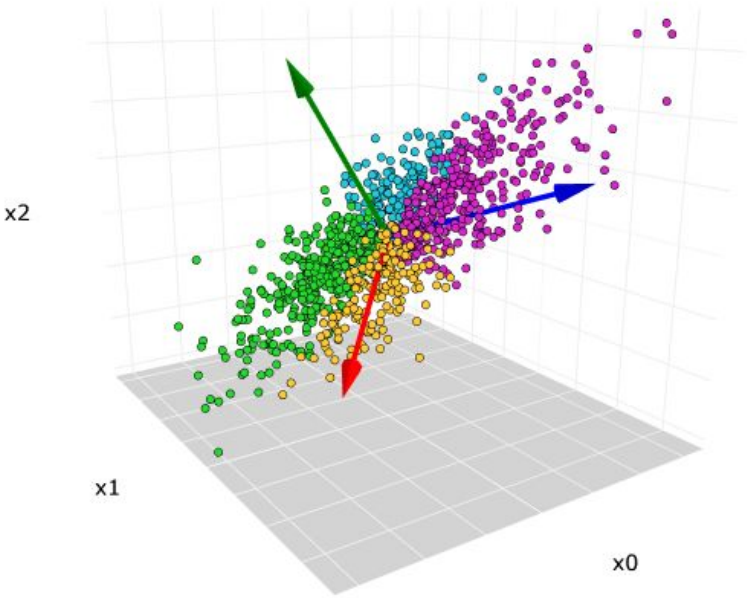

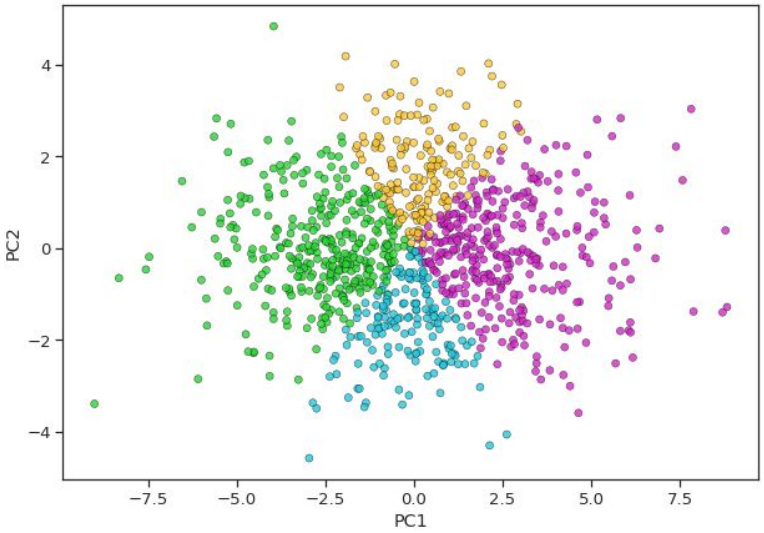

Probability Density Estimation (PDE)

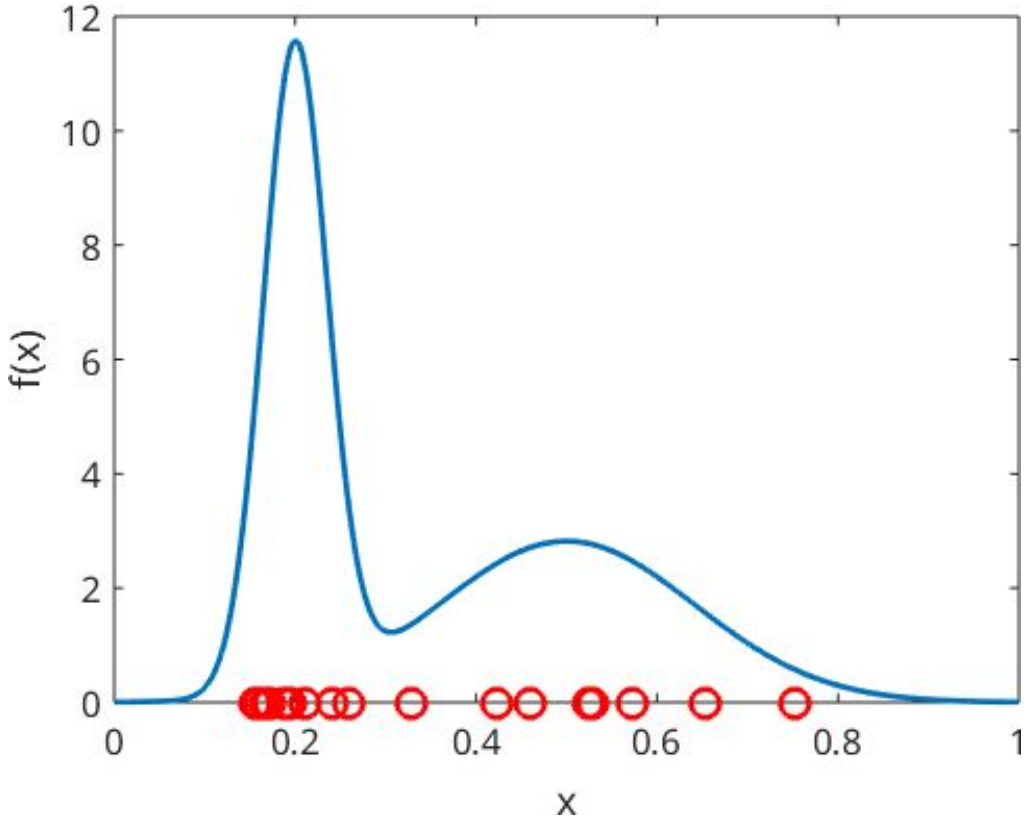

Gaussian Naive Bayes Classifier (GNB)

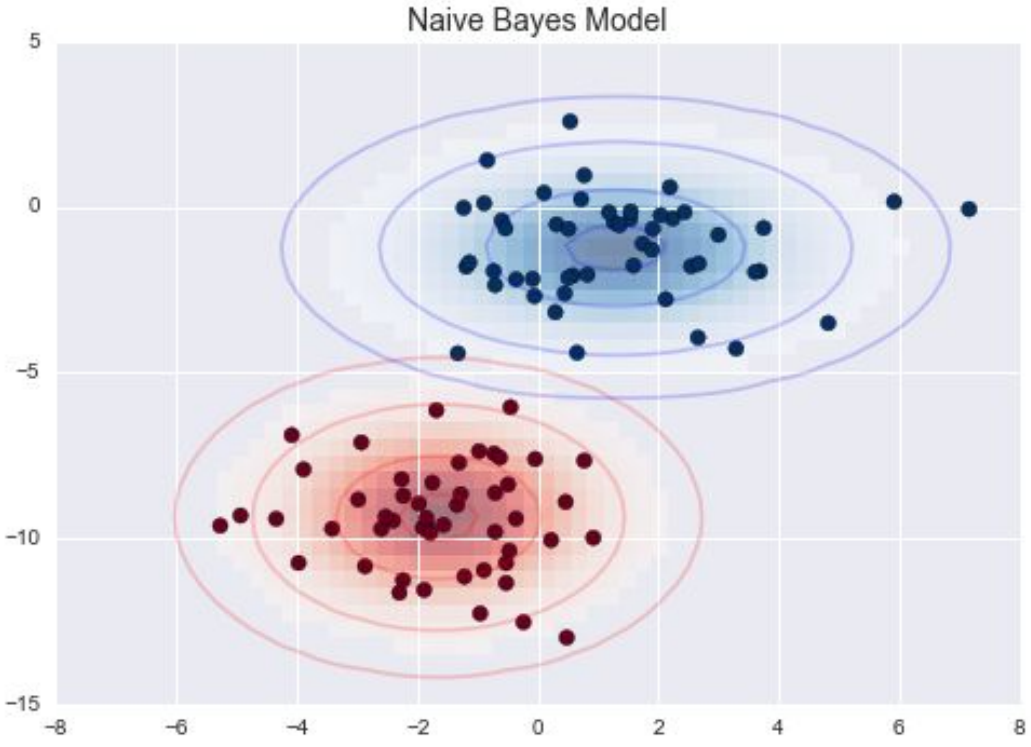

Incremental Setting

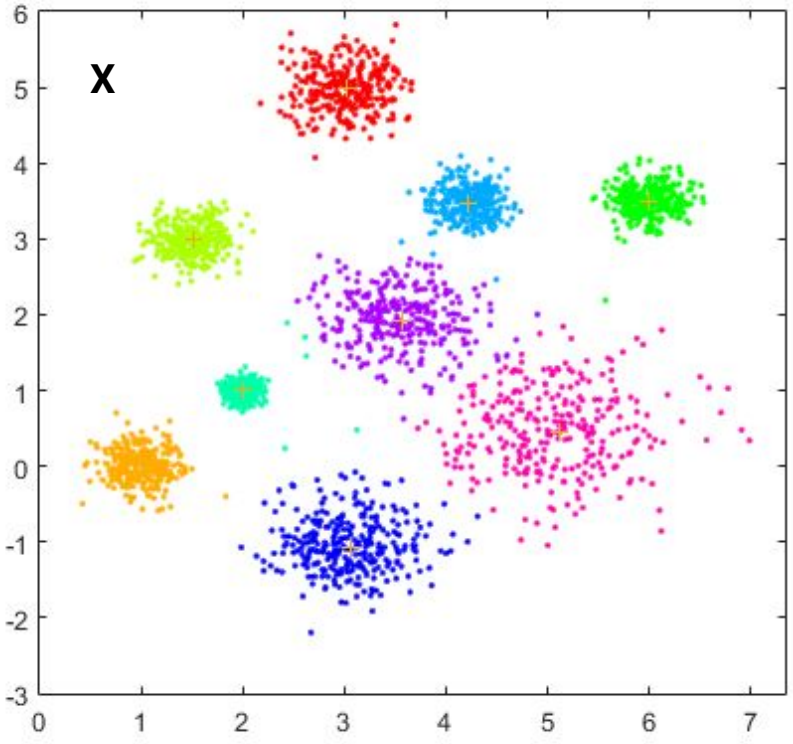

Incremental Setting

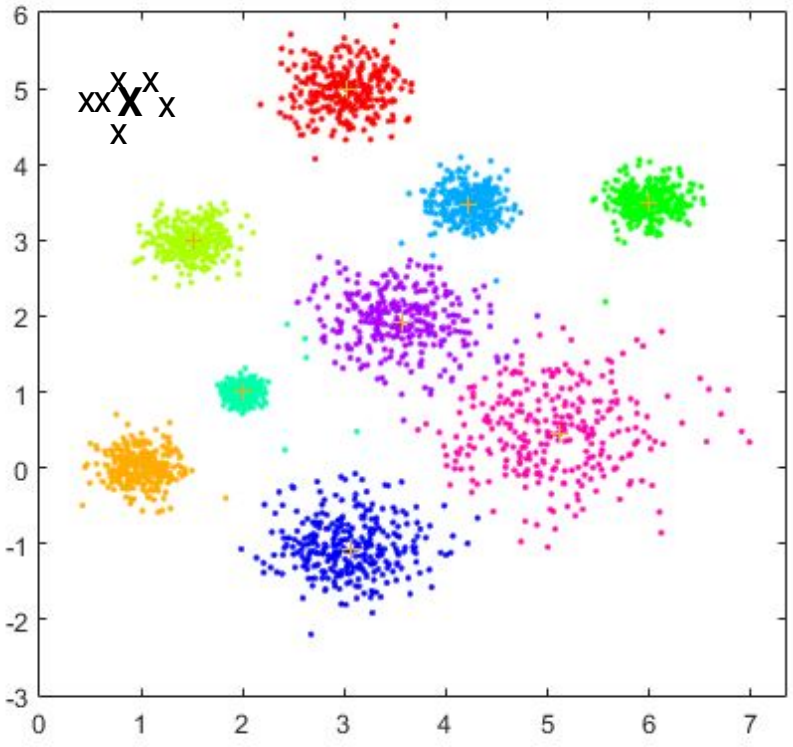

Roadmap

- Implementation in Python Pytorch
- Benchmark: MNIST, ImageNet
- Baseline: Offline mode
- Experiment: iCARL, LwF, PNN, Deep-SLDA
- Performance Evaluation: Average incremental accuracy

MNIST Dataset

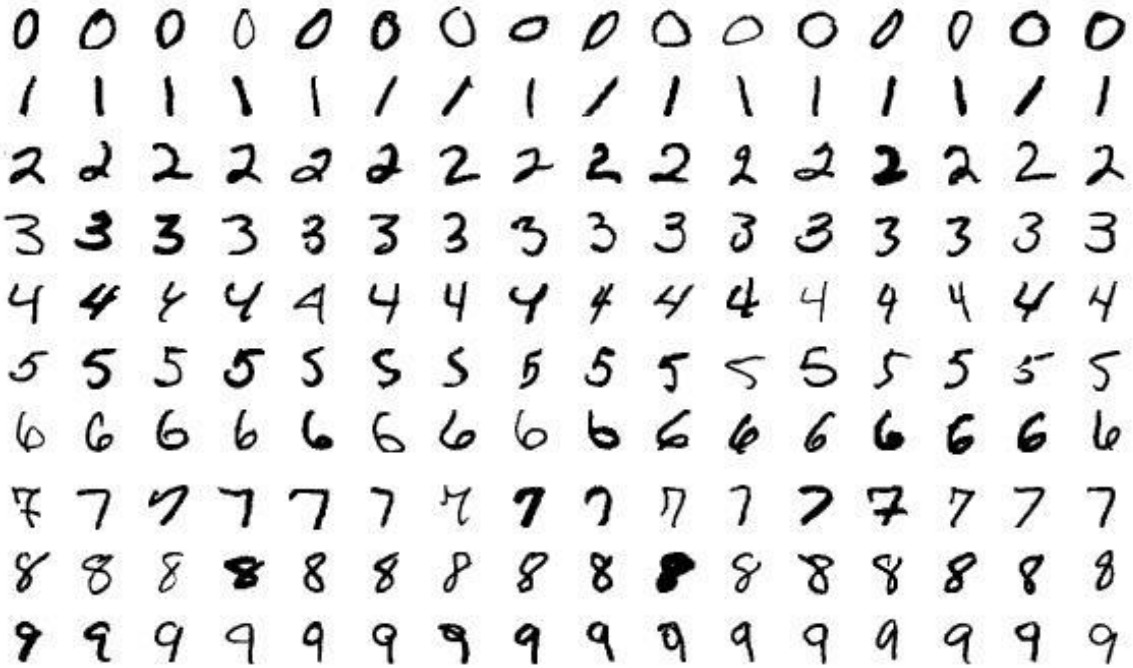

ImageNet

Average Incremental Accuracy

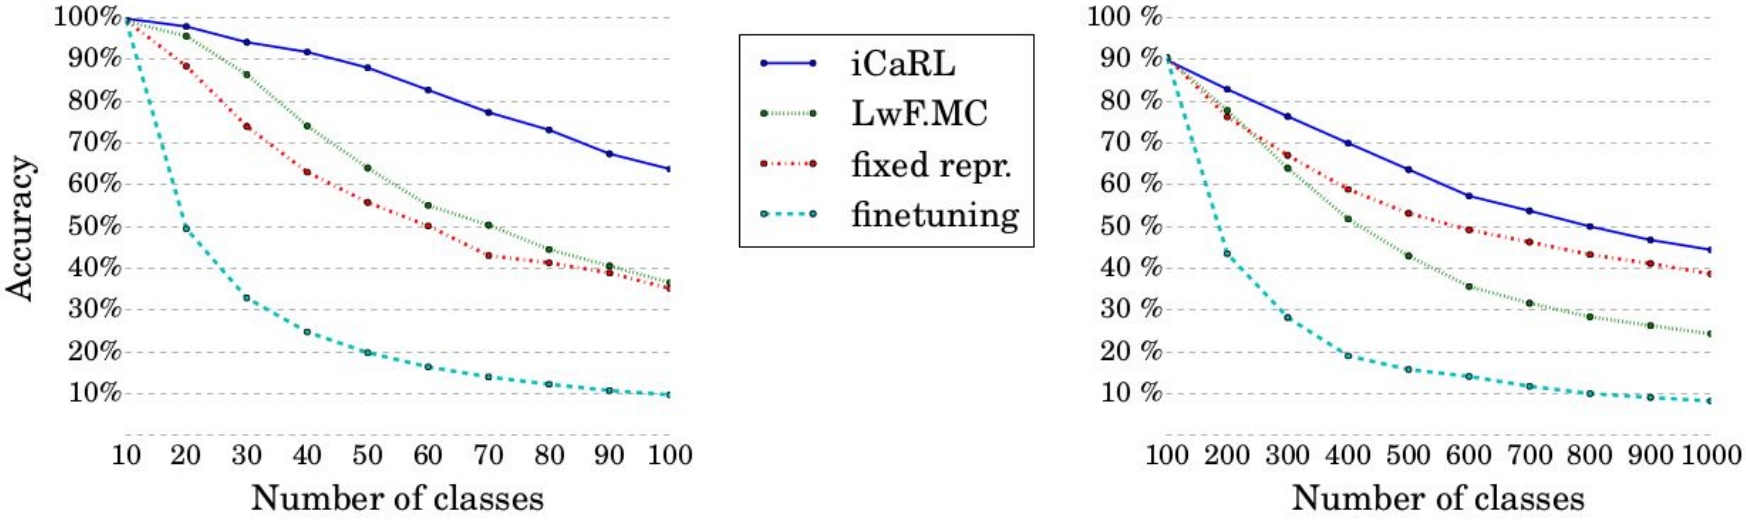

Rebuffi et. al (2017)

References

- Michael McCloskey and Neal Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. Psychology of Learning and Motivation, Vol. 24. Academic Press, 109–165.
- Tyler Hayes and Christopher Kanan. 2020. Lifelong Machine Learning with Deep Streaming Linear Discriminant Analysis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPR-W). 220–221.
- Justin Leo and Jugal Kalita. 2021. Incremental Deep Neural Network Learning using Classification Confidence Thresholding. IEEE Transactions on Neural Networks and Learning Systems (2021).
- Zhizhong Li and Derek Hoiem. 2017. Learning w. Learning without Forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence 40, 12 (2017), 2935–2947.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. 2021. A continual learning survey: Defying forgetting in classification tasks. IEEE transactions on pattern analysis and machine intelligence 44, 7 (2021), 3366–3385.
- Liu, W., Yang, H., Sun, Y., Sun, C. (2018). A Broad Neural Network Structure for Class Incremental Learning. In: Huang, T., Lv, J., Sun, C., Tuzikov, A. (eds) Advances in Neural Networks – ISNN 2018. ISNN 2018. Lecture Notes in Computer Science(), vol 10878. Springer, Cham.
- Shaoning Pang, Seiichi Ozawa, and Nikola Kasabov. Incremental linear discriminant analysis for classification of data streams. IEEE Trans. on Systems, Man, and Cybernetics, part B (Cybernetics), 35(5):905–914, 2005.

Thank you!