Karlsruhe Institute of Technology

Advanced Topics in Continual / Organic Machine Learning

Interactive Systems Lab (ISL)
Institute for Anthropomatics and Robotics (IAR)

Summer 22 Topics

"I'm still learning"
Michelangelo

CostFree Incremental Learning (CF-IL)

- Previous experience is recovering by memory replay paradigm in a single learner
- No need for extracted prior knowledge, external memory
- Evaluated for both Class and Task Incremental Learning

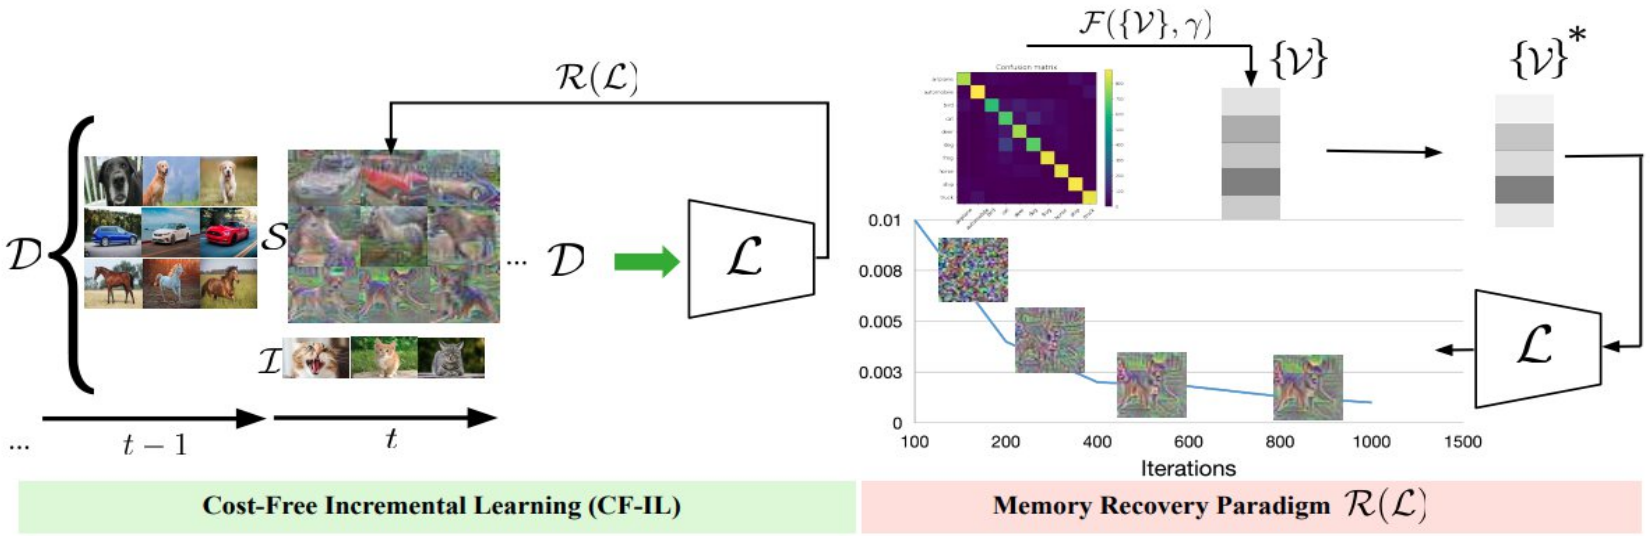

Paper:

PourKeshavarz, Mozhgan, et al. “Looking Back on Learned Experiences for Class/Task Incremental Learning”, ICLR 2022

Incremental Few-Shot Object Detection

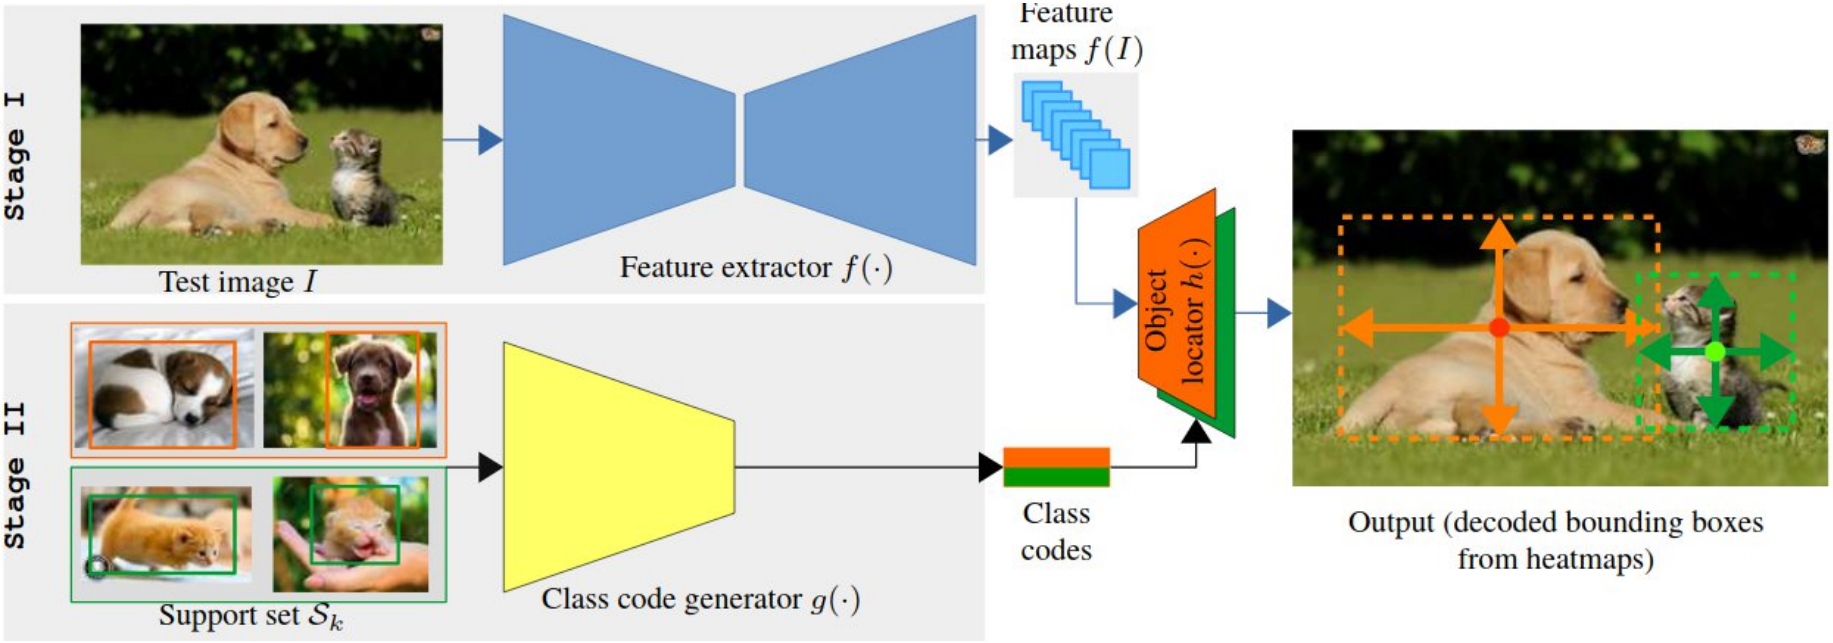

- New will be learned incrementally without revisiting base classes
- Learn to detect new classes using only few images

Paper: Perez-Rua, Juan-Manuel, et al. “Incremental few-shot object detection”, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2020

Efficient Feature Transformations for Discriminative and Generative Continual Learning

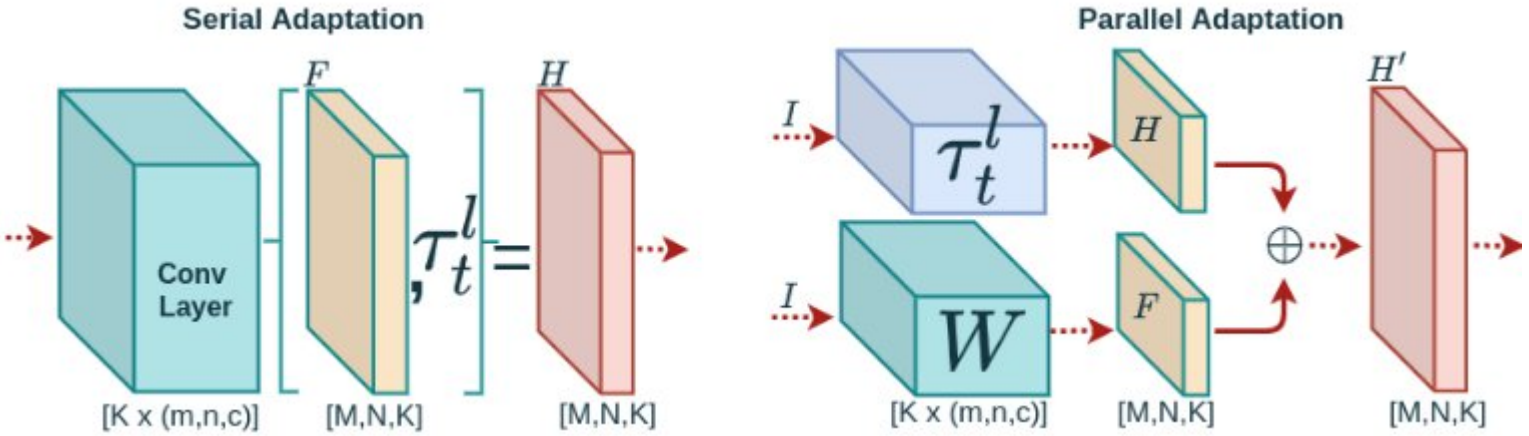

Fig 1. Transform global feature map F to a task-specific feature map H.

- Task-specific feature map transformation
- Parameter-efficient
- Feature distance maximization to improve task prediction

Paper: Verma, Vinay Kumar, et al. “Efficient feature transformations for discriminative and generative continual learning.” CVPR 2021.

Rehearsal Memory for Long Sequences

- Memory-Augmented NNs can handle arbitrarily long sequences
- However, prone to forgetting early content
- Self-Supervised Rehearsal Training circumvents this

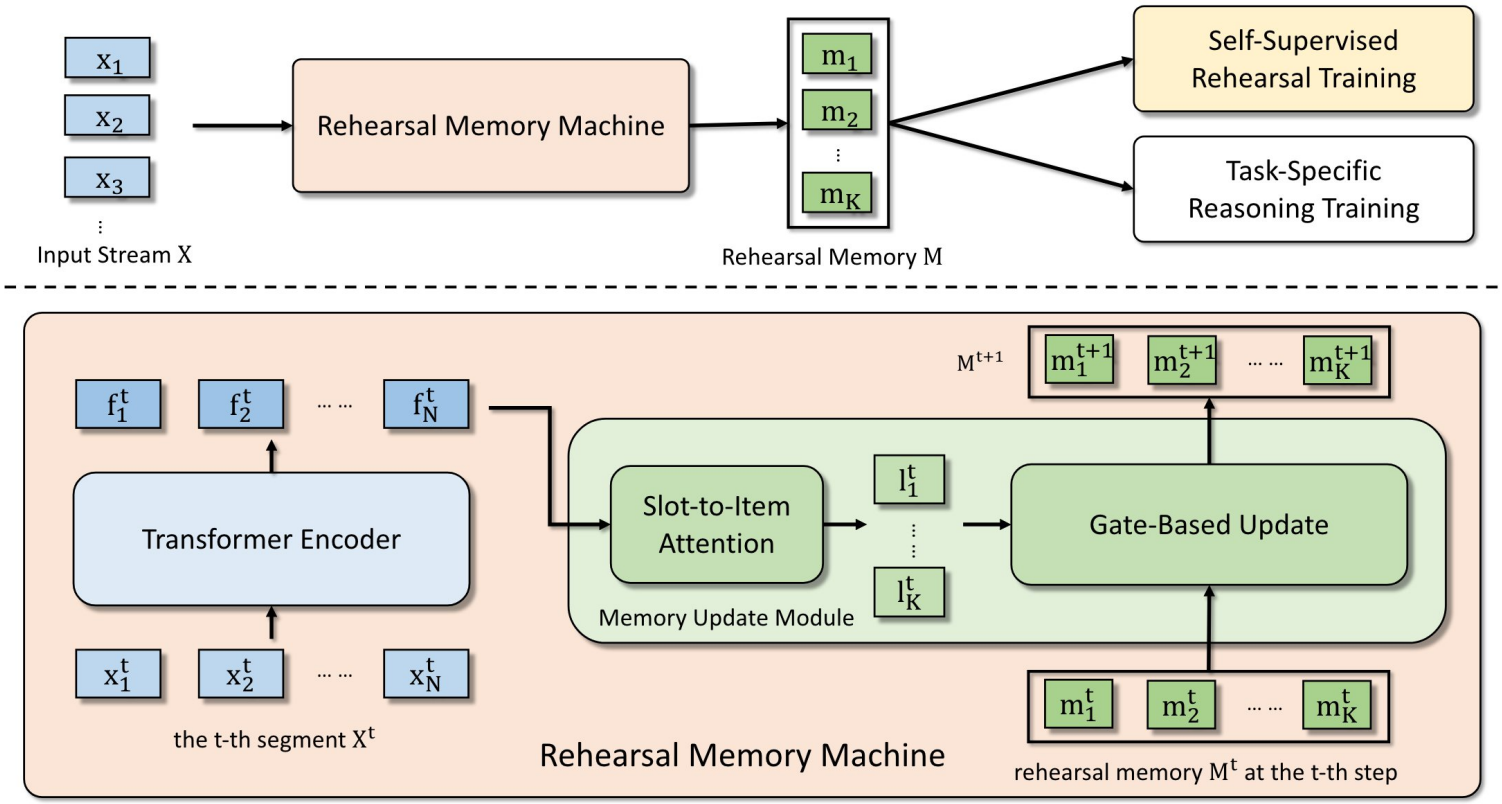

Paper:

- Zhang er al., “Learning to Rehearse in Long Sequence Memorization.” ICML, 2021

Long-Range Transformers – Survey

- Transformer models are SotA for sequence modeling
- Quadratic complexity wrt input length → problem with long sequences
- Various approaches in the literature
  - Hierarchy
  - Sparse Attention
  - Recurrence / Memory

Standard Transformer:

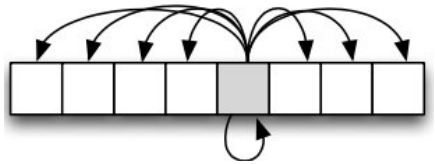

Transformer XL:

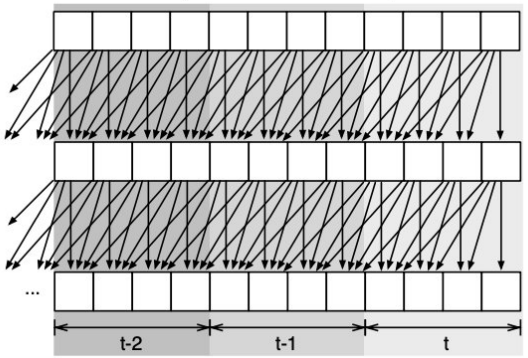

Hierarchical Attention (HIBERT):

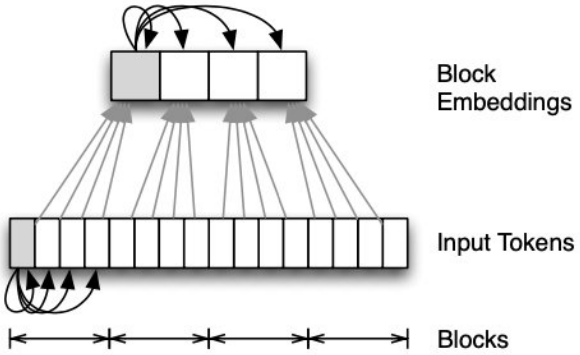

Global-Local Attention (ETC):

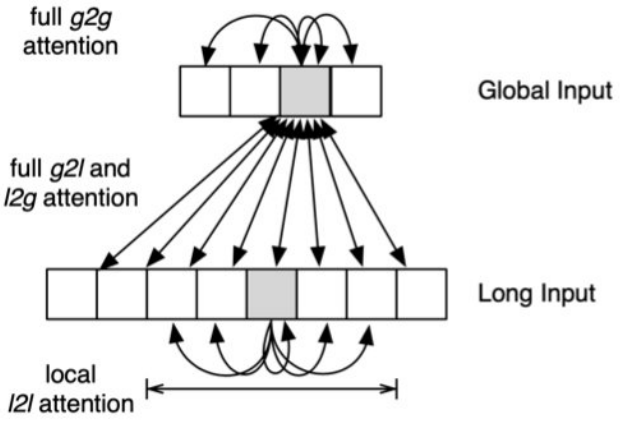

Compressive Transformer:

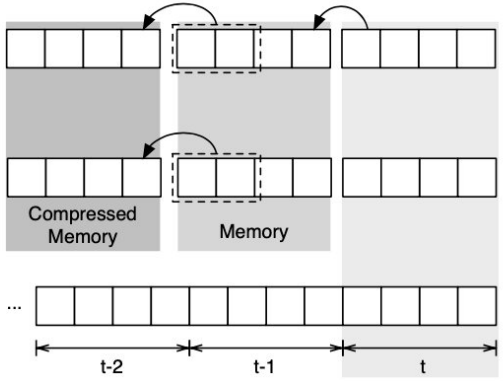

Literature suggestions:

- Tay et al., “Efficient Transformers: A Survey.” arXiv, 2022
- Tay et al., “Long Range Arena : A Benchmark for Efficient Transformers.” ICLR, 2021
- Sukhbaatar et al., “Not All Memories Are Created Equal: Learning to Forget by Expiring.” ICML 2021
- Wu et al., “Memformer: The Memory-Augmented Transformer.” arXiv, 2020
- Zaheer et al., “Big Bird: Transformers for Longer Sequences.” NIPS, 2020
- Zhang et al., “HIBERT: Document Level Pre-Training of Hierarchical Bidirectional Transformers for Document Summarization.” ACL, 2019

Self-Supervised Learning

- A machine learning system that uses supervised learning techniques (e.g. NNs) to learn from automatically labelled data
  - No labelled data needed
  - Explain methods/applications

- Emerging properties in self-supervised vision transformers (Caron et al. 2021)

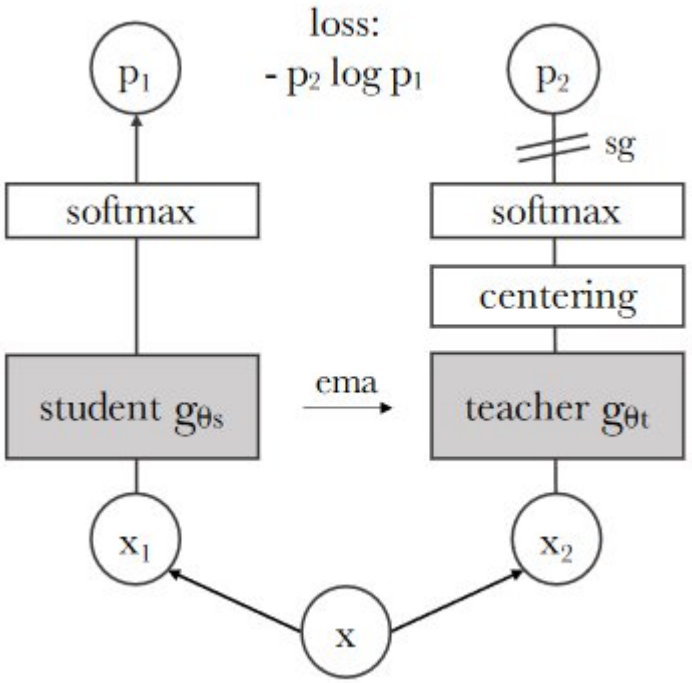

Efficient Conditional GAN Transfer with Knowledge Propagation across Classes

- **Task:** Conditional Image Generation with GAN
- **Method:**
  - GAN transfer method to explicitly propagate the knowledge from the old classes to the new classes
  - Enforce batch normalization (BN) to learn class-specific information of new classes with implicit knowledge sharing among the new ones.

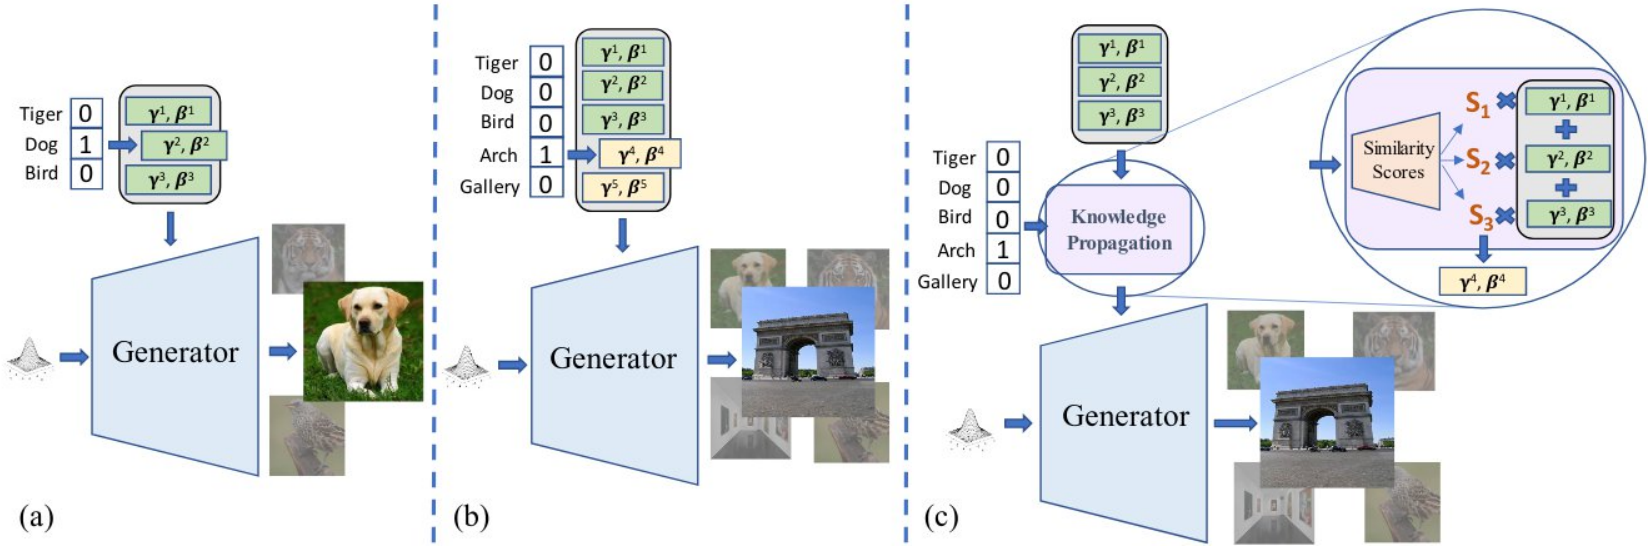

Papers:

- Shahbazi, Mohamad, et al. “Efficient conditional gan transfer with knowledge propagation across classes.” CVPR 2021.

(https://openaccess.thecvf.com/content/CVPR2021/papers/Shahbazi_Efficient_Conditional_GAN_Transfer_With_Knowledge_Propagation_Across_Classes_CVPR_2021_paper.pdf)

Vision and Language Knowledge Distillation

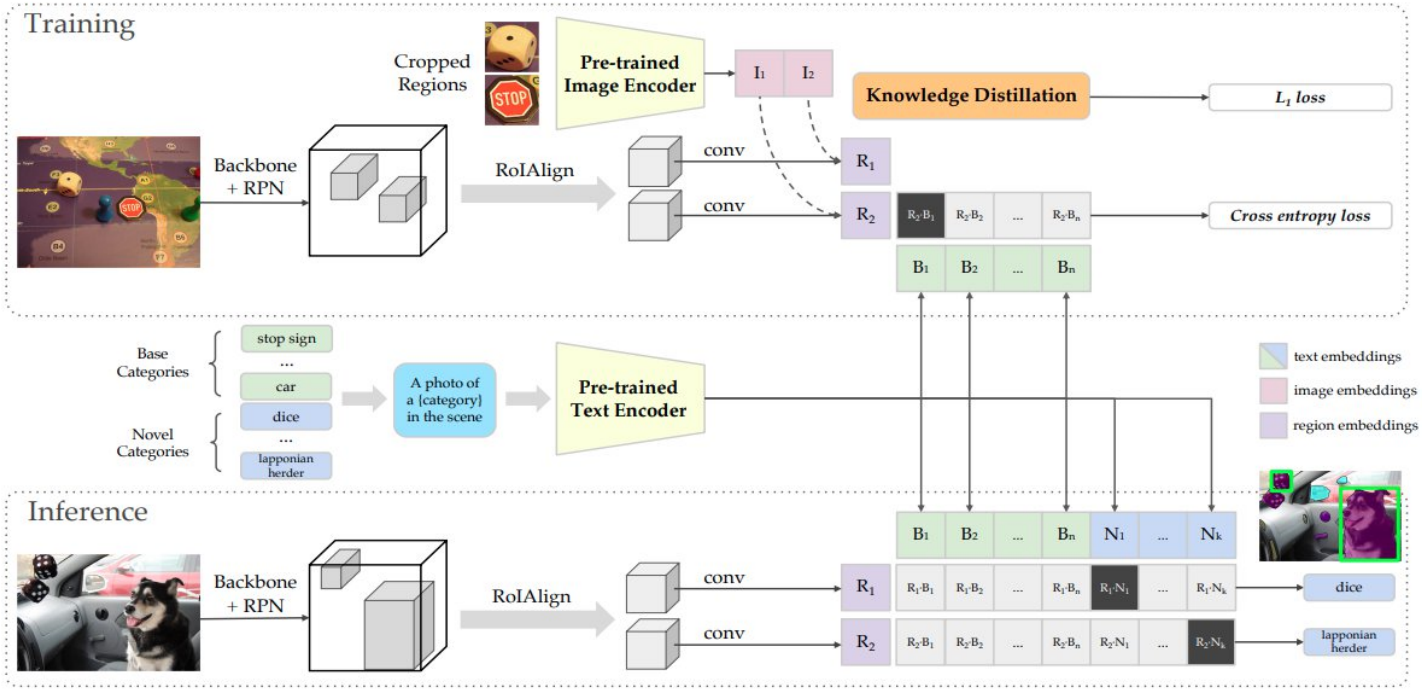

- Objects detection using arbitrary text description
- Append new classes without re-training of the detector
- An open-vocabulary detection method by distilling knowledge from a zero-shot image classification model.

Paper: Gu, Xiuye et al. “Open-Vocabulary Object Detection Via Vision and Language Knowledge Distillation”, ICLR 2022

Human-centric dialog learning

- a supervised trained dialog system can be a good base system
- humans give naturally feedback by using a dialog system (elicitation of laughter, sentiment, ...)
- use the feedback to improve the dialog system
- pay attention to eliminate the risk to learn harmful behaviours

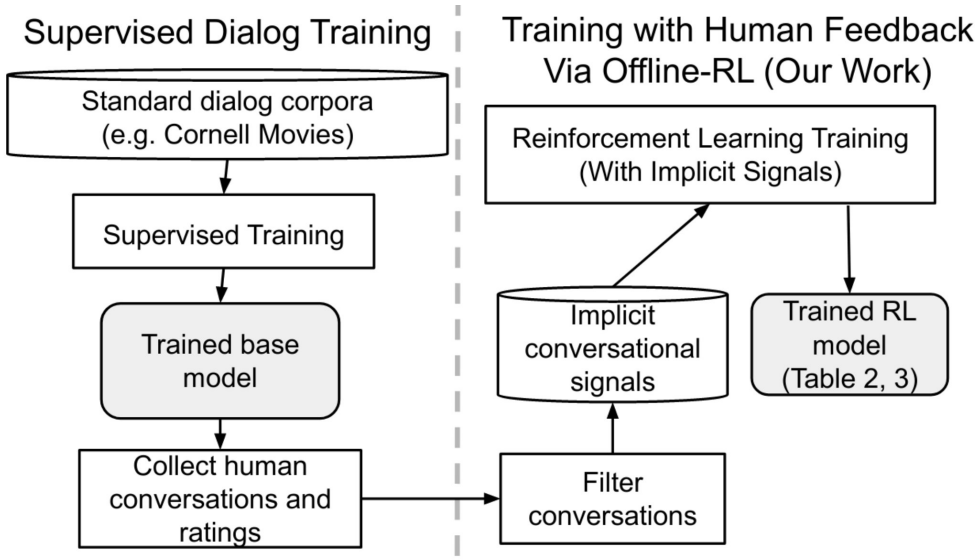

Paper: Natasha Jaques, Judy Hanwen Shen, et al. “Human-centric dialog training via offline reinforcement learning”, EMNLP 2020

Incremental Deep Neural Network Learning Using Classification Confidence Thresholding

- model with its own architecture and initial trained set of data can identify unknown classes during the testing phase and autonomously update itself if evidence of a new class is detected
- threshold approach to prime neural networks for incremental learning to keep accuracies high by limiting forgetting
- lean method to reduce resources used in the retraining

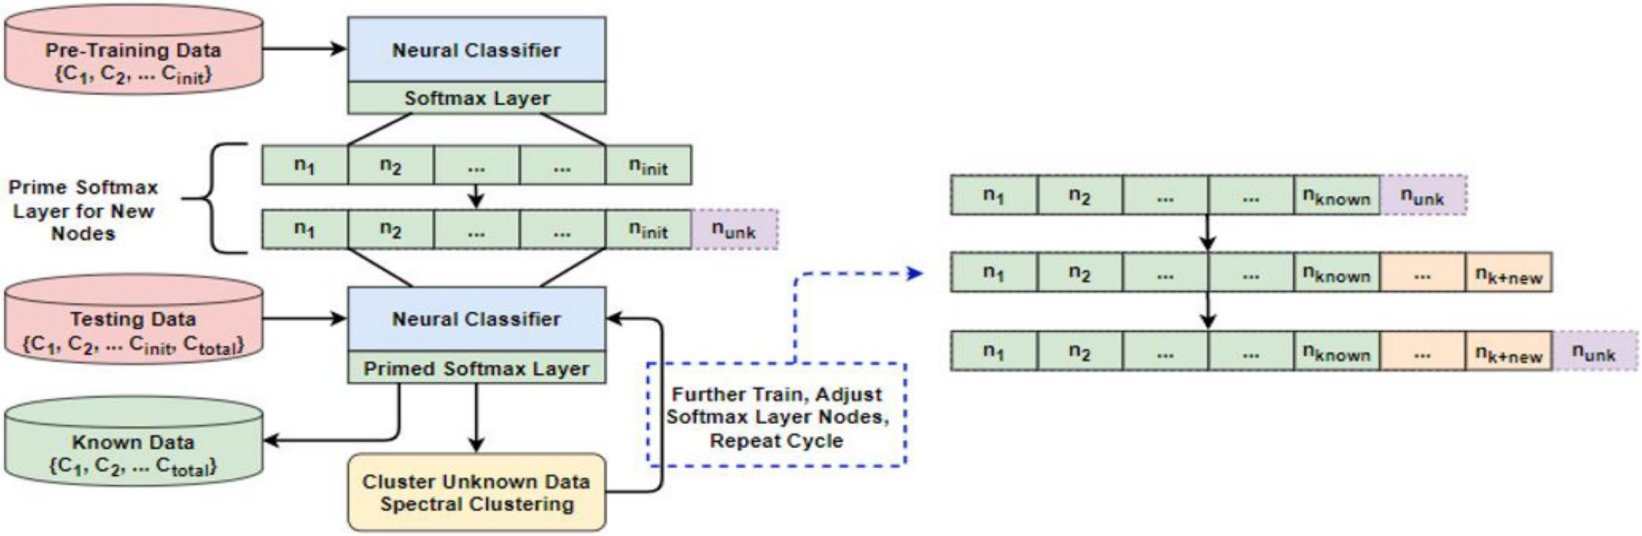

Paper: Justin Leo and Jugal Kalita, “Incremental Deep Neural Network Learning Using Classification Confidence Thresholding”, IEEE Transactions on Neural Networks and Learning Systems, 2021

Zero-Shot Task Generalization in Language Models

- Language models like GPT-3[4] have shown to obtain reasonable zero-shot capabilities
- Since then several works have been done in the area of instruction tuning and prompt engineering
- **Review recent approaches in the area of zero-shot learning with language models such as [1] [2] [3]**

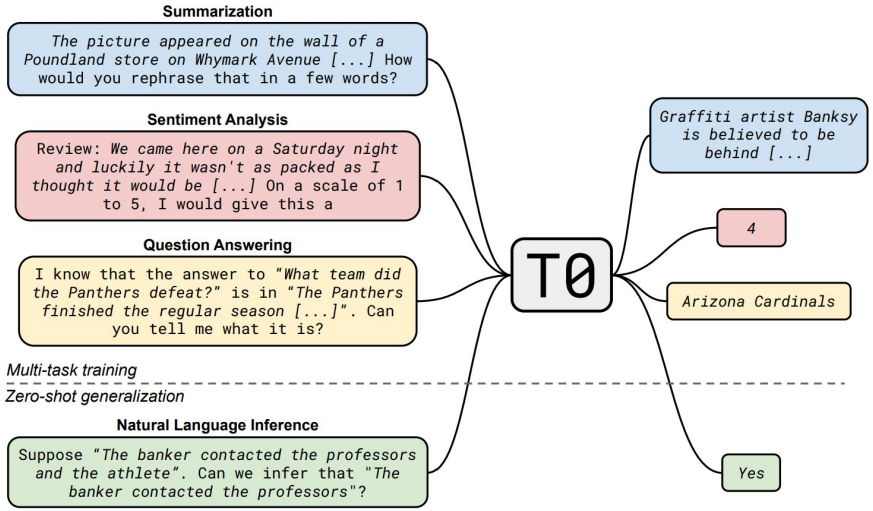

Papers (Suggestions):

- *[1]Multitask Prompted Training Enables Zero-Shot Task Generalization, Sanh et al. (2022)*
- *[2]Finetuned Language Models Are Zero-Shot Learners, Wei et al. (2021)*
- *[3]Adapting language models for zero-shot learning by meta-tuning on dataset and prompt collections, Zhong et al. (2021)*
- *[4]Language models are few-shot learners, Brown et al. (2020)*

(Few-Shot) Fine-Tuning of Pretrained Language Models

- **Review methods for (few-shot) fine-tuning of pretrained language models (PLM)**
- Most current approaches involve prompt engineering and handcrafting
  - For example, PET (Pattern Exploiting Training[4])
- PERFECT[1] promises few-shot learning without handcrafting, by having task-specific adapter layers & multi-token label embeddings

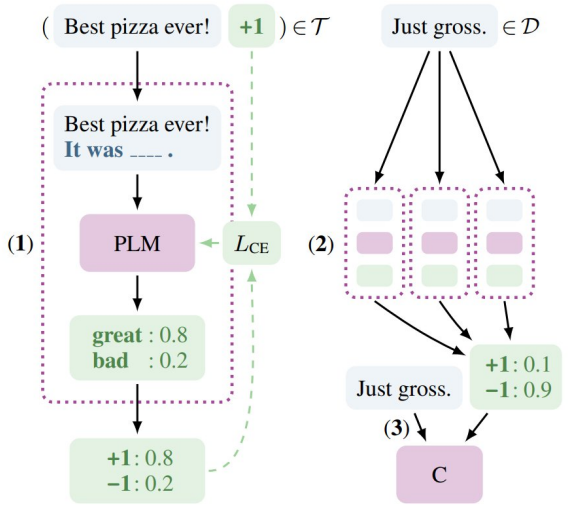

Papers (Suggestions):

- [1]*PERFECT: Prompt-free and Efficient Few-shot Learning with Language Models, Mahabadi et al. (2022)*
- [2]*CROSSFIT : A Few-shot Learning Challenge for Cross-task Generalization in NLP, Ye et al. (2021)*
- [3]*Prefix-Tuning: Optimizing Continuous Prompts for Generation, Li et al. (2021)*
- [4]*Exploiting cloze-questions for few-shot text classification and natural language inference, Schick et al. (2021)*