

Advanced Topics in Continual / Organic Machine Learning

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

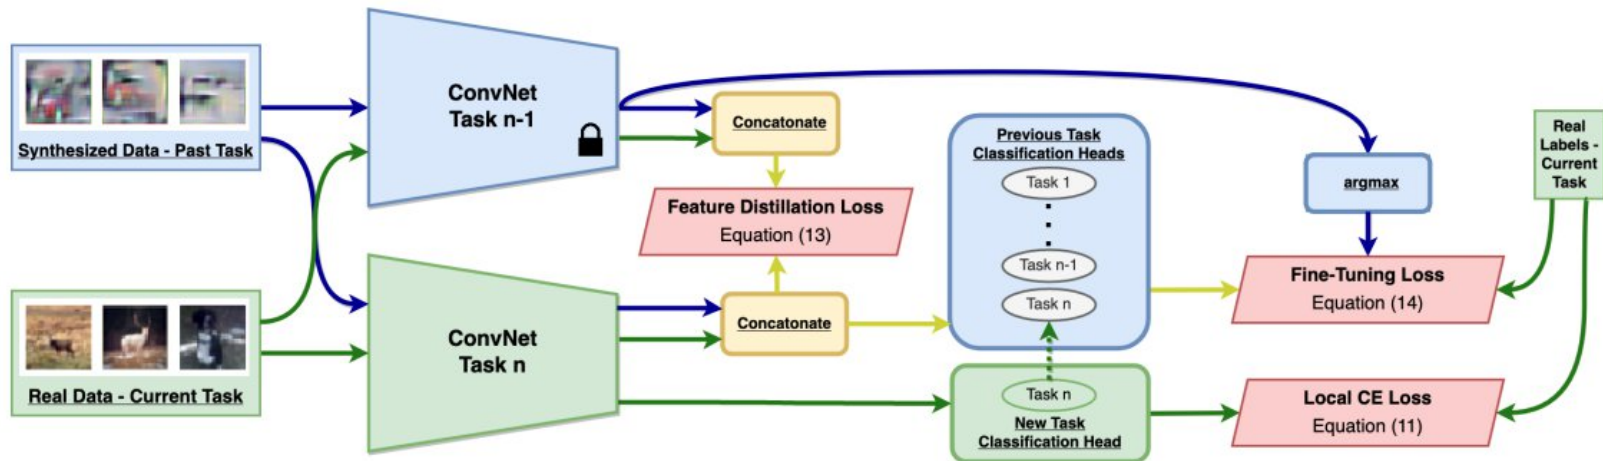
Winter 21-22 Topics

"I'm still learning"
Michelangelo



Data-Free Class Incremental Learning

- Class incremental learning using synthesized data from past tasks distribution.
- Any need to replay mechanism that stores real data from past tasks.
- Try to avoid bias to the new tasks which available real images.



Papers:

- Smith, James, et al. "Always Be Dreaming: A New Approach for Data-Free Class-Incremental Learning", ICCV 2021
- Yin, Hongxu, et al. "Dreaming to Distill: Data-free Knowledge Transfer via DeepInversion.", CVPR 2020.

Semi-Supervised Learning for Understanding Hand-Object Interaction

- Annotation and 3D pose estimation of hand object interactions from a single image are challenging problems.
- Semi-supervised learning with pseudo labels which created with the help of spatial temporal consistency of hand-object videos in the wild.




Papers:

Liu, Shaowei, et al. "Semi-supervised 3d hand-object poses estimation with interactions in time.", *CVPR 2021*.

Hasson, Yana, et al. "Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction.", *CVPR 2020*.

Lifelong Zero-Shot Learning

- Zero-Shot Learning: Recognize unseen classes using semantic attributes
- Lifelong ZSL: Accumulate knowledge across training datasets
- Model: Cross- and Distribution-Aligned VAE + selective retraining + Knowledge Distillation

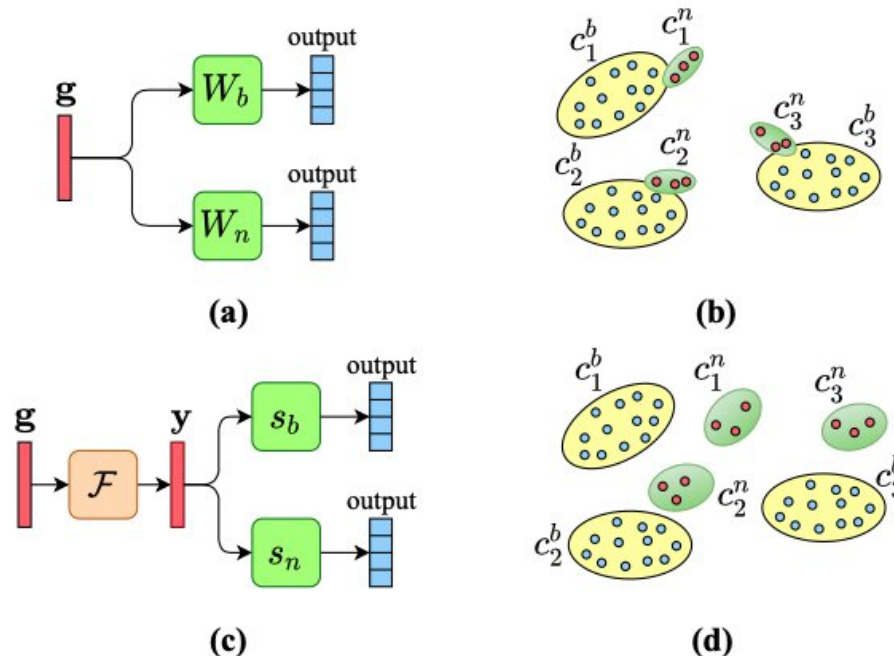


Paper:

- Wei et al. "Lifelong Zero-Shot Learning.", IJCAI-20, <https://doi.org/10.24963/ijcai.2020/77>.

Semantic-aware Knowledge Distillation for Few-Shot Class-Incremental Learning

- Class-Incremental = Continual = Life-Long Learning
- Few-Shot Learning: only very few samples for each new class
- Knowledge Distillation usually adds new weights
⇒ not useful for FSL, too few samples to train new network
- Solution: Use semantic word embeddings, avoid new parameters, fine-tune base network

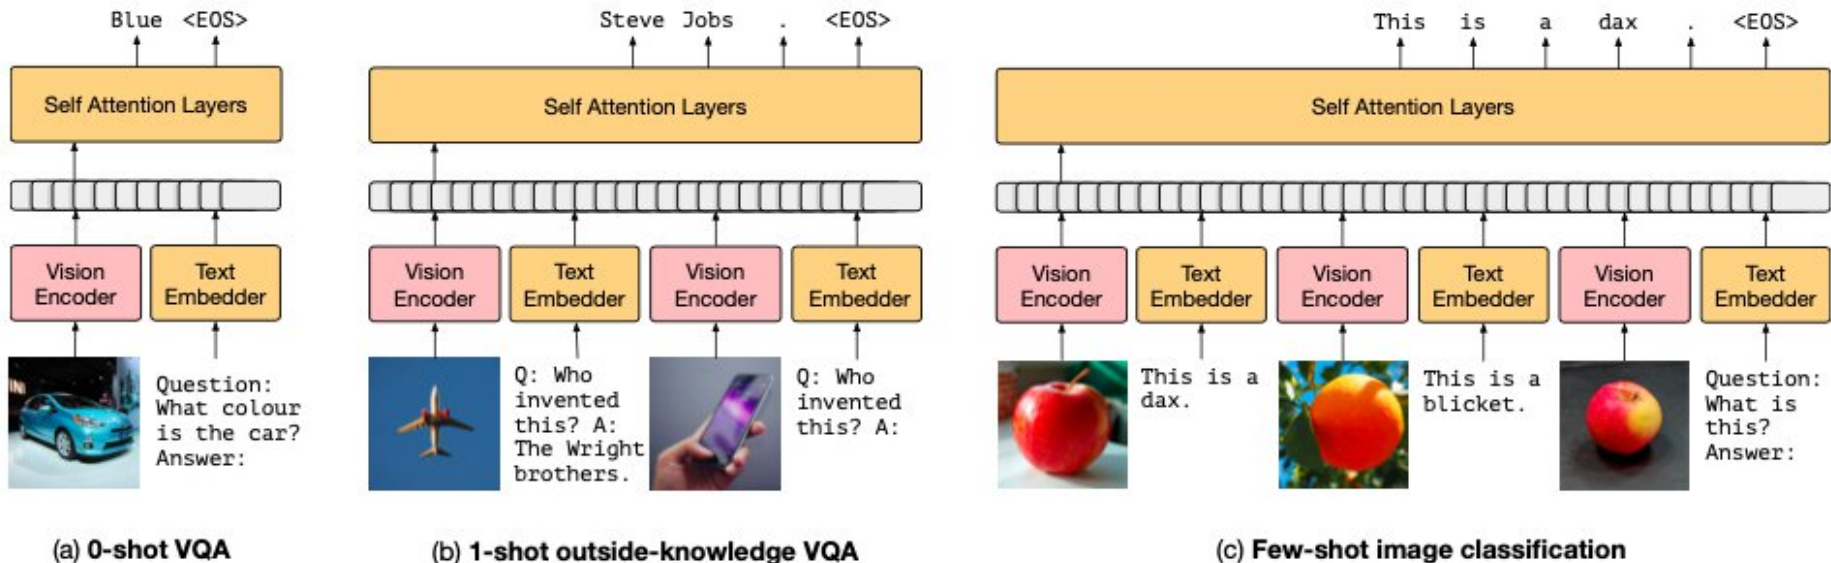


Paper:

- Cheraghian et al. "Semantic-Aware Knowledge Distillation for Few-Shot Class-Incremental Learning.", CVPR 21

Multimodal FSL with Frozen Language Models

- Huge language models (LMs) like GPT-3 show zero/few-shot learning capabilities although only trained to predict next word
- Transfer this to multimodal setting
- Freeze LM, train image embedding on Image Captioning task
- Impressive results:




Paper:

- Tsimpoukelli et al. "Multimodal Few-Shot Learning with Frozen Language Models.", arXiv, Juli 2021

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)





Lifelong Learning for Image Generation with GAN

- **Task:** Lifelong learning for image-conditioned image generation
- **Method:**
 - Factorize convolutional filters into the dynamic base filters
 - Generate these dynamic base filters w/ task specific filter generators
 - Deterministic weight matrix => shared across all tasks







Task 1



Task n-1



(a) Sequential Finetuning

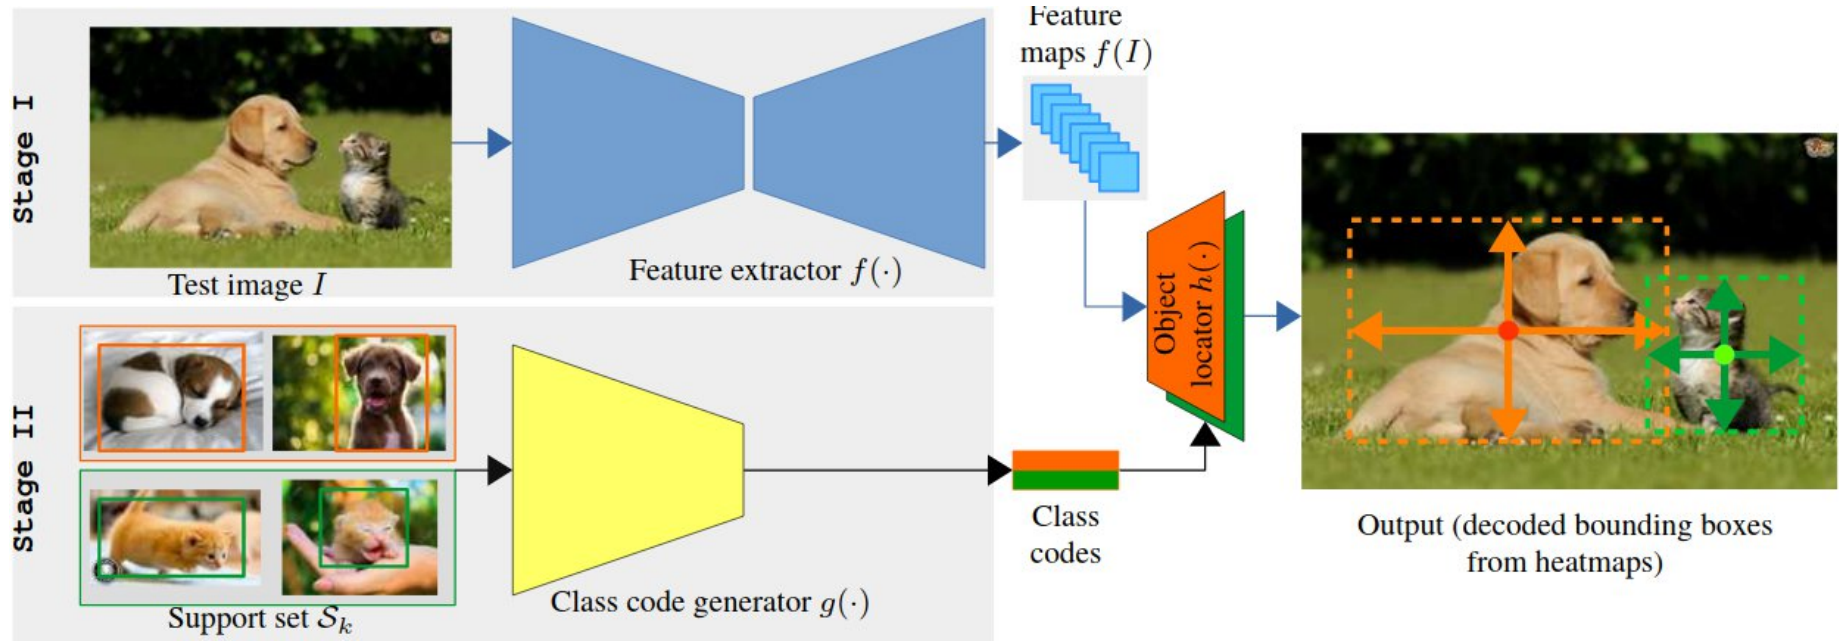


(b) Hyper-Lifelong GAN

Papers:

- Zhai, Mengyao, Lei Chen, and Greg Mori. "Hyper-LifelongGAN: Scalable Lifelong Learning for Image Conditioned Generation", CVPR 2021
 (<https://www2.cs.sfu.ca/~mori/research/papers/zhai-cvpr21.pdf>)

Incremental Few-Shot Object Detection

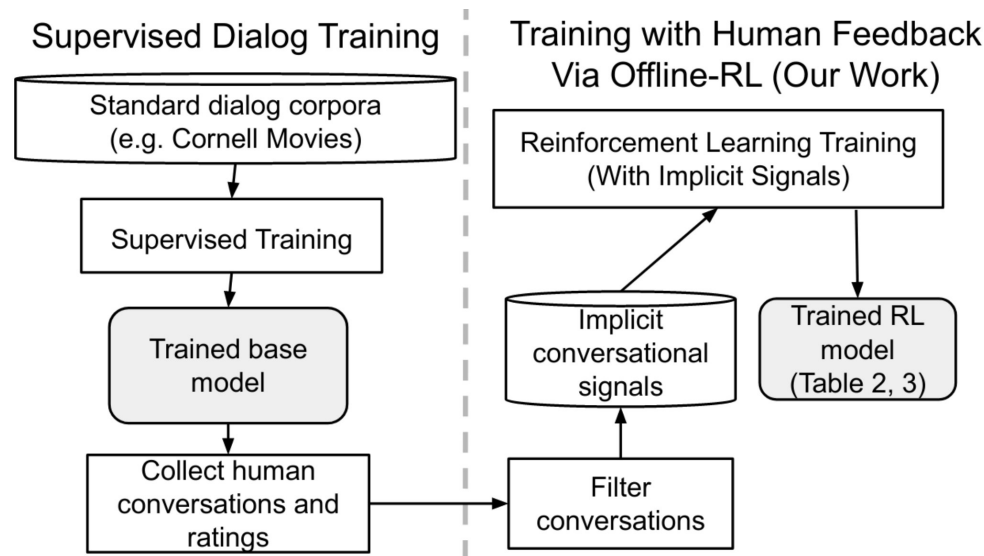


- New will be learned incrementally without revisiting base classes
- Learn to detect new classes using only few images

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

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

Pre-trained model and data poisoning attacks

- pre-trained models are pre-trained on huge datasets (in most cases unsupervised)
- possible to fine-tune to a specific task
- in most NLP tasks: state-of-the-art performance
- however, pre-trained models can have backdoors or are biased

- data can be poisoned so that a model trained with this data can have backdoors or are biased

Papers:

- Weight Poisoning Attacks on Pre-trained Models, Kurita et al., 2020
- You Autocomplete Me: Poisoning Vulnerabilities in Neural Code Completion, Schuster et al, 2020

Not All Memories are Created Equal: Learning to Forget by Expiring

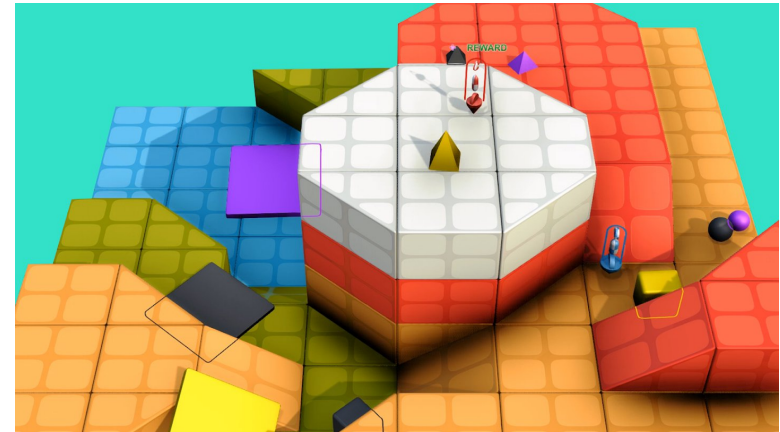
- Equip Transformer models with a memory
- Memories allow to track the state without retraining and to handle longer-term dependencies
- *Expire-Span*¹ aims to retain most important information and expire irrelevant information

Papers:

- *¹Not All Memories are Created Equal: Learning to Forget by Expiring*, Sukhbaatar et al. (2021)
- *Compressive transformers for longrange sequence modelling*, Rae et al. (2020)
- *Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context*, Dai et al. (2019)

Open-Ended Learning Leads to Generally Capable Agents

- RL agents are usually only capable of mastering a single, individual task they were trained on
- Goal: Generalize to a massive, rich space of challenges
- This paper introduces an open-ended learning process, within an iterative process („generations“)
- Impressive results: zero-shot generalisation include good performance on Hide and Seek, Capture the Flag, and Tag



Paper:

Open-Ended Learning Leads to Generally Capable Agents, Stooke et al. (2021)