Karlsruhe Institute of Technology

Advanced Topics in Continual / Organic Machine Learning

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

Summer 21 Topics

"I‘m still learning"
Michelangelo

Class Incremental Semantic Segmentation

- **Task:** Semantic segmentation
- **Method:** 2 modules
  - SegInversion: Use old segmentation model to generate fake images of old classes
  - Half-real half-fake distillation: Use fake data of old classes and real data of new classes to update model

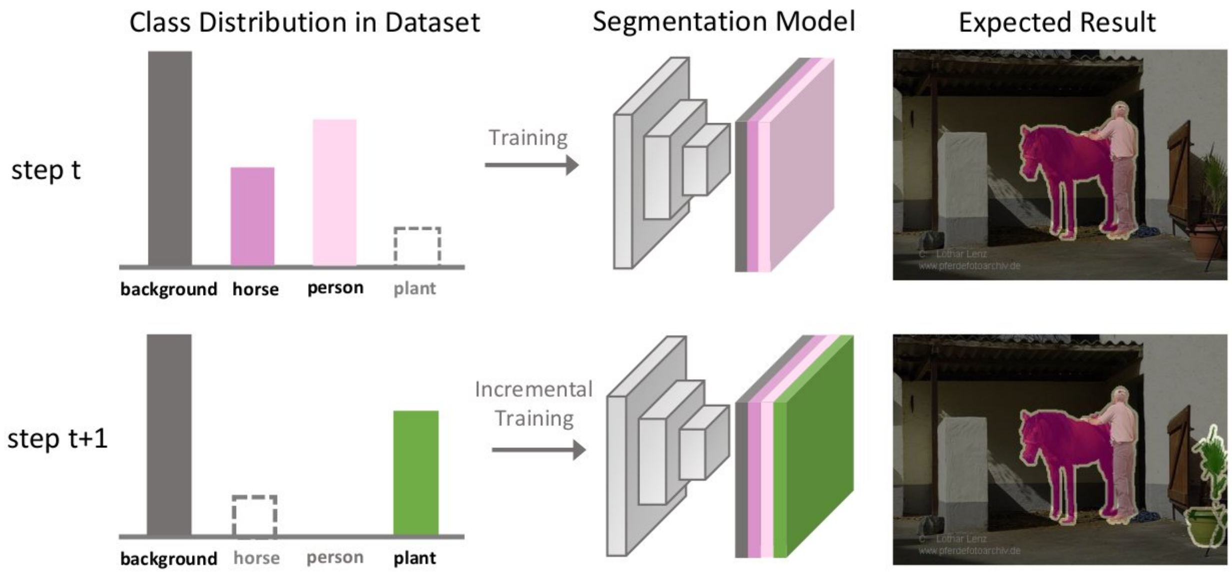

Papers:

- Huang, Z., et al., "Half-real half-fake distillation for class-incremental semantic segmentation, arXiv:2104.00875, 2021. *(https://arxiv.org/abs/2104.00875)*

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

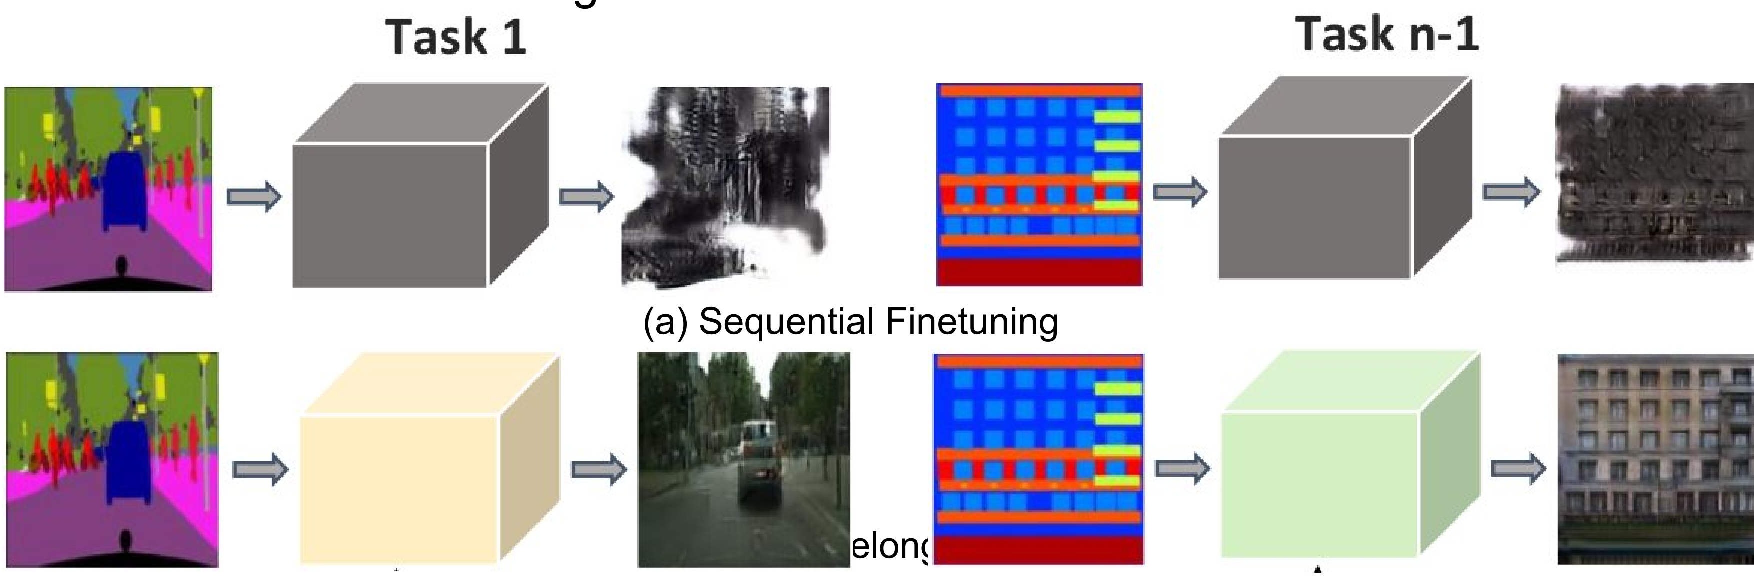

Papers:

- 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)*

Class Incremental Learning

- Avoid from biased model towards old classes and new classes
- Solve the data imbalance problem
- Solve the increasing number of visually similar classes

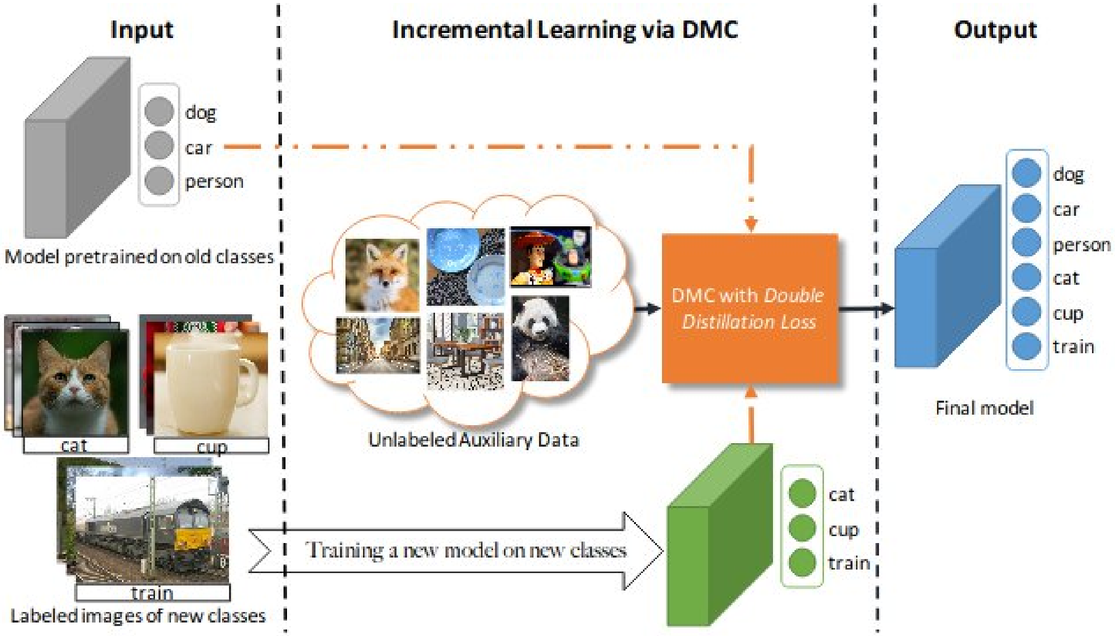

Paper:

- Class-incremental Learning via Deep Model Consolidation, Zhang et al., WACV 2020

Continual Representation Learning (CRL)

Task: Continual Face Recognition & Person Re-id with CRL

- New Continual Learning setting
- Testing with unseen identities
- Neighbourhood selection (NS) and consistency relaxation (CR) for Knowledge Distillation

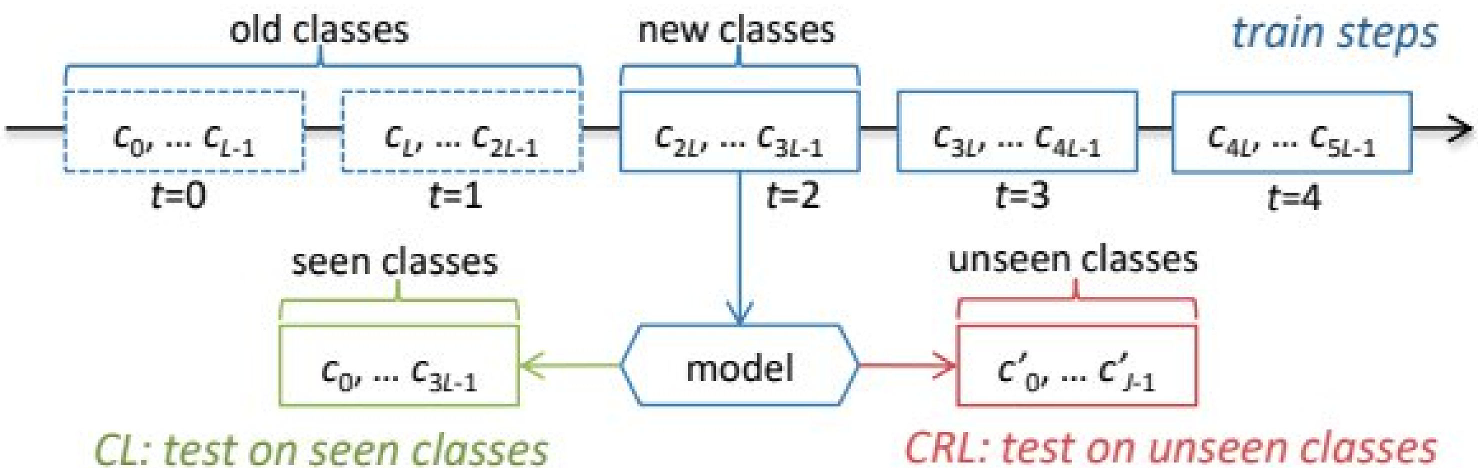

Paper:

- Continual Representation Learning for Biometric Identification, Zhao, Bo, et al., WACV 2021

Personalized Query Rewriting

- Natural language understanding
  - Used in e.g. Amazon Alexa, Google Home and Siri
- Automatic speech recognition (ASR)
  - Error correction
  - Typically independently
- Techniques involved:
  - Memory
  - Pointer-generator network

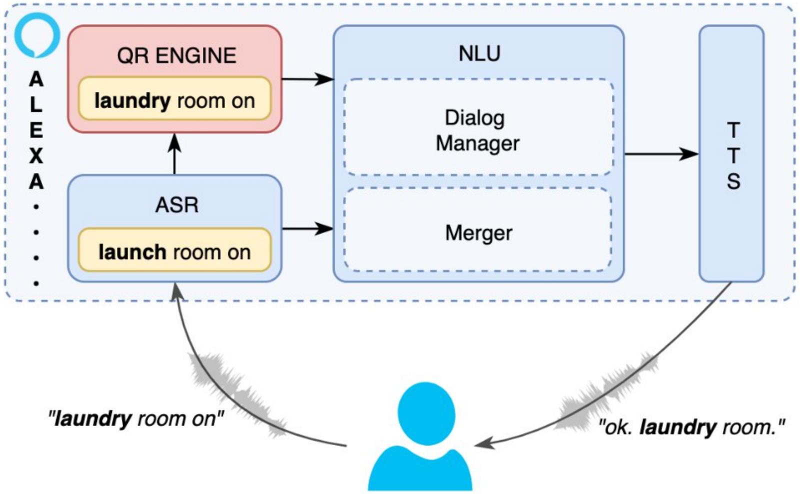

„Personalized Query Rewriting in Conversational AI Agents“, *Roshan-Ghias et al., 2020*

„Personalized Query Rewriting in Conversational AI Agents“, *Roshan-Ghias et al., 2020*

Lifelong Learning E2E Speech Recognition

- Automatic speech recognition (ASR)
  - Training & Deployment
  - No adaptation
  - End-to-end approach (E2E)

- Lifelong learning
  - New training data
  - Gradual improvements
  - Catastrophic forgetting?

- Techniques involved:
  - Elastic Weight Consolidation
  - Knowledge Distillation
  - Episodic memory

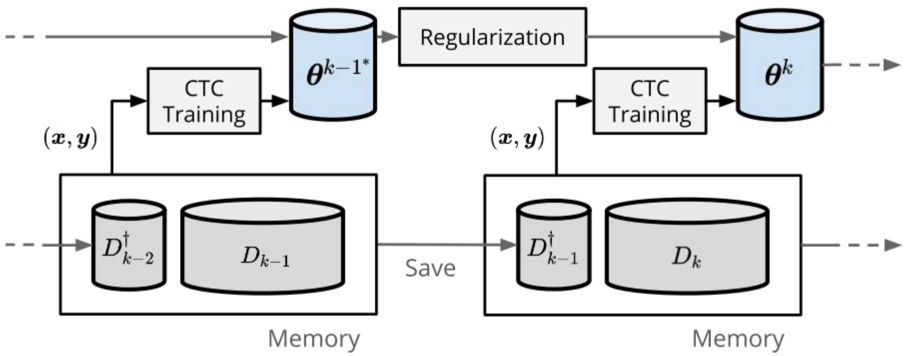

„Towards Lifelong Learning of End-to-end ASR“, *Chang, et al., 2021*

„Towards Lifelong Learning of End-to-end ASR“, *Chang, et al., 2021*

Self-Supervised latent speech representations

- Current ASR systems'
  - Labelled training data
  - Hand crafted features (e.g. Mel-Spectrograms)
- Problem: 7000 spoken languages
- Self-Supervised learning
  - Input: raw audio data
  - Contrastive task training
  - Extract latent speech representation

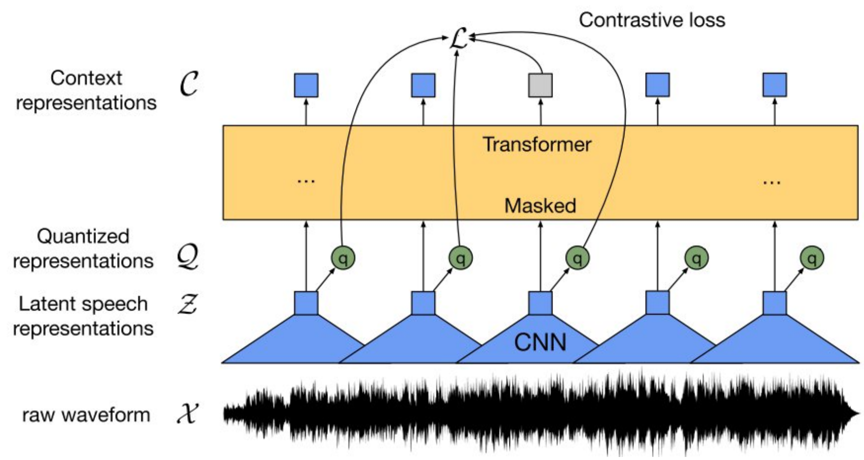

„wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations", *Baevski, et al., 2020*

„wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations", *Baevski, et al., 2020*

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

- Deep Clustering for Unsupervised Learning of Visual Features (Caron et al. 2018)

Perceiver: General Perception with Iterative Attention

- Model build upon a Transformer model (Attention is all you need, Vaswani et. al 2017)
- High dimensional input from multiple modalities

- Perceiver: General Perception with Iterative Attention, Jaegle et al. 2021

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

Improving robustness

- Models are often over-sensitive to small variations (data in production differs from training and validation data)
- Humans are good in handling small variations

Papers:
- Improving Robustness of Task Oriented DialogSystems, Einolghozati al., 2019
- Pretrained Transformers Improve Out-of-Distribution Robustness, Hendrycks et al, 2020

Beyond Maximum Likelihood in NLP

- Language models tend to generate highly repetitive, dull, and incoherent responses
- Currently: Lots of tricks to generate more human-like responses
- Is the likelihood objective itself at fault? Are there better objectives?
  - E.g. Unlikehood Training, Maximum Mutual Information

Input: What are you doing?			
−0.86	I don't know.	−1.09	Get out of here.
−1.03	I don't know!	−1.09	I'm going home.
−1.06	Nothing.	−1.09	Oh my god!
−1.09	Get out of the way.	−1.10	I'm talking to you.

Papers:

- A Diversity-Promoting Objective Function for Neural Conversation Models, Li et al, 2016
- Neural Text Generation with Unlikelihood Training, Welleck et al, 2019
- The Curious Case of Neural Text Degeneration, Holtzman et al, 2019
- Don't Say That! Making Inconsistent Dialogue Unlikely with Unlikelihood Training, Li et al, 2020

Pursuing General-Purpose Algorithms with Reinforcement Learning

- Models are usually trained to master a very specific, narrow task
- RL research, however, seeks to introduce general-purpose algorithms that can be applied in many domains
- E.g., MuZero* combines a tree-based search with a learned model to master Go, Chess, Shogi and Atari games
  - Without knowing the dynamics of the environment

Papers:

- *Mastering Atari, Go, chess and shogi by planning with a learned model, Schrittweiser et al, 2020
- A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play, Silver et al, 2018
- Mastering the game of Go with deep neural networks and tree search, Silver et al, 2016