Karlsruhe Institute of Technology

Advanced Topics in Continual / Organic Machine Learning

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

Course Organizer: F. Irem Eyiokur

"I'm still learning"
Michelangelo

OML

Online Contextualized Few-shot Learning

- Few-shot incremental learning in online/continual setting.
- Classes are organized based in context
- To encode temporal contextual information → RNN
- Prototypical memory to remember prev. learned classes

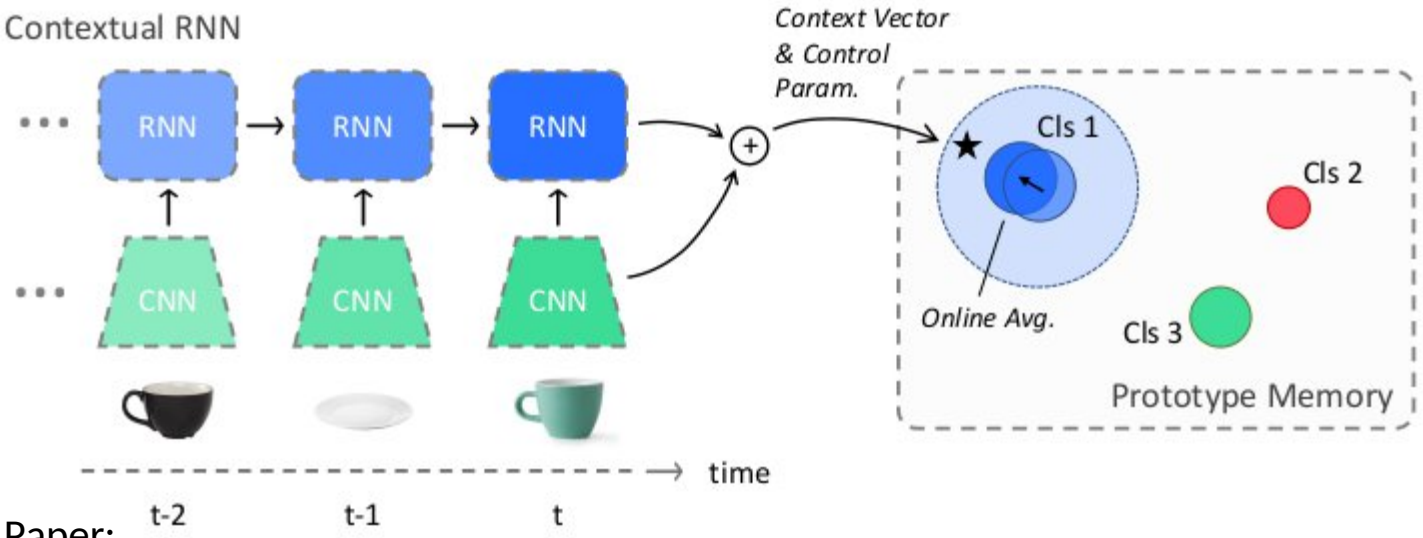

Paper:

Ren, Mengye, et al. “Wandering within a world: Online contextualized few-shot learning”, ICLR 2021

Extension of above paper to unsupervised online learning:

Ren, Mengye, et al. “Online unsupervised learning of visual representations and categories”, arxiv preprint 2022

Sleep-like unsupervised replay reduces catastrophic forgetting in ANNs

Sleep has been hypothesized to play an important role in memory and learning by enabling spontaneous reactivation of previously learned memory patterns.

This hypothesis is implemented as a sleep-like phase called Sleep Replay Consolidation (SRC) in artificial neural networks to protect old memories during new training and alleviate catastrophic forgetting

Paper: Tadros, Timothy et al. “Sleep-like unsupervised replay reduces catastrophic forgetting in artificial neural networks”, Nature Communications, 2022

Online Continual Learning Via Candidates Voting

- Online continual learning to learn sequence of new tasks from data stream
- Selecting candidates from each learned task and then incorporating the corresponding weights instead of single-head classifier
- No need for original data

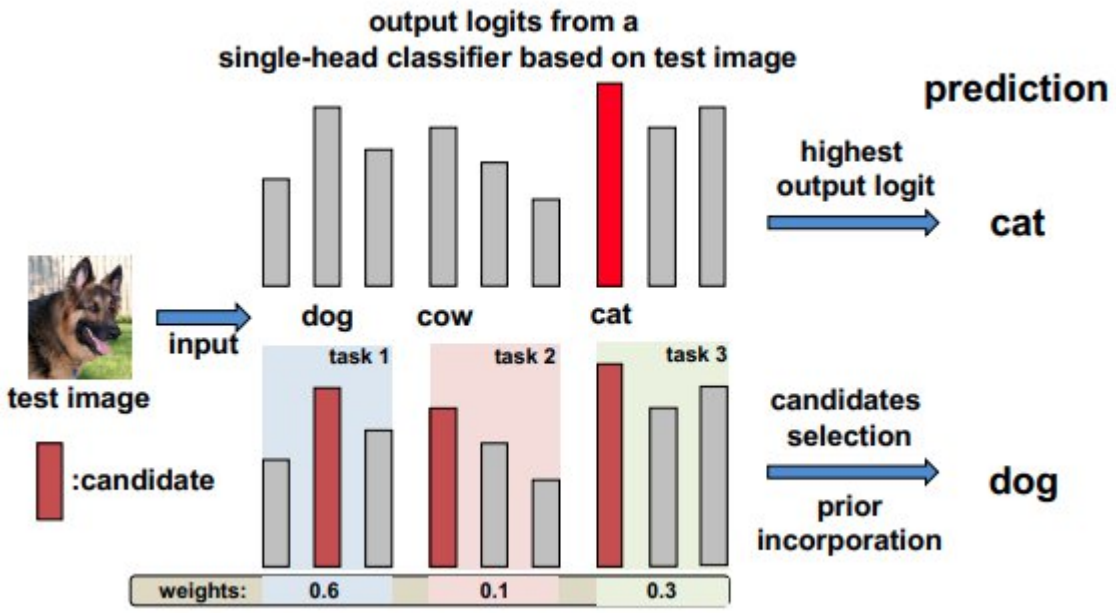

Paper:
Jiangpeng, He, et al. “Online Continual Learning Via Candidates Voting”, WACV 2022

Techniques for In-Context Learning in LLMs

- Prompting / In-Context Learning drives impressive zero-/few-shot performance of LLMs
- Summarize and discuss recent techniques:
  - Chain of Thought Prompting [1]
  - Least-to-Most Prompting [2]
  - Recitation-Augmented Generation [3]
  - Self-Consistency Decoding [4]
- Discuss some theoretical aspects of In-Context Learning [5, 6]

Papers:

1. Jason Wei et al., “Chain of Thought Prompting Elicits Reasoning in Large Language Models,” in *Advances in Neural Information Processing Systems*, 2022
2. Denny Zhou et al., “Least-to-Most Prompting Enables Complex Reasoning in Large Language Models,” in *The Eleventh International Conference on Learning Representations*, 2023
3. Zhiqing Sun et al., “Recitation-Augmented Language Models,” in *The Eleventh International Conference on Learning Representations*, 2023
4. Xuezhi Wang et al., “Self-Consistency Improves Chain of Thought Reasoning in Language Models,” in *The Eleventh International Conference on Learning Representations*, 2023
5. Damai Dai et al., “Why Can GPT Learn In-Context? Language Models Implicitly Perform Gradient Descent as Meta-Optimizers,” in *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2023
6. Sang Michael Xie et al., “An Explanation of In-Context Learning as Implicit Bayesian Inference,” in *International Conference on Learning Representations*, 2022

LLMs in Robotics

- LLMs can be applied to high-level robot control few-shot
  - Open- / Closed-loop planning
  - Natural-Language or Code-Based
- Alternative approach: Train unified model on NL + robotic tasks

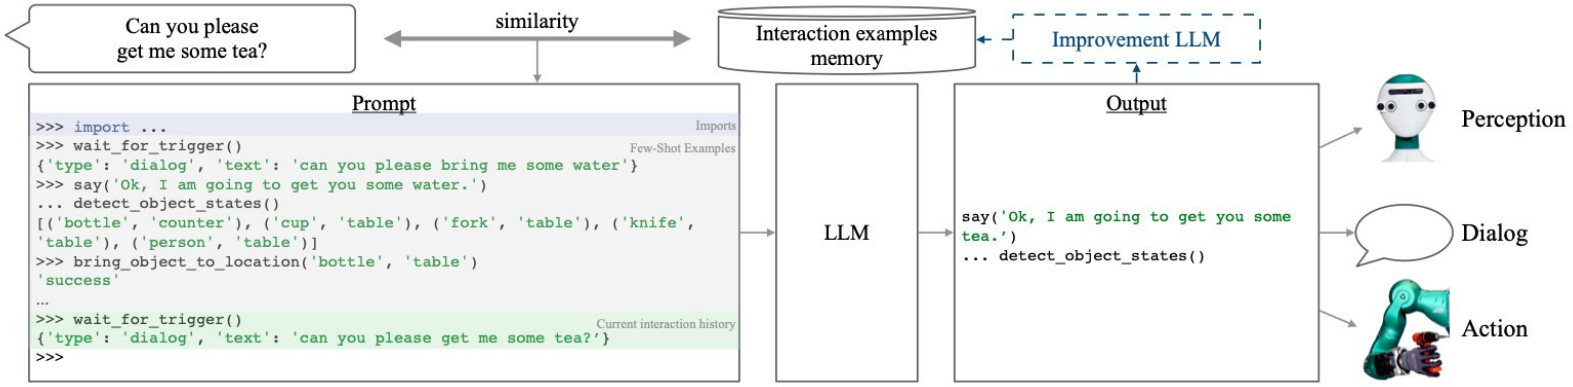

Papers:

- Liang et al., “Code as Policies: Language Model Programs for Embodied Control,” in *Workshop on Language and Robotics at CoRL 2022*
- Allen Z. Ren et al., “Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners,” in *7th Annual Conference on Robot Learning*, 2023
- Michael Ahn et al., “Do As I Can, Not As I Say: Grounding Language in Robotic Affordances,” in *6th Annual Conference on Robot Learning*, 2022
- Wenlong Huang et al., “Inner Monologue: Embodied Reasoning through Planning with Language Models,” in *6th Annual Conference on Robot Learning*, 2022
- Bärmann, Leonard, et al. "Incremental Learning of Humanoid Robot Behavior from Natural Interaction and Large Language Models." *arXiv* (2023)
- Danny Driess et al., “PaLM-E: An Embodied Multimodal Language Model” (arXiv, March 6, 2023)
- Anthony Brohan et al., “RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control,” 2023

Self-Supervised Models are Continual Learners

- **Task:** Self-supervised continual visual representation learning
- **Method:**
  - Benefit from self-supervised loss function
  - Convert it into distillation mechanism for continual learning by using a predictor network
  - Predictor network: Maps the current state of the representations to their past state.

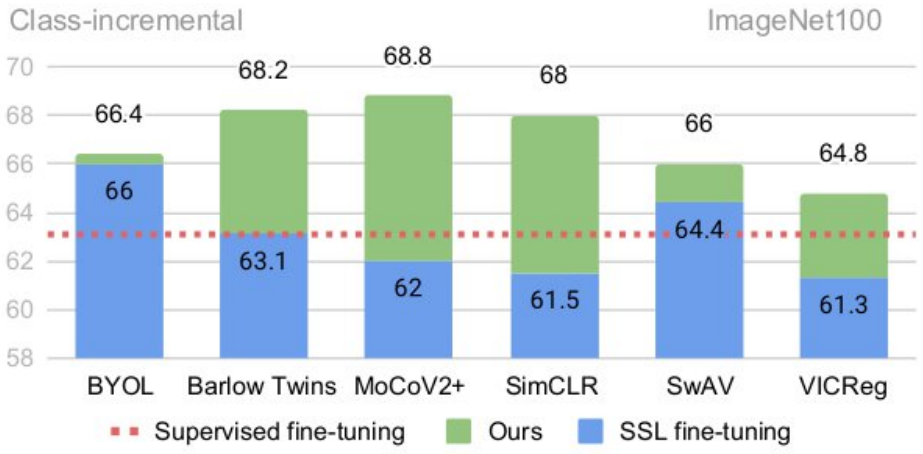

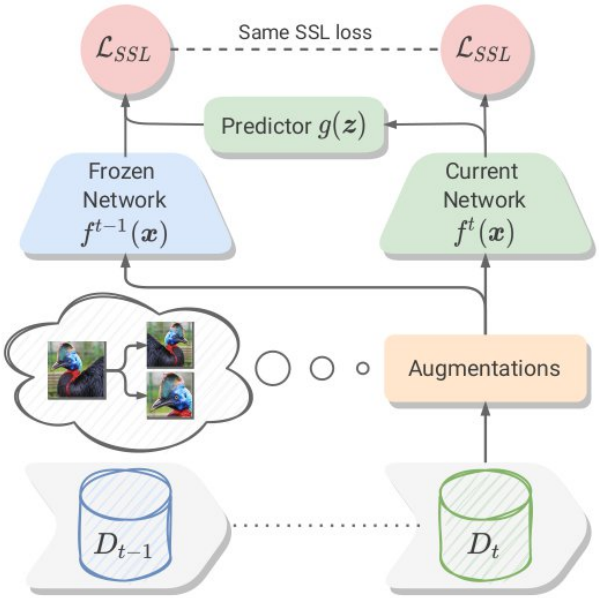

Paper:

- Fini, Enrico, et al. “Self-Supervised Models are Continual Learners” CVPR 2022

Continual Learning with Lifelong Vision Transformer

- Training with sequential data => Existing methods are based on CNN
- Novel attention-based framework Lifelong Vision Transformer (LVT)
- Inter-task attention mechanism: Implicitly absorbs information from the previous tasks
- Slows down the drift of important attention between previous and current tasks.

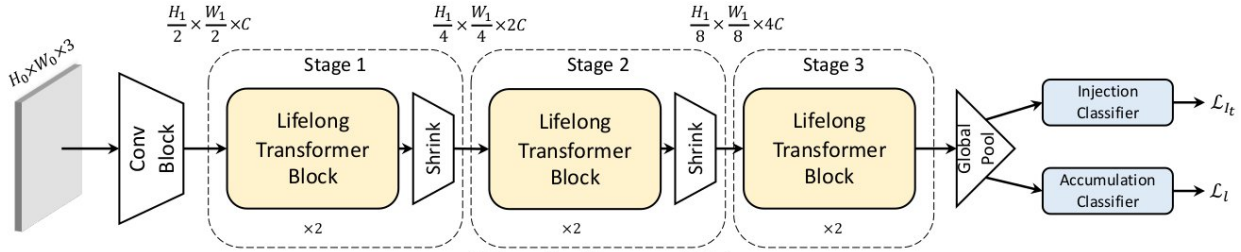

(a) Overall architecture of LVT.

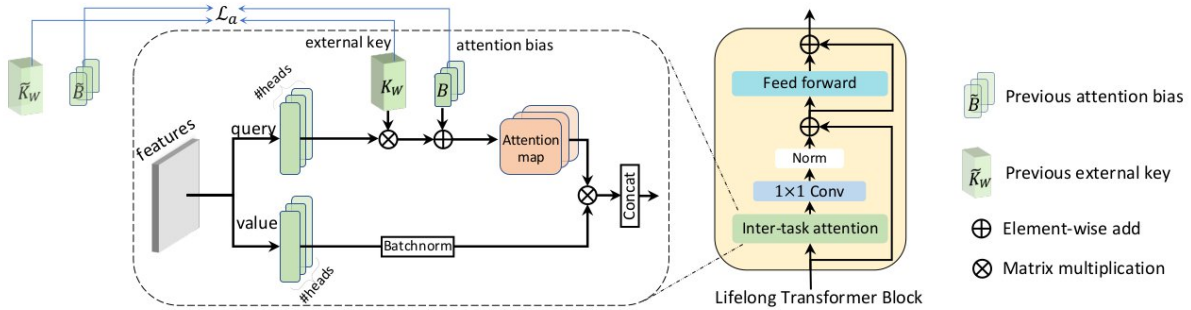

(b) Lifelong transformer block with inter-task attention mechanism.

Paper: Wang, Zhen et al. “Continual Learning with Lifelong Vision Transformer”, CVPR 2022

Online Continual Learning of End-to-End Speech Recognition Models

- Task: Explain Online Gradient Episodic Memory (GEM) method and the application to speech recognition

- Yang, Muqiao, Ian Lane, and Shinji Watanabe. "Online Continual Learning of End-to-End Speech Recognition Models." arXiv preprint arXiv:2207.05071 (2022).

Automatic Prompt Design

- **Review methods for automatic prompt design for pretrained language models (PLM)**
- Prompt design has big impact on the performance of the language model
- Idea: Automatically optimize or search for good prompt designs
- Recent approaches include AutoPrompt[1] (2020), Prefix-Tuning[2] (2021), P-tuning[3] (2021), Prompt Tuning[4] (2022), Automatic Prompt Engineer[5] (2022), Auto-CoT[6] (2023)

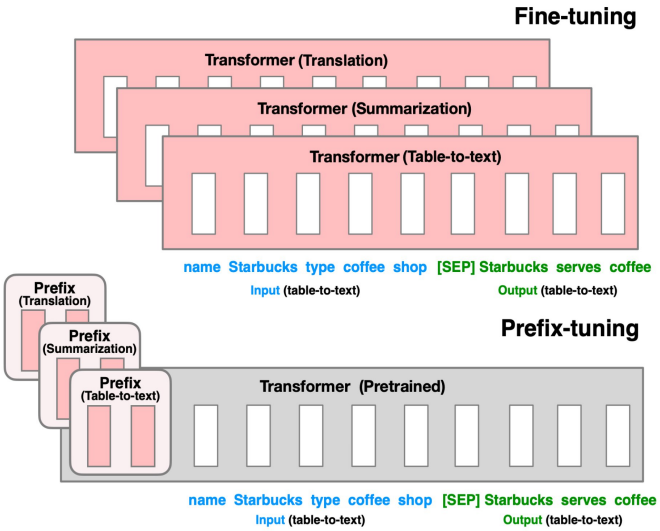

Illustration of Prefix-Tuning from [2]

Papers (Suggestions):

- [1] *AutoPrompt: Eliciting Knowledge from Language Models with ..., Shin et al. (2020)*
- [2] *GPT Understands, Too, Liu et al. (2021)*
- [3] *Prefix-Tuning: Optimizing Continuous Prompts for Generation, Li et al. (2021)*
- [4] *The Power of Scale for Parameter-Efficient Prompt Tuning, Lester et al. (2022)*
- [5] *Large Language Models Are Human-Level Prompt Engineers, Zhou et al. (2022)*
- [6] *Automatic Chain of Thought Prompting in Large Language Models, Zhang et al. (2023)*
- [7] *Automatic Prompt Augmentation and Selection with Chain-of-Thought ..., Shum et al. (2023)*

Variable Few Shot Class Incremental and Open World Learning

- Previous work: Fixed-Few-Shot Class Incremental Learning (number of ways and shots are known, but this is not true in real-world cases)
- Number of ways: how many new classes are added
- Number of shots: how many examples per class are provided
- Variable-Few-Shot Class Incremental Learning (Up-to N-Ways, Up-to K-Shots class incremental learning)
- Detect unknown classes
- Main paper assumes labels for unknown classes, possible solution in real-world: ask user for labels (organic learning like a child)

Paper: Ahmad, Dhamija, Jafarzadeh, Cruz, Rabinowitz, Li, Boult., “Variable Few Shot Class Incremental and Open World Learning”, IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2022

PPT: Pre-trained Prompt Tuning for Few-shot Learning

- pre-trained large language models (PLLMs) have large general domain knowledge
- fine-tuning often too time consuming, hard prompt tuning is a good alternative:
  sentence: I like eating apples
  sentiment: positive
  sentence: The film was horrible
  sentiment:
- alternative: soft prompts that are automatically generated (with backpropagation algorithm)
- in some few-shot learning scenarios prompt tuning is worse than fine-tuning a smaller model
- one solution: adding soft prompts into the pre-training stage to improve performance for these scenarios

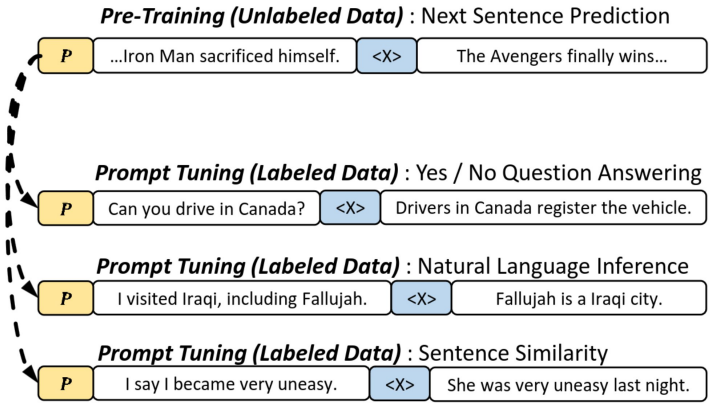

Paper: Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang, “PPT: Pre-trained Prompt Tuning for Few-shot Learning”, Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL), 2022

Continual Learning with Model Merging

- Training Pipeline nowadays:
  1. Pre-train on pre-training objective (e.g. language modelling)
  2. Fine-Tune on *Downstream Task*
- We might continue pre-training the base model
- How should we transfer knowledge and capabilities across trained models?
  - e.g. how to integrate new base model knowledge into the fine-tuned one without redoing the fine-tune?
  - or other scenarios (see figure)
- Matena et al. (2022) present one approach in "*Merging Models with Fisher-Weighted Averaging*"
  - average models in a "knowledge preserving way"
- Your task: explain the scenarios, previous solutions, and the presented approach in this paper