Karlsruhe Institute of Technology

Advanced Topics in Continual / Organic Machine Learning

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

I‘m still learning
Michelangelo

Self-Supervised Learning

- A machine learning system that uses supervised learning techniques (e.g. NNs) to learn from automatically labeled data.

- Why self-supervised learning:
  - Does not need much (human-)labeled data
  - Allows online/continual learning
  - Inspired from how babies learn
    - *The Scientist in the Crib: What Early Learning Tells Us About the Mind*
      Alison Gopnik, Andrew N. Meltzoff and Patricia K. Kuhl
    - *The Development of Embodied Cognition: Six Lessons from Babies*
      Linda Smith and Michael Gasser

- Methods of self-supervised learning

- Applications of self-supervised learning

Zero-shot Learning and Transfer Learning

- How to leverage well-trained machine learning models from rich-data domains/tasks in order to perform in low-resource scenarios, even in extremely zero-shot?

- Why Transfer Learning:
  - Applied to the situations that cannot/hard to get data
  - Finding common relations between tasks

- Investigation of ZS and transfer learning methods used **in NLP**
- Applications of ZS and transfer learning **in NLP**

Personalisation

- Every person speaks or writes their own flavor of their native language, influenced by a number of factors: the content they tend to talk about, their gender, their social status, or their geographical origin.
- When attempting to perform Machine Translation (MT) or Language Generation models in general, these variations have a significant effect on how the system should perform translation, but this is not captured well by standard one-size-fits-all models.
- This is a challenging problem, because modeling each speaker with a personal model leads to a vast number of models and serious data sparsity.
- Factorizing the models is one of the approaches, such as in
  - Extreme Adaptation for Personalized Neural Machine Translation
- An application with concern about data privacy has been mentioned in:
  - Personalized Natural Language Understanding (from Microsoft Corporation)

Survey: techniques in Reinforcement Learning for Dialog

- Sequence-to-Sequence model allows us to build a chat-bot / dialog system
- Generating the response based on the previous sentence

The main problem is that a naive S2S model is not likely to keep track of the flow and suffers from Repetitive Response and Short-sighted conversation decisions.

Reinforcement Learning is an approach that can help these models generalize over rewards that are designed towards longer and more plausible conversations.

These techniques typically involve using an additional discriminator / reward function estimator and combine GAN with Reinforcement Learning to help the model learn better.

Some important recent papers:
Adversarial Learning for Neural Dialogue Generation.
Jiwei Li, Will Monroe, Tianlin Shi, Alan Ritter and Dan Jurafsky. EMNLP 2017.
Learning through Dialogue Interactions by Asking Questions.
Jiwei Li, Alexander Miller, Sumit Chopra, Marc'Aurelio Ranzato, Jason Weston . ICLR 2017
Dialogue Learning With Human-in-the-Loop.
Jiwei Li, Alexander Miller, Sumit Chopra, Marc'Aurelio Ranzato, Jason Weston. ICLR 2017

Integrating new knowledge

Karlsruhe Institute of Technology

- Problems of training the Neural Networks by iterating through all the training data:
    - all training data should be prepared before the training
    - training demands much computation and memory resources
    - training with new data set leads to catastrophic forgetting

- Solution: Continual learning, incremental learning oder long life learning
    - the crucial ability for humans to continually acquire and transfer new knowledge across their lifespans while retaining previously learnt experiences (D. Hassabis, D. Kumaran, C. Summerfield, and M. Botvinick. Neuroscience-inspired artificial intelligence.)
    - some approaches:
        - Deep Generative Replay (DGR) replaces the storage of the previous training data with a Generative Adversarial Network to synthesize training data on all previously learnt tasks
        - store the weights of the model trained on previous tasks, and impose constraints of weight updates on new tasks
        - estimate the importance of the parameters on previous tasks and penalize future changes to the weights on new tasks
        - knowledge distillation
        - use a neural network to learn class-representative prototypes in an embedding space and classify embedded test data by finding their nearest class prototype

Select knowledge to forget

- It is shown that most of the knowledge acquired is of negative value even though it is correct and was acquired solving similar problems. (Markovitch, Shaul, and Paul D. Scott. "The role of forgetting in learning.)
- can be used for different scenarios:
  - forget some wrong trained samples
  - use prior knowledge to train with very few training samples:
    - Learning to Forget for Meta-Learning, Sungyong Baik, Seokil Hong, Kyoung Mu Lee
  - use for generalization
  - Frequently recalled memories are remembered, whilst memories recalled rarely are forgotten (Simpson, Andrew JR. "Use it or Lose it: Selective Memory and Forgetting in a Perpetual Learning Machine.)

Hyperarticulation Speech Dialog

Karlsruhe Institute of Technology

- Problem:
  - Hyperarticulation often occur as a strategy to recover previous recognition errors in spoken dialogue systems.
  - Contrary to this intention a significant performance degradation can be observed at hyperarticulation.

- Papers:
  - Soltau et al. 2000 “Specialized Acoustic Models for Hyperarticulated speech”

Learning new knowledge through generate counterfactual stories

- Counterfactuals can be what-ifs: Would this patient have different blood sugar had they received different medicine?
- Or could be explanations: Why is this bird a cardinal and not a scarlet tanager?
    - Is a kind of what if: If this bird was a scarlet tanager, how would it be different from the picture?
- Frequently applied in medicine and ad placement
- Examples
    - Johannson et al. 2016 “Learning Representations for Counterfactual Inference” - Generic counterfactuals, some focus on medical
    - Hendricks et al. 2018 “Generating Counterfactual Explanations with Natural Language”
    - Qin et al. 2019 “Counterfactual Story Reasoning and Generation”

Defend Adversarial Attacks

- Adversarial attacks try to generate samples that are minimally altered in a way that would be obvious to humans but fools the model
- Generating adversarial examples during training can improve the robustness of models
- Examples:
  - Raghunathan et al. 2018, “Certified Defenses Against Adversarial Examples”
  - Samangouei et al. 2018, “Degense-GAN: Protecting Classifiers Against Adversarial Attacks Using Generative Models”
  - Alzantot et al. 2018, “Generating Natural Language Adversarial Examples”

Self-assessment

- The user or subcomponents should be able to see how certain a model is to its output
- Very important in areas like finding best medical treatment or autonomous driving (where a uncertain (and potential wrong) decision can cost human lives)
- If necessary, a system can ask a question to resolve uncertainty
- Normal confidence values are often imprecise

- Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning (2016, Yarin and Ghahramani)
- Modeling Confidence in Sequence-to-Sequence Models (2019, Niehues and Pham)

Knowledge utilization

- Use external sources like the Internet to give better answers (humans do sometimes the same)
- A Knowledge-Grounded Neural Conversation Model (Ghazvininejad et al., 2017)

Find knowledge in the Internet

- find relevant information in the Internet (Information retrieval)

- Information Retrieval as Statistical Translation (Berger et al., 2017)
- A Study of Smoothing Methods for Language Models Applied to Ad Hoc Information Retrieval (Chengxiang Zhai et al., 20017)

Learning from heterogeneous data

- humans use multiple modalities (vision, sounds, speech, …)
- humans can infer from one modality to another, e.g. if you have never seen/heard the textual description of a formula 1 car, you can describe it after you have seen it

- Modulating early visual processing by language (de Vries et al., 2017)