Course Projects

Jian Tang

tangjianpku@gmail.com

HEC MONTRÉAL

Institut des algorithmes d'apprentissage de Montréal MILA

Overview

- Every team contains at most 4 people and works together on a project
- Goal:
  - Conduct a comprehensive survey on the specific direction
  - Identify a new problem along the direction
- Evaluation
  - We will organize a poster session in the last class
  - The final report is due in two weeks after the course
  - The responsibility of each people in the group should be clearly stated

Research Areas

- **Natural Language Understanding (5)**
  - Sentence representation
  - Reading comprehension
  - Machine translation
  - Summarization
- **Graph & Recommendation (4)**
  - Knowledge graph representation
  - Graph representation
  - Graph generation
  - Recommendation
- **Learning Fundamental (5)**
  - Generative models (VAEs, GANS)
  - Multi-task learning
  - Meta-learning/few-shot learning
  - Transfer learning
  - Adversarial examples

Natural Language Understanding

Machine Reading Comprehension

- Task: given a question and context passage, find the answer from the passage

---

Passage: Tesla later approached Morgan to ask for more funds to build a more powerful transmitter. **When asked where all the money had gone, Tesla responded by saying that he was affected by the Panic of 1901**, which he (Morgan) had caused. Morgan was shocked by the reminder of his part in the stock market crash and by Tesla's breach of contract by asking for more funds. Tesla wrote another plea to Morgan, but it was also fruitless. Morgan still owed Tesla money on the original agreement, and Tesla had been facing foreclosure even before construction of the tower began.

Question: On what did Tesla blame for the loss of the initial money?

Answer: Panic of 1901

- References:
  - Rajpurkar et al. SQuAD: 100,000+ Questions for Machine Comprehension of Text.
  - Wang et al. R-Net: Machine Reading Comprehension with Self-Matching Networks.

Sentence/document representation

- Task: map a sentence/document to a low-dimensional representation
- References:
  - Le et al. Distributed representations of sentences and documents
  - Kim et al. Convolutional neural networks for sentence classification
  - Tang et al. Document modeling with gated recurrent neural networks
  - Lin et al. A structured self-attentive sentence embedding

Neural Machine Translation

- Task: map one sentence in the source language to a sentence in the target language

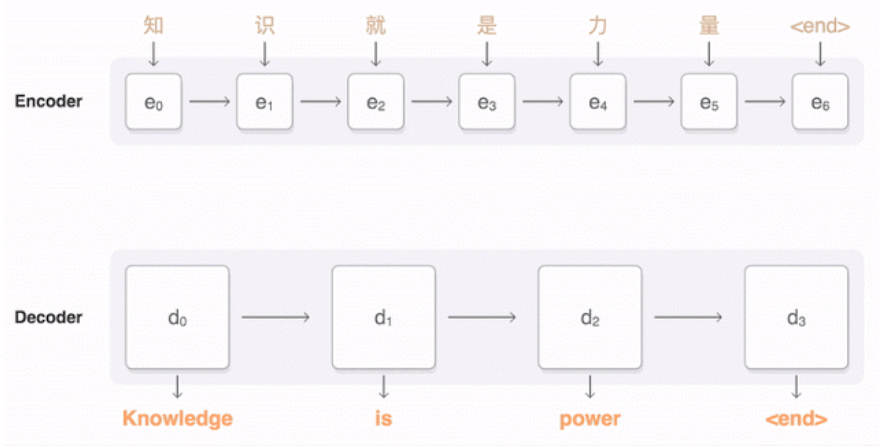

- References:
  - Sutskever et al. Sequence to sequence learning with neural networks.
  - Bahdanau et al. Neural machine translation by joingly learning to align and translate
  - Gehring et al. Convolutional sequence to sequence learning
  - Vaswani et al. Attention is all you need

Neural Text Summarization

Russian Defense Minister Ivanov called **Sunday** *for the creation of a joint front for combating global terrorism.*

Russia calls for joint front against terrorism.

- References
  - Nallapati et al. Abstractive text summarization using sequence-to-sequence RNNs and Beyond.
  - Rush et al. A neural attention model for sentence summarization
  - Paulus et al. A deep reinforced model for abstractive summarization.

Natural Language Inference

- Goal: give a pair of sentences, predict their relationships (entailment, contradiction, neutral)?

A man inspects the uniform of a figure in some East Asian country.	**contradiction** C C C C C	The man is sleeping
An older and younger man smiling.	**neutral** N N E N N	Two men are smiling and laughing at the cats playing on the floor.
A black race car starts up in front of a crowd of people.	**contradiction** C C C C C	A man is driving down a lonely road.
A soccer game with multiple males playing.	**entailment** E E E E E	Some men are playing a sport.
A smiling costumed woman is holding an umbrella.	**neutral** N N E C N	A happy woman in a fairy costume holds an umbrella.

- References:
  - Bowman et al. A large annotated corpus for learning natural language inference.
  - Wang et al. Learning natural language inference with LSTM.
  - Chen et al. Enhanced LSTM for natural language inference.

Graph Representation and Generation

Knowledge Graph Embedding

- Goal: map knowledge graph into low-dimensional spaces

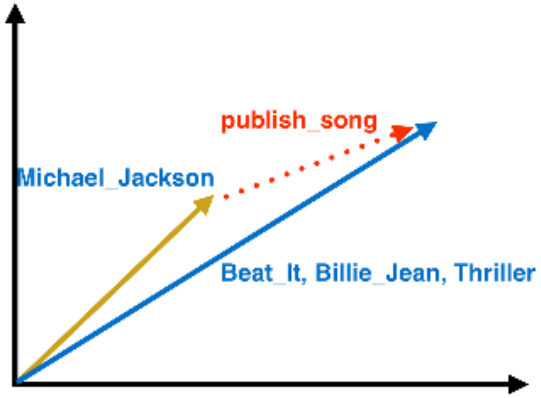

- References:
  - Bordes et al. Translating embeddings for modeling multi-relational data.
  - Lin et al. Learning entity and relation embedding for knowledge graph completion
  - Shen et al. Modeling large-scale structured relationships with shared memory for knowledge base completion
  - https://github.com/thunlp/KB2E

Learning representations of entire graphs

- Goal: map a subgraph/small graph into a low-dimensional vector

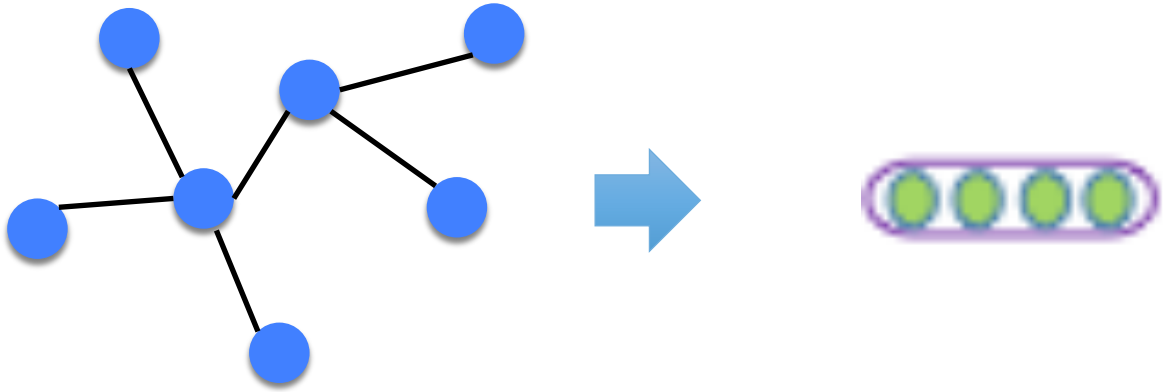

- References:
  - Niepet et al. Learning convolutional neural networks for graphs
  - Dai et al. Discriminative embeddings of latent variables for structured data.
  - Gilmer et al. Neural message passing for quantum chemistry.

Graph Generation

- Generative models such as variational autoencoders (VAEs) and generative adversarial networks (GANs) are widely studied fro generating images, speeches and natural language
- How to generate graphs (e.g., molecules) are under explored

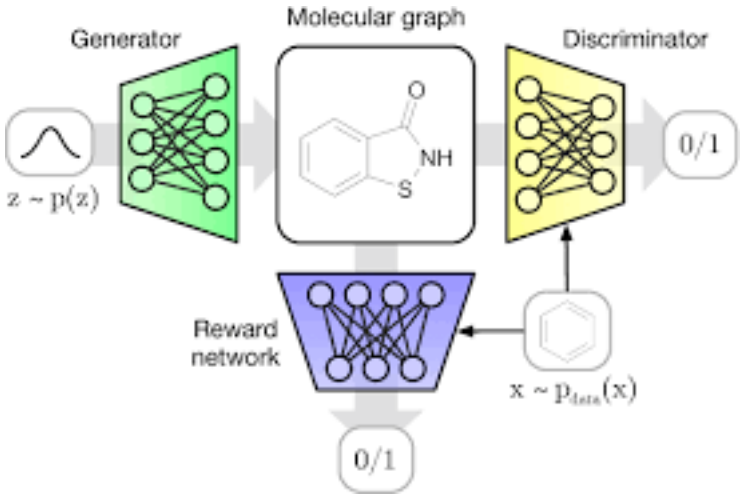

- Nicola De Cao, Thomas Kipf. MolGAN: An implicit generative model for small molecular graphs.
- Jiaxuan You, Bowen Liu, Rex Ying, Vijay Pande, Jure Leskovec. Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation

Deep Learning for Recommendation

- Task: predict the ratings over items given users or suggest relevant items to users

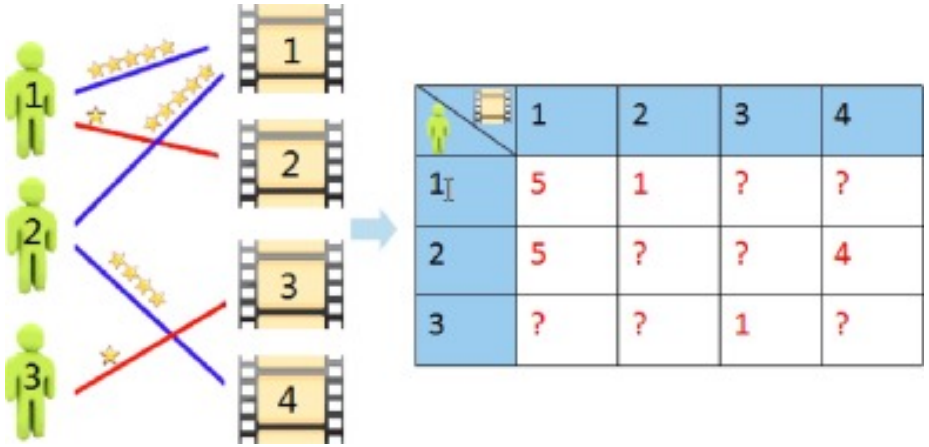

- References:
  - Hidasi et al. Session-based recommendations with Recurrent Neural Networks
  - Covington et al. Deep Neural networks for Youtube recommendations
  - Cheng et al. Wide&Deep learning for recommender systems
  - Salakhutdinov et al. Restricted Boltzmann machines for collaborative filtering.

Learning

Deep Generative Models

- Goal: model the density of the data, or sample from it

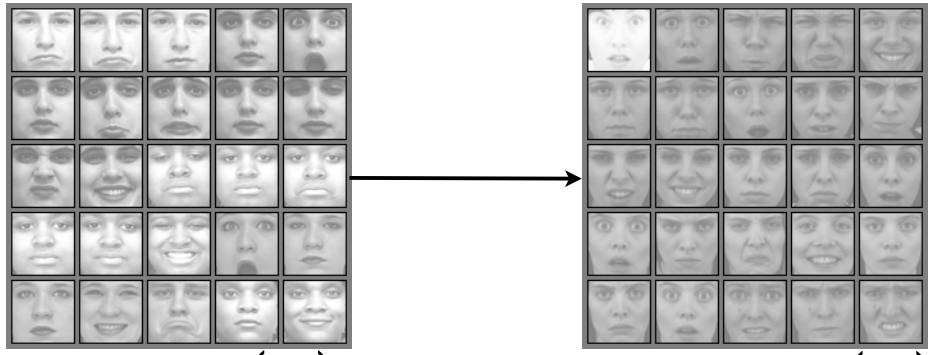

$x \sim p_{data}(x)$ $x \sim p_{model}(x)$

- References
  - Kingma et al. Auto-encoding Variational Bayes
  - Goodfellow et al. Generative adversarial networks.

Multi-task Learning

- Goal: jointly learn multiple tasks

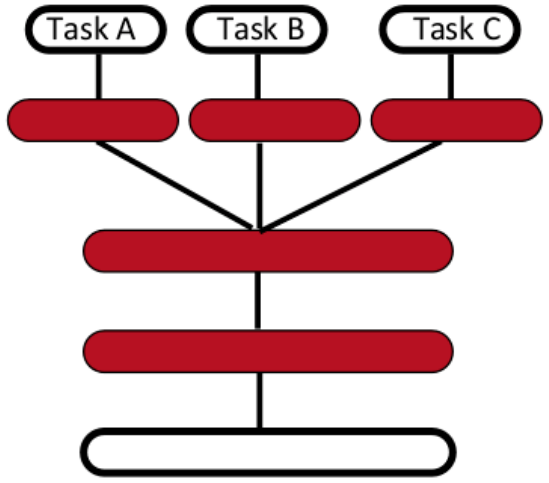

- References:
  - http://sebastianruder.com/multi-task/
  - Liu et al. Deep multi-task learning with shared memory.
  - Kaiser et al. One model to learn them all.

Transfer Learning

- Goal: transfer supervision or knowledge from source tasks to target task

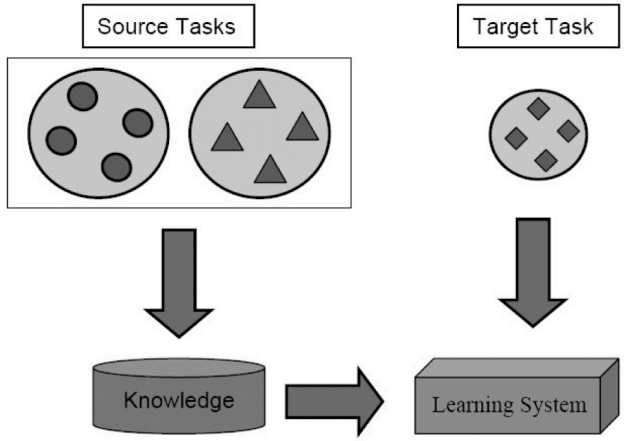

- References:
  - Tzeng et al. Adversarial discriminative domain adaption
  - Csurka et al. Domain adaption for visual applications: a comprehensive survey.

One-shot Learning

- Goal: learning with only a few examples

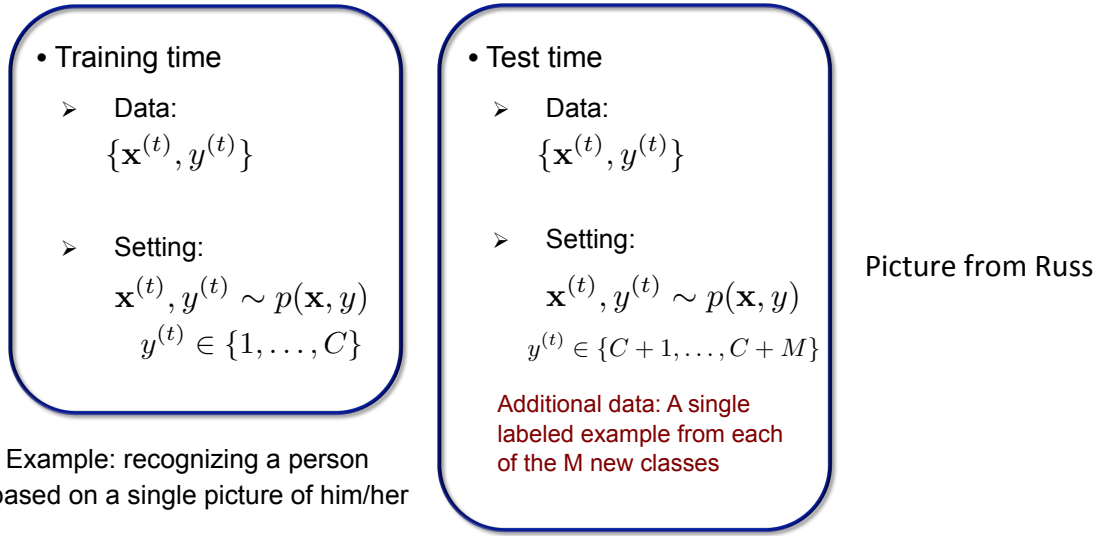

• Example: recognizing a person based on a single picture of him/her

Picture from Russ

- References:
  - Ravi et al. Optimization as a model for few-shot learning. ICLR'2017
  - Santoro et al. One-shot learning with memory-augmented neural networks
  - Vinyals et al. Matching networks for one shot learning.

Adversarial examples

- Many machine learning algorithms are vulnerable to adversarial examples, which misclassify examples that are only slightly different from correctly classified examples.

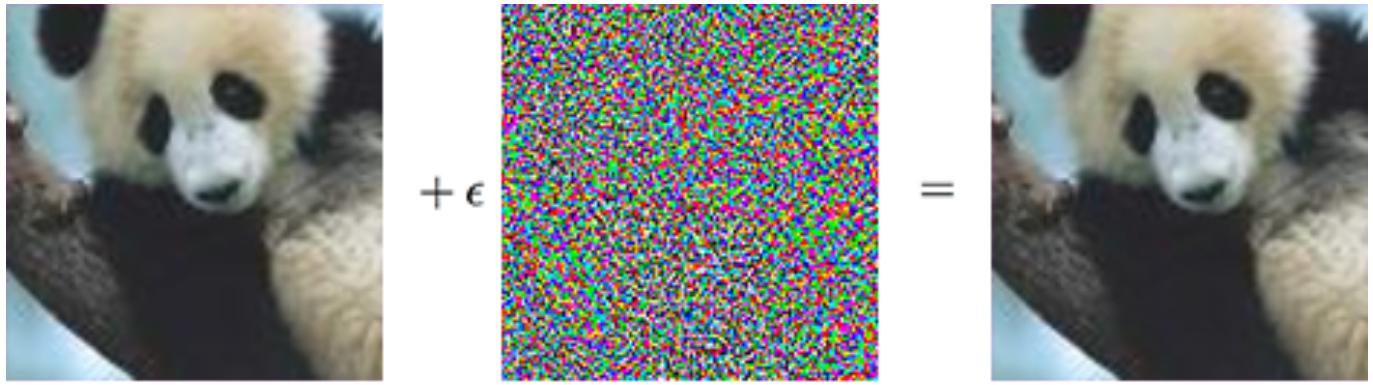

- References
  - Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy. Explaining and harnessing adversarial examples. ICLR 2015
  - Adversarial examples tutorial: https://aaai18adversarial.github.io/

Other Potential Topics

- Deep Learning for computer vision
- Deep Learning for biology
- Deep learning for healthcare
- …

Thanks!