CS-568 Deep Learning

Nazar Khan

PUCIT

Concluding Remarks

Outline

1. What have we covered?
2. What were the general principles?
3. What have we not covered?
4. Research opportunities
5. Acknowledgements

What have we covered?
Theory

What have we covered?
Theory

- 4.3.2 Advanced (Lec. 11)

5. Learning on Images

- 5.1 Convolutional Neural Networks (CNNs) (Lec. 12)
- 5.2 Backpropagation in CNNs (Lec. 13)
- 5.3 Variations and Extensions of CNNs (Lec. 14, 15)

6. Learning on Sequences

- 6.1 Recurrent Neural Networks (RNNs) (Lec. 16)
- 6.2 Backpropagation Through Time (Lec. 17)
- 6.3 Variations and stability of RNNs (Lec. 18)
- 6.4 LSTM and GRU (Lec. 19)
- 6.5 Language Modelling
  - 6.5.1 Text Generation and Language Translation (Lec. 21)
  - 6.5.2 Attention (Lec. 22)
  - 6.5.3 Transformers (Lec. 23, 24)

7. Learning without Calculus

- 7.1 Automatic Differentiation (Lec. 20)

What have we covered?
Theory

8. **Generative Learning**
   - 8.1 Generative Adversarial Networks (Lec. 25)
9. **Learning on Graphs**
   - 9.1 Graph Neural Networks (Lec. 26, 27)

What have we covered?
Recitations

1. Basics of Python (Rec. 1)
2. Using Google Colab (Rec. 1)
3. Basics of NumPy (Rec. 2)
4. Matrix and Vector Calculus (Rec. 3)
5. How to train a Perceptron (Rec. 4)
6. Neural network training tips (Rec. 5)
7. Basics of PyTorch (Rec. 5)
8. Convolution (Rec. 6)
9. CNN in PyTorch (Rec. 7)
10. Transfer Learning (Rec. 8)
11. Sentiment Analysis using an RNN (Rec. 9)
12. Character sequence generation using an RNN (Rec. 10)

What have we covered?
Recitations

13. Language translation using RNN (Rec. 11)
14. Language translation using Attention (Rec. 11)
15. Fake image generation using GAN. (Rec. 12)

What were the general principles?

1. Given enough data, let the machine decide.
2. Let information live.

$$\mathbf{x} + \text{BN}(\text{ReLU}(W\mathbf{x} + \mathbf{b}))$$

3. Use adpative gradient descent.
4. Implement loss function and let automatic differentiation handle the rest.
5. Be cognizant of the domain.

Domain	**Model**
Vectors	Vanilla MLP
Images	CNN
Sequences	RNN, LSTM, Transformer
Graphs	GNN

6. Almost everything boils down to encoding and decoding.
7. DL tries to learn a good representation of knowledge from your data.
8. Become experts of existing DL frameworks.

What have we not covered?

1. Autoencoders
2. Reinforcement Learning
3. Theoretically sound pre-DL concepts.

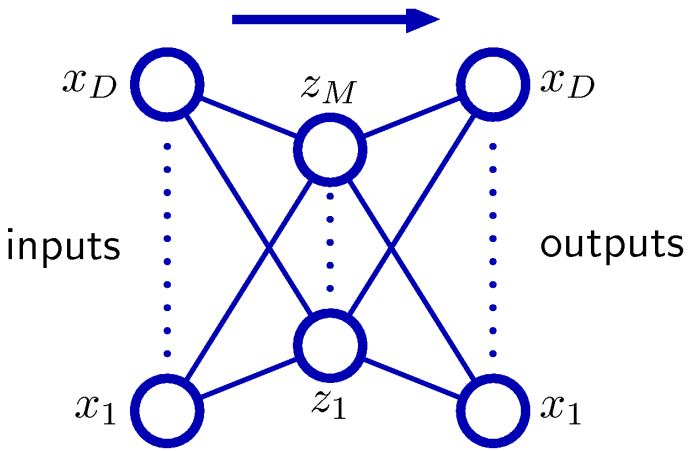

Autoencoder

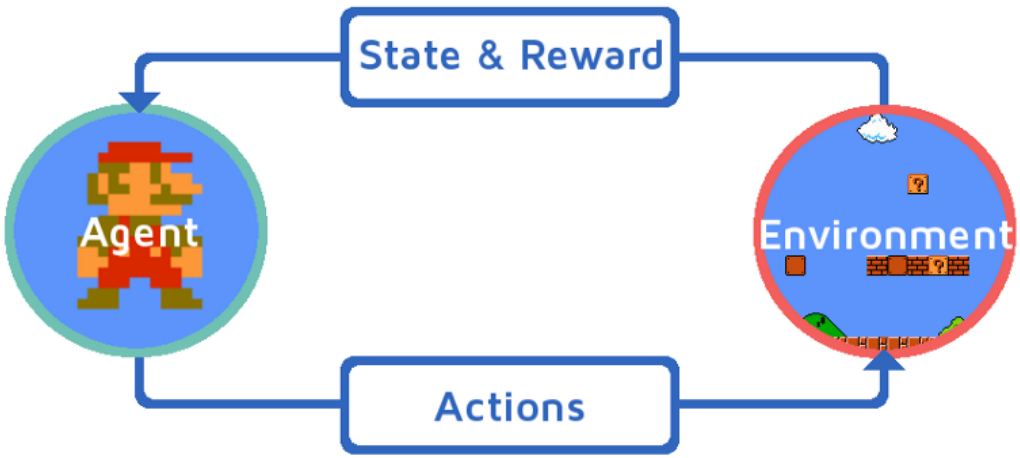

Reinforcement Learning

Research opportunities

- This course has introduced you to a variety of methods and problems in the area of DL.
- This course has *not prepared you fully* for research in DL.
  - That is a whole new ball game.
  - The course project might have nudged the earnest[1] student in the right direction.
  - To prepare yourself for research, the first step is to start devouring research papers.
- In case you are interested, contact us.
- Might not be a bad idea to start your MPhil thesis research right away.

[1] showing sincere and intense conviction

Acknowledgements

- Thank you for your interest in this class.
- The TA: Arbish Akram. If you land a DL job, she deserves your first paycheck![2]
- Members of the CVML group.
- Bhiksha Raj and Christopher Bishop.

Good luck in your future endeavours.
Just remember to **let the machine decide**!

[2] ~~Probably deserves half my pay for this course as well~~.