Introduction to the Course

Deep Learning (ENGG*6600*01)

School of Engineering,
University of Guelph, ON, Canada

Course Instructor: Benyamin Ghojogh
Summer 2023

Introduction of the Instructor and Students

Let us know each other by introducing ourselves!

Feel free to let us know (if you would like):

- Your name
- Your major
- Whether you are studying MEng, MASc, or PhD
- What it your goal for taking this course? What are your expectations from this course?
- How much you know about deep learning?
- How was your vacation between semesters? :)

Introduction of the Course

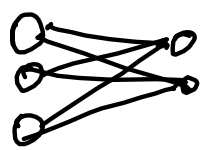

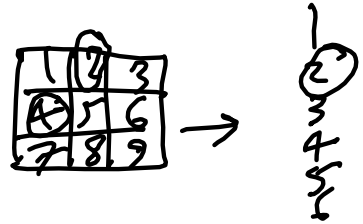

- This course focuses on various topics in deep learning, covering topics from fundamental concepts in neural networks to state-of-the-art deep learning.
- We start with neuron and basic learning methods.
- Then, we cover feed-forward nets and backpropagation.
- Then, we cover regularization techniques such as batch normalization and dropout.
- Then, we cover convolutional neural nets and some well-known networks such as ResNet, AlexNet, etc.
- Then, we cover sequence modeling (useful for NLP and speech processing) including RNN, LSTM, attention, transformers, BERT, and GPT.
- Next, generative models are covered including variational models, generative moment matching nets, GAN, and diffusion models.
- The other covered topics are Boltzmann machines, graph neural nets, deep reinforcement learning, meta-learning, network compression, federated learning, explainable AI, and self-supervised learning.
- Some applications of deep learning, including usage in computer vision, image processing, and NLP, are introduced.
- If time allows, the theory of optimization in networks is also discussed.

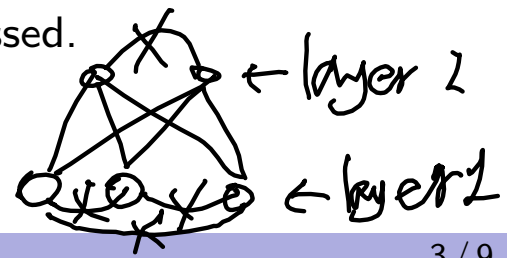

The Tentative Schedule of Weeks

Week 1

Topics: Preliminaries (probability and expectation, the learning model, overfitting, etc), logistic regression, feedforward neural network (perceptron, backpropagation, feed-forward neural network, stochastic gradient descent)

Week 2

Topics: Regularization techniques (overfitting/underfitting, cross validation, weight decay, dataset augmentation, adding noise, early stopping, parameter tying and parameter sharing, bagging, dropout, batch normalization)

Week 3

Topics: Convolutional Networks (CNN, its philosophy, and variants like ResNet, UNet, DenseNet, AlexNet, etc), Deep metric learning, Self-supervised learning, Deep Learning Tuning (choosing the model architecture, choosing the optimizer, and hyper-parameter tuning in networks, e.g., batch size, learning rate, early stopping parameters, etc)

Week 4

Topics: Sequence Modeling (introduction to NLP, RNN, LSTM, GRU, ELMo, etc), Attention mechanism (attention, transformers, BERT, GPT), introduction to Large Language Models (LLM)

Week 5

Topics: Network Compression (knowledge distillation, etc), Restricted Boltzmann Machine (RBM) and deep belief network

Week 6

Topics: Generative Models (variational models, generative moment matching, adversarial learning, GAN, diffusion models)

Week 7

Topics: Midterm exam

Week 8

Topics: Graph Neural Networks (convolution on graphs, GNN vanilla spectral GCN, ChebNet, GCNs, Graph Attention Networks), Few-shot Learning (meta-learning)

Week 9

Topics: Deep Reinforcement Learning

Week 9

Topics: Federated learning, Adversarial attacks

Week 10

Topics: Explainable AI (including shortcut learning and saliency maps)

Week 11

Topics: Theory of neural networks (SAM optimization, optimization landscape, probability bounds, information bottleneck, etc)

Course Materials

- Lecture notes will be provided to you.
- YouTube channel of the course: [Link]
- Our tutorial papers: [Link]
- Additional resource for interested students:
  - Prof. Ali Ghodsi's lectures at the University of Waterloo: [Link]
  - YouTube channel "Carnegie Mellon University Deep Learning": [Link]
- Additional books:
  - Benyamin Ghojogh, Mark Crowley, Fakhri Karray, Ali Ghodsi, "Elements of Dimensionality Reduction and Manifold Learning", Springer, 2023, [Link]
  - Ian Goodfellow, Yoshua Bengio, Aaron Courville, Deep learning, MIT press, 2016, [Link]
  - Simon J.D. Prince, Understanding Deep Learning, to be published by MIT press, [Link]
  - Christopher M. Bishop, "Pattern Recognition and Machine Learning", Springer, 2006, [Link]
- Research articles in the literature

Course's Websites

Introducing the instructor and TA of the course:

- Instructor: Benyamin Ghojogh, email address: bghojogh@uoguelph.ca
- TA: Sepideh Emami Tabrizi, email address: semamita@uoguelph.ca

Course info:

- Classes will be in-person.
- Discussion chats and questions will be in Microsoft Teams group of the course. The students will be added to the Teams group. Please mention (tag) my name and TA's name when you post a question or message in Teams (so it notifies us).
- The course's website is:
  `https://bghojogh.github.io/pages/uoguelph/engg-6600-01-s23/`
- I will probably upload the videos of the classes to my YouTube channel [Click here].
  I will eliminate personal information of students (such as when they introduce themselves) in the videos.

Course Evaluation

- Assignments (20%): Assignments will be posted on CourseLink along with the due dates. They are performed individually. We will probably have several (two to three) assignments.
- The midterm exam (30%): Date will be around week 6. Details to be discussed in class.
- Course project (40%)
  - Date: Week 6 - 11
  - More details will be discussed in class. Report will be electronic submission due in CourseLink.
  - The number of people in each group will be announced in the class.
  - Pick a topic after 6 weeks.
  - Submit the title and proposal/objectives in CourseLink to be checked and approved.
- Group Presentation (10%):
  - Date: Week 11 - 12
  - During class time
- Bonus points: participation in class, participation in the discussions, asking questions, and answering questions.

Course's Goal

- Don't worry much about your marks!
- Focus on understanding the materials of the course.
- Our goal is to learn the important practical and theoretical algorithms in statistical machine learning, so you can use them in both your **industrial projects** and **academic research**.
- About theory and practice:
  - We will learn **some theory** to understand why these methods work.
  - We will also learn how to use the methods **in practice** for practical usage.

Ask Questions!

- Please ask questions whenever you do not understand something.
- Let the class be discussion-based. I do not want to be the sole speaker. We are gonna learn all together.