Introduction to the Course

Deep Learning (ENGG*6600*07)

School of Engineering,
University of Guelph, ON, Canada

Course Instructor: Benyamin Ghojogh
Fall 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 is 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

$\min_x f(x)$ $\frac{\partial f(x)}{\partial x} \overset{\text{set}}{=} 0 \rightarrow x = ?$

- 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 one neuron (Perceptron, ADALINE, logistic regression) and one-layer networks (radial basis function, selforganizing map).
- Then, we cover feed-forward nets, backpropagation, stochastic gradient descent, AdaGrad, RMSProp, and Adam.
- Then, convolutional neural nets and important CNN architectures are introduced.
- Regularization techniques such as weight decay, batch normalization, and dropout are also explained.
- Then, we cover sequence modeling (useful for NLP and speech processing) including RNN, LSTM, attention, transformers, BERT, and GPT.
- Deep metric learning and Siamese network are introduced for data embedding.
- Next, generative models are covered including variational models, generative moment matching nets, GAN, and diffusion models.
- Depending on the time, the other covered topics can be Boltzmann machines, graph neural nets, knowledge distillation for network compression, deep reinforcement learning, meta-learning, federated learning, explainable AI, self-supervised learning, and the theory of optimization in networks.
- Some applications of deep learning, including usage in computer vision, image processing, and NLP, are also introduced.

$\min_x f(x)$

$\frac{\partial f(x)}{\partial x}$ set 0 → $x = \bigcirc$

$\tan(x)$ set 0

$\frac{\sin(x)}{x}$

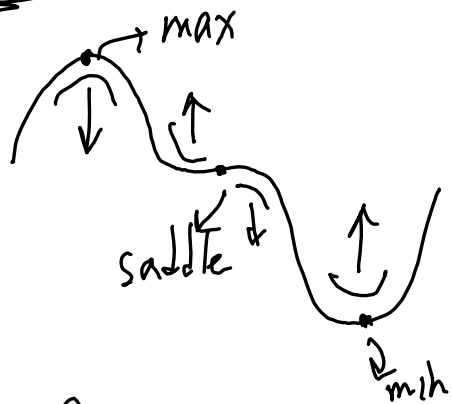

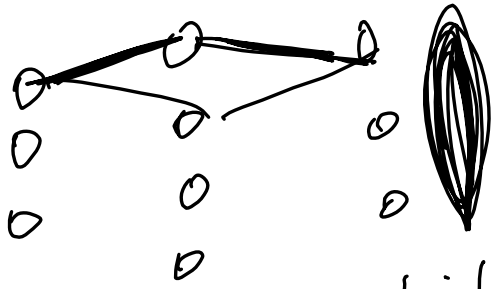

{ interior-point algorithm
barrier methods

$\min f(x)$

s.t. $\bigcirc x \leq 5$

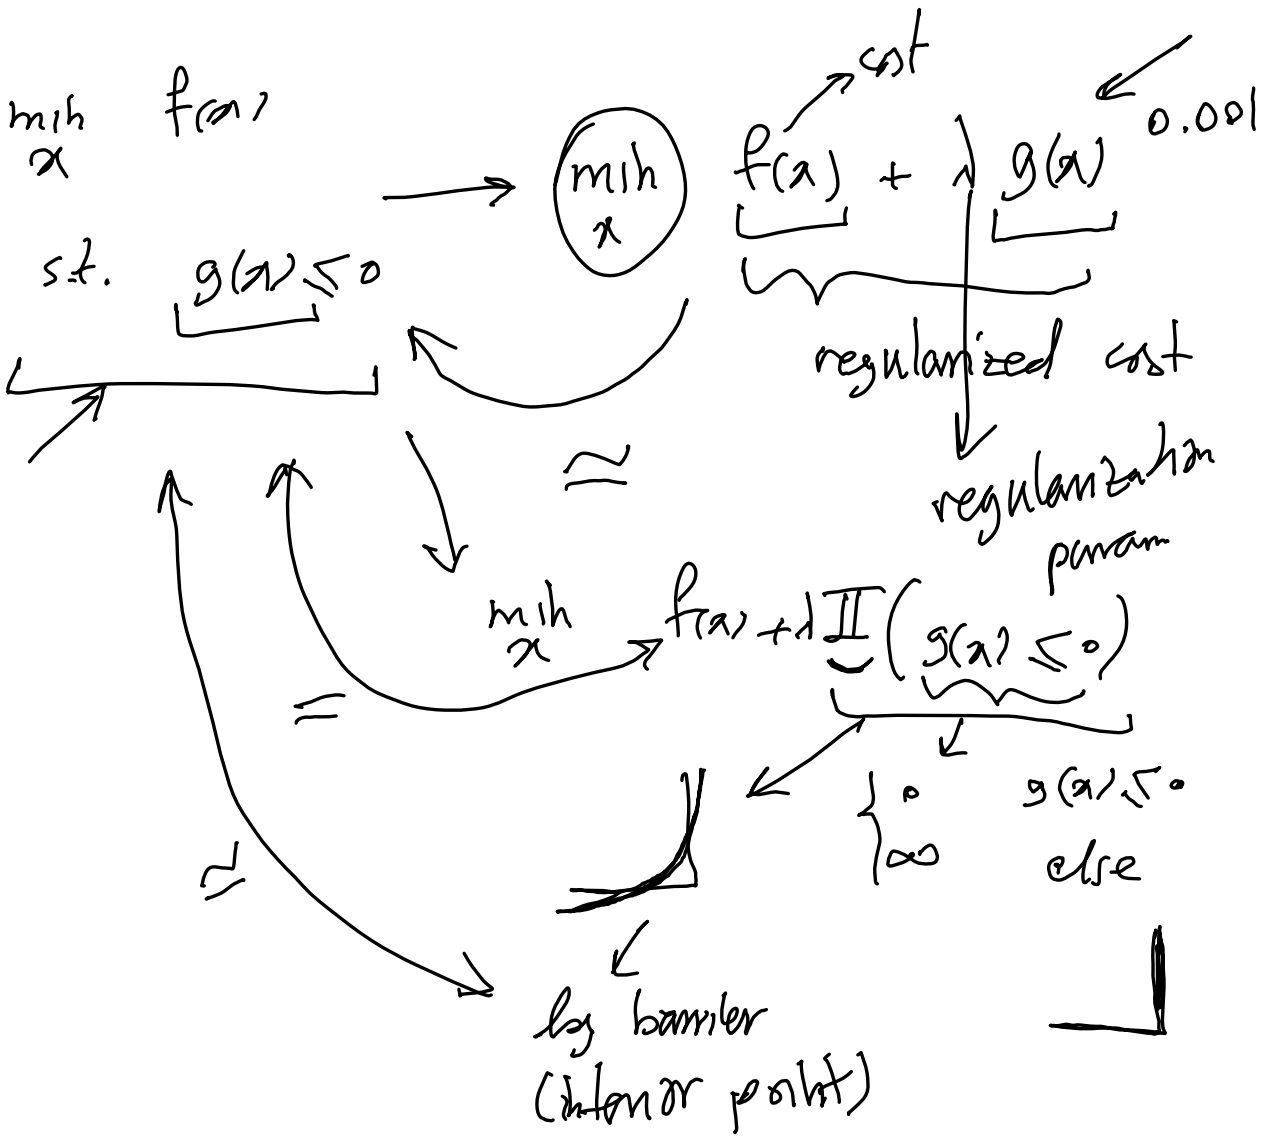
min f(x)
x
s.t. g(x) ≤ 0
min
x
f(x) + λ g(x)
cost
0.001
regularized cost
regularization param
≃
=
≃
min
x
f(x) + λ Π(g(x) ≤ 0)
0
∞
g(x) ≤ 0
else
log barrier
(interior point)

The Tentative Schedule of Weeks

Week 1

Topics: Preliminaries (dataset, learning task, projection, norm, derivative, optimization, eigenvalues), One neuron (Perceptron, ADALINE, logistic regression)

Week 2

Topics: One-layer networks (radial basis function, self-organizing map), fully connected network, activation functions

Week 3

Topics: Backpropagation, stochastic gradient descent, AdaGrad, RMSProp, Adam

Week 4

Topics: Convolutional neural network and important CNN architectures (AlexNet, VGG, Inception, GoogLeNet, U-Net, ResNet, DenseNet)

Week 5

Topics: Regularization in deep learning (wight decay, noise injection, early stopping, dropout)

Week 6

Topics: Recurrent neural network (RNN) and long short term memory (LSTM) network

Week 7

Topics: Midterm exam

Week 8

Topics: Attention, transformer, BERT, and GPT

Week 9

Topics: Deep metric learning, Siamese network, triplet loss

Week 9

Topics: Variational inference and variational autoencoder (VAE)

Week 10

Topics: Generative adversarial network (GAN)

Week 11

Topics: Graph Neural Networks (convolution on graphs, GNN vanilla spectral GCN, ChebNet, GCNs, Graph Attention Networks), Knowledge Distillation (KD)

Week 12

Topics: Group presentations of projects

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 of the course:

- Instructor: Benyamin Ghojogh, email address: bghojogh@uoguelph.ca
- The course does not have any TA in this semester :(

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-6600-07-f23/`
- 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: Assignments will be posted on CourseLink along with the due dates. They are not mandatory but are for you to practice. They are not graded. I will provide the solution keys to you after the due times.
- The midterm exam (50%): 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.