DDLS Research School PhD Course VT2025

Course Title: Training, Validation, Problem Diagnosis, and Troubleshooting of Using Deep Neural Networks in Life Science Applications **(5 hp)**

Instructor: Chun-Biu Li, Stockholm Univ. (cbli@math.su.se)

Course Content:
Deep neural networks (DNNs) constitute the building blocks for many modern machine learning and data science applications, such as regression, classification, dimensionality reduction, generative models, deep reinforcement learning, etc. This course provides a comprehensive training on how to practically train and use DNNs properly for statistically modeling. To be self-contained, an overview of basic concepts of DNNs relevant to life science applications will be included in the course. *Instead of mathematical technicalities, the course will then focus on establishing intuitive but deep understanding of the training and validation processes of DNNs.* Topics covered will include choosing appropriate state-of-the-art network architectures, model hyper-parameters, optimization methods, practices in monitoring the training processes, identifying improper training and troubleshooting, imposing suitable regularizations and their effects, training time complexity and possible speedups, validation of the resulting DNNs and troubleshooting, effective scrutiny of the DNN outputs to facilitate biological interpretations, etc.

Learning Outcomes:
After completing the course, students are expected to be able to
- Explain basic and modern concepts of DNNs
- Select appropriate training, validation methods and software to model a given life science dataset
- Compare strengths and weaknesses of various DNNs architectures, training, regularization and validation schemes
- Perform problem diagnosis and troubleshooting when applying DNNs to real life science problems
- Present analyses and results logically in both written report and oral presentation

Prerequisites:
- The course target students conducting biological studies who may have little or small experience in applying machine learning to analyze and model their data.
- Basic knowledge in linear algebra, calculus, probability and statistics with levels equivalent to those provided in a 1st year bachelor mathematics course for science students
- Have experience in coding (preferable in R, Python or Matlab), and using machine learning/AI libraries and functions

Course Format:
- 2 online meetings (2-3 hours each) + 5-day on-site meetings (*attending of on-site meeting is mandatory*)
- On-site meeting: Theory Section: 9:30-noon, Lunch & Fika: noon-13:30, Practice Section: 13:30-16:00

Location for On-site meetings:
Albano Campus, Stockholm University

Dates of On-site meetings:
2 possible weeks to choose from. Students only need to attend either week below:
1) May 12-16
2) June 9-13

Practice Sections:
Participants divided into groups according to their backgrounds to carry out projects. They are encouraged to bring data/model from their own PhD studies and apply what they learn from the course. They can also pick up existing data/model from well known databases to carry out the analysis. Participants are expected to bring their laptops to the on-site sections.

Assessment:
Project reports, presentations, and seminars, etc.

Tentative Meeting Plan (Theory Sections, 2 online and 5 on-site Lectures):
** The following plan can be subject to change depending on learning progress and interests*

Lecture 1 (Online): Course introduction and some statistical backgrounds
Participants and teachers introduce themselves, course introduction, forming groups, DNNs in various fields of machine learning, cost functions, under-fitting/ over-fitting and generalization.

Lecture 2 (Online): Overview of basic concepts of deep learning
Overview of DNN architectures (feedforward, convolutional, recurrent, gating, residual, U-net, attention, encoder-decoder, etc.), various input formats and output functions, various activation functions, etc.

Lecture 3 (On-site): Continue on basic concepts of deep learning
Overview of DNN architectures (feedforward, convolutional, recurrent, gating, residual, U-net, attention, encoder-decoder, etc.), various input formats and output functions, various activation functions, etc.

Lecture 4 (On-site): Optimization
Stochastic gradient descent and its generalizations, setting hyperparameters, parameter initialization, vanishing and exploding gradients, time complexities and possible speedups.

Lecture 5 (On-site): Regularization
Various regularizations, including norm penalties, data augmentation, noise injections, early stopping, dropout, batch and layer normalizations, etc., and their effects.

Lecture 6 (On-site): Evaluation
Various evaluation techniques, including the monitoring of training/validation curves, parameter updates, and the transformation of representations across layers, error quantifications, model selections, diagnosis of training problems.

Lecture 7 (On-site): Common Problems and Troubleshooting, Guest Speaker(s)
Common problems in training, evaluation, and possible solutions related to discussions in previous lectures and inputs from participants. Include one to two guest speaker(s) (e.g. DDLS fellows) to share difficulties and solutions in applying deep learning to life science studies.

Course References:

1. "Deep Learning" by Goodfellow et al., MIT Press, 2016
2. "Dive into Deep Learning" by Zhang et al., Cambridge University Press, 2023