

GRS MA 752- Mathematical Foundations of Deep Learning

Instructor: Konstantinos Spiliopoulos

Office: 665 Commonwealth Ave, CCDS 438

Email: kspiliop@math.bu.edu

Course web-page: <http://math.bu.edu/people/kspiliop/Spring2026MA752.html>

Meets: Spring 2026, Tuesday-Thursday 11:00-12:15 at 685-725 Comm Ave CAS B27

Office hours: Tuesday-Thursday 10-11

Deep learning is a relatively recent field and mathematical work on deep learning, which is the focus of this class, is even more recent. Key mathematical results will be presented in class and often will be assigned as readings. We will make use of the following book, which grew out of notes from this class as taught in previous years:

- K. Spiliopoulos, R. Sowers, and J. Sirignano, Mathematical Foundations of Deep Learning Models and Algorithms, AMS, 2025: <https://mathdl.github.io/>

Some further recommended readings include

- Lecture notes by Matus Telgarsky: <https://mjt.cs.illinois.edu/dlt/>
- I. Goodfellow, Y. Bengio and A. Courville, Deep Learning, The MIT Press, Cambridge, MA, 2016.
- C. Bishop with H. Bishop, Deep Learning Foundations and Concepts, Springer, 2024

Course Description:

This course is a rigorous introduction to the mathematical foundations of deep learning. In particular, the course will introduce the students to the theory of universal approximation, stochastic gradient type of optimizers and their convergence properties, statistical learning bounds, approximation theory, depth separation results, neural tangent kernel, mean field overparametrized regime, implicit regularization and noisy dynamics, different deep neural network architectures and their properties, transformers, attention mechanism, variational inference, generative modeling, reinforcement learning, Q-learning, several computational aspects such as back propagation, batch normalization and dropout, deep learning for dynamical systems.

The course material will be based on theory, methods (both theoretical and computational) and examples from various scientific disciplines. The class will focus on supervised learning.

Course Prerequisites: Introduction to probability and/or stochastic processes (MA 581 or MA583 or equivalents), Differential Equations (MA226 or MA231 or equivalent), Linear algebra (MA242 or equivalent), basic statistical theory and basic programming skills. Some exposure to real analysis at the undergraduate level will also be useful. PDE's, graduate level probability and statistics will be helpful but not necessary. Students are expected to have the knowledge equivalent to undergraduate level probability or stochastic processes, basic statistical theory, linear algebra, differential equations as well as basic coding skills (ideally in Python).

Tentative Course Syllabus:

Note: I reserve the right to change the course syllabus depending on the progress of the class.

- Week 1: Introduction & Overview
What is deep learning? Some ideas, applications and examples. Overview of course material and course requirements. Review of tasks in machine learning: supervised learning (regression and classification), unsupervised (such as clustering) and reinforcement learning.
- Week 2: Concepts of overfitting and generalization. Linear regression as motivation for neural networks. From the perceptron model to kernels to neural networks.
- Week 3: Shallow and feed forward deep neural networks, universal approximation theorems and their proofs, depth separation results, Barron space and derivation of related bounds.
- Week 4: Gradient flow and gradient descent: conditions for convergence to critical points and minimizers.
Stochastic gradient descent: Robbins-Monro theory and learning rate decay. Effect of regularization. Backpropagation.
- Week 5: Loss landscapes in deep learning: Over-parametrization effects.
Variants of stochastic gradient descent: SGD with momentum and optimality for Nesterov's accelerated gradient descent. RMSprop and ADAM optimizers. Convergence properties.
- Week 6: Practical considerations when training neural networks: Backpropagation, exploding and vanishing gradients, different initialization schemes and their effects, batch-normalization, dropout.
- Week 7: Optimization in shallow neural networks- the linear regime. Neural Tangent Kernel: derivation, convergence theory and properties.
- Week 8: NO class, BU spring break.
- Week 9: Nonlinear overparametrization regime: mean field regime, convergence theory and optimality. Comparisons between the linear and the nonlinear regime. Transitioning from discrete to continuum models.
- Week 10: Deep neural network architectures. Convolutional neural networks. Recurrent neural networks, LSTMs. Transformers and the attention mechanism. Definitions, properties and usage.
- Week 11: Transformers and the attention mechanism. Embeddings, tokenization, decoder transformers, encoder transformers, sequence-to-sequence transformers, large language models. Definitions, properties and usage.
- Week 12: Variational Methods: estimating densities, evidence lower bound, generative adversarial network (GAN), Wasserstein GAN (WGAN).

- Week 13: Diffusion models: encoder-decoder, score matching diffusions, discrete diffusion models, masked diffusion models.
- Week 14: Deep learning and dynamical systems: solving ODEs and PDEs with neural networks, closure models based on neural networks, convergence properties, effect of architectures, applications areas (finance, statistics, etc).
- Week 15: Presentation of class projects.

Assigned work and grading: Registered students will be expected to do regular paper reading, complete a few sets of homework problems (approximately every 2-3 weeks) (75%) and present a final project (25%). Needless to say, you should work on the homework on your own, unless otherwise instructed by me. Late homeworks will not be accepted. The grading policy may change depending on the progress of the class.

Please Note: Students are responsible for knowing the Boston University Academic Conduct Code which is posted at

<https://www.bu.edu/academics/policies/academic-conduct-code/>