# SYSE 6V80: Special Topics – Deep Learning

| | | | |
|---|---|---|---|
| **Instructor:** | Aykut C. Satıcı | **Term:** | Spring 2026 |
| **Email:** | aykut.satici@utdallas.edu | **Course Time:** | Thursday 4:00 – 6:45 p.m. |
| **Room:** | ECSN 3.716 | **Classroom:** | ECSN 2.112 |

**Course Page:** https://elearning.utdallas.edu/ with Teams links.

**Office Hours:** Thursday 11:00 – 12:00 a.m. at ECSN 3.716, or by appointment.

**Prerequisites:** This class requires an understanding and working knowledge of linear algebra and basic probability theory. Exposure to basic optimization theory/algorithms is recommended, but this topic will be reviewed during the course. We will also need a working ability to program in Python.

**Course Description:** This course is designed to equip students with the fundamental methodologies in deep learning, with a practical focus on two domains: (i) image/video classification and generation, and (ii) controlling robotic locomotion and manipulation tasks. The course will cover the training of state-of-the-art neural networks to provide the solution of complex inference tasks.

**Learning Outcomes:**

- Train and evaluate basic artificial neural network architectures,
- Explain how backpropagation works,
- Train and evaluate the convolutional neural network architectures,
- Construct Variational AutoEncoder (VAE) models for representation learning,
- Implement diffusion models for generative learning,
- Experiment with various architectures (e.g. U-Net) to power diffusion models,
- Implement transformer models for representation learning,
- Recognize basic reinforcement learning concepts and algorithms,
- Implement decision-transformers and/or decision-diffusers for robot control.

**Texts and Materials:** There is no required textbook for this course. The following books will be helpful in understanding machine learning concepts. The provided course material may be downloaded during the course; however, these materials are for registered students' use only. Classroom materials may not be reproduced or shared with those not in class or uploaded to other online environments except to implement an approved AccessAbility Resource Center accommodation. Failure to comply with these University requirements is a violation of the Student Code of Conduct.

- M. Nielsen, *Neural Networks and Deep Learning*, Determination Press, 2015.
- G. Sapunov, *Deep Learning with JAX*, Manning, 2024.
  - This GitHub repository holds the code examples.
- S. Prince, *Understanding Deep Learning*, MIT Press, 2023.
- C. Bishop, *Deep Learning: Foundations and Concepts*, Springer, 2023.
- D. V. Godoy, *Deep Learning with PyTorch Step-by-Step*, Independently published, 2022.
- R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2018. http://incompleteideas.net/book/the-book-2nd.html
- M. Morales, *Grokking Deep Reinforcement Learning*, Manning, 2020.

**Class Recordings:** Students are expected to follow appropriate University policies and maintain the security of passwords used to access recorded lectures. Unless the AccessAbility Resource Center has approved the

student to record the instruction, students are expressly prohibited from recording any part of this course. Recordings may not be published, reproduced, or shared with those not in the class or uploaded to other online environments except to implement an approved AccessAbility Resource Center accommodation. Failure to comply with these University requirements is a violation of the Student Code of Conduct.

The instructor may record meetings of this course. These recordings will be made available to all students registered for this class if the intent is to supplement the classroom experience. If the instructor or a UTD school/department/office plans any other uses for the recordings, consent of the students identifiable in the recordings is required prior to such use unless an exception is allowed by law.

**Grading Policy:** Homework (20%), Projects 1 and 2 ($25\% \times 2 = 50\%$), Project 3 (30%).

- Homework must be submitted on the due date and time on eLearning. Extensions may be granted for special circumstances and only when requested at least one day in advance.
- You are responsible for all the information given in class verbally and/or in writing. Any information about the course on the web may be superseded by the information given in the class.

**Important Dates (Tentative):**

| | |
|---|---|
| Project #1 | February 19, 2026 |
| Project #2 | March 19, 2026 |
| Project #3 | April 16, 2026 |

**Project 3:**

- Students will be given the freedom to choose the final project topic out of their interest, subject to the instructor's approval.
- A candidate project topic should address a real life problem and should involve enough to let students demonstrate their proficiency.

**Academic Honesty (Comet Creed):** Academic dishonesty is **not** tolerated. Students are responsible for reading and understanding the student code of conduct. For more information, please see the student code of conduct.

Comet creed was voted by the UT Dallas student body in 2014. It is a standard that Comets choose to live by and encourage others to do the same:

> As a Comet, I pledge honesty, integrity, and service in all that I do.

**Accommodations for Students with Disabilities:** Please review the section within the UT Dallas Syllabus Policies and Procedures webpage.

**UT Dallas Syllabus Policies and Procedures:** Please visit the Syllabus Policies page to view the University's policies and procedures segment of the course syllabus.

Please review the catalog sections regarding the credit/nocredit or pass/fail grading option and withdrawal from class.

**Further Resources:**

- Build a Large Language Model from Scratch
- Introduction to Flow and Diffusion Models

**Tentative Course Outline:**

- Week 1: Introduction
- Week 2: Basic Neural Networks
- Week 3: Basic Machine Learning Concepts
  - Common loss functions
  - Bias-variance trade-off
  - Overfitting and regularization
  - Weight initialization
- Week 4 – 5: Convolutional Neural Networks
- Week 6: Variational AutoEncoders
  - Latent variable models
  - Marginal Log-Likelihood
  - Variational Lower Bound (or Evidence Lower Bound)
- Week 7: Attention and Specialized Architectures
  - Attention mechanism
  - Encoder and decoder
  - U-Net
- Week 8 – 10: Flow Matching and Diffusion
  - Score matching
  - Langevin dynamics
  - Stochastic differential equations
- Week 11 – 13: Transformers
  - Vanilla model
  - Self-attention
  - Vision Transformers (ViT)
  - KV Caching
- Week 14: Basics of Reinforcement Learning
  - Markov Decision Processes
  - Policy evaluation and improvement
  - Deep Q-Networks (DQN), Double DQN
  - REINFORCE, Actor-Critic methods
- Week 15: Offline Reinforcement Learning
  - Decision Transformers
  - Decision Diffusers