

## Overview

Welcome to **ECE324 – Machine Intelligence, Software, and Neural Networks!** This course serves as an introduction to machine learning engineering, with a focus on neural networks. The entire process of developing a machine learning solution, from data collection to software development, as well as ethics in machine learning, will be discussed. Practical techniques in machine learning will be covered, including data augmentation and the use of pre-trained networks. Topics covered will include the fundamentals of neural networks, convolutional neural networks, recurrent neural networks, generative adversarial networks and transformer networks. Students will complete a major hands-on project in machine learning.

## Website &amp; Forum

Website: <https://www.cs.toronto.edu/~guerzhoy/324/>

Forum: <https://piazza.com/utoronto.ca/ece324/>

All course handouts will be posted on the course website. *Students are responsible for reading all announcements on the course forum on Piazza.*

Instructor  
Grading

Instructor	Email	Office	Office Hours
Michael Guerzhoy	<a href="mailto:guerzhoy@cs.toronto.edu">guerzhoy@cs.toronto.edu</a>	BA 2028	Wednesday 6p.m.

The grading scheme for the course is as follows. All due dates are at 23 p.m.

	Worth	Date
Mini-Projects	15%	2 Mini-projects, due tentatively Feb. 13, March 29
Project proposal	5%	Feb. 1
Interim project report	7.5%	March 15
Project presentation	7.5%	last two weeks of class
Project final report	25%	April 13
Midterm	15%	March. 2
Final exam	25%	TBA

Late  
submissions  
References

Everyone has 5 grace days available at the start of the semester. You can use a grace day to avoid a late penalty. Grace days are applied automatically at the first opportunity that they can be used. Late submissions are allowed for up to three days after the deadlines. Late submissions after 3 days are not allowed except for documented medical/personal circumstances.

The following textbooks are not required, but you may find them useful as additional references.

- I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*.  
Available for free at <https://www.deeplearningbook.org/>.
- S. Barocas, M. Hardt, A. Narayanan, *Fairness in Machine Learning: Limitations and Opportunities*.  
Available for free at <https://fairmlbook.org/>.
- H. Daume III, *A Course In Machine Learning*.  
Available for free at <http://ciml.info/>.
- C. Bishop, *Machine Learning and Pattern Recognition*.  
Available for free at <https://www.microsoft.com/en-us/research/people/cmbishop/prml-book/>.
- D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*.  
Available for free at <https://www.inference.org.uk/itprnn/book.pdf>
- Cosma Shalizi, *Advanced Data Analysis from an Elementary point of View*.  
Available for free at <https://www.stat.cmu.edu/~cshalizi/ADAfaEPoV/>

*Topic outline*

The following is a tentative list of topics for the course.

- Statistical inference: maximum likelihood and Bayesian inference
- Causal inference
- Fairness in machine learning: overview of the field; observational measures; fairness and causality
- Word embeddings
- Transformer architectures
- Cognitive science perspectives on language models
- Generative adversarial networks, variational autoencoders, diffusion models
- Graph embeddings and graph neural networks
- Explainability in convolutional networks
- Adversarial examples
- Inductive bias and meta-learning
- Scaling laws for neural networks
- Architectures for autonomous intelligence
- Ethics in machine learning: overview of approaches