

CS 570 – Machine Learning, Winter 2025

Instructor: Michael Kowalczyk

Office: 2222 Jamrich Hall

Office Phone: 227-1600

Class Meetings: 9:00am – 10:40am Monday and Wednesday in room 3102 Jamrich Hall

In-person “walk-in” Office Hours:

Monday and Wednesday, 11:00am – 1:50am

By-appointment Online Office Hours:

Tuesday, Thursday, and Friday, 3:00pm – 5:00pm

Reserve a time by going to ESP in MyNMU; reserve your appointment the day before (at latest).

At your appointed time visit <https://nmu.zoom.us/my/drkow> and use passcode: 147147

Email: mkowalcz@nmu.edu

Course Website: <https://educat.nmu.edu>

Overview:

In this course, we study algorithms for writing computer programs that improve with experience. The basic theory of machine learning will be covered, including concepts such as hypothesis space, bias, overfitting, training sets, and testing sets. We will also dive into the details of specific machine learning techniques, including decision trees, Bayesian classifiers, instance-based learning, and artificial neural networks. The course is focused primarily on supervised learning, though other approaches (semi-supervised, un-supervised, and reinforcement learning) will be explored as well.

Prerequisites:

CS 422 and MA 163 or equivalent.

Required Textbook (available at the NMU Bookstore or through online retailers):

Fundamentals of Machine Learning for Predictive Data Analytics, 2nd edition

John D. Kelleher, Brian Mac Namee, Aoife D’Arcy

The MIT Press, 2020

Equipment:

You will need a computer with a web browser and Internet access. You will also need to install Weka (machine learning software). If for some reason you don’t have a laptop, talk to me as soon as you can, since I plan to have some in-class exercises with them.

Grading:

Grades will be based on assignments, projects, and exams. There will be at least two exams (100 points each). Projects and assignments are weighted based on their size and complexity.

Final Exam Date & Schedule Conflicts:

The final exam will be on Tuesday, April 29th from 8:00am until 9:50am. Any conflicts with the exams (due to religious observances, other coursework, intercollegiate athletics, etc) must be made known to me as soon as you are aware of the conflict.

Course objectives:

The main goals are to 1) understand the details of the most common machine learning algorithms in use today, and 2) understand how to apply them to real-world problems. At the end of the course each student should...

- understand the basic terminology and theory of machine learning,
- know the importance of data preparation and feature engineering,
- know in detail how the most common machine learning techniques work, and the math that drives them,
- have implemented several machine learning algorithms from scratch,
- be familiar with the common approaches to evaluating the performance of generated models, and
- have demonstrated the ability to apply machine learning algorithms to solve real-world problems.

Evaluation of these learning outcomes will be done through assignments, projects, and/or exams.

Academic Conduct:

I strive to work hard with honesty and integrity; I expect my students to do the same. If you have occasion to include something in a project or assignment that you didn't write, clearly indicate what part is not your own authorship.

Disability Services:

If you have a need for disability-related accommodations or services, please inform the Coordinator of Disability Services in the Dean of Students Office at 2001 C. B. Hedgecock Building (227-1700). Reasonable and effective accommodations and services will be provided to students if requests are made in a timely manner, with appropriate documentation, in accordance with federal, state, and University guidelines.