

## Machine Learning for Natural Language Processing

**Instructor:** Dr. Fei Song

Email: [fsong@uoguelph.ca](mailto:fsong@uoguelph.ca)

Office Hours: Tuesdays and Fridays, 2:30 – 4:30 pm (Virtual)

Course Website: <https://courselink.uoguelph.ca/>

### Overview

Natural Language Processing is an interdisciplinary area among Mathematics, Information Theory, Linguistics, and Computer Science. It has been applied successfully to a wide range of problems such as Speech Recognition, Information Retrieval and Extraction, Text Classification, Sentiment Analysis, Automatic Summarization, and Machine Translation. Due to the sequential nature of language data, similar techniques have also been applied to other domains such as Biological and Time Series data analysis. This course provides an introduction to this emerging field, with emphasis on the applications of machine learning techniques. The students should have proficiency in Java and working experience in Python, as well as good understanding of the related mathematics, including Linear Algebra, Probability, and Statistics. In addition to attending lectures, the students are required to review the current literature and present two papers in the class. They are also required to complete two warm-up assignments and an implementation project that applies the related technique(s) they learned to a real world problem. Listed in the following is a set of topics that we intend to cover in the course:

- Introduction to Natural Language Processing
- Language Modeling and N-gram Models
- Information Retrieval Models and Implementations
- Machine Learning for Text Classification and Sentiment Analysis
- Introduction to Deep Learning Models
- Static and contextualized word embeddings
- Unsupervised Machine Learning Methods

### Evaluation

- Warmup Assignments: 30% (2 x 15%)
- Research Presentations: 20% (2 x 10%)
- Term Project: 50%

### Recommended References

Chris Manning and Hinrich Schütz. *Foundations of Statistical Natural Language Processing*. The MIT Press, 1999.

Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Francis Bach. *Deep Learning*. The MIT Press, 2016.

Ian H. Witten, Eibe Frank, Mark A. Hall, and Christopher J. Pal. *Data Mining: Practical Machine Learning Tools and Techniques*. Fourth Edition. Morgan Kaufmann, 2016.

Sebastian Raschka and Vahid Mirjalili. *Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow*. Second Edition. Packt Publishing, 2017.