

First half of the course:

- ▶ Introduction (Aug 27–29)
- ▶ Regular expressions (Sept 3–5)
- ▶ Edit distance (Sept 8)
- ▶ Information theory (Sept 10–12)
- ▶ Language models (Sept 13–22)
- ▶ Parts of speech and HMMs (Sept 22–Oct 3)
- ▶ Parsing (Oct 6–Oct 13)
- ▶ Review (Oct 15)
- ▶ Midterm (Oct 17)

Second half of the course:

- ▶ Introduction (Oct 24–27)
- ▶ Naïve Bayes Classification (Oct 29–Nov 3)
- ▶ Stylometry (Nov 5–Nov 10)
- ▶ Neural nets (Nov 12–14)
- ▶ Word embeddings (Nov 17–19)
- ▶ Machine translation (Nov 21–Dec 3)
- ▶ Large language models (Dec 5–Dec 10)

Machine learning is a form of applied statistics with emphasis on the use of computers to statistically estimate complicated functions.

Goodfellow et al., Deep Learning, 2016. Pg 95.

Machine learning is the science (and art) of programming computers so they can learn from data. [In 1959, Arthur Samuel defined machine learning as the] field of study that gives computers the ability to learn without being explicitly programmed.

Géron, Hands-On Machine Learning, 2019. Pg 2.

[Machine learning is] a set of methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data or to perform other kinds of decision-making under uncertainty.

Murphy, Machine Learning: A Probabilistic Perspective, 2012. Pg 1.

Machine learning main tasks:

- ▶ Regression, where the target type is \mathbb{R}
- ▶ Classification, where the target type is a finite set
 - ▶ Binary classification, where the target is $\{F, T\}$ (or $\{0, 1\}$ or $\{-1, 1\}$...)
- ▶ Density estimation, where the target type is $[0, 1]$.

Other machine learning tasks (see Goodfellow, *Deep Learning*, pg 96–100):

- ▶ Transcription, where the observations are unstructured and the targets are text.
- ▶ Machine translation, where the observations and targets are text.
- ▶ Anomaly detection, where the targets are indicators of whether the observation is atypical.
- ▶ Synthesis and sampling, where there are no observations in deployment, but rather the program produces new observations similar to those in training.
- ▶ Denoising, where the targets are corrected versions of the observations.

Coming up:

- ▶ Do CKY parsing programming assignment (Mon, Oct 27)
- ▶ Take ML basics and bag-of-words quiz (Tues, Oct 28)
- ▶ Read J&M B.(0-8, 10). (Fri, Oct 31)