

# AI & Machine Learning

Introductory materials for artificial intelligence / machine learning applications. These workshops were recently introduced by the Krembil Centre and will continue to expand to address cutting-edge approaches and tools for machine learning.

The CAMH Scientific Computing Cluster supports high performance GPU hardware for machine learning. Accounts can be created by submitting an [S CC Account Request Form](#).

## Introductions:

- [Fundamentals and Neural Networks](#)
- [Machine-Learning Coursera](#)

## AutoML

- [ARXIV - Overview \(PDF\)](#)
- [SKLEARN](#)
- [TPOT - AutoML](#)

## Machine Learning in R:

- [R and Tidyverse](#)
- ['caret' - Classification And REgression Training](#)

## Machine Learning in Python:

- [ML in Python - Notebooks](#)

## Machine Learning Model Evaluation:

- [https://fairlearn.org/v0.10/auto\\_examples/plot\\_intersectional\\_bias.html](https://fairlearn.org/v0.10/auto_examples/plot_intersectional_bias.html)  
*This tutorial introduces learners to the value of intersectionality in the context of mental health, and demonstrates how to carry out an evaluation of machine learning models for intersectional bias, using a case scenario based in mental health and simulated data.*

## Image / Pattern Recognition

- <https://pubs.rsna.org/doi/10.1148/rg.2017160130>

## Natural Language Processing

- [Speech and Language Processing, Daniel Jurafsky](#)
- [NLP - Online Guide](#)
- [GPT-2 Tutorial](#)

## Text Mining with R: A Tidy Approach by Julia Silge and David Robinson

This book/guide is available completely online (<https://www.tidytextmining.com/index.html>) and it offers a nice overview of popular topics in text mining (i.e., sentiment analysis, word frequencies, topic modelling) as well as how to carry out the analyses in R with examples.

## An Introduction to Text Processing and Analysis with R by Michael Clark

This online guide is similar (<https://m-clark.github.io/text-analysis-with-R/intro.html>), and additionally covers strings and word embeddings with sample code/case studies.

This workshop summary on R-Bloggers also provides an introduction to some of these topics using a different R package for analyzing textual data in R (quanteda): <https://www.r-bloggers.com/advancing-text-mining-with-r-and-quanteda/>, and it additionally covers supervised classification, latent semantic analysis, and structural topic models.

## Topic Modelling with BERTopic

This GitHub repository contains code and helpful demonstrations and notebooks for conducting various forms of topic modelling on text using state-of-the-art techniques (i.e., large language models): <https://github.com/MaartenGr/BERTopic>

