

What Next?

Sarath Chandar

Books to read:

Machine Learning:

- Tom Mitchell. Machine Learning
- Christopher Bishop. Pattern recognition and machine learning.
- Trevor Hastie, Robert Tibshirani, Jerome Friedman. The elements of statistical learning.
- Kevin Murphy. Machine learning A probabilistic perspective.

Data Mining:

- Pang-Ning Tan, Michael Steinbach, Anuj Karpatne, Vipin Kumar. Introduction to Data Mining.
- Jiawei Han, Micheline Kamber, Jian Pei. Data Mining: concepts and techniques.
- Jure Leskovec, Anand Rajaraman, Jeffrey Ullman. Mining of Massive Datasets.

Books to read:

Graphical Models:

- Daphne Koller, Nir Friedman. Probabilistic Graphical Models.

Neural Networks:

- Raul Rojas. Neural Networks.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville. Deep Learning.

Reinforcement Learning:

- Rich Sutton and Andrew Barto. Reinforcement Learning: An Introduction.

Artificial Intelligence:

- Stuart Russell, Peter Norvig. Artificial Intelligence: A modern approach.

Next Courses:

- COMP550/559: Natural Language Processing (Fall)
- IFT6269 (UdeM): Probabilistic Graphical Models (Fall)
- IFT6080 (UdeM): Duckietown (Fall)

- COMP767: Reinforcement Learning (Winter)
- IFT6132 (UdeM): Advanced Structured Prediction and Optimization (Winter)
- IFT6135 (UdeM): Learning Representations (Winter)
- IFT6085 (UdeM): Theoretical Principles of Deep Learning (Winter)

Online Courses:

- [Reinforcement Learning](#) by Balaraman Ravindran. (NPTEL)
 - [Reinforcement Learning](#) by David Silver.
 - [Neural Networks](#) by Hugo Larochelle.
 - [Neural Networks](#) by Geoffrey Hinton. (Coursera)
 - [Probabilistic Graphical Models](#) by Daphne Koller. (Coursera)
-
- [Convex Optimization I](#) and [Convex Optimization II](#) by Stephen Boyd. (SEE)
 - [Probability Theory](#) by Krishna Jagannathan. (NPTEL)
 - [Linear Algebra](#) by Gilbert Strang. (MIT OCW)

Thank you!