

Course plan: IFT6266 H2017

Algorithmes d'apprentissage (de représentations)

<https://ift6266h17.wordpress.com>

Professor: Aaron Courville (Aaron.Courville@umontreal.ca)

Teaching Assistant: TBA

Research into statistical learning algorithms have brought important discoveries in the last two decades that have changed the way we think about the problem of making computers more intelligent. An agent is intelligent because it has operational knowledge (not necessarily in an explicit form) that allows it to perform certain tasks or answer questions on some domain. The learning algorithms have the objective that allows the machine to acquire operational knowledge via examples. They seek to uncover the latent structure underlying the observations.

Such trained models have been applied for prediction and decision making and the analysis of data. They are also seeing major use in commercial areas such as “data-mining” and other data analysis-heavy domains. Research into artificial intelligence is also strongly influenced by machine learning and it is these kinds of applications that are the focus of this course (and its practical component).

This course is a **second course** on machine learning following, for example, IFT3395 or IFT6390. Students that have never taken a machine learning course such start with IFT6390 (or the equivalent) before taking IFT6266. The subjects covered by this course (IFT6266) are advanced subjects, and related to a new sub-area of machine learning known as Deep Learning (and/or Representation Learning). Roughly following the new *Deep Learning* textbook of Goodfellow, Bengio and Courville, this course will cover the following subjects:

- Neural Network Fundamentals
 - Deep Feedforward Networks
 - Regularization for Deep Learning
 - Optimization for Training Deep Models
- Advanced neural networks:
 - Convolutional neural networks (with sequential and spatial data)
 - Recurrent networks (in particular the LSTM)
 - Practical Methodology
 - Applications
- Unsupervised Methods and Models
 - Auto-encoders and variants.
- Probabilistic graphical models:
 - Structured Probabilistic Models for Deep Learning
 - Monte Carlo Methods
 - Approximate Inference
 - Deep Generative Models

Pedagogical approach and evaluation: This course is based on multiple sources of pedagogical ma-

Course plan: IFT6266 H2017

Algorithmes d'apprentissage (de représentations)

<https://ift6266h17.wordpress.com>

Professor: Aaron Courville (Aaron.Courville@umontreal.ca)

Teaching Assistant: TBA

material, but will primarily follow the Deep Learning textbook. The lecture notes will draw from the online course of Hugo Larochelle (IFT725, Neural networks) at the University of Sherbrooke and Geoffrey Hinton's Coursera course on Neural Networks for Machine Learning. Each student is responsible for preparing questions before each class (and posting the question on the course website). The discussions in class on these questions will clarify the concepts under study and complement the assigned study material. For each question discussed in class, a student will be assigned the job of transcribing an answer on the course website.

There will be periodic quizzes and a final exam based on the more theoretical side of the material covered in class. Students will exercise the more practical aspects of the material by doing a class project that will involve a large experimental component. Students will be expected to maintain a journal (i.e. a blog visible online by the prof.) that documents the detailed activities and experiments as each student attacks the task given in the project.

The exam will count for 30% of the final grade, quizzes and class participation will make up 30% and the course project will count for 40%.

Schedule:

First class: Thursday, January 5th 2017, 9:30 – 11:30 am.

Day	Time	Location
Monday	14h30-16h30	Z-350 Pav. Claire-McNicoll
Thursday	09h30-11h30	Z-260 Pav. Claire-McNicoll