

IFT6080 – Music and Machine Learning (4 Credits)

Winter 2006

Professor Douglas Eck, 3385 André-Aisenstadt, eckdoug@iro.umontreal.ca

Course Homepage <http://www.iro.umontreal.ca/~pift6080>.

Students must check this page at least once per week for important information.

Meeting Times

Mondays 2:30PM to 4:30PM, 9 January to 10 April, Room Z-300, Pav. C. McNicoll

Wednesdays 2:30PM to 4:30PM, 11 January to 12 April, Room Z-300, Pav. C. McNicoll

Evaluation

Task	Amount
Homework	40%
Weekly Readings	10%
Reading Presentations	10%
Final Project	40%

This is a lab and readings course. There will be no exams.

Homework: (40%) Regular homework assignments dealing with signal processing, music feature extraction and machine learning will be given.

Weekly Readings: (10%) Weekly journal papers, book chapters and tutorial material will be assigned. Careful reading of all material is required.

Reading Presentations: (10%) Each student will present a short summary of one weekly reading and will then lead discussion of the paper.

Final Project: (40%) At mid-term students will submit a 3 to 5 page project proposal for approval. The project must involve an application of machine learning in the domain of music. Students will be furnished with project suggestions and will be shown where to find relevant datasets. Students will present their projects to the class in a 20-30 minute presentation at the end of the term. Example topics:

- Music similarity measures
- Automatic music composition
- Artist / genre classification
- Models of musical timing and rhythm (beat tracking, computer performance, computer/musician synchronization)
- Identification of musical structure (meter, phrase structure, melody)

Goals

The course is structured as a graduate seminar with a focus on reading journal papers and implementing machine learning models relevant to music. Our goal is to explore technological challenges in the domain of music from within the framework of machine learning. Areas of investigation include:

- Signal Processing and Feature Extraction
- Music Similarity Measures
- Music Categorization
- Sequence Generation
- Human-Realistic Music Performance

There is significant programming work in the homeworks and final project. My goal is that the projects deal with open problems in the field. In the Winter 2005 version of this course two projects were transformed into winning contest submissions in Genre Recognition and Note Onset Detection at an international programming contest (“MIREX” at the ISMIR 2005 conference).

Teaching

The primary language for teaching and reading will be English. All graded materials (homeworks, final project, presentation) can be done either in French or English.

Prerequisites

Students will be assumed to have *either* background in machine learning *or* background in music. For those lacking experience in either domain, extra reading will be required. Note that students having background in both machine learning and music are rare; thus students lacking experience in one of these areas should not be discouraged from taking the course.

Texts

Here is a list of recommended texts. There are no required texts for the course. Students with no background in machine learning should read Alpaydin (an undergraduate-level introduction to ML) and Bishop.

- Alpaydin, *Introduction to Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press, 2004.
- Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.
- Gold and Morgan, *Speech and Audio Signal Processing : Processing and Perception of Speech and Music*, Wiley, 1999.
- Hastie, Tibshirani, Friedman, *The Elements of Statistical Learning*, Springer-Verlag, 2001.
- Handel, *Listening, An introduction to the perception of auditory events*, MIT Press, 1993.