

# **CSI 5387: Concept Learning Systems / Machine Learning**

**Instructor:** Nathalie Japkowicz

**e-mail:** [nat@site.uottawa.ca](mailto:nat@site.uottawa.ca)

Objectives of the Course

and

Preliminaries



# Some Information

- ☞ **Instructor:** Dr. Nathalie Japkowicz
- ☞ **Office:** SITE 5-029
- ☞ **Phone Number:** 562-5800 x 6693 (don't rely on it!)
- ☞ **E-mail:** [nat@site.uottawa.ca](mailto:nat@site.uottawa.ca) (best way to contact me!)
- ☞ **Office Hours:**
  - Thursdays, 12:00pm-2:00pm
- ☞ **Extra Seminars:** TAMALE Seminars,
  - Thursdays (invited talks on Machine Learning and Natural Language Processing)
  - See: <http://www.tamale.uottawa.ca> for talk announcements
  - Write to Morvarid Sehatkar ([mseha092@site.uottawa.ca](mailto:mseha092@site.uottawa.ca)) to receive **all announcements by e-mail** (strongly suggested)



# Machine Learning: A Case Study

- ☛ Malfunctioning gearboxes have been the cause for CH-46 US Navy helicopters to crash.
- ☛ Although gearbox malfunctions can be diagnosed by a mechanic prior to a helicopter's take off, what if a malfunction occurs while in-flight, when it is impossible for a human to detect?




- ☛ Machine Learning was shown to be useful in this domain and thus to have the potential of saving human lives!



# How did it Work?

Consider the following common situation:

- ☛ You are in your car, speeding away, when you suddenly hear a “funny” noise.
- ☛ To prevent an accident, you slow down, and either stop the car or bring it to the nearest garage.



- ☛ The in-flight helicopter gearbox fault monitoring system was designed following the same idea. The difference, however, is that many gearbox malfunction cannot be heard by humans and must be monitored by a machine.



# So, Where's the Learning?

- Imagine that, instead of driving your good old battered car, you were asked to drive this truck:




- Would you know a “funny” noise from a “normal” one?
- Well, probably not, since you’ve never driven a truck before!
- While you drove your car during all these years, you effectively learned what your car sounds like and this is why you were able to identify that “funny” noise.



# What did the Computer Learn?

☞ Obviously, a computer cannot hear and can certainly not distinguish between a normal and an abnormal sound.

☞ Sounds, however, can be represented as wave patterns such as this one:  
which in fact is a series of real numbers.



☞ And computers can deal with strings of numbers!

☞ For example, a computer can easily be programmed to distinguish between strings of numbers that contain a “3” in them and those that don’t.



## (Cont' d)

- ☛ In the helicopter gearbox monitoring problem, the assumption is that functioning and malfunctioning gearboxes emit different noises. Thus, the strings of numbers that represent these noises have different characteristics.
- ☛ The exact characteristics of these different categories, however, are unknown and/or are too difficult to describe.
- ☛ Therefore, they cannot be programmed, but rather, they need to be learned by the computer.
- ☛ There are many ways in which a computer can learn how to distinguish between two patterns (e.g., decision trees, neural networks, bayesian networks, etc.) and that is the topic of this course!



# What else can Machine Learning do?

- ☛ Medical Diagnostic (e.g., breast cancer detection)
- ☛ Credit Card Fraud Detection
- ☛ Sonar Detection (e.g., submarines versus shrimps (!) )
- ☛ Speech Recognition (e.g., Telephone automated systems)
- ☛ Autonomous Vehicles (e.g., a vehicle drove unassisted at 70 mph for 90 miles on a public highway. Useful for hazardous missions)
- ☛ Personalized Web Assistants (e.g., an automated assistant can assemble personally customized newspapers)

☛ etc.....



# Required Text Books and Reading Material

- ☛ Ian Witten and Eibe Frank, *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd Edition, Morgan Kaufmann, ISBN 0120884070, 2005.
- ☛ Nathalie Japkowicz and Mohak Shah, *Evaluating Learning Algorithms: A Classification Perspective*, Cambridge University Press, 2011.
- ☛ Introduction to Machine Learning, Nils J. Nilsson (available (free) from the Web)
- ☛ Research papers (available from the Web. Please, see Syllabus for links).
- ☛ The syllabus also lists a number of non-required books that you may find useful. In particular, *Machine Learning*, Tom Mitchell, McGraw Hill, 1997 is the book on which most of my slides are based on.



# Objectives of the Courses:

- ☛ To present a broad introduction of the principles and paradigms underlying machine learning, including discussions and hands-on evaluations of some of the major approaches currently being investigated.
- ☛ To introduce the students to the reading, presenting and critiquing of research papers.
- ☛ To initiate the students to formulating a research problem and carrying this research through.



# Format of the Course:

- ☞ Each week will be devoted to a different topic in the field and a different theme.
- ☞ Part 1 of the lecture will be a presentation (by the lecturer) of the basics concepts pertaining to the weekly topic.
- ☞ Part 2 of the lecture will be a set of presentations (by 1, 2 or 3 students) on recent research papers written on the weekly theme.
- ☞ The last week of the term will be devoted to project presentations.



# What are your interests?

- ☞ For the weekly themes, I usually choose themes which are of current interests to the Machine Learning/Data Mining Community and assign the most important papers recently written on these themes.
- ☞ This year, I have decided to ask you, first, what your interests are. I will then try to select papers according to both your interests and the new developments in the field.



# Course Requirements:

- Weekly paper critiques (1 critique per teams of 2-3 students)
- 1-2 paper presentations
- Assignments (little programming involved as programming packages will be provided)
- Final Project:
  - Project Proposal
  - Project Report
  - Project Presentation

Percent  
of the  
Final  
Grade

20%

30%

50%



# More on Assignments (1):

## ☞ Assignment 1:

- Handed out on: January 24, 2012
- Due on: February 7, 2012
- Topic: Version Spaces, Decision Trees, Evaluation (2 classifiers on 1 domain or 2 classifiers on multiple domains)

## ☞ Assignment 2:

- Handed out on February 14, 2012
- Due on: March 6, 2012
- Topic: Neural Networks, Bayesian Learning, Evaluation (multiple classifiers on multiple domains)



# More on Assignments (2):

## ☛ Assignment 3:

- Handed out on: March 13, 2012
- Due on: March 27, 2012
- Topic: IBL, Rule Learning, SVMs, Classifier Combination



# Project (See Project Description on Course Web site)

- ☛ **Research Project** including a literature review and the design and implementation of a novel learning scheme or the comparison of several existing schemes.
- ☛ **Projects Proposal** (3-5 pages) are due on **February 14**
- ☛ **Project Report** are due on **April 3**
- ☛ **Project Presentations** will take place in **the last week of classes**
- ☛ Suggestions for project topics are listed on the Web site, but you are welcome (and that's even better) to propose your own idea.

**Start thinking about the project early!!!!**