

CS 491/691 Q: Machine Learning Context Learning Systems

Sushil J. Louis
SEM 233 (784-4315)
email: sushil@cs.unr.edu
<http://www.cs.unr.edu/~sushil/>

Description

Machine Learning (3+0) 3 credits

This is a heavy programming course with the overall objective of designing and implementing an integrated system that perceives, reasons, and acts. Survey of sound, vision, and other sensor techniques. Introduction to Artificial Intelligence and Machine Learning. Problem solving, search, and game trees. Knowledge representation. Decision trees, neural-nets, and genetic algorithms. Prerequisite: CS 308. Should have taken AI and OS.

Office Hours

- **Sushil:**

- Mondays: 9:30 - 11:30
- And by appointment (email sushil@cs.unr.edu for an appointment)

TextBook

- **Required:** Patrick H. Winston, *Artificial Intelligence*, 3rd Edition, Addison-Wesley.
- **Recommended:** Tom Mitchell, *Machine Learning*. McGraw-Hill.
- **Recommended:** Ian H. Witten, Eibe Frank, *Data Mining*, Morgan Kaufmann.

Preliminary Syllabus

This course combines AI, Machine Learning, and system building. The idea is to use your knowledge of AI and machine learning to build an integrated system that perceives the world through sensors (cameras, microphones, ...), reasons about the world using AI and ML algorithms, and chooses courses of action to

affect the state of the world. Think of your workstation as a stationary robot with no power, communication, or computational constraints and you will be in the right frame of mind for this course.

You will learn about and use public domain software to implement a useful intelligent system. There is a strong lab component to this course and we will hold a few classes in SEM 252 to go over available software and hardware.

1. Introduction

- What is AI?
- What is Machine Learning?
- What AI-useful software is available?

2. Search

- State spaces, Search, Representations
- DFS, BFS, NDS, Heuristic Search
- Hill Climbing, Beam Search
- Text-to-Speech. Festival speech synthesis system
- Linux usb camera support
- Motion. Detect motion from camera images
- Tree search: Minimax, Alpha-Beta pruning
- CMU Sphinx speech recognition system

3. Machine Learning

- Decision Trees: ID3
- Rule induction
- Neural approaches: Feedforward nets and backpropagation
- Evolutionary approaches: Genetic algorithms and classifier systems

Projects: Find an application that can be improved by knowing something about its context and improve it by designing and implementing a sensing system(s) that provide(s) the needed context. Here is a list of possible applications

- A calendar that reminds you of appointments using a text-to-speech toolkit *at appropriate times*, and uses popups at other times.
- A weather lookup system that tells (text-to-speech) you whenever the weather changes, unless someone is in the room with you - then it waits for the person to leave before speaking.
- Augment xmms (<http://www.xmms.org>) so that it lowers the volume when the phone rings.
- Augment xmms (<http://www.xmms.org>) so that it lowers the volume when someone is speaking.
- A system that warns people standing outside my office that I am about to open my door
- Compile and install CMU Sphinx, a speech understanding system that robustly understands a small vocabulary. See the class web page for links to this software. Use it to execute unix programs like mozilla, xterm, make and to understand “yes” and “no.” You can also work to use it to obtain feedback to better inform your system when it is “appropriate” to use speech and when it is not.

Other projects need my permission.

Homework, Exams, and Projects

There will be a number of programming assignments. No late assignments will be accepted. Use whatever programming language you like but note that I cannot help you much with languages that I do not know. There will be two exams and a final project. Some projects can be done in groups and I have provided the list of possible projects above. Projects will require a written project report and a poster at the end of the semester. This report and poster will be due Monday May 3 at 4:00 p.m. The first exam will be held about 1/3 through the course, the second about 2/3 through the course. Grades will be allocated as shown in the table below.

	Tentative date(s)	percentage
Exams	2/23/04 and 4/5/04	30%
Assignments		30%
Project	5/3/04	40%

This course will use the linux workstations in SEM252 as well as cameras, microphones, and other sensors that will be provided.

Colloquia

From time to time we will have invited speakers to give us a talk on their current research. The class will then meet at the colloquium location. I will notify the class when we have colloquia.

Communications

If I need to communicate with the class as a group I'll post a message on our web page. You are required to check the class web page and your email every day.

In addition to the textbook, there are number of resources available. Our class notes will be available on the WWW from

<http://www.cs.unr.edu/~sushil>

(follow the Machine Learning Link). I will also point out various other places that contain material of interest to us.

Cheating

Cheating is not permitted and will result immediately in a grade of “F.” and/or commencement of administrative proceedings. Please read the section on **Academic Standards** on page 62 of the University Catalog. That section defines cheating and specifies the consequences.