# CS491/591-001: Introduction to Machine Learning

Instructor: Terran Lane
Fall, 2004
Tue/Thu 4:00-5:15
ME218

## 1 Goals and Subjects

This class is intended to introduce you to the study of machine learning, both as it is currently understood and as it has developed over the past twenty years or so. We won't, unfortunately, have time to delve into a number of subjects in nearly as much depth as I would like, but we will attempt to introduce the key principles of ML and gain basic familiarity with the terminology, methodology, algorithms, and mathematical results. The goal is that, by the end of the course, you should know what techniques are available and how/when to apply them (or at least where to look to learn how), be prepared to read the literature independently, and, for graduate students, be ready to start engaging in research in the ML field. We'll be focusing on the statistical approach, which has emerged as the dominant viewpoint in the past ten years or so, though we'll also examine some symbolic and hybrid methods as time permits.

## 2 Textbook

The "official" (primary recommended) text for the class is:

> Witten, I. H and Frank, E., *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*, Morgan Kaufmann, 2000.

It didn't get ordered to the bookstore, but I believe that you can order it through them or your favorite online book retailer.

Some other very useful texts include:

> Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, 2001.
>
> Sutton, R. and Barto, A., *Reinforcement Learning*, MIT Press, 1998. (Available in full text online, if you look around a bit.)
>
> Hand, D., Mannila, H., and Smyth, P., *Principles of Data Mining*, MIT Press, 2001.
>
> Mitchell, T. M., *Machine Learning*, McGraw-Hill, 1997.

I will attempt to make these available on reserve at the library, and may also be able to make electronic excerpts available as well.

# 3 Class Resources (Me, the Web, etc.)

I can be reached most easily by email to `terran@cs.unm.edu`

I will, in general, hold office hours 9:00-11:00 AM on Wednesdays. I'm also available by appointment (send me mail).

Further information on the class (including the content of readings, specifics on homework assignments, etc.) will be made available via the web:

`http://www.cs.unm.edu/~terran/classes/cs591-f04/`

There is a mailing list for this class: `ml-class@cs.unm.edu`. I *strongly* suggest that you subscribe to this list, as it will be used for class administrative updates as well as for some discussion of class topics. (Go to `http://www.cs.unm.edu/cgi-bin/mailman/listinfo/ml-class` to subscribe.)

***Please, please, please,*** if you're having difficulties in the class, for whatever reason, contact me *early*. I really want all of you to enjoy and succeed in this class, but I can't help you do so if you don't talk to me. Preferably early enough to do something about it — two days before the final is probably too late make any difference! I'm happy to talk to you in office hours, or, if you're more comfortable with written communication, by email.

# 4 Assignments, Grading, and Handins

Your grade in this class will be based on four factors:

**(15%)** Participation in discussions — especially on the assigned readings.

**(20%)** Homework assignments.

**(30%)** Exams (2).

**(35%)** A final project.

Each of these will be discussed further below.

## 4.1 Readings

A major part of doing research in any field is reading the peer-reviewed literature, both previous and current. This is really the only way to understand what work has been done, what work is being done now, and what questions the community as a whole finds interesting, important, and difficult. It is vital both in order to move forward in a field and to avoid reinventing the wheel. As it turns out, learning to read research literature is itself a skill (and even a slightly different one for each discipline!). While a substantial part of this class will be based on the text, there will be a number of readings from previous and current literature.

The readings will be announced in class and on the web at least a week in advance. Near the start of the term, the class will be divided into reading groups, each of which will be responsible for reading and discussing the designated papers *in advance* and producing a short (roughly one page) written critique of the readings (due at the *beginning* of the corresponding class). In addition, each group should produce a list of questions and/or observations on the readings that will serve as

a basis for class discussion. The written critiques and questions will contribute to the participation part of your grade.

Note that the readings may or may not be synchronized with class topics and you may or may not have the full background to read the papers. In part, your group members will have different skill sets, so you may be able to help each other through the paper. But you will often find that none of you know enough to fully comprehend a paper. I *do* expect you to take some initiative and do some of your own background reading to fill out your knowledge. You can ask me questions, of course, but you'll also find some great background material in the course text and in the alternate references I list above.

Last note on readings: I have a number of target topics and papers in mind, but I'll welcome suggestions. If there's a topic that you're curious about (and especially if you have a suggested paper or two), I'll be happy to consider it.

## 4.2 Homework Assignments

There will be a small number of homework assignments, which will include a mixture of mathematical exercises, informal analysis, and short programming assignments. Homework is due at the *beginning* of the indicated class session. I encourage you to *discuss* homework problems with your classmates, but each individual is responsible for *doing* and *writing up* her or his own solution to the homework.

## 4.3 Exams

There will be a midterm and a final exam. These will be in-class, individual tests.

## 4.4 The Final Project

The major part of your grade will be based on a final project of substantial content. These are to be individual projects, though in exceptional circumstances I may consider a request to do a group project. The intent of the project is for individuals to explore one or more topics in greater depth than will be possible in the lecture/readings. More detail on the project will be provided later, but for the moment keep an eye out for a topic of interest.

## 4.5 Program Code and Languages

When an assignment (including the final project) includes programming, you should turn in a full copy of your code with the rest of the assignment.

What language you use is entirely up to you, unless I explicitly specify otherwise. I may suggest a language for a particular assignment, but that's advice, not requirement. In general, the motto is: "use the right tool for the job." That said, please honor the *spirit* of the assignment — if the assignment calls for you to *implement* a neural network, for example, *don't* just use the Matlab NN toolbox. Similarly, don't just go find publicly available code for any learning system I specify — your job is to understand how these things work from the inside out, and plugging

in a blackbox library doesn't accomplish that. If you have a question about whether some package/library is acceptable for use in an assignment, *please ask me* before you use it (or before you decide *not* to use it — it doesn't hurt to ask, and you may save yourself a lot of pain).

### 4.6 Electronic Handins

If you wish to submit homeworks or reading group summaries electronically, you are welcome to do so. Please format your submission reasonably nicely and send it in either PDF, PostScript, HTML, or plain ASCII. You may send me MIME attachments, Zip-compressed archives, or tar/gzip archives if you wish. **I DO NOT ACCEPT MS-WORD DOCUMENTS.** If you send me an archive (.zip or .tar.gz), *please* set it up to create a subdirectory for its files, named with your last name and the assignment (e.g., "`lane_hw3`"). This prevents your files from clobbering someone else's files.

Your submission must arrive in my mailbox (as measured by my computer's clock) by the beginning of class on the specified day.

## 5 Other Policies

Now for some class policies that I hope won't be an issue, but are useful to state up front nonetheless.

### 5.1 Late Handins

As stated above, homework and readings critiques are due at the *beginning* of class. Critiques won't be accepted late at all — the point is to be prepared for that day's discussion. Homeworks will be accepted up to one day late for a 50% penalty. The final project will not be accepted late. Extensions will be made only for exceptional circumstances, e.g., illness, family emergency, or alien abduction (all alien abductions must be fully and irrefutably documented). Please contact me as soon as you are able (preferably, in advance) if you need to seek an extension.

### 5.2 Academic Dishonesty (and Legitimate Collaborations)

It shouldn't be necessary to mention this in a senior/graduate-level course, but the issue has come up in the past so I feel that I should state my policies clearly up front.

Dishonest behaviors, including but not limited to plagiarism, copying of another student's work (or providing your own to another), group consultation on individual projects or work, etc., will not be tolerated. My general feeling is that being caught cheating should be more painful than not having done the assignment at all. Therefore, I will generally *at least* assign a negative penalty equal to the full value of the assignment if I discover someone cheating on an assignment. I.e., if an assignment is worth 10% of the final grade, the individual would receive not zero credit for the assignment, but *-10%*. And, of course, if you're clever enough not to get caught, then you're clever enough not to need to cheat in the first place.

That said, this *is* an advanced course and I do want to encourage collaborative thinking and discussion. Therefore, I encourage you to discuss a number of the class assignments. Specifically:

- The reading assignments are intended to be group discussion work. I encourage (and even *require*) you to discuss them among yourselves before coming to class. The corresponding critique and list of questions is also a group assignment, and should include the names of everyone who participated. (Please, if you didn't participate in developing critique, don't put your name on it.)

- I am happy for you to *discuss* homework assignments and the final project, though each individual is to write his or her own programs, documentation, reports, mathematical results, perform independent experiments, etc. Program code and/or experimental results may not be shared.

- Exams are, of course, completely individual work.

Most importantly, if you have a question about whether some specific collaboration is permissible, *please ask me before you do it*.

# 6 A Final (Initial?) Note

I look forward to working with you this semester. This is a thrilling field to me that has undergone dramatic change even in the past few years. I hope that you'll find it as fascinating as I do!

# 7 Miscellanea

- Staples are your friend.

- The spell checker is also a friend, albeit a frequently deceptive one. Rely on it with caution.

- Gauss was a real person, so things derived from his name should be captialized. E.g., Gaussian distribution, Gauss-Jordan elimination, Gauss integral, etc. Similarly, Markov, Laplace, Dirichlet, etc. should all be properly capitalized.

- Plots should always be fully labeled — both axes, and a title indicating content. If the plot contains more than one data set, please distinguish the different data sets by symbol or line type (e.g., '+' vs 'o' or ':' vs '–') and include a legend that names each. Finally, be sure to distinguish *discrete* from *continuous* data in your plots. E.g., if you have a sample of 10 measured points, do *not* simply draw a curve through them. At the very least, indicate the location of each point with a discrete symbol, such as a cross or square, superimposed on the curve. Better to omit the curve altogether as it implies the presence of unmeasured points. (Unless, of course, you have good reason to draw such a curve, and make it clear that it's interpolated rather than measured.)

- I expect proper grammar, spelling, punctuation, etc., on all assignments. And it wouldn't hurt to brush up on it in email too. `;-)` Being a computer hacker does not excuse you from natural language skills!

# Schedule

The following schedule is possibly optimistic, but is definitely tentative and subject to revision, depending on how we progress. Ideally, we'll touch on all of the topics given here (and maybe some additional), but we may dwell longer on some, in the interest of improved understanding, at the cost of neglecting others.

**Weeks 1–2** Introduction; basic concepts and definitions; examples; empirical methodology.

**Weeks 3–5** Introductory classification and regression; decision trees; classification by linear machines; support vector machines; linear regression.

**Weeks 6–8** Introduction to generative classification; Bayes' rule; principles of Maximum Likelihood and MAP; graphical models.

**Weeks 9–10** Unsupervised learning; mixture models; expectation-maximization; graphical models with hidden variables.

**Week 11** Time series models; Markov chains; hidden Markov models.

**Weeks 12–14** Reinforcement learning and planning; Markov decision processes; POMDPs; Q-learning; $E^3$, Reinforce.

**Week 15** Final project presentations.

**Final Exam Week** The final exam will be in ME218 on Tues, Dec 14 at 5:30-7:30 PM.