CS 5804: Introduction to Artificial Intelligence

Graduate level course in AI with coverage of search, knowledge representation and reasoning, machine learning (paradigms, models, and algorithms), use of knowledge in learning, and AI applications. The emphasis of the course is on recent developments in AI (i.e., beyond monkeys and bananas), especially contributions that forged novel connections among diverse areas, or addressed problems of significant impact. I have had to make some tough calls in deciding which material to include (exclude), so the course is not intended to be comprehensive. The goal is to emphasize certain thematic issues that recur in AI systems and applications.

Meeting Times	MW 5:30-6:45pm, McBryde 233
Instructor	Naren Ramakrishnan, 1-8451, Torg 2160L
	naren@cs.vt.edu, http://www.cs.vt.edu/∼naren
Office Hours	Tuesdays 1-5pm, or walk in any time.
Listserv	CS5804_91718@listserv.vt.edu
	(yes, the name is rather long winded.)
Course Web Page	http://courses.cs.vt.edu/∼cs5804

If you are unable to make the above times and need to meet with us, you can setup an alternative time via email. If you need adaptations or accommodations because of a disability (learning disability, attention deficit disorder, psychological, or physical), if you have emergency medical information to share with the instructor, or if you need special arrangements in case the building must be evacuated, please meet with the instructor ASAP.

Pre-requisites: There are really no formal pre-requisites beyond graduate student standing. You are expected to have basic knowledge of probability, statistics, must have taken undergraduate courses in CS theory (algorithms, NP-completeness), and must not be averse to math. Knowledge of propositional and predicate logic or experience with PROLOG will be beneficial. I do not expect that you have taken an undergraduate course in AI.

Evaluation: There will be 8 homeworks, which will involve a mix of theoretical problems, programming assignments, and questions that will focus on your (surprise!) writing skills. The topical content of the programming assignments will usually be language-independent, so you will be free to use your favorite platform/language. No late submissions will be accepted. There will be a midterm exam and a final (both closed book and closed notes). Detailed breakdown: homeworks (50%), midterm exam (20%), final (30%).

All homeworks are designed by the instructor. In addition, the instructor grades both the exams individually. If you have an exam or homework that you feel has been graded incorrectly, please contact us, and we can discuss a re-grading if appropriate.

Keeping in Touch: Please use the listserv actively for discussions and exchanging ideas. Since it is created automatically by a central university system, any student registered in CS 5804 will be added to the mailing list. If you do not receive a test mail from the instructor by the end of the first week of classes, ensure that your email address is properly recorded in the university system.

Workload: The course moves at a very fast pace! I assume that this is your most intensive course this semester and that you relish 2-3 sleepless nights per week (I did, when I was a graduate student). Most assignments involve a fair amount of design, so plan your schedule accordingly.

Electronic Accounts and Programming: You are expected to have accounts on the CS graduate lab network or some other equivalent facility. Familiarity with high-level programming is expected, in an operating system of your choice. You are also expected to be a good coder, and to choose appropriate data structures and algorithm design strategies. Programs in AI can get unwieldly and sound choice of data structures can be useful in completing the homeworks.

Book: There is really no textbook that covers all relevant themes in the required detail. Being graduate students, you are expected to be able to cull ideas from research papers and form your own conceptual model of what modern AI is about. In short, no babysitting. Still, Russell and Norvig's *AIMA* is a classic book in coverage and will help you to keep track of the larger picture. We will follow a suggested sequence from the book and at periodic times, drill down into research literature to understand ideas in detail. We will also cover most of the material from the *RL* book (*Reinforcement Learning: An Introduction*, by Sutton and Barto). The *NDP* book (see below) is recommended. The formal list of readings (this is bound to change in the future) is:

AIMA S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, Upper Saddle River, NJ, 1995.

RL R.S. Sutton and A.G. Barto, *Reinforcement Learning*, MIT Press, 1998.

NDP D.P. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, Sep 1996.

P1 V. Kumar, Algorithms for Constraint Satisfaction: A Survey, *AI Magazine*, Vol. 13, No. 1, pages 32-44, Spring 1992.

P2 B. Selman, H. Levesque, and D. Mitchell, A New Method for Solving Hard Satisfiability Problems, in *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI'92)*, pages 440-446, July 1992.

P3 G. Tesauro, Temporal Difference Learning and TD-Gammon, *Communications of the ACM*, Vol. 38, No. 3, pages 58-68, March 1995.

P4 W. Zhang and T.G. Dietterich, A Reinforcement Learning Approach to Job-Shop Scheduling, in *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1114-1120, 1995.

P5 D. Subramanian, P. Druschel, and J. Chen, Ants and Reinforcement Learning: A Case Study in Routing in Dynamic Networks, In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 832-839, 1997.

(These papers will be linked in from the course web page. Please stay tuned to the course web page for additions and changes!)