UNDERGRADUATE ARTIFICIAL INTELLIGENCE DAY

Presentations 4:00-6:00pm BA1130

4:15 - 4:35pm: Knowledge and Representation	**Sheila McIlraith and Tyler Lu**
4:35 - 4:55pm: Computer Vision	**David Fleet and Fernando Flores-Mangas**
4:55 - 5:15pm: Computational Biology	**Michael Brudno and Orion Buske**
5:15 - 5:35pm: Machine Learning	**Rich Zemel and Danny Tarlow**
5:35 - 5:55pm: Computational Linguistics	**Gerald Penn and Katie Fraser**

Reception 6:00-7:30pm BA3200

Talk to graduate students: ask them about their research, courses, grad school, etc. And, there's free food!

Interested in A.I.? Where to go from here:

Get involved in research

Note: graduate schools look closely at undergraduate research experience when choosing who to accept, so you might want to start looking for research opportunities early on. Or, perhaps you're just keen on learning something new, seeing what a career in research would be like, and exploring the opportunities out there:

- Apply for a paid summer internship (NSERC USRA) with one of the CS profs (applications for this year will be out mid-March)
- Not eligible for NSERC? (e.g. International Student), apply for the University of Toronto Excellence Awards (UTEA): applications for this year will be out mid-March, similar to NSERC
- Enrol in a project course (CSC494/495): find out how by visiting the Computer Science Undergraduate Office (or e-mail: ug@cs.toronto.edu) OR e-mail the professor of your choice and ask how you can participate in their research!

The A.I. faculty at UofT are doing some fascinating work. Don't be afraid to approach them and talk about it!

Join the Undergraduate Artificial Intelligence Group (UAIG)

The Undergraduate Artificial Intelligence Group aims to bring together University of Toronto students interested in the various sub-fields of Artificial Intelligence. Weekly meetings serve as a breeding ground for ideas and help to stimulate discussion. The theme of the group differs from year to year, depending on the interests of students involved.
For information about upcoming meetings join the mailing list by emailing uoft.uaig@gmail.com or check out UAIG's website, www.uaig.wikispaces.com.

Attend other A.I. related talks

Liked this seminar, and looking for more like it? There are tons of A.I.- related seminars that take place all throughout the year – you just have to know where to look:

- Distinguished Lecture Series (open to everyone): all talks are at 11:00 a.m., in the Bahen Centre, Room 1180. Check the DCS website for 2012-2013 dates (http://web.cs.toronto.edu/news/lectures.htm)
- Research in Action Showcase: the CS department showcases its projects at the beginning of April (stay tuned)
- Mailing Lists: ask to get added if you want to get e-mails about upcoming seminars:

Group	*Seminar schedule*	*Contact*
Knowledge and Representation	**Wed, 12pm-2pm**	**edelisle@cs.toronto.edu**
Computer Vision	**Fri, 11am-12pm**	**http://www.cs.toronto.edu/vis/seminars/**
Computational Linguistics	**Wed, morning**	**cl-mo@cs.toronto.edu**
Machine Learning	**Mon/Thurs, 11am-12pm**	**dtarlow@cs.toronto.edu**
Computational Biology	**Wed, 10am**	**misko@cs.toronto.edu**

Also, many faculty members have their own individual group meetings. If you are interested in attending those, email the professor!

A.I. related courses offered by the Department of Computer Science

Note: you don't have to be in third or fourth year to take these courses – as long as you have the prerequisites (or can get them waived), you're all set!

CSC321 Introduction to Neural Networks and Machine Learning: The first half of the course is about supervised learning for regression and classification problems and will include the perceptron learning procedure, backpropagation, and methods for ensuring good generalisation to new data. The second half of the course is about unsupervised learning methods that discover hidden causes and will include K-means, the EM algorithm, Boltzmann machines, and deep belief nets.

CSC411 Machine Learning and Data Mining: An introduction to methods for automated learning of relationships on the basis of empirical data. Classification and regression using nearest neighbour methods, decision trees, linear models, and neural networks. Clustering algorithms. Problems of overfitting and of assessing accuracy. Problems with handling large databases.

CSC412 Probabilistic Learning and Reasoning: An introduction to probability as a means of representing and reasoning with uncertain knowledge. Qualitative and quantitative specification of probability distributions using probabilistic graphical models. Algorithms for inference and probabilistic reasoning with graphical models. Statistical approaches and algorithms for learning probability models from empirical data. Applications of these models in artificial intelligence and machine learning.

CSC310 Information Theory: Measuring information. The source coding theorem. Data compression using ad hoc methods and dictionary-based methods. Probabilistic source models, and their use via Huffman and arithmetic coding. Noisy channels and the channel coding theorem. Error correcting codes, and their decoding by algebraic and probabilistic methods.

CSC320H1 Introduction to Visual Computing: A unified introduction to image synthesis and image analysis aimed at students with an interest in computer graphics, computer vision or the visual arts. Focus on three major topics: (1) visual computing principles - computational and mathematical methods for creating, capturing, analyzing and manipulating digital photographs (raster algorithms, image acquisition, basic image processing, image warping, anti-aliasing); (2) digital special effects - applying these principles to create special effects found in movies and commercials; (3) visual programming - using C/C++ and OpenGL to create graphical user interfaces for synthesizing and manipulating photographs.

CSC420 Introduction to Image Understanding: Introduction to basic concepts in computer vision. Extraction of image features at multiple scales. Robust estimation of model parameters. Multiview geometry and reconstruction. Image motion estimation and tracking. Object recognition. Topics in scene understanding as time permits.

CSC418 Computer Graphics: Identification and characterization of the objects manipulated in computer graphics, the operations possible on these objects, efficient algorithms to perform these operations, and interfaces to transform one type of object to another. Display devices, display data structures and procedures, graphical input, object modelling, transformations, illumination models, primary and secondary light effects; graphics packages and systems. Students, individually or in teams, implement graphical algorithms or entire graphics systems.

CSC384 Introduction to Artificial Intelligence: Theories and algorithms that capture (or approximate) some of the core elements of computational intelligence. Topics include: search; logical representations and reasoning, classical automated planning, representing and reasoning with uncertainty, learning, decision making (planning) under uncertainty. Assignments provide practical experience, both theory and programming, of the core topics.

CSC486 Knowledge Representation and Reasoning: Representing knowledge symbolically in a form suitable for automated reasoning, and associated reasoning methods: first-order logic, entailment, the resolution method, Horn clauses, procedural representations, production systems, description logics, inheritance networks, defaults and probabilities, tractable reasoning, abductive explanation, the representation of action, planning.

CSC330 Logical Specifications: Logic and its use as a declarative language in computer science. Syntax and semantics of propositional and predicate calculus. Proving entailment and non-entailment rigorously. Formal derivations. Satisfiability. Applications, including information systems, program verification, artificial intelligence, software engineering. Computational tools, including Prolog. Other logics.

CSC485 Computational Linguistics: Computational linguistics and the understanding of language by computer. Possible topics include: augmented context-free grammars; chart parsing, ,statistical parsing; semantics and semantic interpretation; ambiguity resolution techniques; discourse structure and reference resolution. Emphasis on statistical learning methods for lexical, syntactic and semantic knowledge.

CSC401 Natural Language Computing: Introduction to techniques involving natural language and speech in applications such as information retrieval, extraction, and filtering; intelligent Web searching; spelling and grammar checking; speech recognition and synthesis; and multi-lingual systems including machine translation. N-grams, POS-tagging, semantic distance metrics, indexing, on-line lexicons and thesauri, markup languages, collections of on-line documents, corpus analysis. PERL and other software.

Additional note: you don't have to be a graduate student to take some amazing graduate-level courses. As long as you have the necessary background, you can petition to get into a graduate course (find out how by talking to the Computer Science Undergraduate Office). The graduate-level course descriptions are available here:
http://web.cs.toronto.edu/program/gc/2011-2012_Course_Descriptions.htm