

# Doina Precup

School of Computer Science  
3480 University st.  
McGill University  
Montreal, Quebec, H3A 2A7

dprecup@cs.mcgill.ca  
<http://www.cs.mcgill.ca/~dprecup>  
Fax: (514) 398-3883  
Phone : (514) 398-6443

---

**RESEARCH INTERESTS** Artificial Intelligence; Machine Learning; Reinforcement Learning; Markov Decision Processes; Planning and Scheduling; Reasoning Under Uncertainty; Applications of Machine Learning and Artificial Intelligence.

**CURRENT POSITION** Tenure-track assistant professor, School of Computer Science, McGill University.  
Since July 2000.

**EDUCATION**

**University of Massachusetts, Amherst**  
Ph.D. in Computer Science  
Dissertation: *Temporal Abstraction in Reinforcement Learning*  
Advisor: Dr. Richard S. Sutton May 2000

**University of Massachusetts, Amherst**  
M.Sc. in Computer Science, GPA 4.0/4.0 February 1997

**Technical University Cluj-Napoca, Romania**  
Advanced Studies (M.S.) in Computer Science and Engineering, GPA 10/10 July 1995

**Technical University Cluj-Napoca, Romania**  
B.Sc. in Computer Science and Engineering, Magna cum Laudae, GPA 10/10 July 1994

**HONORS**

Graduate Fellowship, University of Massachusetts, Amherst, 1997-1998 (one of 3 recipients of the highest level of funding)

Fulbright scholarship for study towards a M.S. degree, 1995-1997

Magna cum Laudae (highest) distinction for B.S. graduation

Merit scholarship from the Soros foundation, 1994

Tiriac scholarship from Mercedes Benz, 1994-1995

Merit scholarship from the Romanian government, 1990-1995

Several prizes received as a high school and college student in Romania, in the national computer science, mathematics and physics contests (including first prize and best paper)

award in the national computer science contest for high school students, 1986, second prize in the national physics contest for high school students, 1985, a special prize in the national physics contest for high school students, 1987, and a special prize in the national mathematics contest for college students, 1990).

High school valedictorian, Informatics High School, Cluj-Napoca, Romania

## RESEARCH EXPERIENCE

**McGill University, School of Computer Science** Montreal, Canada  
*Assistant Professor* Since July 2000

I am conducting research in supervised learning and reinforcement learning. A complete list of the research grants and students supervised is included below.

**AT&T Labs - Research** Florham Park, NJ  
*Summer Manager* May 1999 - September 1999

I elaborated new algorithms for multi-step off-policy learning. I also worked on applying reinforcement learning to an interactive robot, which uses different input modalities (speech recognition, text, vision) and different output modalities (speech, screen display, web page display) and learns human preferences for different modes of interaction.

**AT&T Labs - Research** Florham Park, NJ  
*Research Consultant* September 1998 - December 1998

I worked on studying the theoretical properties and testing new algorithms for learning about temporally extended actions in reinforcement learning.

**Computer Science Department, University of Massachusetts Amherst**  
*Research Assistant, Adaptive Networks Laboratory* September 1996 - August 1998

I participated in the elaboration of a new framework for describing temporally extended actions in reinforcement learning. In particular, I studied ways in which such actions can be modeled, and I designed new algorithms for learning the models and values of such actions.

## GRANTS

*Abstraction in Reinforcement Learning.*  
NSERC Individual Research Grant, \$27,000(CAD) per year, 2001–2005.

*Hierarchy and Structure in Reinforcement Learning.*  
FCAR Nouveau Chercheur Individual Grant, \$15,000(CAD) per year, 2001–2004.

*Temporal Abstraction in Reinforcement Learning.*  
Research Development Fund from the Faculty of Graduate Studies and Research, McGill University, 2000–2001, in the amount of \$18,333(CAD).

*Temporal Abstraction in Reinforcement Learning.*  
Startup grant from the School of Computer Science, McGill University, 2000–2001, in the amount of \$31,667(CAD).

*Trajectory Simulator for Mobile Robots Guided by Fuzzy Logic.*  
Research Project funded by the MADIR foundation, 1994, in the amount of \$800(U.S.). The research was completed during the same year (joint work with Tudor Precup and Catalin Sipos), and received a Certificate of Achievement from MADIR. The results have been published in several papers, mentioned above.

TEACHING  
EXPERIENCE**School of Computer Science, McGill University***Assistant Professor*

Since July 2000

Courses:

- *308-424: Topics in Artificial Intelligence.*  
I taught this undergraduate class in the Fall 2000 semester (enrollment: 70 students), and I will teach it again in the Fall 2001 semester.
- *308-766: Machine Learning.*  
This is a new graduate-level course that I introduced at McGill during the Winter 2001 semester (enrollment: 7 students) and I will teach it again in the Fall 2001 semester (listed 308-761A). The course has been approved by McGill as 308-652 (starting Fall 2002).
- *308-526: Probabilistic Reasoning in Artificial Intelligence.*  
This is a new course offered for graduate and advanced undergraduate students, which has been approved by McGill University. I proposed this class, and I will teach it for the first time in the Spring 2002 semester.

**Computer Science Department, University of Massachusetts at Amherst***Teaching Assistant*

Spring 2000

I graded assignments and taught the discussion section of CMPSCI 320, "Introduction to software engineering", a core course for the undergraduate curriculum. The course was taught in Java and covered basic topics in software engineering and object-oriented program design. The programming component of the course is covered solely in the discussion section. I was the only one designing the content of this section, providing programming examples and assignments for the course. We covered a large variety of topics, from basic concepts such as inheritance and modularity, to advanced Java features, such as RMI, user interfaces and threads.

**Computer Science Department, University of Massachusetts at Amherst***Teaching Associate*

Fall 1998

I was a lecturer for one of the sections of CMPSCI 187, "Programming with Data Structures", a core course for the undergraduate curriculum. The course was taught in Java for the first time, so in addition to the usual teaching and grading duties, I participated in developing the curriculum, designing programming examples, assignments and exams.

**Computer Science Department, University of Massachusetts at Amherst***Guest lecturer*

Fall 1998

I gave a couple of lectures during the reinforcement learning class, taught by Prof. Andrew G. Barto.

**Computer Science Department, Technical University Cluj-Napoca**

Romania

*Lecturer*

March 1995 - July 1995; position was on hold until July 2000

I taught laboratory sections for three courses: Artificial Intelligence, Programming Languages, and Compilers. Besides lecturing, my duties included grading of assignments and projects. I also advised a graduation (B.Sc.) project in 1995.

**PROGRAMMING Transactions and Database Systems Ltd**

Cluj-Napoca, Romania

EXPERIENCE *Analyst/programmer*

1991-1995

I developed database applications for a science library and for human resource management for several industrial clients.

STUDENT  
SUPERVISION

I am currently supervising the following graduate students:

1. Bohdana Ratitch (co-supervision with Prof. Denis Therien).  
PhD candidate (passed her proposal defense on December 2 2000).  
Thesis title: Study of Reinforcement Learning Methods with Generalization Capabilities.
2. Danielle Azar (co-supervision with Prof. Sue Whitesides and Prof. Haroun Sahraoui).  
PhD candidate (passed her proposal defense in December 2000).  
Thesis title: Genetic Algorithms for Building and Improving Software Quality Estimation Models
3. Francois Rivest (co-supervision with Prof. Thomas Schultz, Psychology Department).  
MSc student.  
Thesis title: Knowledge-Based Methods for Constructing Neural Networks.
4. Ron Singer.  
MSc student  
Thesis title: Information Extraction from Text using Hidden Markov Models.
5. Yannick Daoudi (co-supervision with Prof. Michael Hallett)  
MSc student  
Thesis title: Neural Network Based Prediction System for Classifying Golgi Resident Proteins.
6. John Boadway  
Incoming MSc student  
Thesis title: Reinforcement Learning for Intelligent Trading Agents

I have also supervised several undergraduate students:

1. David Airapetian, NSERC summer intern (Summer 2001).  
Project title: Reinforcement learning applied to function inlining in Java.
2. John Boadway, research assistant (Summer 2001).  
Project title: Reinforcement learning applied to autonomous trading agents.
3. Muhammad Haroon Saeed, research assistant (Summer 2001).  
Project title: Reinforcement learning in multi-agent systems.
4. Mani Ghasemlou, honors student (Summer 2001).  
Project title: ALife simulation with perceptron agents..
5. Pablo Samuel Castro, honors student (Winter 2001).  
Project title: Computer music improvisation using reinforcement learning.
6. Martin Stolle, honors student (Winter 2001).  
Project title: Discovering macro-actions in reinforcement learning.
7. Adeline Asgaraly, honors student (Fall 2000).  
Project title: Reinforcement Learning for Air Traffic Control.

PUBLICATIONS **Refereed Journal Publications**

Sutton, R. S., Precup, D., Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. In *Artificial Intelligence*, vol 112, pp.181-211. An earlier version appeared as Technical Report UM-CS-1998-74, Department of Computer Science, University of Massachusetts, Amherst, MA 01003.

Precup, D., Utgoff, P.E. (To appear). Classification using  $\Phi$ -machines and constructive function approximation. Accepted for publication in *Machine Learning*, pending revisions.

**Refereed Conferences**

Letia, I.A., Precup, D., Craciun, F. (2001). “Developing Collaborative Golog Agents by Reinforcement Learning”. To appear in *Proceedings of the Thirteenth IEEE International Conference on Tools with Artificial Intelligence (ICTAI-2001)*. IEEE Computer Society.

Precup, D., Sutton, R.S., Dasgupta, S. (2001). “Off-Policy Temporal-Difference Learning with Function Approximation”. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, pp. 417–424. Morgan Kaufman.

Precup, D., Sutton, R. S., Singh, S. (2000) “Eligibility Traces for Off-Policy Policy Evaluation”. In *Proceedings of the Seventeenth Conference on Machine Learning (ICML 2000)*, pp. 759-766. Morgan Kaufman.

Sutton, R. S., Singh, S., Precup, D., Ravindran, B. (1999). “Improved Switching among Temporally Abstract Actions”. In *Advances in Neural Information Processing Systems 11 (Proceedings of NIPS’98)*, pp. 1066-1072. MIT Press.

Sutton, R. S., Precup, D., Singh, S. (1998). “Intra-Option Learning about Temporally Abstract Actions”. In *Proceedings of the Fifteenth International Conference on Machine Learning, ICML’98*, pp. 556-564. Morgan Kaufman.

Precup, D., Utgoff, P.E. (1998). “Classification using  $\Phi$ -machines and Constructive Function Approximation”. In *Proceedings of the Fifteenth International Conference on Machine Learning, ICML’98*, pp.439-444. Morgan Kaufman. An earlier version appeared as Technical Report UM-CS-1997-005, Department of Computer Science, University of Massachusetts, Amherst, MA 01003.

Precup, D., Sutton, R. S., Singh, S. (1998). “Theoretical Results on Reinforcement Learning with Temporally Abstract Options”. In *Machine Learning: ECML-98. 10th European Conference on Machine Learning, Chemnitz, Germany, April 1998. Proceedings*, pp. 382-393. Springer Verlag.

Precup, D., Sutton, R. S. (1998). “Multi-Time Models for Temporally Abstract Planning”. In *Advances in Neural Information Processing Systems 10 (Proceedings of NIPS’97)*, pp. 1050-1056. MIT Press.

Moss, J. E. B., Utgoff, P. E., Cavazos, J., Precup, D., Stefanovic, D., Brodley, C. E., Scheeff, D. T. (1998). “Learning to Schedule Straight-Line Code”. In *Advances in Neural Information Processing Systems 10 (Proceedings of NIPS’97)*, pp.929-935. MIT Press.

Precup, D., Sutton, R. S. (1997) “Exponentiated Gradient Methods for Reinforcement Learning”. In *Proceedings of the 14th International Conference on Machine Learning, ICML'97*, pp.272-277. Morgan Kaufman.

Letia, I.A., Precup, D. (1995). “Knowledge Transfer when Learning a Second Programming Language”. In *Proceedings of the 6th IFIP World Conference on Computers in Education*, pp.97-106. Chapman and Hall.

### **Refereed Workshops and Symposia**

McGeoch, C.C., Precup, D., Cohen, P.R. (1997). “How to Find Big-Oh in Your Data Set (and How Not To)”. In *Advances in Intelligent Data Analysis: Reasoning about Data. Proceedings of the Second International Symposium on Intelligent Data Analysis, IDA-97*, pp. 41-52. Springer Verlag.

Precup, D., Sutton, R. S., Singh, S. (1997). “Planning with Closed-Loop Macro Actions”. In *Working Notes of the AAAI Fall Symposium '97 on Model-directed Autonomous Systems*, pp. 70-76.

Precup, D., Sutton, R. S. (1997) “Multi-Time Models for Reinforcement Learning”. In *Proceedings of the ICML'97 Workshop on Modeling in Reinforcement Learning*.

Precup, D., Precup, T. (1995). “Trajectory Simulation and Optimization for Fuzzy Controlled Mobile Robot”. In *Proceeding of the 3rd IFAC/IFIP/IFORS Workshop, Intelligent Manufacturing Systems, IMS'95, Bucharest, Romania. Preprints*, pp. 39-43. Also accepted for publication in the post-prints, edited by Elsevier Science.

Precup, D. (1995). “An Analogical Model of the Student for Learning a Second Programming Language”. In *Proceedings of the International Workshop on Intelligent Computer Communication ICC'95, Cluj-Napoca, Romania*, pp. 123-130.

Precup, T., Handra-Luca, V., Precup, D., Sipos, C. (1994). “Robot Simulation Using Object-Oriented Programming”. In *Proceedings of the International Conference on Technical Informatics CONTI'94, Timisoara, Romania, Vol. 2*, pp.42-51.

### **Book Chapters**

Utgoff, P.E., Precup, D. (1998). “Constructive function approximation.” In *Feature extraction, construction, and selection: A data-mining perspective*, Motoda and Liu (Eds.), pp. 219–235. Kluwer.

### **Unrefereed Publications**

McGovern, A., Precup, D., Ravindran, B., Singh, S., Sutton, R. S. (1998). “Hierarchical Optimal Control of MDPs”. In *Proceedings of the Tenth Yale Workshop on Adaptive and Learning Systems*, pp.186-191.

Maties, V., Precup, T., Precup, D., Sipos, C. (1994). “Simulation of a Fuzzy-Guided Mobile Robot in a Static Environment”. In *Proceedings of the National Conference on Systems Theory, Robotics and Automatic Control SINTES 7, Craiova, Romania*, pp.203-208.

Precup, D., Sipos, C., Precup, T. (1994). "Simulation of Automatic Fuzzy Controllers for Robot Guidance". In *Proceedings of the Basis of Electronics Workshop, Cluj-Napoca, Romania*, pp.58-63.

### Technical reports which do not overlap with previous publications

Perkins, T.J., Precup, D. (1999). "Using Options for Knowledge Transfer in Reinforcement Learning". *Technical Report UM-CS-99-34*, Department of Computer Science, University of Massachusetts, Amherst, MA 01003.

Utgoff, P.E., Precup, D. (1997). "Relative Value Function Approximation". *Technical Report UM-CS-1997-003*, Department of Computer Science, University of Massachusetts, Amherst, MA 01003.

### Other Documents

Temporal Abstraction in Reinforcement Learning, PhD Dissertation, University of Massachusetts, Amherst, MA, 2000.

Using Analogic Reasoning for Natural Language Semantic Processing, "Advanced studies in computer science" MScCSE graduation project), Technical University Cluj-Napoca, Romania, 1995.

Student Modeling for Learning a Second Programming Language, BScCSE graduation project, Technical University Cluj-Napoca, Romania, 1995.

## TALKS

### Invited Talks - Universities and Research Centers

*Off-Policy Temporal Difference Learning with Function Approximation*  
Department seminar, Department of Computer Science and Operations Research, University of Montreal (Host: Yoshua Bengio). March 27, 2001

*Reinforcement Learning in Artificial Intelligence*  
McGill Cognitive Science / TUNGS Seminar. March 16, 2001

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
NEC Research, Princeton, NJ (Host: Eric Baum). May 3, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
University of California at San Diego (Host: Charles Elkan). April 21, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
Washington University, St. Louis (Host: Catalin Roman). April 14, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
University of Colorado, Boulder (Host: Clayton Lewis). April 4, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
University of Michigan at Ann Arbor (Host: Michael Wellman). March 30, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
Gatsby Research Institute, London, UK (Host: Peter Dayan). March 23, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
University of Pittsburgh (Host: Martha Pollack). March 13, 2000

*Options: A Framework for Temporal Abstraction in Reinforcement Learning*  
McGill University, Montreal (Host: Denis Therien). February 17, 2000

*Eligibility Traces for Off-Policy Learning*  
AI and Statistics Seminar, AT&T Research Labs August 28, 1999

*Intra-option Learning for Temporally Abstract Options*  
AT&T Research Labs (Host: Richard S. Sutton) November 16, 1998

*Between MDPs and Semi-MDPs: Learning, Planning, and Representing Knowledge at Multiple Temporal Scales*  
Istituto Politecnico di Milano, Italy (Host: Marco Colombetti) April 6, 1998

*Between MDPs and Semi-MDPs: Learning, Planning, and Representing Knowledge at Multiple Temporal Scales*  
IDSIA, Lugano, Switzerland (Host: Juergen Schmidhuber) April 3, 1998

*Between MDPs and Semi-MDPs: Learning, Planning, and Representing Knowledge at Multiple Temporal Scales*  
MIT AI Lab Student Seminar March 3, 1998

*Multi-time Representations of Temporally Abstract Knowledge*  
Institute for Information Sciences, USC (Host: Wei-Min Shen) September 16, 1997

### **Invited Talks at Conferences and Workshops**

*Temporal Abstraction in Reinforcement Learning*  
UAI-2000 workshop “Beyond MDPs: Representations and Algorithms” organized by Michael Kearns, Michael Littman and Satinder Singh. June 30, 2000

### **Presentations at Conferences and Workshops**

*Eligibility Traces for Off-Policy Policy Evaluation*  
The 17th International Conference on Machine Learning, ICML’00 June 30, 2000



<i>Eligibility Traces for Off-Policy Policy Evaluation</i>	1999 New England Spring Symposium on Reinforcement Learning	April 23, 1999
<i>Intra-option Learning for Temporally Abstract Options</i>	The 15th International Conference on Machine Learning, ICML'98	July 24, 1998
<i>Classification Using <math>\Phi</math>-machines and Constructive Function Approximation</i>	The 15th International Conference on Machine Learning, ICML'98	July 24, 1998
<i>Between MDPs and Semi-MDPs: Learning, Planning, and Representing Knowledge at Multiple Temporal Scales</i>	1998 New England Spring Symposium on Reinforcement Learning	May 1, 1998
<i>Theoretical Results on Reinforcement Learning with Temporally Abstract Options</i>	European Conference on Machine Learning, Chemnitz, Germany	April 22, 1998
<i>Multi-time models for Temporally Abstract Planning</i>	Poster at Neural Information Processing Systems (NIPS'97)	December 3, 1997
<i>Reinforcement Learning with Temporally Abstract Behaviors</i>	1997 New England Fall Reinforcement Learning Workshop	November 21, 1997
<i>Planning with Closed-loop Macro-actions</i>	AAAI Fall Symposium on Model Directed Autonomous Systems	November 10, 1997
<i>Multi-time Models for Reinforcement Learning</i>	ICML'97 Workshop on Modeling in Reinforcement Learning	July 12, 1997
<i>Exponentiated Gradient Descent for Reinforcement Learning</i>	The 14th International Conference on Machine Learning, ICML'97	July 10, 1997
<i>Exponentiated Gradient Descent for Supervised and Reinforcement Learning</i>	Active Learning Workshop, MIT	August 1996

### Professional Service

#### SERVICE

Co-organizer (along with Yoshua Bengio and Balasz Kegl) of a workshop on machine learning techniques for learning probabilities, to be held at the Centre de Recherches Mathematiques in May 2003.

Co-organizer (along with Peter Stone) of a workshop on learning agents, held on May 29 2001, in Montreal, as part of the Agents'2001 conference.

Co-organizer (with Ronald Parr, Leslie Kaelbling and Tom Dietterich) of the NIPS'98 Workshop on Abstraction and Hierarchy in Reinforcement Learning.

Reviewer for NIPS'2001.

Reviewer for NSERC research grants (2001).

Regular reviewer for submissions to Journal of Artificial Intelligence (JAIR), Machine Learning, Artificial Intelligence.

Beta reader for the textbook Machine Learning by Tom Mitchell.

### **McGill Service**

Organizer of the Computer Science Colloquium, 2001-2002 (partly funded by a grant from the Faculty of Graduate Studies and Research).

Member of the undergraduate committee in the School of Computer Science (since August 2000).

Representative, PSEAL Advisory Board (since August 2001).

### **University of Massachusetts Service**

Senator (elected position), Graduate Student Senate, University of Massachusetts, Amherst. 1999.

Graduate student representative (elected position), Computer Science Department, University of Massachusetts, Amherst. 1998.

### **PERSONAL**

Born on January 28, 1971 in Cluj-Napoca, Romania. Lived in the U.S. between 1995-2000, during graduate studies. Currently living in Montreal, Canada.

Citizenship: Romanian.

Foreign languages: Romanian (maternal), French (fluent), Italian (intermediate), German (beginner).

### **REFERENCES**

Available upon request.