

Amir-massoud Farahmand

Principal Research Scientist
Mitsubishi Electric Research Laboratories (MERL)

Webpage: <http://academic.SoloGen.net> ([Google Scholar Profile](#))

Research Goal

- Understanding the principles required to design *reinforcement learning* and *adaptive situated agents*. These agents interact with their environment, collect data, and use data not only to predict, which is the common objective in machine learning, but also to control the environment, with the goal of maximizing their long-term performance.
- Solving hard industrial problems using principles developed to design adaptive situated agents.

Research Interests

Theory: Reinforcement Learning (high-dimensional problems, regularized algorithms, model learning, representation learning and deep RL, learning from demonstration, inverse optimal control); Machine Learning (statistical learning theory, nonparametric algorithms, time series, non-i.i.d. processes, manifold learning, online learning); Large-scale Optimization; Evolutionary Computation

Applications: Robotics (uncalibrated visual-servoing, behavior-based architectures); Smart air conditioning systems; PDE control; Hybrid vehicles energy management system, Fault detection/prognostics for time series

Education

- **PhD in Computing Science**, University of Alberta, September 2011
 - Dissertation: *Regularization in Reinforcement Learning*
 - Supervisors: Csaba Szepesvári and Martin Jägersand
- **Master of Science in Electrical Engineering** (with specialization in Control Theory), University of Tehran, 2005
 - Thesis: *Learning and Evolution in Hierarchical Behavior-based Systems*
- **Bachelor of Science in Electrical Engineering** (with specialization in Telecommunication Systems), K.N. Toosi University of Technology, 1998-2002.

Work Experience

2016 April – Present **Mitsubishi Electric Research Laboratories: Principal Research Scientist**

2014 December – 2016 March **Mitsubishi Electric Research Laboratories: Member of Research Staff**

2014 March – 2014 October **Carnegie Mellon University: Postdoctoral Fellow** working with J. Andrew Bagnell
Topics:

- Maximum entropy inverse optimal control for high-dimensional problems
- Learning positive functions as sum of squares

- 2011 August – 2014 February** **McGill University: Postdoctoral Fellow** working with Doina Precup
Topics:
- Classification-based reinforcement learning
 - Learning from demonstrations
 - Sample-based approximate regularizers
- 2005 September – 2011 July** **University of Alberta: Research Assistant** working with Csaba Szepesvári and Martin Jägersand
Topics:
- Regularization in reinforcement learning
 - Manifold-adaptive methods in machine learning
 - Application of machine learning in robotic visual-servoing
- 2008 (Fall)** **University of Alberta: Teaching Assistant**
Course: Probabilistic graphical models (Russell Greiner)
- 2003 – 2005** **University of Tehran: Research Assistant** working with Majid Nili Ahmadabadi
Topic: Learning and evolution in hierarchical behavior-based systems
- 2001 – 2002** **K. N. Toosi University of Technology: Research Assistant** working with M.S. Abrishamian
Topic: Calculating resonant frequencies of a metallic cavity using Finite Element Method
- 2001 (Winter-Spring)** **K. N. Toosi University of Technology: Teaching Assistant**
Courses:
- Signals and Systems (M.S. Abrishamian)
 - Telecommunication Systems (E. Kalantari)
- 2001 (Summer)** **Iran Telecommunication Research Center: Intern** working with E. Geranpayeh
Topic: Analyzing the effect of metallic electrode with buffer layer on dielectric waveguides

Honors and Achievements

- International Conference on Machine Learning (ICML) Best Reviewer Award, 2015.
- NSERC Postdoctoral Fellowship, Natural Sciences and Engineering Research Council of Canada (NSERC), \$80000, 2012–2014.
- PhD Outstanding Thesis Award for the period of 2011–2012, Department of Computing Science, University of Alberta.
- NSERC Pre-approved Industrial Research and Development Fellowship (IRDF) Candidate, Natural Sciences and Engineering Research Council of Canada (NSERC), 2012.
- Nomination for Canadian Artificial Intelligence Association (CAIAC) AI Doctoral Dissertation Award, Department of Computing Science, University of Alberta, 2011.
- Finalist for Research Award, Department of Computing Science, University of Alberta, 2007.
- PhD Academic Achievement Award, Department of Computing Science, University of Alberta, 2006.
- Best Presentation Award of “*Evolving Learning Systems*” Technical Session in IEEE Congress on Evolutionary Computation (CEC) for presenting “*Hybrid Behavior Co-evolution and Structure Learning in Behavior-based Systems*”, Vancouver, Canada, 2006.
- Provost Doctoral Entrance Award, University of Alberta, \$16000, 2005–2006.
- Honored M.S. Thesis, University of Tehran, 2005.
- Bronze medal in Iran’s National Physics Olympiad, 1997.

Publications

Journal Papers (refereed)

1. A.M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Policy Iteration for Nonparametric Function Spaces**,” Journal of Machine Learning Research (JMLR), Vol. 17, No. 139, 2016.
2. A.M. Farahmand, Doina Precup, André M.S. Barreto, and Mohammad Ghavamzadeh, “**Classification-based Approximate Policy Iteration**,” IEEE Transactions on Automatic Control, Vol. 60, No. 11, 2015.
3. A. M. Farahmand, Csaba Szepesvári, “**Regularized Least-Squares Regression: Learning from a β -mixing Sequence**,” Journal of Statistical Planning and Inference (JSPI), Vol. 142, Issue 2, pp. 493–505, 2012.
4. A. M. Farahmand and Csaba Szepesvári, “**Model Selection in Reinforcement Learning**,” Machine Learning Journal (MLJ), Vol. 85, No. 3, pp. 299–332, 2011.
5. A. M. Farahmand, Majid Nili Ahmadabadi, Babak N. Araabi, Caro Lucas, “**Interaction of Culture-based Learning and Cooperative Co-evolution and its Application to Automatic Behavior-based System Design**,” IEEE Transactions on Evolutionary Computation, Vol. 14, No. 1, pp. 23-57, 2010. (Impact Factor: 2.426)
6. M. J. Yazdanpanah, E. Madanian, and A. M. Farahmand, “**Channel Assignment in Cellular Communications using a New Modification on Hopfield Networks**,” Iranian Journal of Science and Technology, Transaction B: Engineering, Vol. 29, No. B4, 2005.

Submitted

7. A.M. Farahmand, Daniel Nikovski, Yuji Igarashi, and Hiroki Konaka, “**Truncated Approximate Dynamic Programming**,” Submitted to the Journal of Artificial Intelligence Research (JAIR), 2016. [In the first round of revisions]

Conference Papers (refereed)

Note: Conferences are the major publication venue in computer science. The top conferences in machine learning and artificial intelligence, such as Neural Information Processing Systems (NIPS), International Conference on Machine Learning (ICML), AAAI Conference on Artificial Intelligence (AAAI), and Artificial Intelligence and Statistics (AISTATS) are very competitive and have acceptance rate between 20-30%. Moreover, International Conference on Robotics and Automation (ICRA) and International Conference on Intelligent Robots and Systems (IROS) are the top conferences in robotics and Conference on Decision and Control (CDC) and American Control Conference (ACC) are the top venues for control engineering.

8. A.M. Farahmand, Sepideh Pourazarm, Daniel Nikovski, “**Random Projection Filter Bank for Time Series**,” Accepted at the Advances in the Neural Information Processing Systems (NIPS), 2017. (21% acceptance rate)
9. A.M. Farahmand, André M.S. Barreto, Daniel Nikovski, “**Value-Aware Model Learning for Reinforcement Learning**,” In the Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017. (32% acceptance rate)
10. A.M. Farahmand, Saleh Nabi, and Daniel Nikovski, “**Deep Reinforcement Learning for Partial Differential Equation Control**,” In the Proceedings of the American Control Conference (ACC), 2017.
11. Sepideh Pourazarm, A. M. Farahmand, Daniel Nikovski, “**Fault Detection and Prognosis of Time Series Data with Random Projection Filter Bank**,” Accepted at the Annual Conference of the Prognostics and Health Management (PHM) Society, 2017.
12. Mouhacine Benosman and A.M. Farahmand, “**Gaussian Process-based Parameter Identification for Dynamical Systems**,” In the Proceedings of the World Congress of the International Federation of Automatic Control (IFAC), 2017.
13. A.M. Farahmand, Saleh Nabi, Piyush Grover, Daniel Nikovski, “**Learning to Control Partial Differential Equations: Regularized Fitted Q-Iteration Approach**,” In the Proceedings of the IEEE Conference on Decision and Control (CDC), December 2016.
14. Mouhacine Benosman, A.M. Farahmand, Meng Xia, “**Learning-Based Modular Indirect Adaptive Control for a Class of Nonlinear Systems**,” In the Proceedings of American Control Conference (ACC), July 2016.

15. A.M. Farahmand, Daniel Nikovski, Yuji Igarashi, Hiroki Konaka, “**Truncated Approximate Dynamic Programming with Task-Dependent Terminal Value**,” In Proceedings of AAAI Conference on Artificial Intelligence, February 2016. (26% acceptance rate)
16. De-An Huang, A. M. Farahmand, Kris M. Kitani, J. Andrew Bagnell, “**Approximate MaxEnt Inverse Optimal Control and its Application for Mental Simulation of Human Interactions**,” In Proceedings of AAAI Conference on Artificial Intelligence, 2015. (27% acceptance rate)
17. Philip Bachman, A. M. Farahmand, Doina Precup, “**Sample-based Approximate Regularization**,” In Proceedings of International Conference on Machine Learning (ICML), 2014. (25% acceptance rate)
18. Beomjoon Kim, A. M. Farahmand, Joelle Pineau, Doina Precup, “**Learning from Limited Demonstrations**,” In the Proceedings of Advances in the Neural Information Processing Systems (NIPS), 2013. (25% acceptance rate -- 3.6% spotlight presentation)
19. Mahdi Milani Fard, Yuri Grinberg, A.M. Farahmand, Joelle Pineau, Doina Precup, “**Bellman Error Based Feature Generation using Random Projections on Sparse Spaces**,” In the Proceedings of Advances in the Neural Information Processing Systems (NIPS), 2013. (25% acceptance rate)
20. A. M. Farahmand and Doina Precup, “**Value Pursuit Iteration**,” In the Proceedings of Advances in Neural Information Processing Systems (NIPS), 2012 (25% acceptance rate).
21. A. M. Farahmand, “**Action-Gap Phenomenon in Reinforcement Learning**,” In the Proceedings of Advances in the Neural Information Processing Systems (NIPS), 2011. (22% acceptance rate -- 5% spotlight presentation)
22. A. M. Farahmand, Remi Munos, Csaba Szepesvári, “**Error Propagation for Approximate Policy and Value Iteration**,” In the Proceedings of Advances in Neural Information Processing Systems (NIPS – 23), Vancouver, Canada, December 2010. (24% acceptance rate)
23. Azad Shademan, A. M. Farahmand, and Martin Jägersand, “**Robust Jacobian Estimation for Uncalibrated Visual Servoing**,” In the Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Anchorage, Alaska, USA, May 2010.
24. A. M. Farahmand, Azad Shademan, Martin Jägersand, and Csaba Szepesvári, “**Model-based and Model-free Reinforcement Learning for Visual Servoing**,” In the Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Kobe, Japan, May 2009.
25. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Fitted Q-Iteration for Planning in Continuous-Space Markovian Decision Problems**,” In the Proceedings of the American Control Conference (ACC), St. Louis, Missouri, USA, pp. 725–730, June 2009.
26. Azad Shademan, A. M. Farahmand, Martin Jägersand, “**Towards Learning Robotic Reaching and Pointing: An Uncalibrated Visual Servoing Approach**,” Sixth Canadian Conference on Computer and Robot Vision (CRV), Kelowna, British Columbia, Canada, 2009.
27. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Policy Iteration**,” In the Proceedings of the Advances in Neural Information Processing Systems (NIPS – 21), Vancouver, Canada, December 2008. (24% acceptance rate)
28. A. M. Farahmand, Csaba Szepesvári, and Jean-Yves Audibert, “**Manifold-Adaptive Dimension Estimation**,” In Proceedings of International Conference on Machine Learning (ICML), pp. 265–272, 2007. (29% acceptance rate)
29. A. M. Farahmand, Azad Shademan, and Martin Jägersand, “**Global Visual-Motor Estimation for Uncalibrated Visual Servoing**,” In Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1969-1974, 2007.
30. A. M. Farahmand, Majid Nili Ahmadabadi, Caro Lucas, and Babak N. Araabi, “**Hybrid Behavior Co-evolution and Structure Learning in Behavior-based Systems**,” In Proceedings of IEEE Congress on Evolutionary Computation (CEC), Vancouver, Canada, 2006. (Chosen as the best presentation of the “Evolving Learning Systems” technical session)
31. A. M. Farahmand and Mohammad javad Yazdanpanah, “**Channel Assignment using Chaotic Simulated Annealing Enhanced Hopfield Neural Network**,” In Proceedings of International Joint Conference on Neural Networks (IJCNN), Vancouver, Canada, 2006.
32. Mohammad G. Azar, Majid Nili Ahmadabadi, A. M. Farahmand, and Babak N. Araabi, “**Learning to Coordinate Behaviors in Soft Behavior-based Systems using Reinforcement Learning**,” International Joint Conference on Neural Networks (IJCNN), Vancouver, Canada, 2006.

33. A. M. Farahmand and M. J. Yazdanpanah, “**Locally Optimal Takagi-Sugeno Fuzzy Controllers,**” In Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference (CDC-ECC), pp. 4095-4099, Seville, Spain, December 2005.
34. A. M. Farahmand, Majid Nili Ahmadabadi, and Babak N. Araabi, “**Behavior Hierarchy Learning in a Behavior-based System using Reinforcement Learning,**” In Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Sendai, Japan, 2004.

Book Chapters (refereed)

35. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Fitted Q-Iteration: Application to Bounded Resource Planning,**” in Recent Advances in Reinforcement Learning, 8th European Workshop, EWRL 2008, Revised and Selected Papers, Springer, LNCS 5323, pp. 55—68, 2008.

Workshop Papers or Extended Abstracts (lightly refereed)

36. Mouhacine Benosman and A.M. Farahmand, “**Towards Stability in Learning-based Control: A Bayesian Optimization-based Adaptive Controller,**” The 3rd Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2017 (Extended Abstract).
37. A.M. Farahmand, André M.S. Barreto, Daniel Nikovski, “**Value-Aware Loss Function for Model Learning in Reinforcement Learning,**” The 13th European Workshop on Reinforcement Learning (EWRL), December 2016.
38. Mouhacine Benosman and A.M. Farahmand, “**Bayesian Optimization-based Modular Indirect Adaptive Control for a Class of Nonlinear Systems,**” IFAC International Workshop on Adaptation and Learning in Control and Signal Processing, June 2016.
39. J. Andrew Bagnell and A. M. Farahmand, “**Learning Positive Functions in a Hilbert Space,**” NIPS Workshop on Optimization for Machine Learning, December 2015.
40. De-An Huang, A. M. Farahmand, Kris M. Kitani, and J. Andrew Bagnell, “**Approximate MaxEnt Inverse Optimal Control,**” The 2nd Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), June 2015.
41. Beomjoon Kim, A. M. Farahmand, Joelle Pineau, and Doina Precup, “**Approximate Policy Iteration with Demonstration Data,**” First Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2013 (Extended Abstract).
42. A. M. Farahmand, Doina Precup, André M.S. Barreto, and Mohammad Ghavamzadeh, “**CAPi: Generalized Classification-based Approximate Policy Iteration,**” First Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2013 (Extended Abstract).
43. A. M. Farahmand, Doina Precup, and Mohammad Ghavamzadeh, “**Generalized Classification-based Approximate Policy Iteration,**” Tenth European Workshop on Reinforcement Learning (EWRL 2011), Edinburgh, Scotland, June 2012.
44. A. M. Farahmand and Csaba Szepesvári, “**BErMin: A Model Selection Algorithm for Reinforcement Learning Problems,**” Workshop on New Frontiers of Model Order Selection, Advances in Neural Information Processing Systems (NIPS), December 2011.
45. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularization in Reinforcement Learning,**” Multidisciplinary Symposium on Reinforcement Learning (MSRL-2009), Montreal, QC, Canada, June 2009.
46. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Robot Learning with Regularized Reinforcement Learning,**” Workshop on Regression in Robotics: Approaches and Applications, Robotics: Science and Systems Conference (RSS-2009), Seattle, WA, June 2009.
47. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Policy Iteration,**” Eighth European Workshop on Reinforcement Learning (EWRL 2008), Villeneuve d'Ascq, France, July 2008.
48. A. M. Farahmand, Csaba Szepesvári, and Jean-Yves Audibert, “**Towards Manifold-Adaptive Learning,**” Workshop on Topology Learning, Advances in Neural Information Processing Systems (NIPS), 2007.

Theses and Dissertation

49. **PhD Dissertation: Regularization in Reinforcement Learning**, Department of Computing Science, University of Alberta. Defended on September 14, 2011 (Supervisor: Csaba Szepesvári and Martin Jägersand – Examining Committee: Peter L. Bartlett, Michael Bowling, Alexander Melnikov, Dale Schuurmans, Richard S. Sutton)
50. **M.S. Thesis: Learning and Evolution in Hierarchical Behavior-based Systems**, Department of Electrical and Computer Engineering, University of Tehran, 2005. (Advisors: Majid Nili Ahmadabadi, Babak N. Araabi, Caro Lucas) (in Persian)
51. **B.S. Thesis: Calculating Resonant Frequencies of a Metallic Cavity using Finite Element Method**, K. N. Toosi University of Technology, 2002. (Advisor: Mohammad-Sadegh Abrishamian) (in Persian)

Working Papers

52. A. M. Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor, “**Regularized Fitted Q-Iteration Algorithm**,” Working paper (Journal paper).

Technical Reports (Selected)

53. Kota Hara, Ming-Yu Liu, Oncel Tuzel, and A.M. Farahmand, “**Attentional Network for Visual Object Detection**,” arXiv:1702.01478, 2017.
54. A.M. Farahmand, Doina Precup, André M.S. Barreto, and Mohammad Ghavamzadeh, “**Classification-based Approximate Policy Iteration: Experiments and Extended Discussions**,” 2014. (A shorter version of this work is published in IEEE Transactions on Automatic Control, 2015).
55. A. M. Farahmand, Majid Nili Ahmadabadi, and Babak N. Araabi, “**Behavior and Hierarchy Development in Behavior-based Systems using Reinforcement Learning**,” Technical Report, University of Tehran, 2005.
56. A. M. Farahmand, Caro Lucas, and Babak N. Araabi, “**Chaos Control Survey**,” A Technical Report for the Graduate Seminar Course, University of Tehran, 2004 (in Persian).
57. A. M. Farahmand and Mohammad javad Yazdanpanah, “**A Class of Nonlinear Controllers for Synchronization of Chaotic Semipassive Systems**,” Technical Report, University of Tehran, 2003.
58. A. M. Farahmand, Ramin Pashai, and Ezatollah Geranpayeh, “**Effect of Metallic Electrode and Buffer Layer on Dielectric Waveguides**,” Technical Report, Iran Telecommunication Research Center (ITRC), 2001.

Patents

1. Daniel Nikovski and A. M. Farahmand, “**Method and System for Selecting Power Sources in Hybrid Electric Vehicles**,” US9637111, Filing date: Jun 9, 2015, Issue date: May 2, 2017.

Note: Four other patent applications has been filed in 2016.

Talks

Invited Talks

- *Industrial Applications of Reinforcement Learning*, University of Saskatchewan, October 2016.
- *Truncated Approximate Dynamic Programming with Task-Dependent Terminal Value*, University of Alberta, January 2016.
- *Reinforcement Learning and Learning from Demonstrations: A Unified Approach*
 - Mitsubishi Electric Research Laboratories, September 2014.
 - Microsoft Research (Redmond), October 2014.
- *Sample-based Approximate Regularization*, Carnegie Mellon University, September 2014.

- *Solving High-Dimensional Reinforcement Learning Problems while Avoiding the Curse of Dimensionality*, Concordia University, February 2014.
- *Beyond Value-based Regularities in Sequential Decision-Making Problems*, Carnegie Mellon University, October 2013.
- *Regularities in Reinforcement Learning Problems*, Department of Computer and Information Sciences, University of Alabama at Birmingham, March 2013.
- *Classification-based Approximate Policy Iteration*, 7th Barbados Workshop on Reinforcement Learning Workshop, Bellairs Institute, Barbados, April 2012 (Host: Doina Precup).
- *Regularities in Sequential Decision-Making Problems: A Research Program*, Department of Computer Science, McGill University, May 2010 (Host: Doina Precup).
- *Regularities in Sequential Decision-Making Problems*, Department of Electrical and Computer Engineering, University of Tehran, January 2010 (Host: Majid Nili Ahmadabadi).
- *Manifold Learning*, Applied Mathematics Institute (AMI) Seminar, Department of Mathematics, University of Alberta, Fall 2007 (Host: Suneeta Vardarajan).
- *Learning and Evolution in Behavior-based Systems*, Institute of Physics and Mathematics (IPM) – School of Cognitive Sciences’ monthly seminar on Machine Learning, Winter 2005 (Host: Majid Nili Ahmadabadi).
- *Hierarchical Reinforcement Learning*, Institute of Physics and Mathematics (IPM) - School of Cognitive Sciences’ monthly seminar on Machine Learning, Summer 2004 (Host: Majid Nili Ahmadabadi).
- *Behavior-based Systems*, Institute of Physics and Mathematics (IPM) – School of Cognitive Sciences’ monthly seminar on Machine Learning, Winter 2004 (Host: Majid Nili Ahmadabadi).

Conference Presentations (Selected)

- *Deep Reinforcement Learning for Partial Differential Equation Control*, American Control Conference (ACC), May 2017.
- *Learning to Control Partial Differential Equations: Regularized Fitted Q-Iteration Approach*, IEEE Conference on Decision and Control (CDC), December 2016.
- *Value-Aware Loss Function for Model Learning in Reinforcement Learning*, European Workshop on Reinforcement Learning (EWRL), December 2016.
- *Truncated Approximate Dynamic Programming with Task-Dependent Terminal Value*, AAAI Conference on Artificial Intelligence, February 2016.
- *BErMin: A Model Selection Algorithm for Reinforcement Learning Problems*, Workshop on New Frontiers of Model Order Selection, Advances in Neural Information Processing Systems (NIPS), Sierra Nevada, Spain, December 2011. (Watch on [Videolectures](#))
- *Action-Gap Phenomenon in Reinforcement Learning*, 25th Annual Conference on Advances in Neural Information Processing Systems (NIPS), Granada, Spain, December 2011 (Spotlight presentation). (Watch on [Videolectures](#))
- *Manifold-Adaptive Dimension Estimation*, International Conference on Machine Learning (ICML), Corvallis, Oregon, USA, 2007.
- *Hybrid Behavior Co-evolution and Structure Learning in Behavior-based Systems*, IEEE Congress on Evolutionary Computation (CEC), Vancouver, Canada, 2006. (***Chosen as the best presentation of the “Evolving Learning Systems” technical session***)

Teaching Experience

Lectures/Short Courses

- Guest Lecturer: *Statistical Learning Theory* module for the graduate-level Machine Learning course (COMP-652), McGill University, Fall 2012. (2 sessions; Invited by Doina Precup)
- *Control Theory*, Department of Computing Science, University of Alberta, Summer 2006. (1 session)
- Guest Lecturer: *Hierarchical Reinforcement Learning* module for the graduate-level Distributed AI course, University of Tehran, Spring 2004. (1 session; Invited by Majid Nili Ahmadabadi)

- *Fuzzy Logic*, Student Branch of IEEE, University of Tehran, Summer 2003. (3 sessions)
- *Neural Networks and Evolutionary Algorithms*, K.N. Toosi Univ. of Tech., 2000. (several sessions)

Teaching Assistantship

- Probabilistic Graphical Models (for Russell Greiner), Fall 2008. (homework grading; MatLab tutorial lecture; help sessions; 20 students)
- Signals and Systems (for Mohammad sadegh Abrishamian), Spring 2001. (homework and quizzes design and grading, problem-solving sessions; ~30 students)
- Telecommunication Systems I (for Esmael Kalantari), Spring 2001. (problem-solving sessions; ~10 students)

Supervisory Experience

- Yangchen Pan (PhD Student – MERL). Topic: *Reinforcement Learning for High-Dimensional Action Spaces* (2017 June – August)
- Sepideh Pourazarm (Postdoctoral Researcher – MERL). Topic: *Random Projection Filter Bank for Prognostics* (2017 January – June)
- Jordi Grau-Moya (PhD Candidate Intern – MERL). Topic: *Deep Least-Squares Temporal Difference Learning* (2017 February – April)
- De-An Huang (MS student – Carnegie Mellon University). Topic: *Approximate MaxEnt Inverse Optimal Control* (2014 March – October)
- Beomjoon Kim (MS student – McGill University). Topic: *Learning from Demonstration* (2013 April – December)
- Howard Huang (undergraduate – McGill University). Topic: *Empirical evaluation of the CAPI algorithm* (2012 May – August)

Grants

- *Reinforcement Learning with Big Data*, INRIA International Associate Team with McGill University Proposal (2013–15), PI from McGill University.

Professional Activities and Services

Peer Reviewing

Journal Reviewer (35 papers)

- Journal of Machine Learning Research (JMLR) (2007, 2008, 2010 (2), 2011, 2013, 2014 (3), 2015, 2017)
- Machine Learning Journal (2015, 2016)
- Biometrika (2015, 2016(2))
- IEEE Transactions on Neural Network and Learning Systems (2011, 2015)
- IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) (2009, 2010)
- IEEE Transactions on Automatic Control (2011, 2012 (4), 2013)
- Journal of Artificial Intelligence Research (JAIR) (2016)
- Artificial Intelligence Journal (AIJ) (2007)
- Automatica (2009)
- International Journal of Robust and Nonlinear Control (2012)
- AI Communication (2009)
- IEEE Transactions on Vehicular Technology (2009)
- Autonomous Robots (2008)
- IEEE Transactions on Robotic and Automation (2007)

- International Journal of Neural Systems (2006)

Conference Program Committee/Reviewer (147 papers)

- International Conference on Machine Learning (**ICML**) (2006: co-reviewer for a paper; 2010: PC for 7 papers; 2012: PC for 6 papers; 2013: PC for 7 papers; 2014: PC for 8 papers; 2015: PC for 4 papers, best reviewer award; 2016: PC for 2 papers, 2017: PC for 4 papers)
- Advances in Neural Information Processing Systems (**NIPS**) (2007: co-reviewer for 2 papers; 2008: PC for 7 papers; 2012: PC for 8 papers; 2014: PC for 5 papers, 2015: PC for 8 papers; 2016: PC for 6 papers; 2017: PC for 6 papers)
- Artificial Intelligence and Statistics (**AISTATS**) (2015: 4 papers, 2016: 5 papers)
- AAAI Conference on Artificial Intelligence (**AAAI**) (2016: 5 papers)
- International Joint Conference on Artificial Intelligence (**IJCAI**) (2011: one paper; 2013: PC for 3 papers; 2015: PC for a paper)
- International Joint Conference on Neural Networks (**IJCNN**) (2007: reviewer for 3 papers)
- Conference on Learning Theory (**COLT**) (2007: co-reviewer for 2 papers; 2009: co-reviewer for a paper; 2010: reviewer for a paper)
- International Conference on Algorithmic Learning Theory (**ALT**) (2014: reviewer for a paper)
- European Conference on Machine Learning (**ECML**) (2009: reviewer for 2 papers, 2012: reviewer for a paper)
- Algorithmic Learning Theory (2013: reviewer for a paper)
- European Workshop on Reinforcement Learning (**EWRL**) (2012: PC for 2 papers; 2015: PC for a paper; 2016: PC for 3 papers)
- Reinforcement Learning and Decision Making (**RLDM**) (2015: reviewer for a paper, 2017: reviewer for 6 papers)
- International Conference on Robotics and Automation (**ICRA**) (2009: co-reviewer for a paper)
- IEEE International Conference on Intelligence Robots and Systems (**IROS**) (2008: reviewer for a paper; 2009: reviewer for a paper; 2010: reviewer for a paper)
- IEEE Conference on Decision and Control (**CDC**) (2009: reviewer for a paper; 2014: reviewer for a paper; 2016: reviewer for a paper)
- American Control Conference (**ACC**) (2006: reviewer for a paper; 2009: reviewer for 2 papers; 2015: reviewer for a paper)
- World Congress of the International Federation of Automatic Control (**IFAC**) (2017: reviewer for a paper)
- European Control Conference (**ECC**) (2007: reviewer for a paper)
- IEEE Multi-Conference on Systems and Control (**MSC**) (2008: reviewer for a paper)
- IEEE Joint International Conference on Control Applications (**CCA**), Symposium on Computer-Aided Control System Design (**CACSD**), and International Symposium on Intelligent Control (**ISIC**). (2006: reviewer for a paper)
- Canadian Conference on Electrical and Computer Engineering (**CCECE**) (2012: reviewer for a paper)
- German Conference on Artificial Intelligence (**KI**) (2010: PC for a paper)
- Computer Society of Iran Computer Conference (2016)
- Knowledge, Skill, and Behavior Transfer in Autonomous Robots (KSBT) - AAAI 2015 Workshop (Program Committee for 2 papers)
- Autonomously Learning Robots (NIPS 2014 Workshop): Reviewer for two papers

Organizing Professional Activities

- Chair: *Machine Learning* session at the American Control Conference (ACC), May 2017.
- Co-chair: *Machine Learning* session at the 55th IEEE Conference on Decision and Control (CDC), December 2016.
- Co-chair: *Sequential Decision Making with Big Data* workshop at the 28th Conference on Artificial Intelligence (AAAI-14), July 2014.
- Organizing Teatime Talks (Weekly presentations) - McGill University (Summer 2012)
- Organizing reading groups at the School of Computer Science - McGill University

- Random Projection in Machine Learning (Fall 2011)
- Online Learning (Winter 2012)
- Organizing reading groups at the Department of Computing Science - University of Alberta
 - Compressive Sampling Reading Group (Fall 2008)
 - Reinforcement Learning and Function Approximation (Winter 2007 – Fall 2007)
- Organizing several short courses at the University of Alberta, University of Tehran, and K. N. Toosi University of Technology (see **Teaching Experience – Short Courses**)
- Scientific Chair of the 1st Soft Computing workshop in K.N. Toosi Univ. of Technology (2001)

Professional Associations

- Association for the Advancement of Artificial Intelligence (AAAI), 2016 – present
- Institute of Electrical and Electronics Engineering (IEEE) – Student member since 1999 until 2010 (with interruptions)

Other Activities

- Writing short stories (in Persian)
- Publishing **Anti Memoirs** blog since 2002 – one of the earliest Persian blogs (weblog.SoloGen.net)
- Co-founder of the first Persian online newspaper, Aftab, in 1998 [discontinued]
- Elected officer for Scientific Student Society of E.E. Dept. of K.N. Toosi University of Technology (Chair: 2001-2002; Research Vice-chair: 2000-2001)
- Co-founder of the Σαω scientific group (2000 – present)
- Publishing a scientific blog **Thesilog** (thesilog.SoloGen.net) since 2004 (in English) [inactive]

References

Csaba Szepesvári

Professor, Department of Computing Science, University of Alberta
 webpage: <http://www.ualberta.ca/~szepesva>

Doina Precup

Associate Professor, School of Computer Science, McGill University
 webpage: <http://www.cs.mcgill.ca/~dprecup>

J. Andrew Bagnell

Associate Professor, Robotics Institute, Carnegie Mellon University
 webpage: <http://robotwhisperer.org/>

Joelle Pineau

Associate Professor, School of Computer Science, McGill University
 webpage: <http://www.cs.mcgill.ca/~jpineau/>

Mohammad Ghavamzadeh

Google DeepMind & INRIA Lille – Team SequeL
 webpage: <http://chercheurs.lille.inria.fr/~ghavamza/>

Note: Please coordinate with me before contacting any of the references.