

Alessandro Lazaric

Curriculum Vitæ

INRIA Lille-Nord Europe
40, Avenue Halley
59650 Villeneuve d'Ascq, France
☎ + 33 3 59 57 79 03
✉ alessandro.lazaric@inria.fr
🌐 researchers.lille.inria.fr/~lazaric



Academic Positions

- 2012 **Junior Researcher (CR1)**, INRIA – Team *SequeL*, Lille, France.
- 2010 **Junior Researcher (CR2)**, INRIA – Team *SequeL*, Lille, France.
- 2008 **Postdoctoral Researcher**, INRIA – Team *SequeL*, Lille, France.
Supervisor: Prof. Rémi Munos.
- 2008 **Postdoctoral Researcher**, *Politecnico di Milano*, Milan, Italy.
Supervisor: Prof. Andrea Bonarini.

Education

- 2005 **PhD Computer Engineering**, *Politecnico di Milano*, Milan, Italy.
PhD thesis: “Knowledge Transfer in Reinforcement Learning”. *Supervisors*: Prof. Andrea Bonarini, Dr. Marcello Restelli. *Examiner*: Peter Stone (The University of Texas at Austin).
- 2003 **MSc Computer Engineering**, *University of Illinois at Chicago (UIC)*, Chicago, IL, USA.
Master thesis: “Learning with multiple state aggregation representations”. *Master advisor*: Prof. Piotr Gmytrasiewicz.
- 1999 **MSc Computer Engineering (Cum Laude)**, *Politecnico di Milano*, Milan, Italy.
Master thesis: “A Multi-Resolution Reinforcement Learning Approach to the Efficient Solution of Complex Problems”. *Supervisor*: Prof. Andrea Bonarini.

Awards

- INRIA award for scientific excellence (prime d'excellence scientifique) (2013-2016)
- Winner of a Ph.D. scholarship granted by the Italian Ministry of Education, University and Research by ranking second among the Ph.D. applicants at Politecnico di Milano (Italy)
- M.S. with honors. Politecnico di Milano (Italy)

Languages

Italian **Native**
English **Fluent**
French **Fluent**

Teaching and Supervision

Lecturer

- 2013 **Reinforcement Learning**, *ENS Cachan – Master 2 MVA*, Cachan, France.
- 2013 **Introduction to Reinforcement Learning**, *Ecole Centrale de Lille – Option Décision et Analyse de Données*, Lille, France.

2012
Advanced Topics in Machine Learning, *Politecnico di Milano – PhD Course*, Milan, Italy.
Focused on statistical learning theory, online learning, and multi-armed bandit.

Teaching Assistant

2008
Computer Science III, *Politecnico di Milano*, Milan, Italy.
Lecturer: Prof. Marcello Restelli

2007
Principles of Computer Science, *Politecnico di Milano*, Milan, Italy.
Lecturer: Prof. Andrea Bonarini

2006
Software Engineering, *Politecnico di Milano*, Milan, Italy.
Lecturer: Prof. Andrea Morzenti

2005
Principles of Computer Science (C), *Politecnico di Milano*, Milan, Italy.
Lecturer: Prof. Matteo Matteucci

2005
Artificial Intelligence and Robotics, *Politecnico di Milano*, Milan, Italy.
Lecturer: Prof. Marcello Restelli

PhD Supervision

2014
Daniele Calandriello, *Large-scale Graph Learning*, Université de Lille, France.
Co-supervisor: Michal Valko, Philippe Preux

2014
Marc Abeille, *Reinforcement Learning for Trading*, ENS Cachan, France.
Co-supervisor: Rémi Munos

2012
Marta Soare, *Exploration–exploitation in Linear Bandit*, Université de Lille, France.
Co-supervisor: Rémi Munos

2011
Amir Sani, *Risk in Machine Learning*, Université de Lille, France.
Co-supervisor: Rémi Munos

PhD Examiner

2014
Mahdi Milani Fard, *McGill University*, Montreal, Canada.
Supervisor: Prof. Joelle Pineau

MsC Supervision

2014
Daniele Calandriello, *Sparse Multi-task Reinforcement Learning*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli. **Best MsC thesis award at AI*IA 2014.**

2011
Francesco Trovò, *Regret Estimation for Multi Slot Incentive Compatible Multi Armed bandit*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Nicola Gatti, Erkki Oja. *Francesco is currently a PhD student at Politecnico di Milano.*

2007
Mario Quaresimale, *Analysis of Opponent Modeling for improving Reinforcement Learning Agents in the Kuhn Poker Game*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli

2007
Emanuele Venneri, *Macro-actions Optimization in Multi-Task Reinforcement Learning*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli

2007
Eliseo Ferrante, *Transfer of Knowledge in Reinforcement Learning through Policy-based Proto- Value Functions.*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli. *Eliseo is currently a postdoctorate researcher at “Socio-ecology and Social Evolution” laboratory in K.U. Leuven*

2007
Alessandro Prete, *Learning Driving Tasks in a Racing Game Using Reinforcement Learning*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli

- 2007 **Francesco Recchia and Guglielmo Roelsfema**, *Bound Algorithms for Learning Optimal Policies using State Aggregation*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2007 **Nicola Parolini and Giorgio Pandolfi**, *Multiagent Reinforcement Learning in Extended Form Games: the Bilateral Bargaining*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2007 **Mario Lozza**, *Reinforcement Learning in Texas Hold-em Poker*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2007 **Matteo Lazzarotto**, *Factored Representation for Intrinsically Motivated Reinforcement Learning*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2006 **Simone Tognetti**, *Analysis and Application of Algebraic Transformations in Partially Symmetric Environments for Learning Speed Improvement*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2005 **Francesco Carsana e Ilario Ceruti**, *Multi-resolution Approximation in Hierarchical Reinforcement Learning for Complex Environment*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- 2005 **Patrick Vitali**, *Autonomous Development for Artificial Agents through Self-Regulated Learning*, Politecnico di Milano, Milan, Italy.
Co-supervisor: Marcello Restelli
- MsC Examiner**
- 2012 **Matteo Pirotta, Alessio Pecorino**, *Safe Policy Iteration: A Monotonically Improving Approximate Policy Iteration Approach*, Politecnico di Milano, Milan, Italy.
Supervisor: Marcello Restelli.
- 2008 **Sofia Ceppi**, *Characterization of Equilibria in Energy Markets Using Decentralized Auctions*, Politecnico di Milano, Milan, Italy.
Supervisor: Nicola Gatti.

Project Funding

- 2015
2018 **EduBand (Educational Bandit)**, Associate Team with Carnegie Mellon University, Funding: 50,000 €.
Collaborators: Emma Brunskill, Michal Valko
- 2014
2018 **ExTra-Learn (Extraction and Transfer of Knowledge in Reinforcement Learning)**, ANR-JCJC, Funding: 250,000 €.
- 2011
2015 **ComLACS (Composing Learning for Artificial Cognitive Systems)**, EU FP7.
PI: John Shawe-Taylor. Partners: University College London (UCL), University of Bristol (UB), Royal Holloway, University of London, Radboud Universiteit Nijmegen, Technische Universitat Berlin, Montanuniversitat Leoben, Inria, Technische Universität Darmstadt.
- 2009
2013 **Explo-RA (Exploration-Exploitation for Efficient Resource Allocation, applications to optimization, control, learning and games)**, ANR.
PI: Rémi Munos. Partners: Inra, HEC Paris, Les Ponts, Paris 5, Paris 8.

Organization of Scientific Events

- 2015 **12th European Workshop on Reinforcement Learning (EWRL)**, Chair, Lille, France.
Co-chair: Mohammad Ghavamzadeh, Rémi Munos.

2009

ICML Workshop on “On-line Learning with Limited Feedback”, *Chair*, Montreal, Canada.

Co-chair: J.-Y. Audibert, P. Auer, R. Munos, D. Ryabko, C. Szepesvari.

Tutorials and Invited Talks

2014

Invited Talk on “Transfer in Reinforcement Learning”, 30 minutes de science, INRIA, Lille, France.

2014

Invited Talk on “Approximate Dynamic Programming Meets Statistical Learning Theory”, GDR “Controle des décollements” and GDR MOSAR, Paris, France.

Host: L. Mathelin.

2013

Tutorial on “Statistical Learning Theory in Reinforcement Learning and Approximate Dynamic Programming”, *Lecturer*, EWRL-13, Dagstuhl, Germany.

Co-organizers: M. Ghavamzadeh.

2013

Invited Talk on “Transfer of Samples from Multiple MDPs”, Carnegie Mellon University, Pittsburgh, USA.

Host: E. Brunskill.

2012

Invited Talk on “Recent Advancements on Multi-Armed Bandit”, Politecnico di Milano, Milan, Italy.

Host: M. Restelli.

2012

Tutorial on “Statistical Learning Theory in Reinforcement Learning and Approximate Dynamic Programming”, *Lecturer*, ICML-12, Edinburgh, UK.

Co-organizers: M. Ghavamzadeh.

2011

Invited Talk on “Transfer of Samples from Multiple MDPs”, Université de Liège, Liège, Belgium.

Host: D. Ernst.

2010

Invited Talk on “Analysis of a Classification-based Policy Iteration Algorithm”, University of Alberta, Edmonton, Canada.

Host: Cs. Szepesvari.

2010

Tutorial on “Reinforcement Learning and Beyond”, *Lecturer*, AAMAS-10, Toronto, Canada.

Co-organizers: K. Verbeeck, J. Steven, P. Vrancx, M. E. Taylor, F. S. Melo.

2009

Tutorial on “Reinforcement Learning and Beyond”, *Lecturer*, AAMAS-09, Budapest, Hungary.

Co-organizers: K. Verbeeck, J. Steven, P. Vrancx, M. E. Taylor, F. S. Melo.

2009

Invited Talk on “Transfer in Reinforcement Learning: It’s All About Similarity”, Barbados Reinforcement Learning Workshop, McGill’s Bellairs Institute, Barbados.

2009

Tutorial on “Transfer Learning for Reinforcement Learning Domains”, *Lecturer*, ECML-09, Bled, Slovenia.

Co-organizers: M. E. Taylor.

2008

Invited Talk on “Knowledge Transfer in Reinforcement Learning”, IRIDIA, Bruxelles, Belgium.

Host: C. Pinciroli, M. Dorigo.

Responsibilities

- Program Committee and Reviewer: ICML (2009–), NIPS (2009–), ECML (2010), AAMAS (2011–2013), AAAI (2010, 2011, 2014, 2015), ICRA (2010), IJCAI (2009, 2011, 2013, 2015), COLT (2010–2014), ALT (2010, 2014), RLDM (2015), JFPDA (2013, 2014), UAI (2012), CDC (2012), various workshops.

- Journal Reviewer: JMLR, MLJ, JAIR, JOTA, IEEE Information Theory, IEEE TAC, JAAMAS.
- Member of the selection committee for junior researcher positions at INRIA Lille (2015).
- Member of the selection committee for PhD and postdoctorate positions at INRIA Lille (2014, 2015).
- Member of the selection committee for “delegation” positions at INRIA Lille (2014).
- Responsible for the Activity Report of Sequel (2011, 2012).
- Responsible for the reading group on “Multi-armed Bandit” at Sequel (2009).

Journal Publications

2012

Lazaric, A., M. Ghavamzadeh, and R. Munos. “Finite-Sample Analysis of Least-Squares Policy Iteration”. In: *Journal of Machine Learning Research (JMLR)*.

2012

Lazaric, Alessandro and Rémi Munos. “Learning with stochastic inputs and adversarial outputs”. In: *Journal of Computer and System Sciences (JCSS)* 78.5, pp. 1516–1537.

2009

Bonarini, Andrea, Alessandro Lazaric, Francesco Montrone, and Marcello Restelli. “Reinforcement distribution in fuzzy Q-learning”. In: *Fuzzy Sets and Systems* 160.10, pp. 1420–1443.

Conference Publications

2014

Calandriello, D., A. Lazaric, and M. Restelli. “Sparse Multi-task Reinforcement Learning”. In: *Proceedings of the Twenty-Eighth Annual Conference on Neural Information Processing Systems (NIPS'14)*.

2014

Gheshlaghi-Azar, M., A. Lazaric, and E. Brunskill. “Resource-efficient Stochastic Optimization of a Locally Smooth Function under Correlated Bandit Feedback”. In: *Proceedings of the Thirty-First International Conference on Machine Learning (ICML'14)*.

2014

Sani, A., G. Neu, and A. Lazaric. “Exploiting easy data in online optimization”. In: *Proceedings of the Twenty-Eighth Annual Conference on Neural Information Processing Systems (NIPS'14)*.

2014

Soare, M., A. Lazaric, and R. Munos. “Best-arm Identification in Linear Bandit”. In: *Proceedings of the Twenty-Eighth Annual Conference on Neural Information Processing Systems (NIPS'14)*.

2013

Gheshlaghi-Azar, M., A. Lazaric, and E. Brunskill. “Regret Bounds for Reinforcement Learning with Policy Advice”. In: *Proceedings of the European Conference on Machine Learning (ECML'13)*.

2013

Gheshlaghi-Azar, M., A. Lazaric, and E. Brunskill. “Sequential Transfer in Multi-arm Bandit with Finite Set of Models”. In: *Proceedings of the Twenty-Seventh Annual Conference on Neural Information Processing Systems (NIPS'13)*.

- 2012
Gabillon, Victor, Mohammad Ghavamzadeh, and Alessandro Lazaric. "Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence". In: *Proceedings of the Twenty-Sixth Annual Conference on Neural Information Processing Systems (NIPS'12)*.
- 2012
Gatti, Nicola, Alessandro Lazaric, and Francesco Trovò. "A Truthful Learning Mechanism for Multi-Slot Sponsored Search Auctions with Externalities (Extended Abstract)". In: *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*.
- 2012
Geist, Matthieu, Bruno Scherrer, Alessandro Lazaric, and Mohammad Ghavamzadeh. "A Dantzig Selector Approach to Temporal Difference Learning". In: *Proceedings of the Thirtieth International Conference on Machine Learning (ICML-2012)*.
- 2012
Ghavamzadeh, Mohammad and Alessandro Lazaric. "Conservative and Greedy Approaches to Classification-based Policy Iteration". In: *Proceedings of the 26th Conference on Artificial Intelligence (AAAI'12)*.
- 2012
Sani, Amir, Alessandro Lazaric, and Rémi Munos. "Risk Averse Multi-Arm Bandits". In: *Proceedings of the Twenty-Sixth Annual Conference on Neural Information Processing Systems (NIPS'12)*.
- 2012
Valko, Michal, Mohammad Ghavamzadeh, and Alessandro Lazaric. "Semi-Supervised Inverse Reinforcement Learning". In: *Proceedings of the 10th European Workshop on Reinforcement Learning*.
- 2011
Carpentier, Alexandra, Alessandro Lazaric, Mohammad Ghavamzadeh, Rémi Munos, and Peter Auer. "Upper-confidence-bound Algorithms for Active Learning in Multi-armed Bandits". In: *Proceedings of the 22nd International Conference on Algorithmic Learning Theory (ALT'11)*.
- 2011
Gabillon, Victor, Alessandro Lazaric, Mohammad Ghavamzadeh, and Sebastien Bubeck. "Multi-bandit Best Arm Identification". In: *Proceedings of the Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS'11)*, pp. 2222–2230.
- 2011
Gabillon, Victor, Alessandro Lazaric, Mohammad Ghavamzadeh, and Bruno Scherrer. "Classification-based Policy Iteration with a Critic". In: *Proceedings of the Twenty-Eighth International Conference on Machine Learning (ICML-2011)*.
- 2011
Ghavamzadeh, M., A. Lazaric, R. Munos, and M. Hoffman. "Finite-Sample Analysis of Lasso-TD". In: *Proceedings of the Twenty-Eighth International Conference on Machine Learning (ICML-2011)*.
- 2011
Lazaric, Alessandro and Marcello Restelli. "Transfer from Multiple MDPs". In: *Proceedings of the Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS'11)*.
- 2010
Ghavamzadeh, M., A. Lazaric, O. A. Maillard, and R. Munos. "LSTD with Random Projections". In: *Advances in Neural Information Processing Systems*.

- 2010
Lazaric, Alessandro and Mohammad Ghavamzadeh. "Bayesian Multi-Task Reinforcement Learning". In: *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-2010)*.
- 2010
Lazaric, Alessandro, Mohammad Ghavamzadeh, and Rémi Munos. "Analysis of a Classification-based Policy Iteration Algorithm". In: *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-2010)*.
- 2010
Lazaric, Alessandro, Mohammad Ghavamzadeh, and Rémi Munos. "Finite-Sample Analysis of LSTD". In: *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-2010)*.
- 2010
Maillard, Odalric-Ambrym, Alessandro Lazaric, Mohammad Ghavamzadeh, and Rémi Munos. "Finite-Sample Analysis of Bellman Residual Minimization". In: *Proceedings of the Second Asian Conference on Machine Learning (ACML-2010)*.
- 2009
Lazaric, A. and R. Munos. "Hybrid Stochastic-Adversarial On-line Learning". English. In: *Proceedings of the Twenty-Second Annual Conference on Learning Theory (COLT-2009)*. Montreal, Canada.
- 2008
Bonarini, A., C. Caccia, A. Lazaric, and M. Restelli. "Batch Reinforcement Learning for Controlling a Mobile Wheeled Pendulum Robot". In: *IFIP International Conference on Artificial Intelligence (IFIP-AI-2008)*. Milan, Italy, pp. 151–160.
- 2008
Ferrante, E., A. Lazaric, and M. Restelli. "Transfer of task representation in reinforcement learning using policy-based proto-value functions". In: *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*. Estoril, Portugal: International Foundation for Autonomous Agents and Multiagent Systems, pp. 1329–1332.
- 2008
Lazaric, A., M. Quaresimale, and M. Restelli. "On the usefulness of opponent modeling: the Kuhn Poker case study". In: *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-08)*. Estoril, Portugal, pp. 1345–1348.
- 2008
Lazaric, A., M. Restelli, and A. Bonarini. "Reinforcement Learning in Continuous Action Spaces through Sequential Monte Carlo Methods". In: *Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems (NIPS-2008)*. Ed. by J.C. Platt, D. Koller, Y. Singer, and S. Roweis. Vancouver, Canada, pp. 833–840.
- 2008
Lazaric, A., M. Restelli, and A. Bonarini. "Transfer of samples in batch reinforcement learning". In: *Proceedings of the Twenty-Fifth Annual International Conference on Machine Learning (ICML-2008)*. Ed. by Andrew McCallum and Sam Roweis. Helsinki, Finland, pp. 544–551.
- 2007
Bonarini, A., A. Lazaric, and M. Restelli. "Piecewise Constant Reinforcement Learning for Robotic Applications". In: *Proceedings of the Fourth International Conference on Informatics in Control, Automation and Robotics (ICINCO)*. Angers, France, pp. 214–221.

- 2007
Bonarini, A., A. Lazaric, and M. Restelli. "Reinforcement Learning in Complex Environments Through Multiple Adaptive Partitions". In: *Proceedings of the Tenth Congress of the Italian Association for Artificial Intelligence (AI*IA)*. Rome, Italy, pp. 531–542.
- 2007
Lazaric, A., E. Munoz de Cote, N. Gatti, and M. Restelli. "Reinforcement Learning in Extensive Form Games with Incomplete Information: the Bargaining Case Study". In: *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-2007)*. Honolulu, Hawai'i, pp. 216–218.
- 2006
Bonarini, A., A. Lazaric, and M. Restelli. "Self-Development Framework for Reinforcement Learning Agents". In: *Proceedings of the Fifth International Conference on Development and Learning (ICDL)*. Bloomington, USA.
- 2006
Munoz de Cote, E., A. Lazaric, and M. Restelli. "Learning to cooperate in multi-agent social dilemmas". In: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-2006)*. Hakodate, Japan, pp. 783–785.

Chapters in Books and Collections

- 2011
Busoniu, L., A. Lazaric, M. Ghavamzadeh, R. Munos, R. Babuska, and B. De Schutter. "Least-squares methods for policy iteration". In: *Reinforcement Learning: State of the Art*. Ed. by M. Wiering and M. van Otterlo. Springer.
- 2011
Lazaric, A. "Transfer in Reinforcement Learning: a Framework and a Survey". In: *Reinforcement Learning: State of the Art*. Ed. by M. Wiering and M. van Otterlo. Springer.
- 2009
Gatti, N., A. Lazaric, and M. Restelli. "Towards Automated Bargaining in Electronic Markets: a Partially Two-Sided Competition Model". In: *Agent-Mediated Electronic Commerce X and Trading Agent Competition VI*. Lecture Notes in Computer Science. Berlin, Germany: Springer.
- 2008
Lazaric, A., E. Munoz de Cote, F. Dercole, and M. Restelli. "Bifurcation Analysis of Reinforcement Learning Agents". In: *Adaptive Agents and Multi-Agent Systems III*. Ed. by K. Tuyls, A. Nowe, Z. Guessoum, and D. Kudenko. Vol. 4865. Lecture Notes in Artificial Intelligence. Berlin, Germany: Springer-Verlag, pp. 129–144.
- 2006
Bonarini, A., A. Lazaric, and M. Restelli. "Incremental Skill Acquisition for Self-motivated Learning Animats". In: *Proceedings of the Ninth International Conference on Simulation of Adaptive Behavior (SAB-06)*. Vol. 4095. Lecture Notes in Artificial Intelligence. Berlin: Springer Verlag, pp. 357–368.

Workshop Publications

- 2010
Gabillon, Victor, Alessandro Lazaric, and Mohammad Ghavamzadeh. "Rollout Allocation Strategies for Classification-based Policy Iteration". In: *ICML 2010 Workshop on Reinforcement Learning and Search in Very Large Spaces*.

- 2009
Lazaric, A. and R. Munos. "Hybrid Stochastic-Adversarial On-line Learning". English. In: *ECML 2009 Workshop on Learning from non-IID data: Theory, Algorithms and Practice*. Bled, Slovenia.
- 2008
Gatti, N., A. Lazaric, and M. Restelli. "Towards Automated Bargaining in Electronic Markets: a Partially Two-Sided Competition Model". In: *Proceedings of the AAMAS Workshop on Agent-Mediated Electronic Commerce (AMEC-2008)*. Lisbon, Portugal, pp. 29–42.
- 2007
Lazaric, A., D. Loiacono, A. Prete, M. Restelli, and P.L. Lanzi. "Learning Driving Tasks in TORCS Using Reinforcement Learning". In: *NIPS 2007 Workshop on Machine Learning and Games*.
- 2007
Munoz de Cote, E., A. Lazaric, M. Restelli, and A. Bonarini. "A Learning Approach to Dynamic Coalition Formation". In: *Proceedings of the AAMAS Workshop on Adaptive and Learning Agents (ALAg-07)*. Honolulu, Hawa'i.
- 2006
Bonarini, A., A. Lazaric, and M. Restelli. "Learning in Complex Environments through Multiple Adaptive Partitions". In: *Proceedings of the ECAI Workshop on Planning, Learning and Monitoring with Uncertainty and Dynamic Worlds*. Riva del Garda, Italy.
- 2006
Bonarini, A., A. Lazaric, and M. Restelli. "Learning Reusable Skills through Self-Motivation". In: *Proceedings of the ICML Workshop on Structural Knowledge Transfer for Machine Learning*. Pittsburgh, USA.
- 2005
Bonarini, A., A. Lazaric, E. Munoz de Cote, and M. Restelli. "Improving Cooperation among Self-Interested Reinforcement Learning Agents". In: *Proceedings of the ECML Workshop on Reinforcement Learning in Non-Stationary Environments*. Porto, Portugal.
- 2005
Bonarini, A., A. Lazaric, and M. Restelli. "LEAP: Learning Entities Adaptive Partitioning". In: *Proceedings of the NIPS Workshop on Reinforcement Learning Benchmarks and Bake-offs II*.
- 2005
Bonarini, A., A. Lazaric, and M. Restelli. "Learning Optimal Policies with State Aggregation". In: *Proceedings of the Seventh European Workshop on Reinforcement Learning (EWRL)*.
- 2005
Bonarini, A., A. Lazaric, and M. Restelli. "Yahtzee: a Large Stochastic Environment for RL Benchmarks". In: *Proceedings of the NIPS Workshop on Reinforcement Learning Benchmarks and Bake-offs II*.