

# Anna Harutyunyan

---

CONTACT	Artificial Intelligence Lab Vrije Universiteit Brussel Pleinlaan 2 B-1050 Brussels Belgium	aharutyu@vub.ac.be anna.harutyunyan.net +32 2 629 33 08
EDUCATION	<b>PhD, Computer Science: Artificial Intelligence</b> Vrije Universiteit Brussel, Brussels, Belgium Keywords: <b>reinforcement learning</b> , reward shaping, off-policy learning Advisor: Ann Nowé	<b>2013 - present</b>
	<b>MSc, Computer Science: Algorithms</b> Oregon State University, Oregon USA Thesis: “Maximum Flow in Planar Digraphs” Advisor: Glencora Borradaile Grade: Magna Cum Laude	<b>2011 - 2012</b>
	<b>BSc, Computer Science, Mathematics</b> Utah State University, Utah USA Grade: Magna Cum Laude	<b>2007 - 2010</b>
AWARDS AND HONORS	<b>FoCAS Best Paper Award</b> , ALA Workshop at AAMAS For the paper “Off-Policy Shaping Ensembles in Reinforcement Learning”	<b>2015</b>
	<b>Graduate Research Fellowship</b> During MSc studies at Oregon State University	<b>2010 - 2011</b>
	<b>CRA Distributed Research Experience for Undergraduates Grant</b> Toward a 10-week summer research internship at Oregon State University	<b>2010</b>
	<b>Dean’s List Award for Outstanding Scholastic Achievements</b> During BSc studies at Utah State University	<b>2008 - 2010</b>
	<b>CRLA/ITPC Regular Tutoring Certificate (Mathematics)</b> Academic Learning Center at Utah State University	<b>2009</b>
	<b>ISEP Study Abroad Scholarship</b> With placement at the University of Chester, England	<b>2009</b>
	<b>Four-Year Full Academic Scholarship, Utah State University</b> Granted by the Huntsman Foundation	<b>2007 - 2010</b>

TALKS	<i>Potential-based reward shaping as a tool to safely incorporate auxiliary information</i> Sequel seminar, INRIA Lille	2017
	<i>Off-Policy TD Learning from Returns</i> Reasoning and Learning Lab, McGill University	2016
	<i>Off-Policy Shaping Ensembles in Reinforcement Learning</i> 14th Adaptive Learning Agents (ALA) Workshop at AAMAS	2015
	<i>Parallel Reward Shaping Architectures</i> 9th Barbados Workshop on Reinforcement Learning	2015
	<i>Expressing Arbitrary Reward Functions as Potential-Based Advice</i> 29th Conference on Artificial Intelligence (AAAI)	2015
	<i>Reward Shaping Architectures in Reinforcement Learning</i> Reasoning and Learning Lab, McGill University	2014
	<i>Boundary-to-Boundary Flows in Planar Graphs</i> 24th International Workshop on Combinatorial Algorithms (IWOCA)	2013
	<i>An Adaptive Approach to Finding Maximum Flow in Planar Graphs</i> REU Seminar, Oregon State University	2012
TEACHING	<b>Courses at Vrije Universiteit Brussel</b>	
	• Multi-Agent Learning Seminar (MSc)	2013 - 2017
	• Theory of Computation (MSc)	2014 - 2016
	<b>Courses at Oregon State University</b>	
	• Advanced Algorithms (MSc)	Fall 2011
	• Algorithm Analysis (BSc)	Winter 2012
	<b>Thesis Students Advised</b>	
	• Timothy Verstraeten (MSc) Title: “Modeling Exoskeleton-Assisted Human Motion Using Gaussian Processes” Graduated with the highest distinction.	2014 - 2015
ACADEMIC SERVICE	<b>Organization</b>	
	Autonomous Learning Agents (ALA) Workshop at AAMAS 2017	
	The Future of Interactive Learning Machines (FILM) Workshop at NIPS 2016	
	<b>Program Committee</b>	
	AAAI 2017, IJCAI 2016-2017, ECAI 2016, ALA Workshop at AAMAS 2015-2016	
	<b>Reviewing and sub-reviewing</b>	
	NIPS 2016, AAMAS 2014-2015, ICML 2014, IJCNN 2014, UKCI 2014, Benelearn 2014, Humanoids 2014, LATIN 2014, SAC 2014, ...	
RESEARCH VISITS AND ACADEMIC TRAINING	Deep Learning Summer School University of Montreal, Montreal, Canada	2016

Prof. Doina Precup  
McGill University, Montreal, Canada **Fall 2014**  
(1 month)

Worked on variants of the Horde of shapings architecture.

PROFESSIONAL  
EXPERIENCE

**Research Intern** **Fall 2015**  
Google DeepMind, London, UK

Worked with Rémi Munos on designing new off-policy reinforcement learning algorithms. Part of this internship work resulted in the ALT 2016 and NIPS 2016 publications.

**Summer Research Intern** **Summer 2010**  
Computer Research Association (CRA-W), Corvallis, Oregon USA

Developed an experimental algorithm with Prof. Glencora Borradaile at Oregon State University for solving a class of 2D alignment problems, which later, for its application to spreadsheets, led to a Google Research Award, and an academic publication (VL/HCC-12).

**Quality Assurance Engineer** **2010**  
Department of Computer Science, Utah State University

Maintained and tested a database with two other students. Responsibilities included QA, data entry, and designing automated test suites.

**Senior Mathematics Tutor** **2008 - 2010**  
Academic Learning Center, Utah State University

Tutored the entire undergraduate mathematics curriculum from pre-algebra to differential equations.

TECHNICAL  
SKILLS

Matlab, Python, Lua, C/C++, Java, L<sup>A</sup>T<sub>E</sub>X, bash, version control

LANGUAGES

Native: Russian and Armenian  
Fluent: English  
Elementary proficiency: French, Dutch

PUBLICATIONS

**Journal**

- [1] K. Tanghe, **Harutyunyan, A.**, E. Aertbelien, F. De Groote, J. De Schutter, P. Vrancx, and A. Nowé. Predicting Seat-Off and Detecting Start-of-Assistance Events for Assisting Sit-to-Stand with an Exoskeleton. *Robotics and Automation Letters, IEEE*, 1(2):792 – 799, 2016.

**Conference**

- [2] Rémi Munos, Tom Stepleton, **Anna Harutyunyan**, and Marc G. Bellemare. Safe and efficient off-policy reinforcement learning. In *Proceedings of Neural Information Processing Systems (NIPS)* (To Appear), 2016.
- [3] **Anna Harutyunyan**, Marc G. Bellemare, Tom Stepleton, and Rémi Munos. Q( $\lambda$ ) with off-policy corrections. In *Proceedings of Algorithmic Learning Theory (ALT)* (To Appear), 2016.
- [4] Tim Brys, **Anna Harutyunyan**, Halit Bener Suay, Sonia Chernova, Matthew E. Taylor, and Ann Nowé. Reinforcement Learning from Demonstration through Shaping. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

- [5] **Anna Harutyunyan**, Sam Devlin, Peter Vrancx, and Ann Nowé. Expressing Arbitrary Reward Functions as Potential-Based Advice. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 2652–2658, 2015.
- [6] Tim Brys, **Anna Harutyunyan**, Matthew E. Taylor, and Ann Nowé. Policy Transfer using Reward Shaping. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 181–188, 2015.
- [7] Tim Brys, **Anna Harutyunyan**, Peter Vrancx, Matthew E. Taylor, Daniel Kudenko, and Ann Nowé. Multi-Objectivization of Reinforcement Learning Problems by Reward Shaping. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 2315–2322, 2014.

#### Conference (Short)

- [8] **Anna Harutyunyan**, Tim Brys, Peter Vrancx, and Ann Nowé. Multi-Scale Reward Shaping via an Off-Policy Ensemble. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1641–1642, 2015.
- [9] **Anna Harutyunyan**, Tim Brys, Peter Vrancx, and Ann Nowé. Off-Policy Shaping Ensembles in Reinforcement Learning. In *Proceedings of the Twenty-First European Conference on Artificial Intelligence (ECAI)*, pages 1021–1022, 2014.

#### Refereed Workshop and Symposium

- [10] Timothy Verstraeten, Roxana Radulescu, Yannick Jadoul, Tom Jaspers, Robrecht Conjaerts, Tim Brys, **Anna Harutyunyan**, Peter Vrancx, and Ann Nowé. Human-Guided Ensemble Learning in StarCraft. In *Autonomous Learning Agents (ALA) Workshop at AAMAS*, 2016.
- [11] **Anna Harutyunyan**, Tim Brys, Peter Vrancx, and Ann Nowé. Off-Policy Reward Shaping with Ensembles. In *Autonomous Learning Agents (ALA) Workshop at AAMAS*, 2015.
- [12] Glencora Borradaile and **Anna Harutyunyan**. Boundary-to-boundary flows in planar graphs. In *Proceedings of the International Workshop on Combinatorial Algorithms (IWOCA)*, volume 8288 of *Lecture Notes in Computer Science*, pages 67–80. Springer, 2013.
- [13] Glencora Borradaile and **Anna Harutyunyan**. Maximum st-flow in directed planar graphs via shortest paths. In *Proceedings of the International Workshop on Combinatorial Algorithms (IWOCA)*, volume 8288 of *Lecture Notes in Computer Science*, pages 423–427. Springer, 2013.
- [14] **Anna Harutyunyan**, Glencora Borradaile, Chris Chambers, and Christopher Scaffidi. Planted-model evaluation of algorithms for identifying differences between spreadsheets. In *Proceedings of the IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pages 7–14, 2012.

#### Miscellaneous

- [15] **Anna Harutyunyan**, Tim Brys, Peter Vrancx, and Ann Nowé. Shaping Mario with Human Advice (Demonstration). In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1913–1914, 2015.

#### Thesis

- [16] Anna Harutyunyan. *Maximum Flow in Planar Digraphs*. Oregon State University, Corvallis OR, USA, November 2012.