

Florent Delgrange

Researcher in Reinforcement Learning and World Models
Designing principled RL methods that enable agents to plan, adapt, and
improve reliably by leveraging world models

✉ florent.delgrange@ai.vub.ac.be

🆔 0000-0003-2254-0596

🌐 [florentdelgrange](https://florentdelgrange.com)

🌐 delgrange.me

Work Experience

- 2024 – Today **Postdoctoral Researcher**, *AI Lab, VUB – Vrije Universiteit Brussel*, Belgium, since September 2024
Reinforcement learning and world models for reliable, adaptive agents, with Ann Nowé and Guillermo A. Pérez.
- Developed **RL** methods leveraging **world models** to reliably improve agents' **adaptability**.
 - Designed **compositional methods** that combine learning agents with **reactive** and **program synthesis** to support scalable, high-performing agent behavior in large-scale environments.
 - Ongoing work on agents that combine **foundation models** for task acquisition, **RL** for **test-time adaptation**, and **agentic orchestration** for task decomposition.
 - **Temporary guest professor**: lecturer of the course *Theory of Computation* at VUB.
- 2025 **Visiting Researcher**, *Delft University of Technology*, The Netherlands, 2 months
Worked with Prof. Frans A. Oliehoek on **reinforcement learning**, **abstractions**, and **world models** in the Sequential Decision Making group.
- 2020 – 2024 **Doctoral Researcher**, *AI Lab, VUB & University of Antwerp (UAntwerpen)*, Belgium
Formal verification and representation learning for deep reinforcement learning, supervised by Ann Nowé (VUB) and Guillermo A. Pérez (UAntwerpen).
- Introduced **latent abstraction** methods in learned world models, through Wasserstein/variational auto-encoders, to **enable model checking** of deep RL policies.
 - Extended this work to **representation learning** in RL for **partially observable** environments.
- 2018 – 2019 **Scientific Researcher**, *RWTH Aachen University & UMONS – Université de Mons*, Germany and Belgium, 1 year
Many-sided synthesis for stochastic systems, with Joost-Pieter Katoen and Mickael Randour.
- Developed novel **reactive synthesis** approaches for stochastic systems and **multi-objective MDPs**.
- 2017 **Internship**, *Nokia*, Antwerp, Belgium, 3 months
Machine learning to detect, identify, and troubleshoot several impairments impacting data transmission technologies (such as DSL lines).

Education

- 2020 – 2024 **Doctor of Science**, *Computer Science*, VUB & UAntwerpen, Belgium
- **Title of the dissertation**: Activating Formal Verification of Deep Reinforcement Learning Policies by Model Checking Bisimilar Latent Space Models.
- 2016 – 2018 **Master of Science**, *Computer Science*, UMONS, Belgium
- **Grade**: *Magna Cum Laude* (Ranking: 1/15).
 - **Master's thesis**: Multi-objective synthesis in Markov decision processes, 19/20.
 - **Master's project**: Stochastic shortest path in Markov decision processes, 18.5/20.

- 2013 – 2016 **Bachelor of Science, Computer Science**, UMONS, Belgium
- **Grade:** *Cum Laude*.
 - **Research internship:** UML statecharts visualizer for Sismic, a statechart library for Python, in the software engineering lab, 1 month.

Research Interests

- Core Reinforcement learning and representation learning for reliable, adaptive agents.
- Methods World models, formal verification, abstraction, and reactive/program synthesis for certifiable decision-making.
- Themes Foundation models for task acquisition, test-time adaptation and generalization, agentic orchestration, multi-objective decision making, uncertainty and partial observability, and deep generative modeling.

Honors and Awards

- 2026 **Best Blue Sky Paper Award**, The 25th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2026, [verified computing research association certificate](#)
- 2025 **Top Reviewer**, *The Thirty-Ninth Annual Conference on Neural Information Processing Systems*, NeurIPS 2025
- 2025 **International mobility grant for a long stay abroad**, FWO Research Foundation – Flanders, for visiting Prof. Frans Oliehoek at TU Delft
- 2025 **VUB Research Council backup mandate**, for being ranked in the top 20% of applicants for the FWO junior fellowship (1308 applicants), 1 year
- 2025 **Best Poster Award at BeNeRL**, Belgium-Netherlands Workshop on Reinforcement Learning, Eindhoven University of Technology
- 2025 **Distinguished PC Member**, The 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025, top 5%
- 2019 **Best Paper Award (CONCUR 2019)**, The 30th International Conference on Concurrency Theory, Amsterdam, The Netherlands
- 2019 **Best MIMA Poster Award (Mathematics, Information technology, Modeling and Applications)**, MdC 2019: Mardi des Chercheurs, UMONS & Université Polytechnique Hauts-de-France
- 2018 **Award for the best Master's thesis in Computer Science**, UMONS
- 2016 **Research Initiation Grant**, UMONS, in the Software engineering lab, 1 grant per department.

Teaching

- 2024 – 2025 **Theory of Computation**, *Lecturer*, VUB
CS: Master (3 ECTS, 28 hours).
- 2021 – 2024 **Theory of Computation**, *Teaching Assistant*, VUB
CS: Master. Exercise sessions (12 hours / year).
- 2018 – 2019 **Formal methods for system design**, *Teaching Assistant*, UMONS
CS and Maths: Master. Exercise sessions (15 hours).

2018 **University certificate in artificial intelligence (“Hands on AI”)**, *Teaching Assistant*, UMONS

Non-university persons, engineers and science students: Master and PhD. Reinforcement learning practical sessions (12 hours).

2018 **Computer programming and algorithms II**, *Student teaching assistant*, UMONS
CS and Maths: 1st year. Practical sessions (15 hours) in Java 8.

2017 **Computer programming and algorithms I**, *Student teaching assistant*, UMONS
CS, Maths and Physics: 1st and 2nd year. Practical sessions (20 hours) in Python 3.

Other

Aug *Student Tutor*, UMONS
2013–2014 Gave computer science lectures to university students retaking exams.

Students

Bachelor Thesis, VUB, *CS: Bachelor*

Olivier Delattre (2020 – 2021), Cédric Simon Segers (2020 – 2021), Abdoullah El Yachouti (2021 – 2022)

Master’s project & Scientific reading and writing, UMONS, *CS: Master*

Dimitri Waelkens (2018 – 2019)

Academic Service and Community Engagement

Program committee

Conferences AAAI (2023), AAMAS (2024, 2025: **distinguished PC member**), and NeurIPS (2024, 2025: **top reviewer**)

Workshops EWRL (2023), ALA (2025), and MODeM (2024)

Reviewer

Conferences RP (2019), SETTA (2019), and AAMAS (2020, 2023)

Journals LMCS (2024)

Organization

2023 **The Sixteenth European Workshop on Reinforcement Learning**, *Organizing committee*, VUB

Local organization of EWRL 2023 at VUB, attracting 235 researchers.

2026 **The Adaptive and Learning Agents Workshop**, *Organizing committee*, AAMAS 2026 Workshop, Paphos, Cyprus

Publications

Peer-reviewed conference proceedings

AAMAS 2026 **Foundation World Models for Agents that Learn, Verify, and Adapt Reliably Beyond Static Environments**, *Florent Delgrange*, In Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems, May 25 – 29, 2026, Paphos, Cyprus, **Best Blue Sky Paper Award**; DOI: 10.65109/WCEI7331; [verified computing research association certificate](#)

- ICLR 2026 **Deep SPI: Safe Policy Improvement via World Models**, *Florent Delgrange, Raphaël Avalos, Willem Röpke*, The Fourteenth International Conference on Learning Representations, 2026
- ALA 2026 **CTDE2: Continuous Training Discrete Execution**, *Nicolas Rowies, Florent Delgrange, Ann Nowé, Diederik M. Roijers*, Proceedings of the Adaptive and Learning Agents Workshop at AAMAS 2026, Paphos, Cyprus
PDF
- AAMAS 2025 **Composing Reinforcement Learning Policies, with Formal Guarantees**, *Florent Delgrange, Guy Avni, Anna Lukina, Christian Schilling, Ann Nowé, Guillermo A. Pérez*, AAAI Track. In Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems, Detroit, Michigan, USA, May 19 – 23, IFAAMAS, 10 pages, 2025
- ICLR 2024 **The Wasserstein Believer: Learning Belief Updates for Partially Observable Environments through Reliable Latent Space Models**, *Raphael Avalos[†], Florent Delgrange[†], Ann Nowé, Guillermo A. Pérez, Diederik M. Roijers*, The Twelfth International Conference on Learning Representations, 2024, [†]*equal contribution*;
Also appears in the following peer-reviewed international workshop proceedings:
- The Sixteenth European Workshop on Reinforcement Learning (EWRL 2023), and
 - The Adaptive and Learning Agents Workshop at AAMAS 2023 (ALA 2023).
- ICLR 2023 **Wasserstein Auto-encoded MDPs: Formal Verification of Efficiently Distilled RL Policies with Many-sided Guarantees**, *Florent Delgrange, Ann Nowé, Guillermo A. Pérez*, The Eleventh International Conference on Learning Representations, 2023
Also appears in the peer-reviewed proceedings of The Sixteenth European Workshop on Reinforcement Learning (EWRL 2023)
- AAAI 2022 **Distillation of RL Policies with Formal Guarantees via Variational Abstraction of Markov Decision Processes**, *Florent Delgrange, Ann Nowé, Guillermo A. Pérez*, Proceedings of the AAAI Conference on Artificial Intelligence Vol. 36 No. 6: AAAI-22 Technical Tracks 6, 6497-6505, 2022
- TACAS 2020 **Simple Strategies in Multi-Objective MDPs**, *Florent Delgrange**, *Joost-Pieter Katoen, Tim Quatmann, Mickael Randour*, Tools and Algorithms for the Construction and Analysis of Systems - 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I
- CONCUR 2019 **Life is Random, Time is Not: Markov Decision Processes with Window Objectives**, *Thomas Brihaye, Florent Delgrange**, *Youssef Oualhadj, Mickael Randour*, 30th International Conference on Concurrency Theory, August 27-30, 2019, Amsterdam, the Netherlands, **Best Paper Award**
- [Peer-reviewed journals](#)
- NCAA **A Framework for Flexibly Guiding Learning Agents**, *Mahmoud Elbarbari, Florent Delgrange, Ivo Vervlimmeren, Kyriakos Efthymiadis, Bram Vanderborght, Ann Nowé*, Neural Computing and Applications, Special Issue on Adaptive and Learning Agents, 2022

LMCS **Life is Random, Time is Not: Markov Decision Processes with Window Objectives**, *Thomas Brihaye, Florent Delgrange**, *Youssef Oualhadj, Mickael Randour*, Logical Methods in Computer Science, Volume 16, 2020

[Other peer-reviewed international workshop proceedings](#)

ALA 2025 **Integrating RL and Planning through Optimal Transport World Models**, *Willem Röpke, Raphaël Avalos, Roxana Rădulescu, Ann Nowé, Diederik M Roijers, Florent Delgrange*, Proceedings of the Adaptive and Learning Agents Workshop, 2025

EWRL 2024 **Controller Synthesis from Deep Reinforcement Learning Policies**, *Florent Delgrange, Guy Avni, Anna Lukina, Christian Schilling, Guillermo A. Pérez, Ann Nowé*, Seventeenth European Workshop on Reinforcement Learning, 2024
Also appears in the following peer-reviewed international workshop proceedings:

- The PRL workshop series on *Bridging the Gap Between AI Planning and Reinforcement Learning* at AAI 2025 (PRL 2025),
- The Multi-Objective Decision Making Workshop at ECAI 2024 (MODeM 2024).

ALA 2023 **WAE-PCN: Wasserstein-autoencoded Pareto Conditioned Networks**, *Florent Delgrange[†], Mathieu Reymond[†], Ann Nowe, Guillermo A. Pérez*, Proceedings of the Adaptive and Learning Agents Workshop, 2023, [†]*equal contribution*.

* *alphabetical order*.

Programming skills

Languages Python, Java, C++, OCaml, Scala
ML libs JAX, TensorFlow[-agents, -probability].
Tools Model checkers PRISM and STORM

Attended events

Talks

- **Foundation World Models for Agents that Learn, Verify, and Adapt Reliably Beyond Static Environments**
 - AAMAS 2026: 25th International Conference on Autonomous Agents and Multiagent Systems (**Best Blue Sky Paper Award**, *Paphos, Cyprus*)
- **Composing Reinforcement Learning Policies, with Formal Guarantees**
 - SYNT 2025: 14th workshop on synthesis (CAV 2025 workshop, *Zagreb, Croatia*)
 - AAMAS 2025: 24th International Conference on Autonomous Agents and Multiagent Systems (*Detroit, USA*)
 - the PRL workshop series on *Bridging the Gap Between AI Planning and Reinforcement Learning* (AAAI 2025 workshop, *Philadelphia, USA*)
- **The Wasserstein Believer: Learning Belief Updates for Partially Observable Environments through Reliable Latent Space Models**
 - BNAIC/BeNeLearn 2024: Joint International Scientific Conferences on AI and Machine Learning (2024, *Utrecht, The Netherlands*)
 - Invited talk at the University of Liège when visiting Damien Ernst
- **Activating Formal Verification of Deep RL Policies by Model Checking Bisimilar Latent Space Models**

- CFV seminar (2024, *ULB – Université Libre de Bruxelles*)
- Flanders AI Day 2024 (*Ghent, Belgium*)
- Dagstuhl seminar on Stochastic Games (2024)
- LiVe: Learning in Verification Workshop (ETAPS 2024 satellite event, *Luxembourg*)
- **Wasserstein-autoencoded Pareto Conditioned Networks**
 - ALA: Adaptive and Learning Agents Workshop (2023, AAMAS workshop, *London, UK*)
- **Wasserstein Auto-encoded MDPs: Formal Verification of Efficiently Distilled RL Policies with Many-sided Guarantees**
 - BNAIC/BeNeLearn 2022: Joint International Scientific Conferences on AI and Machine Learning (2022, *Mechelen, Belgium*)
- **Distillation of RL Policies with Formal Guarantees via Variational Abstraction of Markov Decision Processes**
 - Logic and AI workshop (2022, AI Flanders Day, *Leuven*)
 - Invited talk at TU Delft, when visiting Anna Lukina
 - iVerif: Workshop on artificial intelligence and Verification (2021, FSTTCS Virtual Pre-Conference Workshop)
- **Life is Random, Time is not: Markov Decision Processes with Window Objectives.**
 - CONCUR (2019, *Amsterdam, The Netherlands*)
 - MoRe: Multi-objective Reasoning in Verification and Synthesis (2019, LICS workshop, *SFU, Vancouver, Canada*)
 - UnRAVeL: invited talk (2019, *RWTH Aachen, Germany*)
 - CFV seminar (2019, *ULB – Université Libre de Bruxelles*).

Posters

- **Deep SPI: Safe Policy Improvement via World Models**
 - ICLR 2026: The Fourteenth International Conference on Learning Representations (2026, *Rio de Janeiro, Brazil*)
- **Continuous Training Discrete Execution**
 - ALA 2026: Adaptive and Learning Agents Workshop (AAMAS workshop, *Paphos, Cyprus*)
- **Composing Reinforcement Learning Policies, with Formal Guarantees**
 - BeNeRL 2025 (**Best Poster Award**, *TU Eindhoven, The Netherlands*)
 - BNAIC/BeNeLearn 2024 (*Utrecht, The Netherlands*)
 - EWRL 2024: Seventeenth European Workshop on Reinforcement Learning (*Toulouse, France*)
 - Multi-Objective Decision Making Workshop at ECAI 2024 (*Santiago de Compostela, Spain*)
- **The Wasserstein Believer: Learning Belief Updates for Partially Observable Environments through Reliable Latent Space Models**
 - BNAIC/BeNeLearn 2024 (*Utrecht, The Netherlands*)
 - BeNeRL 2024: the Belgium-Netherlands workshop on Reinforcement Learning (2024, *Amsterdam, The Netherlands*)
 - ICLR 2024: The Twelfth International Conference on Learning Representations (2024, *Vienna, Austria*)
 - EWRL: European Workshop on Reinforcement Learning (2023, *Brussels, Belgium*)

- ALA 2023 (*London, UK*)
- **Wasserstein-autoencoded Pareto Conditioned Networks**
 - ALA 2023 (*London, UK*)
- **Wasserstein Auto-encoded MDPs: Formal Verification of Efficiently Distilled RL Policies with Many-sided Guarantees**
 - EWRL 2023 (*Brussels, Belgium*)
 - ICLR 2023: The Eleventh International Conference on Learning Representations (2023, *Kigali, Rwanda*)
 - BNAIC/BeNeLearn 2022 (*Mechelen, Belgium*)
- **Distillation of RL Policies with Formal Guarantees via Variational Abstraction of Markov Decision Processes**
 - Highlights of logic, games and automata (2022, conference, *Paris, France*)
 - BeNeRL: the Belgium-Netherlands workshop on Reinforcement Learning (2022, *Leiden University*)
 - AAAI-2022 (2022, *Virtual*)
 - iVerif (2021, *FSTTCS Virtual Workshop*)
 - AI Flanders Research Days (2021, *Virtual, Belgium*)
- **Life is Random, Time is not: Markov Decision Processes with Window Objectives.**
 - Theory and Algorithms in Graph and Stochastic Games (2019, *Mons, Belgium*)
 - MdC: Mardi des Chercheurs (**Best Poster Award**, 2019, *UMONS*).

Full list

2026

- **ICLR 2026** (conference, *Rio de Janeiro, Brazil*)
- **AAMAS 2026** (conference, **Best Blue Sky Paper Award**, *Paphos, Cyprus*)
- **ALA 2026: Adaptive and Learning Agents Workshop** (AAMAS workshop, organizer, *Paphos, Cyprus*)

2025

- **SYNT 2025: 14th workshop on synthesis** (CAV 2025 workshop, *Zagreb, Croatia*)
- **SAIV 2025: The 8th International Symposium on AI Verification** (symposium, *Zagreb, Croatia*)
- **BeNeRL 2025** (workshop, *TU Eindhoven, the Netherlands*)
- **AAMAS 2025** (conference, *Detroit, USA*)
- **PRL 2025: Bridging the Gap Between AI Planning and Reinforcement Learning** (AAAI workshop, *Philadelphia, USA*)

2024

- **LiVe 2024: Learning in Verification** (ETAPS workshop, *Luxembourg*)
- **ICLR 2024** (conference, *Vienna, Austria*)
- **Schloss Dagstuhl**, seminar on Stochastic Games (*Germany*)
- **BeNeRL 2024** (workshop, *Amsterdam, The Netherlands*)
- **AI Flanders Research Day** (*Gent, Belgium*)
- **MODeM 2024** (ECAI 2024 workshop, *Santiago de Compostela, Spain*)
- **EWRL 2024** (workshop, *Toulouse, France*)
- **Meta's GenAI Forum** (workshop, by invitation only, *Brussels, Belgium*)
- **BNAIC/BeNeLearn 2024** (conference, *Utrecht, The Netherlands*)

2023

- **ICLR 2023** (conference, *Kigali, Rwanda*)
- **AAMAS 2023** (conference, *London, UK*)
- **EWRL 2023** (workshop, *Brussels, Belgium*)

2022

- **BNAIC/BeNeLearn 2022** (conference, *Mechelen, Belgium*)
- **EWRL 2022**: The 15th European Workshop on Reinforcement Learning (workshop, *Milano, Italy*)
- **Highlights of Logic, Games, and Automata** (conference, *Paris, France*)
- **AI Flanders Research Day** (*Leuven, Belgium*)
- **BeNeRL: the Belgium-Netherlands workshop on Reinforcement Learning** (workshop, *Leiden, the Netherlands*)
- **AAAI-2022** (conference, *Virtual*)

2021

- **iVerif** (FSTTCS Virtual Workshop)
- **AAMAS: 20th International Conference on Autonomous Agents and Multiagent Systems** (*Virtual*)
- **ALA: Adaptive and Learning Agents** (AAMAS Virtual Workshop)
- **AI Flanders Research Days** (*Online, Belgium*)

2019

- **CONCUR: 30th International Conference on Concurrency Theory** (conference, *Amsterdam, The Netherlands*)
- **LearnAut: Learning and Automata** (LICS workshop, *SFU, Vancouver, Canada*)
- **MoRe** (LICS workshop, *SFU, Vancouver, Canada*)
- **FMAI: Formal Methods and AI** (workshop, *IRISA Rennes, France*)
- **SYNT Camp** (ETAPS workshop, *Prague, Czech Republic*)
- **LiVe: Learning in Verification** (ETAPS workshop, *Prague, Czech Republic*)
- **Theory and Algorithms in Graph and Stochastic Games** (workshop, *Mons, Belgium*)
- **MdC: Mardi des Chercheurs** (*UMONS, Mons, Belgium*)

2018

- **Highlights of Logic, Games, and Automata** (*conference, Berlin, Germany*)
- **MOVEP: Modeling and verification of parallel processes** (*summer school, ENS Paris-Saclay, Cachan, France*)
- **MoRe** (FLoC workshop, *University of Oxford, UK*)
- **Logic and Learning** (FoPSS summer school, *University of Oxford, UK*)
- **Logic and Learning** (workshop, *The Alan Turing Institute, London, UK*)

Miscellaneous

Pop. science

- Pint of Science, *organization of the event, Mons, 2019.*
- Spring of science: "How Mario learns to play", *popular introduction to reinforcement learning, Mons, 2019.*
- Math & science days: "It's your turn to play!", *popular talk about game theory, UMONS, 2019.*

Activities and
Societies

During studies: member of the Computer Science Club, General Association of Science Students, council of the Faculty of Science, and council of the Department of Computer Science, *UMONS, 2013 – 2018.*