

CURRICULUM VITAE — GERGELY NEU

Date of birth: December 24, 1984
Address: DTIC, Universitat Pompeu Fabra
C/ Roc Boronat 138
08018 Barcelona, Spain
Web: <http://cs.bme.hu/~gergo>
E-mail: gergely.neu@gmail.com

Appointments:

10/2018–	Research assistant professor (untenured*), Artificial Intelligence group Department of Information and Communication Technologies Universitat Pompeu Fabra Barcelona, Spain
10/2015– 09/2018	Postdoctoral researcher, Artificial Intelligence group Department of Information and Communication Technologies Universitat Pompeu Fabra Barcelona, Spain
09/2013–08/2015	Postdoctoral researcher, SequeL team, INRIA Lille – Nord Europe Villeneuve d’Ascq, France
09/2008–06/2013	PhD Student, Department of Computer Science and Information Theory Budapest University of Technology and Economics Budapest, Hungary

Academic Qualifications:

2013	PhD. in Computer Science Thesis title: <i>Online learning in non-stationary Markov decision processes.</i> Advisors: András György (MTA SZTAKI, Hungary), Csaba Szepesvári (University of Alberta, Canada) and László Györfi (Budapest University of Technology and Economics, Hungary) Budapest University of Technology and Economics Budapest, Hungary
2008	MSc. in Electrical Engineering Thesis title: <i>Inverse reinforcement learning and its application to learning to parse.</i> Advisors: Csaba Szepesvári (University of Alberta, Canada) and László Györfi (Budapest University of Technology and Economics, Hungary) Budapest University of Technology and Economics Budapest, Hungary

PhD students

2017–2022	Julia Okhovskaya (UPF $\rightsquigarrow$ assistant professor at TU Delft)
2017–2022	Joan Bas-Serrano (UPF $\rightsquigarrow$ research engineer at SocialPoint)
2019–2024	Germano Gabbianelli (UPF $\rightsquigarrow$ research engineer at a stealth startup)
2021–2025	Nneka Okolo (UPF $\rightsquigarrow$ postdoc at MIT)
2021–	Antoine Moulin (UPF)
2021–	Ludovic Schwartz (UPF)
2023–	Baptiste Abélès (UPF)
2021–	Luca Viano (EPFL, co-supervised with Volkan Cevher)
2021–	Maximillian Müller (U. Tübingen, co-supervised with Matthias Hein)
2021–	Lukas Zierahn (U. Milan, co-supervised with Nicolò Cesa-Bianchi)
2024–	Adrian Müller (ETH Zürich, co-supervised with Niao He)

* Employed on a temporary “distinguished researcher” contract since 2022. See attached contract.

Postdocs:

2019–2020	Ciara Pike-Burke (UPF $\rightsquigarrow$ lecturer at Imperial College London)
2020–2021	Mastane Achab (UPF $\rightsquigarrow$ founder at DeepGambit, an AI startup)
2021–2023	Matteo Papini (UPF $\rightsquigarrow$ assistant professor at Politecnico di Milano)
2023–	Eugenio Clerico (UPF)
2024–	Hamish Flynn (UPF)

Visiting positions:

11/2022–12/2022	Visiting researcher at the Simons Institute for the Theory of Computing for the program on Data-Driven Decision Processes, Berkeley CA, USA
08/2020–12/2020	Visiting researcher at the Simons Institute for the Theory of Computing for the program on Theory of Reinforcement Learning, Berkeley CA, USA (virtual)
06/2018–09/2018	Visiting researcher at Google Brain, Zürich, Switzerland
02/2011–06/2011	Visiting PhD student at the University of Alberta, Edmonton AB, Canada
08/2009–10/2009	Visiting PhD student at the University of Alberta, Edmonton AB, Canada
02/2008–06/2008	Visiting MSc student at the University of Alberta, Edmonton AB, Canada

Research Grants and Awards:

10/2021–09/2026	ERC Starting Grant ($\approx$ €1,500,000 individual funding)
10/2019	Bosch Young AI Researcher Award ($\approx$ €50,000 individual funding)
02/2019	Google Faculty Research Award 2018 ($\approx$ €25,000 individual funding)
10/2018–09/2021	“la Caixa” Junior Leader fellowship ($\approx$ €300,000 individual funding)
10/2015–09/2018	UPFellows Fellowship (Marie Curie COFUND programme, $\approx$ €150,000 individual funding)
09/2013–08/2014	ERCIM Alain Bensoussan Fellowship (Marie Curie COFUND programme, $\approx$ €40,000 individual funding)
12/2020–	ELLIS Scholar
06/2019	Outstanding reviewer award at ICML 2019 (awarded to top 5% of reviewers)
07/2016	Outstanding reviewer award at ICML 2016 (awarded to top 1% of reviewers)
07/2016	Outstanding PC member award at IJCAI 2016 (awarded to top 10% of reviewers)
10/2011	PhD student of the year at MTA SZTAKI
04/2011	MTA SZTAKI youth prize for excellent research activity
10/2010	PhD student of the year at MTA SZTAKI
10/2009	PhD student of the year at MTA SZTAKI
11/2009	Pro Scientia Gold Medal (awarded biannually to the two best Computer Science student-researchers in Hungary)
06/2009	First prize at the National Scientific Students’ Associations Conference (Hungary)
09/2007–06/2008	Scholarship of the Hungarian Republic (awarded to the top 0.8% of students)
06/2007	Second prize at the National Scientific Students’ Associations Conference (Hungary)
09/2006–06/2007	Scholarship of the Hungarian Republic (awarded to the top 0.8% of students)

Publications in Peer-Reviewed Journals*:

-
- J6*. Gábor Lugosi, Mihalis Markakis and **Gergely Neu** (2024): *On the hardness of learning from censored and nonstationary demand*. In *INFORMS Journal on Optimization*: Volume 6, Issue 2, pp. 63–83.
- J5. **Gergely Neu** and Gábor Bartók (2016): *Importance weighting without importance weights: An efficient algorithm for combinatorial semi-bandits*. In *Journal of Machine Learning Research*: Volume 17, Issue 154, pp. 1–21.
- J4*. Luc Devroye, Gábor Lugosi and **Gergely Neu** (2015): *Random-walk perturbations for online combinatorial optimization*. In *IEEE Transactions on Information Theory*: Volume 61, pp. 4099–4016.
- J3. **Gergely Neu**, András György, Csaba Szepesvári and András Antos (2014): *Online Markov decision processes under bandit feedback*. In *IEEE Transactions on Automatic Control*: Volume 59, pp. 676–691.
- J2. András György and **Gergely Neu** (2014): *Near-optimal rates for limited-delay universal lossy source coding*. In *IEEE Transactions on Information Theory*: Volume 60, pp. 2823–2834.
- J1. **Gergely Neu** and Csaba Szepesvári (2009): *Training parsers by inverse reinforcement learning*. In *Machine Learning*: Volume 77, Issue 2, pp. 303–337.

* Authors listed in alphabetical order for papers marked with *.

Publications in Peer-Reviewed Conference Proceedings* :

- C49*. **Gergely Neu** and Nneka Okolo (2025): Offline RL via Feature-Occupancy Gradient Ascent. In AI & Statistics (AISTATS).
- C48*. Baptiste Abélès, Eugenio Clerico, **Gergely Neu** (2025): Online-to-PAC generalization bounds under graph-mixing dependencies. In AI & Statistics (AISTATS).
- C47*. Baptiste Abélès, Eugenio Clerico, **Gergely Neu** (2025): Generalization bounds for mixing processes via delayed online-to-PAC conversions. In Algorithmic Learning Theory (ALT).
- C46*. Sergio Calo, Anders Jonsson, **Gergely Neu**, Ludovic Schwartz, Javier Segovia-Aguas (2024): Bisimulation Metrics are Optimal Transport Distances, and Can be Computed Efficiently. In Neural Information Processing Systems (NeurIPS).
- C45*. **Gergely Neu**, Matteo Papini and Ludovic Schwartz (2024): Optimistic Information-Directed Sampling. In Conference on Learning Theory (COLT).
- C44*. **Gergely Neu**, Nneka Okolo (2024): Dealing with Unbounded Gradients in Stochastic Saddle-Point Optimization. In International Conference of Machine Learning (ICML).
- C43*. Germano Gabbianelli, **Gergely Neu**, Nneka Okolo, Matteo Papini (2024): Offline Primal-Dual Reinforcement Learning for Linear MDPs. In AI&Statistics (AISTATS).
- C42*. **Gergely Neu**, Julia Olkhovskaya, Sattar Vakili (2024): Adversarial Contextual Bandits Go Kernelized. In Algorithmic Learning Theory (ALT).
- C41*. Germano Gabbianelli, **Gergely Neu**, Matteo Papini (2024): *Importance-Weighted Offline Learning Done Right*. In Algorithmic Learning Theory (ALT).
- C40*. Germano Gabbianelli, **Gergely Neu**, Nneka Okolo, Matteo Papini (2024): *Offline Primal-Dual Reinforcement Learning for Linear MDPs*. In AI&Statistics (AISTATS).
- C39. Julia Olkhovskaya, Jack Mayo, Tim van Erven, **Gergely Neu**, Chen-Yu Wei (2023): *First-and Second-Order Bounds for Adversarial Linear Contextual Bandits*. In Neural Information Processing Systems (NeurIPS).
- C38*. Antoine Moulin and **Gergely Neu** (2023): Optimistic Planning by Regularized Dynamic Programming. In International Conference of Machine Learning (ICML), pp. 25337-25357.
- C37. Lukas Zierahn, Dirk van der Hoeven, **Gergely Neu**, Nicolò Cesa-Bianchi (2023): Nonstochastic Contextual Combinatorial Bandits. In AI&Statistics (AISTATS), pp. 8771-8813.
- C36*. **Gergely Neu**, Nneka Okolo (2023): Efficient Global Planning in Large MDPs via Stochastic Primal-Dual Optimization. In Algorithmic Learning Theory (ALT), pp. 1101-1123.
- C35*. Germano Gabbianelli, **Gergely Neu**, Matteo Papini (2023): Online learning with off-policy feedback. In Algorithmic Learning Theory (ALT), pp. 620-641.
- C34. Luca Viano, Angeliki Kamoutsis, **Gergely Neu**, Igor Krawczuk and Volkan Cevher (2022): *Proximal Point Imitation Learning*. In Neural Information Processing Systems (NeurIPS).
- C33*. **Gergely Neu**, Julia Olkhovskaya, Matteo Papini and Ludovic Schwartz (2022): *Lifting the Information Ratio: An Information-Theoretic Analysis of Thompson Sampling for Contextual Bandits*. In Neural Information Processing Systems (NeurIPS).
- C32*. **Gergely Neu** and Gábor Lugosi (2022): *Generalization bounds via convex analysis*. In Conference on Learning Theory (COLT), pp. 3524–3546.
- C31*. **Gergely Neu** and Julia Olkhovskaya (2021): *Online learning in MDPs with linear function approximation and bandit feedback*. In Neural Information Processing Systems (NeurIPS), pp. 10407–10417.

* Authors listed in alphabetical order for papers marked with *.

- C30. **Gergely Neu**, Gintare Karolina Dziugaite, Mahdi Haghifam and Daniel M. Roy (2021): *Information-theoretic generalization bounds for stochastic gradient descent*. In Conference on Learning Theory (COLT), pp. 3526–3545.
- C29. Fan Lu, Prashant G. Mehta, Sean P. Meyn and **Gergely Neu** (2021): *Convex Q-Learning*. In American Control Conference (ACC), pp. 4749–4756.
- C28*. Joan Bas-Serrano, Sebastian Curi, Andreas Krause and **Gergely Neu** (2021): *Logistic Q-Learning*. In AI&Statistics (AISTATS), pp. 3610–3618.
- C27*. **Gergely Neu** and Ciara Pike-Burke (2020): *A Unifying View of Optimism in Episodic Reinforcement Learning*. In Neural Information Processing Systems (NeurIPS), pp. 1392–1403.
- C26*. **Gergely Neu** and Julia Olkhovskaya (2020): *Efficient and robust algorithms for adversarial linear contextual bandits*. In Conference on Learning Theory (COLT), pp. 3049–3068.
- C25*. **Gergely Neu** and Nikita Zhivotovskiy (2020): *Fast rates for online prediction with abstention*. In Conference on Learning Theory (COLT), pp. 3030–3048.
- C24*. Joan Bas-Serrano and **Gergely Neu** (2020): *Faster saddle-point optimization for solving large-scale Markov decision processes*. In Conference on Learning for Dynamics and Control (L4DC), pp. 413–423.
- C23*. Nicole Mücke, **Gergely Neu** and Lorenzo Rosasco (2019): *Beating SGD saturation with tail-averaging and minibatching*. In Neural Information Processing Systems (NeurIPS), pp. 12568–12577.
- C22. Carlos Riquelme, Hugo Penedones, Damien Vincent, Hartmut Maennel, Sylvain Gelly, Timothy Mann, André Barreto and **Gergely Neu** (2019): *Adaptive temporal-difference learning for policy evaluation with per-state uncertainty estimates*. In Neural Information Processing Systems (NeurIPS), pp. 11872–11882.
- C21*. Wojciech Kotłowski and **Gergely Neu** (2019): *Bandit principal component analysis*. In Conference on Learning Theory (COLT).
- C20*. Gábor Lugosi, **Gergely Neu** and Julia Olkhovskaya (2019): *Online influence maximization with local observations*. In Algorithmic Learning Theory (ALT), pp. 557–580.
- C19*. **Gergely Neu** and Lorenzo Rosasco (2018): *Iterate averaging as regularization for stochastic gradient descent*. In Conference on Learning Theory (COLT), pp. 3222–3242.
- C18*. Nicolás Cesa-Bianchi, Claudio Gentile, Gábor Lugosi and **Gergely Neu** (2017): *Boltzmann exploration done right*. In Neural Information Processing Systems (NeurIPS), pp. 6287–6296.
- C17. **Gergely Neu** and Vicenç Gómez (2017): *Fast rates for online learning in Linearly Solvable Markov Decision Processes*. In Conference on Learning Theory (COLT), pp. 1567–1588.
- C16*. Tongliang Liu, Gábor Lugosi, **Gergely Neu** and Dacheng Tao (2017): *Algorithmic stability and hypothesis complexity*. In International Conference of Machine Learning (ICML), pp. 2159–2167.
- C15. Tomáš Kocák, **Gergely Neu**, Michal Valko (2016): *Online learning with Erdős-Rényi side-observation graphs*. In Uncertainty in Artificial Intelligence (UAI), pp. 339–346.
- C14. Tomáš Kocák, **Gergely Neu**, Michal Valko (2016): *Online learning with noisy side observations*. In AI&Statistics (AISTATS), pp. 1186–1194.
- C13. **Gergely Neu** (2015): *Explore no more: Improved high-probability regret bounds for non-stochastic bandits*. In Neural Information Processing Systems (NeurIPS), pp. 3150–3158.
- C12. **Gergely Neu** (2015): *First-order regret bounds for combinatorial semi-bandits*. In Conference on Learning Theory (COLT), pp. 1360–1375.
- C11. Amir Sani, **Gergely Neu** and Alessandro Lazaric (2014): *Exploiting easy data in online optimization*. In Neural Information Processing (NeurIPS), pp. 810–818.

- C10. Tomáš Kocák, **Gergely Neu**, Michal Valko and Rémi Munos (2014): *Efficient learning by implicit exploration in bandit problems with side observations*. In Neural Information Processing (NeurIPS), pp. 613–621.
- C9. **Gergely Neu** and Michal Valko (2014): *Online combinatorial optimization with stochastic decision sets and adversarial losses*. In Neural Information Processing (NeurIPS), pp. 2780–2788.
- C8. Alexander Zimin and **Gergely Neu** (2013): *Online learning in episodic Markov decision processes by Relative Entropy Policy Search*. In Neural Information Processing (NeurIPS), pp. 1583–1591.
- C7. **Gergely Neu** and Gábor Bartók (2013): *An efficient algorithm for learning with semi-bandit feedback*. In Algorithmic Learning Theory (ALT), pp. 234–248.
- C6*. Luc Devroye, Gábor Lugosi and **Gergely Neu** (2013): *Prediction by random-walk perturbation*. In Conference on Learning Theory (COLT), pp. 460–473.
- C5. **Gergely Neu**, András György and Csaba Szepesvári (2012): *The adversarial stochastic shortest path problem with unknown transition probabilities*. In AI&Statistics (AISTATS), pp. 805–813.
- C4. András György and **Gergely Neu** (2011): *Near-optimal rates for limited-delay universal lossy source coding*. In IEEE International Symposium on Information Theory (ISIT), pp. 2344–2348.
- C3. **Gergely Neu**, András György, Csaba Szepesvári and András Antos (2010): *Online Markov decision processes under bandit feedback*. In Neural Information Processing Systems (NeurIPS), pp. 1804–1812.
- C2. **Gergely Neu**, András György and Csaba Szepesvári (2010): *The online loop-free stochastic shortest path problem*. In Conference on Learning Theory (COLT), pp. 231–243.
- C1. **Gergely Neu** and Csaba Szepesvári (2007): *Apprenticeship learning using inverse reinforcement learning and gradient methods*. In Uncertainty in Artificial Intelligence (UAI), pp. 295–302.

Selected invited talks:

04/2025	BIRS Workshop on Advances in Stochastic Control and RL, Banff, AB, Canada
02/2025	MFO Workshop on Overparametrization, Regularization, Identifiability and Uncertainty in Machine Learning, Oberwolfach, Germany
09/2024	Workshop on Mathematical Aspects of Learning Theory, CRM, Barcelona, Spain
02/2024	Algorithmic Learning Theory (ALT 2024), San Diego, California, USA
05/2023	The 7th London Symposium on Information Theory (LSIT 2023), London, UK
05/2023	Royal Statistical Society (discussant at discussion meeting), London, UK
09/2022	The 15th European Workshop on Reinforcement Learning (EWRL 2022), Milan, Italy
10/2022	ELLIS workshop on Interactive Learning and Interventional Representations, Schwarzwald, Germany
06/2022	ELLIS workshop on Theory, Algorithms and Computations of Modern Learning Systems, Arenzano, Italy
12/2021	BIRS Workshop on Mathematical Statistics and Learning, Banff, AB, Canada
11/2021	CWI Lectures on Decision Making under Uncertainty, Amsterdam, the Netherlands
09/2020	Theory of Reinforcement Learning Boot Camp, Simons Institute, Berkeley, CA, USA (virtual)
07/2020	ICML workshop on Theory of Reinforcement Learning, Vienna, Austria (virtual)
03/2020	ELLIS workshop on Interactive Learning and Interventional Representations, Oberwolfach, Germany (virtual)
09/2019	DALI 2019 workshop on Theory, Algorithms and Computations of Modern Learning Systems, San Sebastián, Spain
12/2018	Tutorial at the Greek Stochastics Kappa workshop, Athens, Greece
10/2018	The 14th European Workshop on Reinforcement Learning (EWRL 2018), Lille, France
05/2018	Workshop on “Modern Challenges in Learning Theory”, Montreal, Canada
06/2017	Learning theory workshop, Foundations of Computational Mathematics, Barcelona, Spain
04/2016	DALI 2016 workshop on Reinforcement Learning, Sestri Levante, Italy
12/2015	NeurIPS 2015 workshop “Learning faster from easy data II”, Montreal, Canada
04/2015	DALI 2015 workshop on Learning Theory, La Palma, Spain
08/2013	The 11th European Workshop on Reinforcement Learning (EWRL 2013), Dagstuhl, Germany

Teaching:

2022–	Reinforcement Learning at Universitat Pompeu Fabra
2022–	Optimization Techniques at Universitat Pompeu Fabra
2020–	Stochastic Models and Optimization at Barcelona Graduate School of Economics
2016–2020	Machine Learning at Universitat Pompeu Fabra (selected lectures)
2006–2011	TA for various courses at Budapest University of Technology and Economics (Probability Theory, Information Theory, Queuing Theory, Electromagnetic Theory)

Program committees and grant panels:

Program chair	The 36th Annual Conference on Learning Theory (COLT 2023) The 31st International Conference on Algorithmic Learning Theory (ALT 2020)
Senior area chair	International Conference on Machine Learning (2025)
PC member / Area chair	Neural Information Processing Systems (NeurIPS 2018-) Conference on Learning Theory (COLT 2019-) Algorithmic Learning Theory (ALT 2015-)
Associate editor	Journal of Machine Learning Research (JMLR, 2023-), Transactions of Machine Learning Research (TMLR, 2021-2025), Machine Learning Journal (MLJ, 2021-2025)
Journal reviewer	Annals of Statistics, Journal of the Royal Statistical Society Series B, Journal of Machine Learning Research (JMLR), Transactions of Machine Learning Research (TMLR), Machine Learning Journal (MLJ), Journal of AI Research (JAIR), AI Journal (AIJ), IEEE Trans. on Information Theory (IEEE T-IT), IEEE Trans. on Automatic Control (IEEE TAC), Automatica, Mathematical Programming, Games and Economic Behavior, Operations Research, Mathematics of OR, Theoretical Computer Science, IEEE Signal Processing Letters (IEEE SPL), IEEE Trans. on Networking (IEEE TNET), Neurocomputing
Conference reviewer	International Conf. on Machine Learning (ICML 2010–), Neural Information Processing Systems (NeurIPS 2012–), Conf. on Learning Theory (COLT 2012–), Algorithmic Learning Theory (ALT 2011–), AI & Statistics (AISTATS 2013–), IEEE International Symposium on Information Theory (ISIT 2012–), American Control Conf. (ACC 2014–), International Joint Conf. on Artificial Intelligence (IJCAI 2015–2017), International Conf. on Learning Representations (ICLR 2019)
External grant reviewer	European Research Council (several ERC consolidator and starting grant proposals since 2019), German Research Foundation (panel member for several DFG priority programs and clusters of excellence, since 2023), Natural Sciences and Engineering Research Council of Canada (NSERC, several Discovery grants since 2022), National Research, Development and Innovation Office of Hungary (NFKI grant 2021), Austrian Science Fund (several FWF grants since 2017), Netherlands Organisation for Scientific Research (several NWO Veni / Vidi / Vici grants since 2018), Israel Science Foundation (several ISF grants between 2019-2023, and several more rejected invites since October 2023)

Organization and community roles:

General chair	The 13th European Workshop on Reinforcement Learning (EWRL 2016)
Local organization chair	The 29th International Conference on Algorithmic Learning Theory (ALT 2018)
Sponsorship chair	Association for Algorithmic Learning Theory (AALT 2021–2023)
Workshop chair	Data, Learning and Inference workshop (DALI 2020)
Steering committees	Association for Computational Learning (member / director, 2020–2023) Association for Algorithmic Learning Theory (member, 2020–2022, 2023-2026) European Workshop on Reinforcement Learning (acting chair of SC, 2022–) ELLIS PhD and Postdoc Program (member, 2023-)
Founding member	ELLIS Unit Barcelona (2023-)
Co-founder and organizer	RL Theory Virtual Seminars (2020-)
Social chair	Neural Information Processing Systems (NeurIPS 2019)

PhD committees:

2024	Gellért Weisz (University College London, UK)
2023	Fredrik Hellström (Chalmers University of Technology, Gothenburg, Sweden)
2022	Sarah Perrin (University of Lille, France)
2021	Matteo Papini (Politecnico di Milano, Italy)
2020	Mastane Achab (Telecom Paris, France)