



IRINA RISH

Full Professor at Université de Montréal
Core Member at Mila – Quebec Artificial Intelligence Institute
6666, rue St-Urbain, bureau 200
Montréal (QC) H2S 3H1, Canada

irina.rish@mila.quebec
<http://irina-rish.com/>
1-914-720-4877

CITIZENSHIP AND LANGUAGES

USA citizen. **Russian:** Native Language. **English:** Superior. **German:** Intermediate. **French:** Novice.

RESEARCH INTERESTS

Machine learning, continual learning, out-of-distribution generalization, robustness, large-scale foundation models, neural scaling laws, sparse models, computational neuroscience.

SUMMARY

[Irina Rish](#) is a Full Professor in the Computer Science and Operations Research Department at the [Université de Montréal \(UdeM\)](#), where she leads the [Autonomous AI Lab](#), and a core faculty member of [MILA - Quebec AI Institute](#). She holds [Canada Excellence Research Chair \(CERC\)](#) and a [CIFAR Chair](#). Dr. Rish completed her MSc and PhD in AI at the University of California, Irvine, and also holds an MSc in Applied Mathematics from Moscow Gubkin Institute. Irina is the recipient of the INCITE 2023 and several ALCC compute grants by the US Department of Energy and lead an [INCITE project on Scalable Foundation Models](#) on Summit & Frontier supercomputers at the [Oak Ridge Leadership Computing Facility](#), focusing on developing open-source AI Foundation Models. She is currently a CSO at [42.com](#) (multimodal foundation models for finance), and an advisor at [Nolano.ai](#) (memory- and inference efficient time-series foundation models).

Irina's extensive research career spans multiple AI domains, from automated reasoning and probabilistic inference in graphical models, to machine learning, sparse modeling, and neuroscience-inspired AI. Irina's current research endeavors concentrate on continual learning, out-of-distribution generalization, robustness; and understanding neural scaling laws and emergent behaviors (w.r.t. both capabilities and alignment) in foundation models - a vital stride towards achieving maximally beneficial Artificial General Intelligence (AGI). She teaches [courses on AI scaling and alignment](#), and runs [Neural Scaling & Alignment workshop series](#).

Before joining UdeM in 2019, Irina was a research scientist at the IBM T.J. Watson Research Center, where she worked on various projects at the intersection of neuroscience and AI, and led the Neuro-AI challenge. She received IBM Eminence & Excellence Award and IBM Outstanding Innovation Award (2018), IBM Outstanding Technical Achievement Award (2017), and IBM Research Accomplishment Award (2009). She holds 64 patents, over 130 research papers, several book chapters, three edited books, and a monograph on Sparse Modeling.

EDUCATION

PhD	Computer Science (Artificial Intelligence) Information and Computer Science Department, University of California, Irvine PhD committee: Rina Dechter (chair), Sandra Irani, Padhraic Smyth. Thesis: Efficient reasoning in graphical models.	July 1999
MS	Computer Science (Artificial Intelligence) Information and Computer Science Department, University of California, Irvine GPA 3.91/4.	June 1994
MS/BS	Applied Mathematics Moscow Gubkin Institute, Applied Math Department, Moscow, Russia Thesis: A rule-based expert system for geophysical applications. GPA 4.98/5.	June 1991

IRINA RISH

SKILLS

Research and Teaching (25+ years). **Programming:** MATLAB (15+ yrs), C/C++ (15+ yrs), Java (1 yr), PROLOG (3 yrs), Pascal (2 yrs), FORTRAN (2 yrs), assembly languages (3 yrs).

HONORS AND AWARDS

IBM Eminence & Excellence Award	2018
IBM Outstanding Innovation Award	2018
IBM Outstanding Technical Achievement Award	2017
IBM 10 th invention plateau award (63 granted patents)	2015
IBM Research Accomplishment Award	2009
Best paper award at ECML-2003	2003
IBM Manager's Choice Awards (4)	2015, 2016, 2017, 2018

RESEARCH FUNDING AND COMPUTE AWARDS

Compute awards:

- 2025-2026** Compute award (Frontier supercomputer): U.S. Department of Energy ALCC Award for "Real-Time Adaptive Disruption Forecasting" at OLCF (Oak Ridge Leadership Computing Facility).
- 2024-2025** Compute award (Frontier supercomputer): U.S. Department of Energy ALCC Award for "Continual Learning of Foundation Models for Scientific Imaging" at OLCF (Oak Ridge Leadership Computing Facility).
- 2023** Compute award (Summit supercomputer): U.S. Department of Energy INCITE Award for "Scalable Foundation Models for Transferrable Generalist AI" at OLCF (Oak Ridge Leadership Computing Facility).

Research funding:

- 2025-2029** Simons Foundation grant to support a collaboration on the "Physics of Learning and Neural Computation" (\$108,816 USD/approx \$155,000 CAD per year, for 4 years)
- 2025-2027** Korea National AI Research Lab (KAIST) partnership grant (\$110,000 CAD per year, for 2 years)
- 2025-2026** Industry partnerships: CapitalOne (\$50,000 CAD), Softmax.ai (\$150,000 USD/approx \$213,000 CAD)
- 2025** Foresight Institute grant to support a project on "Detecting AI collusion by inferring coordination through scalable multi-agent models" (150,000 USD/approx \$213,000 CAD)
- 2024-2027** Canada CIFAR AI chair (Extended)
- 2019-2024** Canada CIFAR AI chair (325,000 CAD/year for 5 years)
- 2020-2027** Canada Excellence Research Chair (CERC) in Autonomous AI: \$10 million from the Government of Canada over seven years, and nearly \$25 million CAD from organizations such as the Quebec Artificial Intelligence Institute (Mila), and the Université de Montréal, as well as several other public and private organizations)
- 2021-2024** Fonds de recherche du Québec - Nature et technologies (FRQNT); Études de la prise de décision dynamique en milieu naturel et artificiel: Une approche interdisciplinaire et intégrative. \$190,500 CAD; Principal Investigator.
- 2021-2023** IVADO; Neocortical mechanisms for continual learning. \$249,970 CAD; Co-Investigator

IRINA RISH

- 2020-2022** IVADO; Ultrasound classification of chronic liver disease with deep learning. \$225,000 CAD; Co-Investigator
- 2020-2022** IVADO; A neurodevelopmentally-informed computational model of flexible human decision-making. \$200,000 CAD; Co-Investigator

RESEARCH EXPERIENCE

- 2022-now** Full professor in the Department of Computer Science and Operations Research at the Université de Montréal, a core faculty member at Mila (Quebec Artificial Intelligence Institute), Canada CIFAR AI Chair and Canada Excellence Research Chair (CERC) in Autonomous AI.
- 2019-2022** Associate professor in the Department of Computer Science and Operations Research at the Université de Montréal, a core faculty member at Mila (Quebec Artificial Intelligence Institute), Canada CIFAR AI Chair and Canada Excellence Research Chair (CERC) in Autonomous AI.
- 2017 - 2019** Research Staff Member (RSM) at the IBM Research and project lead on Brain-inspired AI/Neuro-AI project. Research area: Artificial Intelligence and Machine Learning. Research focus: optimization in deep learning, continual (lifelong) learning, online decision making, dialog generation, representation learning, nonlinear dynamical models, brain imaging data analysis. Initiated and led research projects on the above topics, jointly with IBM colleagues and with external collaborators at various Universities. Supervised students and postdocs; served on the PhD thesis committee of several graduate students. Initiated and maintained external collaborators with various Universities and with internal IBM collaborators. Supervised students and postdocs; served on the PhD thesis committee of several graduate students. Actively published in AI conferences and neuroscience journals.
- 2008 - 2016** Research Staff Member (RSM) at the Computational Biology Center (CBC), IBM T.J. Watson Research Center. Research area: machine learning for neuroscience and computational psychiatry. Research focus: sparse predictive models, interpretable machine learning, recurrent deep neural networks; mental state recognition and diagnosis of psychiatric and neurological disorders from brain imaging data (functional MRI, EEG), wearable sensor data (wearable EEG), speech and text. Designed and implemented MATLAB tools for neuroimaging analysis, analyzed multiple datasets, developed machine-learning methods for discovering statistical biomarkers (features) associated with various mental conditions, including schizophrenia, cocaine addiction, pain, and others, actively published in both AI and neuroscience venues, maintained external and internal collaborations, supervised interns and postdocs.
- 1999 - 2007** Research Staff Member (RSM) at IBM T.J. Watson Research Center (Distributed Computing). Research in the areas of machine learning and probabilistic reasoning, graphical models (Bayesian networks) and approximation algorithms, dimensionality reduction and sparse predictive models, active learning, matrix factorization/blind-source separation. Applications: problem diagnosis and performance analysis of large-scale distributed systems (“autonomic computing”). Project lead on Self-Aware Distributed Systems (2004-2007). The project resulted in development of the Real-Time Active Inference and Learning (RAIL) tool for real-time problem diagnosis in distributed systems using active probing and Bayesian inference.
- 1992 – 99** Research Assistant at Computer Science Department, University of California, Irvine. Research areas: constraint satisfaction/optimization, Bayesian networks, decision-theoretic planning. Approximation algorithms for probabilistic inference. Application domains: medical diagnosis, probabilistic decoding. Developed a C++ library of probabilistic inference algorithms for Bayesian networks. C/C++ (SUN/Unix and Windows NT). Development and experimental evaluation of CSP algorithms for propositional satisfiability. Analysis of empirical complexity distributions, statistical models.

IRINA RISH

Planning under uncertainty. Markov Decision Processes. Robustness in planning and scheduling. Scheduling of power generating units in a power plant (EPRI project).

Molecular biology and cancer research. Developed and implemented a probabilistic model of mutagenesis in yeast and successfully tested it on real data (joint work with the Department of Biochemistry, UCI).

1989 - 91 Diploma project at Moscow Oil and Gas Institute. Development and implementation of a rule-based expert system for geophysical applications. C/Prolog (MSDOS).

INDUSTRY EXPERIENCE

2025 - 2026 Scientific Director/CSO at [42.com](https://www.42.com)

2023 - 2025 Scientific Director/CSO at [Nolano.ai](https://www.nolano.ai)

1999 – 2019 Research Staff Member (RSM) at IBM T.J. Watson Research Center (see above).

1998 Summer internship at Rockwell Science Center, Palo Alto (RPAL). Research and software development for military situation assessment problems. Probabilistic modeling using Bayesian networks. Visual C++, Visual Basic, MapObjects (Windows NT).

1995 – 96 Project manager assistant, Continuous Software Corp., Irvine, CA. Unix shell programming. Object-oriented programming.

1994 – 95 Software engineer at Orbit Technology, Irvine, CA. CAD/CAM systems for manufacturing. Visual C++ (Windows NT). Developed/maintained a subsystem for visualization of metal sheet folding and unfolding processes.

1989 – 92 Software engineer at the Central Geophysical Expedition, Moscow, Russia. Developed and implemented: an expert system for classification and decision making in geophysical data processing; resource manager and GUI for a client/server geophysical system. C/C++, PROLOG.

1986 – 89 Software engineer, Moscow Oil & Gas Gubkin Institute (Moscow, Russia), Department of Applied Mathematics. System programmer in a support group. UNIX system administration. Participated in development and improvement of C-compilers for PDP-11 compatible computers. C/UNIX, assembly language. Participated in implementing of a cross-system for Intel8085 on PDP-11 w/ Unix and VAX/Unix. Languages: C, assembly language Intel 8085. Developed application software packages for oil & gas industry. IBM370, assembly language, FORTRAN, PASCAL.

TEACHING EXPERIENCE

2022-Present Full professor in the Department of Computer Science and Operations Research at the Université de Montréal, a core faculty member at Mila (Quebec Artificial Intelligence Institute), Canada CIFAR AI Chair and Canada Excellence Research Chair (CERC) in Autonomous AI.

2019-2022 Associate professor at University of Montreal, Dept. of Computer Science and Operations Research: courses taught: Continual Learning, winter semester.

2011 Adjunct professor at Columbia University, Electrical Engineering Department; course taught: Sparse Modeling, spring semester.

2007 Adjunct professor at Columbia University, Computer Science Department; course taught: Learning and Empirical Inference, spring semester.

2003 Adjunct professor at Columbia University, Electrical Engineering Department; course taught: Statistical Pattern Recognition, spring semester.

2002 Adjunct professor at Columbia University, Electrical Engineering Department; course taught: Statistical Machine Learning, spring semester.

1997 Instructor at UCI Extension. Classes taught: Introduction to Client/Server Computing.

1996-97 Instructor at Chapman University. Classes taught: Data Structures in Pascal, Data Structures in C++.

1992-95 Teaching Assistant, Information and Computer Science Department, UCI; responsibility for discussion sections, test and homework preparation, grading, occasional lectures. Classes taught: Introduction to Artificial Intelligence, Principles of Operating Systems, Compilers and Interpreters, Design and Analysis of Algorithms, Belief Networks.

RECENT INVITED TALKS

- June 2026: Beyond Scaling: Toward Continual and Adaptive Intelligence, a talk at the Mila's Partners' Symposium.
- June 2026: [Scaling, Transfer and Continual Learning in Foundation Models](#), a keynote talk at the 7th International Conference on the Mathematics of Neuroscience and AI (NM2026).
- June 2026: [Beyond Scaling: Toward Continual and Adaptive Intelligence](#), an invited talk at the Foresight Vision Weekend United Kingdom 2026, Foresight Institute, London, UK.
- February 2026: [From Scaling Curves to Smarter Models: Prediction, Mechanisms, and Reasoning](#), an invited talk at the [Information Theory and Applications Workshop \(ITA\)](#).
- January 2026: [Reasoning: the Effects of Data Pruning, Representation Capacity, and Latent Trajectories](#), an invited talk at the [IVADO Reasoning Workshop](#).
- December 2025: [Scaling, Robustness and Continual Learning in Foundation Models: Recent Advances from the Autonomous AI Lab](#), an invited talk at the Amii AI Seminar Series, University of Alberta, Edmonton, Canada.
- December 2025: [From Continual Pretraining to Cross-Modal Transfer](#), an invited talk at the [NeurIPS 2025 CCFM Workshop](#).
- November 2025: Scaling, Robustness and Continual Learning, ADDIROUM event, Université de Montréal, Montréal, Canada.
- October 2025: Scaling, Robustness, and Continual Learning in Foundation Models: Recent Advances from the Autonomous AI Lab. Capital One Research Spotlight Series, online.
- September 2025: CERC in Autonomous AI: 2024–2025 Research Summary and Future Plans. CERC-AAI annual research presentation, Montréal, Canada.
- Dec 17, 2024: [Towards Time Series Foundation Models](#), an invited talk at FSC-BOK-BIS Joint Conference: AI, Finance and Central Banking: Opportunities, Challenges, and Policy Responses.
- Dec 14, 2024: [Continual learning in the era of foundation models](#), a keynote at the [NeurIPS 2024 workshop on Scalable Continual Learning for Lifelong Foundation Models](#).
- Nov 21, 2024: [Scale to AGI \(some recent developments in scaling laws\)](#), a talk at the [Curve conference](#).
- Oct 22, 2024: [Rethinking Intelligence from Dynamical Systems Perspective](#), an invited talk at [MAIN conference](#), Montreal.
- Sept 11, 2024: keynote at [AutoML 2024](#).
- June 13, 1:30pm-3:00pm EDT, [Emergent Behaviors & Foundation Models in AI](#), an invited talk at [LARGE LANGUAGE MODELS : SCIENCE AND STAKES](#) (June 3 – June 14 2024)
- April 25, 2024: [Recent Advances in Foundation Models: Continual Training, Time-Series and VLM Evals](#), an invited talk at the [Symposium on AI and LLMs](#), CUNY, NYC.
- April 24, 2024: [Towards AGI: Predicting Scaling Behaviors and Training Large Foundation Models on Supercomputers](#), an invited talk at the [World Summit AI Americas 2024](#), Montreal.
- April 18, 2024: [Scaling to AGI](#), an invited talk at the University of Chicago [Data Science Institute \(DSI\) Distinguished Speaker Series](#).
- April 13, 2024: [The AI Scaling Revolution and the Future of Intelligence](#), a keynote talk at the [MTACConf2024](#), Provo, Utah.
- April 5, 2024: [Lag-LLama: Towards Time-Series Foundation Models](#), an online guest lecture at the University of Warsaw
- March 25, 2024: [Continual Learning of Foundation Models](#), an invited talk at the [NSF Open-Source Generative AI Workshop](#), NYC
- March 18, 2024: [On Scaling Laws, Foundation Models and HPC](#), an invited talk at [The Future of Discovery in the Age of Human-AI Collaboration](#), [Michigan Institute for Data Science \(MIDAS\)](#)
- March 12, 2024: Towards AGI: Predicting Scaling Behaviors and Training Large Foundation Models on Supercomputers, an invited talk at Microsoft Research, Montreal.
- March 7, 2024: [Scaling Laws and Emergent Behaviors in Foundation Models](#), an invited talk at Statistical Physics Meets Machine Learning II session, [March Meeting 2024 - American Physical Society](#), Minneapolis.

- Feb 26, 2024: [AI 4 Psychology and Psychology 4 AI: Towards Better Alignment Among Humans and Machines](#), a keynote talk at AAAI-2024 [Machine Learning for Cognitive and Mental Health \(ML4CMH\) workshop](#).
- Feb 7, 2024: *On Scaling Laws and Foundation Models: an Overview of Our INCITE 2023 Efforts*, an invited talk at the Trillion Parameter Consortium (TPC) Seminar Series.
- Dec 11, 2023: *Complex systems view of large-scale AI systems*, an invited talk at the [Alignment Workshop](#), New Orleans, LA, USA.
- Nov 17, 2023: *Divide and Conquer: Scaling AI via Federated Learning and Distributed Alternatives to Backpropagation*, an invited talk at the [The Second Workshop on Federated and Privacy Preserving AI for HPC, Denver, CO](#).
- October 20, 2023: [From AI Systems to Systems \(Neuro-\)AI?](#) A talk at the UNIQUE Scientific Retreat, Tremblant.
- August 29, 2023: *The Large-scale AI Revolution: Breakthroughs and Challenges*, a keynote talk at the Smoky Mountains Computational Sciences and Engineering Conference (SMC2023) Knoxville, Tennessee.
- August 25, 2023: *Continual Learning at Scale*, a keynote talk at the Second Conference on Lifelong Learning Agents - CoLLAs 2023.
- July 29, 2023: *Backpropagation Alternatives and Scalable AI*, a plenary talk at the ICML 2023 workshop on Localized Learning.
- June 24, 2023: *Scaling Laws and Emergent Behaviors*, in the Human-Level AI workshop, Boston.
- June 18, 2023: *On Scaling Laws, Emergent Properties and AI Democratization efforts*, a plenary talk at AGI-23, 16th Annual AGI Conference in Stockholm, Sweden.
- May 23, 2023: *Recent Advances in Foundation Models: Scaling Laws, Emergent Behaviors, and AI Democratization*, a plenary talk at the Cosmic Connections: A ML X Astrophysics Symposium at Simons Foundation.
- March 29, 2023: *Advances in Foundation Models and Scaling Laws*, a keynote talk at EDGeS workshop, San Francisco (Burlingame).
- March 8, 2023: *On Scaling Laws, Emergent Properties and AI Democratization efforts*, a keynote talk at the Predictable AI workshop (<https://www.predictable-ai.org/>)
- Feb 7, 2023: *Towards Genuinely Open AI: The Role of Supercomputers*, an invited talk at the The 5th R-CCS International Symposium | RIKEN.
- Nov 7 2022: *Computational Psychology & Psychological Computation*, an invited talk at the AI Helps Ukraine - Charity Conference.
- Oct 17, 2022: *Towards General and Robust AI at Scale*, an invited talk at the Fields Medal Symposium.
- Oct 4, 2022: *Towards General and Robust AI at Scale*, an invited talk at the Morgan Stanley ML Seminar, New York.
- May 5, 2022: *Towards General and Robust AI at Scale*, a keynote talk at the World Summit AI Americas, Montreal.
- April 7, 2022: *Towards General and Robust AI at Scale*, an invited talk at the NAISys workshop (From Neuroscience to Artificially Intelligent Systems), Cold Spring Harbor, NY.
- March 14, 2022: *Towards General and Robust AI*, an invited talk at the symposium on Statistical Physics Meets Machine Learning, at the APS March meeting.
- Feb 2, 2022: *Towards General and Robust AI at Scale*, an invited talk at the AI Sweden NLP Seminar
- Jan 27, 2022: *Towards General and Robust AI at Scale*, an invited talk at the Institute of Advanced Research in AI.
- Jan 25, 2022: *Towards General and Robust AI at Scale*, an invited talk at the Accelerated Data Analytics and Computing (ADAC) Institute workshop.
- Aug 26, 2021: *On Neural Scaling Laws*, an invited talk in IBM Research AI seminar series.
- July 30, 2021: *Tutorial on Lifelong Learning* at DLRL Summer School.
- July 2021: *Continual Learning with Deep Architectures*, a tutorial at ICML 2021.
- Spring 2021: *Reward Processing Biases in Humans and RL Agents*, 2021 Reinforcement Learning Virtual School.
- Spring 2021: *Lifelong learning and Robustness*, Mila partners symposium.

- Spring 2021: *Beyond Backprop: Different Approaches to Credit Assignment in Neural Nets*, invited talk at Neuro Inspired Computational Elements (NICE) Conference 2021.
- December 2020: *Bringing AI and Neuroscience Together*, an invited talk at Machine Learning in Science & Engineering (MLSE 2020).
- October 2020: *Dialogue Modeling Via Hash Functions: Applications to Psychotherapy*, an invited talk at WeCNLP 2020.
- October 2020: *Towards Broad and Robust AI*, CBDS Distinguished Speaker Series Online, Yale U.
- Sept 2020: *Towards Less Artificial Intelligence*, a plenary talk at Montreal AI Symposium.
- July 2020: *Lifelong Learning: Towards Broad and Robust AI*, an invited talk at the 4th Lifelong Learning workshop at ICML 2020.
- April 2020: *AI against COVID-19 @* webinar by Islamic World Educational, Scientific and Cultural Organization (ICESCO) & AI Civic Forum (AlgoraLab, The Future Society, Mila and UdeM).
- Jan 2020: *Dialog Modeling via Hash Functions*, Tea talk @ Mila.
- Dec 2019: *Where Artificial Intelligence and Neuroscience Meet*, FutureTech podcast.
- Dec 2019: *Beyond Backprop: Online Alternating Minimization with Auxiliary Variables*, MILA workshop on Neural Backpropagation,
- June 2019: *Towards the Synergy Between AI and Neuroscience*, talk at University of Montreal.
- May 2019: *AI and Neuroscience: Bridging the Gap*, invited talk, New York University (NYU).
- April 2019: *AI in Neuroscience*. invited talk, National Institute of Health (NIH).
- March 2019: *AI and Neuroscience: Bridging the Gap*, **plenary talk**. Neuro-Inspired Computing Elements (NICE) conference.
- December 2018: *Modeling Brain Dynamics: van der Pol, LSTMs and beyond*, invited talk, Montreal AI & Neuroscience (MAIN) conference.
- November 2018: *AI for Neuroscience and Neuroscience for AI*, invited talk, MLconf 2018, San Francisco.
- November 2018: *AI and Neuroscience*, invited talk, Helsinki Distinguished Lecture Series on Future Information Technology, Aalto University, Aalto, Finland.
- September 2018: *AI for Mental Health*, **plenary talk** at Brazilian Congress of Innovation in Psychological Assessment and Neuroscience, Porto Alegre, Brazil.
- July 2017, *Brain and AI*, **plenary talk** at SAI conference, London.
- March 2017: *Learning About the Brain and Brain-Inspired Learning*, invited talk at MLconf, NYC.
- December 2016: *Learning About the Brain: Neuroimaging and Beyond*. **Plenary talk** at NIPS-2016, Barcelona, Spain.
- November 2015: *Sparse Models and Beyond*, invited talk at MLconf 2015, San Francisco.

PROFESSIONAL ACTIVITIES

- Editorial board member of AI Journal
- IEEE TPAMI Associate Editor
- Program Committee: **senior area chair** (NIPS-2017, NIPS-2018, ICML-2018), **area chair** (IJCAI-2015, ICML-2015, ICML-2016, NIPS-2010), **workshop chair** (UAI-2015, ICML-2012), **tutorial chair** (UAI-2012), program committee member over multiple years (ICML, KR, IJCAI, UAI, AAAI).
- Reviewer: Constraints Journal, Journal of AI Research, Journal of Machine Learning research (JMLR), Machine Learning Journal.
- Evaluation committees:
 - Evaluated INRIA teams affiliated with the theme "Optimisation, machine learning, and statistical methods" (March 2014)
- Tutorials:
 - Human Brain Mapping (HBM) 2015: Machine learning in neuroscience and biology
 - ICML-2010, Haifa, Israel: Sparse Modeling (with G. Grabarnik)
 - ECML-2006, Berlin, Germany: Autonomic Computing
 - AAAI-2000, Austin, TX: Approximation Techniques for Automated Inference (with R. Dechter)
 - IC-AI'2000, Las Vegas: Advances in Bayesian Learning
 - AAAI-98, Madison, WI: Principles and Strategies of Automated Inference (with R. Dechter)

IRINA RISH

- Workshops organized:
 - NIPS-03: Robust Communication Dynamics in Complex Networks
 - NIPS-05: Workshop on Value of Information in Inference, Learning and Decision-Making
 - NIPS-06: Workshop on Novel Applications of Dimensionality Reduction
 - ECML-06: Autonomic Computing: A New Challenge for Machine Learning
 - ICML-08: Sparse Optimization and Variable Selection: Theory, Algorithms and Applications
 - NIPS-08: New Directions in Statistical Learning for Meaningful and Reproducible fMRI Analysis
 - NIPS-10: Practical Applications of Sparse Modeling: Open Issues and New Directions
 - NIPS-11: Machine Learning and Interpretation in Neuroimaging (MLINI-2011)
 - NIPS-12: Machine Learning and Interpretation in Neuroimaging (MLINI-2012)
 - NIPS-13: Machine Learning and Interpretation in Neuroimaging (MLINI-2013)
 - NIPS-14: Machine Learning and Interpretation in Neuroimaging (MLINI-2014)
 - NIPS-15: Machine Learning and Interpretation in Neuroimaging (MLINI-2015)
 - NIPS-16: Machine Learning and Interpretation in Neuroimaging (MLINI-2016)
 - [Neural Scaling Laws workshop series](#): 5 workshops from Oct 2021 to Aug 2023

RESEARCH SUPERVISION: STUDENTS AND POSTDOCS

Name (last, first)	Role	Start date	End date
INTERNS			
Ho, Mai-Thi	intern	1/2022	7/2022
Cecchi, Lucas	intern	6/2021	9/2021
Yohav Vinnikov	intern	1/2021	6/2021
Mathieu, Samuel	intern	5/2022	8/2022
Runwal, Bharat	intern	12/2022	3/2023
MASTER'S STUDENTS			
Bayat, Reza	Master's student	09/2021	12/2024
Bhagwatkar, Rishika	Master's student	09/2023	08/2026
Chandra Raparthy, Sharath	Master's student	09/2020	08/2022
Darabi, Amin	Master's student	09/2023	08/2026
Francis, Sreya	Master's student	01/2020	04/2021
Gagnon-Audet, Jean-Christophe	Master's student	01/2020	08/2023
Gauthier, Shanel	Master's student	09/2020	09/2021
Gupta, Kshitij	Master's student	09/2021	05/2024

IRINA RISH

Haggi Mani, Parviz	Master's student	01/2020	08/2021
Hassen, Nadhir	Master's student	09/2020	04/2022
Huard, Charles	Master's student	09/2020	09/2020
Khorasani, Arian	Master's student	07/2021	08/2024
Laleh, Touraj	Master's student	01/2020	06/2023
Lemay, David	Master's student	09/2024	04/2026
Mansouri Tehrani, Sayed Mohammadamin	Master's student	09/2020	06/2023
Memarian, Amin	Master's student	09/2019	12/2022
Mishra, Sparsha	Master's student	10/2022	09/2024
Misra, Diganta	Master's student	10/2021	03/2024
Mocanu, Remus	Master's student	09/2020	08/2023
Moiescu-Pareja, Gabriela	Master's student	09/2021	12/2024
Normandin, Fabrice	Master's student	09/2019	12/2021
Piche, Marc-André	Master's student	03/2020	04/2023
Samsami, Mohammad Reza	Master's student	09/2022	08/2025
Simard, Alexandre	Master's student	09/2021	08/2023
Spinney, Sean	Master's student	09/2021	07/2023
Tenison, Irene	Master's student	09/2020	06/2022
Vaidhya, Tejas	Master's student	09/2022	12/2024
Whipps, David	Master's student	09/2022	08/2024
Yasafova, Daria	Master's student	09/2023	08/2027
PhD STUDENTS			
Abrevaya, German	PhD student	06/2022	08/2022
Anokhin, Ivan	PhD student	09/2023	08/2028
Arefin, Md Rifat	PhD student	09/2021	08/2027

IRINA RISH

Ashok, Arjun	PhD student	01/2023	12/2029
Caballero, Ethan	PhD student	09/2022	04/2030
Darvishi Bayazi, Mohammad Javad	PhD student	09/2020	12/2024
Dehgan, Arthur	PhD student	01/2020	12/2024
Dufort-Labbe, Simon	PhD student	01/2024	04/2026
Faramarzi, Mojtaba	PhD student	01/2021	12/2026
Ibrahim, Adam	PhD student	09/2023	08/2024
Islah, Nizar	PhD student	01/2023	12/2026
Jain, Arnav Kumar	PhD student	01/2022	08/2026
Joseph, Sonia	PhD student	09/2022	08/2026
Legate, Gwendoline	PhD student	06/2023	12/2026
Liu, Michelle	PhD student	09/2020	08/2022
McCracken, Gavin	PhD student	05/2023	06/2027
Nest, Timothy	PhD student	06/2021	04/2027
Pourrahi, Motahareh	PhD student	09/2023	08/2026
Ramezani, Mahta	PhD student	01/2020	08/2028
Riemer, Matthew	PhD student	09/2020	08/2026
Roger, Alexis	PhD student	09/2024	08/2027
Romascanu, Andrei Mircea	PhD student	09/2023	08/2028
Subbaraj, Gopeshh Raaj	PhD student	01/2021	12/2026
Thérien, Benjamin	PhD student	09/2023	08/2028
Tikeng Notsawo, Pascal Junior	PhD student	09/2023	08/2028
Williams, Andrew	PhD student	09/2023	08/2028
POSTDOCTORAL FELLOWS AND RESEARCH ASSOCIATES			
Ahuja, Kartik	Postdoctoral fellow	04/2020	08/2022

IRINA RISH

Bashivan, Pouya	Postdoctoral fellow	04/2020	08/2020
Brennan, Connor	Postdoctoral fellow	03/2021	04/2024
Fournier, Quentin	Postdoctoral fellow	01/2023	01/2024
Girotti, Manuela	Postdoctoral fellow	10/2020	08/2021
Haggi Mani, Parviz	Research Associate	04/2025	Present
Humane, Prateek	Master's student / Postdoctoral fellow	09/2023	12/2031
Lehnert, Lucas	Postdoctoral fellow	06/2021	05/2023
Lesort, Timothée	Postdoctoral fellow	03/2021	08/2023
Lobacheva, Ekaterina	Postdoctoral fellow	04/2024	05/2025
Nobandegani, Ardavan	Research Associate	01/2023	08/2024
Puelma-Touzel, Maximilien	Postdoctoral fellow / Research Associate	12/2020	09/2024
Shahtalebi, Soroosh	Postdoctoral fellow	02/2021	06/2021

PUBLICATIONS

Thesis

- I. Rish. Efficient Reasoning in Graphical Models, PhD thesis, 1999.

Monograph

- I. Rish and G. Grabarnik, Sparse Modeling: Theory, Algorithms and Applications. Chapman and Hall/CRC Machine Learning and Pattern Recognition, Dec. 2014.

Edited Books (3)

1. **I. Rish**, G. Langs, L. Wehbe, G. Cecchi, K. Kevin Chang and B. Murphy, editors. *Machine Learning and Interpretation in Neuroimaging: 4th International Workshop*, MLINI 2014, held at NIPS 2014, Montreal, QC, Canada, December 13, 2014, Revised Selected and Invited Contributions. Springer, 2014.
2. **I. Rish**, G. Cecchi, A. Lozano, A. Niculescu-Mizil, editors. *Practical applications of sparse modeling*. MIT Press, 2014.
3. G. Langs, **I. Rish**, M. Grosse-Wentrup, B. Murphy, editors. *Machine Learning and Interpretation in Neuroimaging: 1st International Workshop*, MLINI 2011, held at NIPS 2011, Sierra Nevada, Spain, December 16-17, 2011, Revised Selected and Invited Contributions. Springer, 2012.

Book chapters (4)

1. J. Honorio, D. Samaras, **I. Rish** and G. Cecchi. Variable-Selection in Gaussian MRFs. In *Log-Linear Models, Extensions and Applications*. Edited by Alexander Aravkin et al., MIT Press, 2018.
2. **I. Rish** and G. Cecchi. Functional Network Disruptions in Schizophrenia. In *Biological Networks and Pathway Analysis*, edited by Y. Nikolsky and T. Tatarinova, Springer, 2017.
3. **I. Rish** and G. Grabarnik, Sparse Signal Recovery with Exponential-Family Noise. In *Compressed Sensing and Sparse Filtering*. Edited by Carmi, Avishy Y., Mihaylova, Lyudmila S., Godsill, Simon J. Springer, 2014.

4. M. Carroll, G. Guillermo, **I. Rish**, R. Garg, M. Baliki, and A. Apkarian. Reliability Estimation and Enhancement via Spatial Smoothing in Sparse fMRI Modeling. In *Practical Applications of Sparse Modeling*, pp. 123-150, MIT Press, 2014.

Publications (195)

1. Sihui Wei, Gavin McCracken, Gabriela Moisesescu-Pareja, Harley Wiltzer, Doina Precup, Irina Rish, Jonathan Love, "[Deep neural networks divide and conquer dihedral multiplication](#)", The Forty-Third International Conference on Machine Learning (ICML), 2026.
2. Benjamin Thérien, Xiaolong Huang, Aaron Defazio, Irina Rish, Eugene Belilovsky, "[MuLoCo: Muon is a Practical Inner Optimizer for DiLoCo](#)", The Forty-Third International Conference on Machine Learning (ICML), 2026.
3. Hiroki Naganuma, Shagun Gupta, Youssef Briki, Ioannis Mitliagkas, Irina Rish, Parameswaran Raman, Hao-Jun Michael Shi, "[Adaptive Batch Sizes Using Non-Euclidean Gradient Noise Scales for Stochastic Sign and Spectral Descent](#)", The Forty-Third International Conference on Machine Learning (ICML), 2026.
4. Matthew Riemer, Tommaso Tosato, Maximilian Puelma Touzel, Amin Memarian, Guillaume Dumas, Glen Berseth, Irina Rish, "[Collusion Risks Among AI Reasoning Agents Justify Certification Requirements for Making Market Decisions](#)" (position paper), The Forty-Third International Conference on Machine Learning (ICML), 2026
5. Benjamin Thérien, Charles-Étienne Joseph, Boris Knyazev, Edouard Oyallon, Irina Rish, and Eugene Belilovsky, "[muLO: Compute-Efficient Meta-Generalization of Learned Optimizers](#)", The Fourteenth International Conference on Learning Representations (ICLR), 2026.
6. Tommaso Tosato, Saskia Helbling, Yorguin-Jose Mantilla-Ramos, Mahmood Hegazy, Alberto Tosato, David John Lemay, Irina Rish, Guillaume Dumas, "[Persistent Instability in LLM's Personality Measurements: Effects of Scale, Reasoning, and Conversation History](#)", in the Alignment Track of The 40th Annual AAAI Conference on Artificial Intelligence (AAAI'26).
7. Alexis Roger, Prateek Humane, Zhenghan Tai, Gwen Legate, Andrei Mircea, Vasilii Feofanov, Irina Rish, "[Forecasting Emerges from Auto-Regressive Pretraining: Latent Predictive Structure in Language Models](#)", in the Forecasting as a New Frontier of Intelligence Workshop at ICML 2026 (oral).
8. Prateek Humane, Alexis Roger, Zhenghan Tai, Gwen Legate, Andrei Mircea, Vasilii Feofanov, Irina Rish, "[Feature Geometry of Language Models Transfer Across Modalities to Time Series](#)", in the Mechanistic Interpretability Workshop at ICML 2026.
9. Zhenghan Tai, Alexis Roger, Prateek Humane, Gwen Legate, Andrei Mircea, Vasilii Feofanov, Irina Rish, "[Language Pretraining Gives Structured Forecasters a Sequential Prior](#)", in the 2nd ICML Workshop on Foundation Models for Structured Data (FMSD) at ICML 2026.
10. Alexis Roger, Prateek Humane, Zhenghan Tai, Gwen Legate, Andrei Mircea, Vasilii Feofanov, Irina Rish, "[Reusable Low-Rank Subspaces Explain Why Cross-Modal Transfer Adapts with Tiny Updates](#)", in the 2nd Workshop on Connecting Low-rank Representations in AI (CoLoRAI) at ICML 2026.
11. Andrei Mircea, Ildus Sadrudinov, Irina Rish, Ekaterina Lobacheva, "[Shared Gradient Discovery and Superposition: Learning Dynamics of Generalization in LLMs](#)", in the ICLR 2026 Workshop on Scientific Methods for Understanding Deep Learning (Sci4DL 2026).
12. Gopeshh Subbaraj, Istabrak Abbes, Artem Zhohus, Matthew Riemer, Irina Rish, Sarath Chandar, "[Emergent Reasoning via Recursive Latent Reinforcement Pretraining](#)", in the ICLR 2026 Workshop LLM Reasoning.
13. Shahrad Mohammadzadeh, Jacob Chmura, Ivan Anokhin, Jacob-Junqi Tian, Mandana Samiei, Taz Scott-Talib, Irina Rish, Doina Precup, Reihaneh Rabbany, Nishanth Anand, "[AIF-GEN: Open-Source Platform and Synthetic Dataset Suite for Reinforcement Learning on Large Language Models](#)", in LLA Workshop @ ICLR 2026.
14. Anatole Jacquin de Margerie, Alexis Roger, Irina Rish, "[Image Tiling for High-Resolution Reasoning: Balancing Local Detail with Global Context](#)", in the Reproducible AI Workshop at AAAI 2026

15. Alexis Roger, Gwen Legate, Kashif Rasul, Yuriy Nevmyvaka, Irina Rish, "[Small Vocabularies, Big Gains: Pretraining and Tokenization in Time Series Models](#)", in the Artificial Intelligence for Time Series Analysis (AI4TS) Workshop at AAAI 2026.
16. Benjamin Thérien, Charles-Étienne Joseph, Zain Sarwar, Ashwinee Panda, Anirban Das, Shi-Xiong Zhang, Stephen Rawls, Sambit Sahu, Eugene Belilovsky, Irina Rish, "[Continual Pre-training of MoEs: How robust is your router?](#)", Transactions on Machine Learning Research (TMLR), 2025.
17. Javad Bayazi, Mohammad Sajjad Ghaemi, Irina Rish, and Jocelyn Faubert, "Critical Role of EEG Signals in Assessment of Sex-Specific Insights in Neurological Diagnostics via Machine Learning Approach", Nature Scientific Reports, 2025.
18. Rishika Bhagwatkar, Syrielle Montariol, Angelika Romanou, Beatriz Borges, Irina Rish, Antoine Bosselut, "[CAVE : Detecting and Explaining Commonsense Anomalies in Visual Environments](#)", in Proc of EMNLP 2025.
19. Andrei Mircea, Ekaterina Lobacheva, Supriyo Chakraborty, Nima Chitsazan, Irina Rish, "[Understanding Language Model Scaling Laws in Terms of Training Dynamics via Loss Deceleration and Zero-Sum Learning](#)", in ACL 2025.
20. Tejas Vaidhya, Ayush Kaushal, Vineet Jain, Francis Couture-Harpin, Prashant Shishodia, Majid Behbahani, Irina Rish, Yuriy Nevmyvaka, "[Scaling Laws and Efficient Inference for Ternary Language Models](#)", in ACL 2025.
21. Paul Janson, Vaibhav Singh, Paria Mehrbod, Adam Ibrahim, Irina Rish, Eugene Belilovsky, Benjamin Thérien, "[Beyond Cosine Decay: On the effectiveness of Infinite Learning Rate Schedule for Continual Pre-training](#)", in CoLLAs 2025 (oral).
22. Gwen Legate, Irina Rish, Eugene Belilovsky, "[Warming Up for Zeroth-Order Federated Pre-Training with Low Resource Clients](#)", in CoLLAs 2025.
23. Istabrak Abbes, Gopeshh Subbaraj, Matthew Riemer, Nizar Islah, Tsuguchika Tabaru, Hiroaki Kingetsu, Sarath Chandar, Irina Rish, "[Revisiting Replay and Gradient Alignment for Continual Pre-Training of Large Language Models](#)", in CoLLAs 2025.
24. Tianyu Zhang, Andrew Robert Williams, Phillip Wozny, Kai-Hendrik Cohrs, Koen Ponse, Marco Jiralerspong, Soham Rajesh Phade, Sunil Srinivasa, Lu Li, Yang Zhang, Prateek Gupta, Erman Acar, Irina Rish, Yoshua Bengio, Stephan Zheng. "[AI for Global Climate Cooperation: Modeling Global Climate Negotiations, Agreements, and Long-Term Cooperation in RICE-N](#)", The Forty-Second International Conference on Machine Learning (ICML), 2025.
25. Andrew Robert Williams*, Arjun Ashok*, Étienne Marcotte, Valentina Zantedeschi, Jithendaraa Subramanian, Roland Riachi, James Requeima, Alexandre Lacoste, Irina Rish, Nicolas Chapados, Alexandre Drouin. "[Context is Key: A Benchmark for Forecasting with Essential Textual Information](#)", The Forty-Second International Conference on Machine Learning (ICML), 2025.
26. Ayush Kaushal, Tejas Vaidhya, Arnab Kumar Mondal, Tejas Pandey, Aaryan Bhagat, Irina Rish. "[Spectra: Surprising Effectiveness of Pretraining Ternary Language Models at Scale](#)", The Thirteenth International Conference on Learning Representations (ICLR), 2025 (spotlight).
27. Md Rifat Arefin, Gopeshh Subbaraj, Nicolas Gontier, Yann LeCun, Irina Rish, Ravid Shwartz-Ziv, Christopher Pal. "[Seq-VCR: Preventing Collapse in Intermediate Transformer Representations for Enhanced Reasoning](#)", The Thirteenth International Conference on Learning Representations (ICLR), 2025.
28. Matthew Riemer, Gopeshh Subbaraj, Glen Berseth, Irina Rish. "[Enabling Realtime Reinforcement Learning at Scale with Staggered Asynchronous Inference](#)", The Thirteenth International Conference on Learning Representations (ICLR), 2025.
29. Arnab Kumar Jain, Harley Wiltzer, Jesse Farebrother, Irina Rish, Glen Berseth, Sanjiban Choudhury. "[Non-Adversarial Inverse Reinforcement Learning via Successor Feature Matching](#)", The Thirteenth International Conference on Learning Representations (ICLR), 2025.
30. Ivan Anokhin, Rishav Rishav, Matthew Riemer, Stephen Chung, Irina Rish, Samira Ebrahimi Kahou. "[Handling Delay in Reinforcement Learning Caused by Parallel Computations of Neurons](#)", The Thirteenth International Conference on Learning Representations (ICLR), 2025.
31. Adam Ibrahim, Benjamin Thérien, Kshitij Gupta, Mats L. Richter, Quentin Anthony, Timothée Lesort, Eugene Belilovsky, Irina Rish, "[Simple and Scalable Strategies to Continually Pre-train Large Language Models](#)", Transactions on Machine Learning Research (TMLR), 2024.

32. Germán Abrevaya, Mahta Ramezani-Panahi, Jean-Christophe Gagnon-Audet, Pablo Polosecki, [Irina Rish](#), Silvina Ponce Dawson, Guillermo Cecchi, Guillaume Dumas, "[Effective Latent Differential Equation Models via Attention and Multiple Shooting](#)". Transactions on ML Research (TMLR), 2024.
33. Mohammad-Javad Darvishi-Bayazi, Mohammad Sajjad Ghaemi, Timothee Lesort, Md Rifat Arefin, Jocelyn Faubert, [Irina Rish](#), "[Amplifying Pathological Detection in EEG Signaling Pathways through Cross-Dataset Transfer Learning](#)", Computers in Biology and Medicine, Vol 169, February 2024.
34. Maurice Weber, Daniel Y Fu, Quentin Gregory Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adam, "[RedPajama: an Open Dataset for Training Large Language Models](#)", NeurIPS 2024 Datasets and Benchmarks Track (spotlight).
35. Connor Brennan, Andrew Robert Williams, Omar G. Younis, Vedant Vyas, Daria Yasafova, Irina Rish, "[Using Unity to Help Solve Reinforcement Learning](#)", NeurIPS 2024 Datasets and Benchmarks Track.
36. Rishika Bhagwatkar, Shravan Nayak, Pouya Bashivan, Irina Rish, "[Improving Adversarial Robustness in Vision-Language Models with Architecture and Prompt Design](#)", in Proc of EMNLP 2024.
37. Alessio Mora, Irene Tenison, Paolo Bellavista, Irina Rish. "[Knowledge distillation for federated learning: a practical guide](#)." Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, Survey Track. pp 8188-8196.
38. Rifat Arefin, Yan Zhang, Aristide Baratin, Francesco Locatello, Irina Rish, Dianbo Liu, and Kenji Kawaguchi. "[Unsupervised Concept Discovery Mitigates Spurious Correlations](#)". In Proc of ICML 2024. *arXiv preprint arXiv:2402.13368* (2024).
39. Mohammad-Javad Darvishi-Bayazi, Mohammad Sajjad Ghaemi, Timothee Lesort, Md Rifat Arefin, Jocelyn Faubert, Irina Rish, "[Amplifying Pathological Detection in EEG Signaling Pathways through Cross-Dataset Transfer Learning](#)", in CoLLAs 2024 - Journal Track.
40. Tommaso Tosato, Mahmood Hegazy, David Lemay, Mohammed Abukalam, Irina Rish, Guillaume Dumas, "[LLMs and Personalities: Inconsistencies Across Scales](#)", in NeurIPS 2024 workshop on NeurIPS 2024 Workshop on Behavioral ML.
41. Andrei Mircea, Ekaterina Lobacheva, Supriyo Chakraborty, Nima Chitsazan, Irina Rish. "[Language model scaling laws and zero-sum learning](#)", in NeurIPS 2024 workshop on Scientific Methods for Understanding Deep Learning.
42. Benjamin Thérien, Charles-Étienne Joseph, Boris Knyazev, Edouard Oyallon, Irina Rish, Eugene Belilovsky, "[muLO: Compute-Efficient Meta-Generalization of Learned Optimizers](#)", in NeurIPS 2024 workshop on Optimization for Machine Learning.
43. Arjun Ashok, Andrew Robert Williams, Étienne Marcotte, Valentina Zantedeschi, Jithendaraa Subramanian, Roland Riachi, James Requeima, Alexandre Lacoste, Irina Rish, Nicolas Chapados, Alexandre Drouin, "[Context is Key: A Benchmark for Forecasting with Essential Textual Information](#)", in NeurIPS 2024 workshop on Time Series in the Age of Large Models.
44. Mohammad Javad Darvishi Bayazi, Hena Ghonia, Roland Riachi, Bruno Aristimunha, Arian Khorasani, Md Rifat Arefin, Sylvain Chevallier, Guillaume Dumas, Irina Rish, "[Introducing Brain Foundation Models](#)", in NeurIPS 2024 workshop on Time Series in the Age of Large Models.
45. Mohammad-Javad Darvishi-Bayazi, Md Rifat Arefin, Jocelyn Faubert, Irina Rish, "[VFA: Vision Frequency Analysis of Foundation Models and Human](#)", In ICML 2024 Workshop on Foundation Models in the Wild.
46. Arnav Kumar Jain, Harley Wiltzer, Jesse Farebrother, Irina Rish, Glen Berseth, Sanjiban Choudhury, "[Revisiting Successor Features for Inverse Reinforcement Learning](#)". In ICML 2024 Workshop on Model of Human Feedback for AI Alignment.
47. Notsawo Jr, Pascal, Hattie Zhou, Mohammad Pezeshki, Irina Rish, and Guillaume Dumas. "[Predicting Grokking Long Before it Happens: A look into the loss landscape of models which grok](#)". In the ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models (ME-FoMo).
48. Justine Gehring*, Nizar Islah*, Diganta Misra*, Massimo Caccia, Irina Rish. "[GitChameleon: Breaking the version barrier for code generation models](#)". In Proc of the 4th DMLR workshop at ICLR 2024.
49. Karolis Jucys*, George Adamopoulos*, Mehrab Hamidi, Stephanie Milani, Mohammad Reza Samsami, Artem Zholus, Sonia Joseph, Blake Richards, Irina Rish, and Ozgur Simsek. "Interpretability in Action: Exploratory Analysis of VPT, a Minecraft Agent". In the ICML 2024 Workshop on Mechanistic Interpretability.

50. Matthew Riemer*, Gopeshh Subbaraj*, Glen Berseth, Irina Rish. "Realtime Reinforcement Learning: Towards Rapid Asynchronous Deployment of Large Models". In the ICML 2024 Workshop on Aligning Reinforcement Learning Experimentalists and Theorists and the RLC 2024 Workshop on Finding the Frame.
51. Maximilian Puelma Touzel, Amin Memarian, Matthew D Riemer, Andrei Mircea, Andrew Robert Williams, Elin Ahlstrand, Lucas Lehnert, Rupali Bhati, Guillaume Dumas, Irina Rish. "Scalable Approaches for a Theory of Many Minds". ICML 2024 Workshop on Agentic Markets.
52. Ivan Anokhin, Rishav, Stephen Chung, Irina Rish, Samira Ebrahimi Kahou. "Handling Delay in Reinforcement Learning Caused by Parallel Computations of Neurons". In the ICML 2024 Workshop on Aligning Reinforcement Learning Experimentalists and Theorists and the RLC 2024 Workshop on Deployable RL.
53. Andrei Mircea, Ekaterina Lobacheva, Irina Rish. "Gradient Dissent in Language Model Training and Saturation." In High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning.
54. Bayazi, Mohammad Javad Darvishi, Andrew Law, Sergio Mejia Romero, Sion Jennings, Irina Rish, and Jocelyn Faubert. "[Neural efficiency in an aviation task with different levels of difficulty: Assessing different biometrics during a performance task.](#)" *Journal of Vision* 23, no. 9 (2023): 5647-5647.
55. Mohammad-Javad Darvishi-Bayazi, Andrew Law, Sergio Mejia Romero, Sion Jennings, Irina Rish, Jocelyn Faubert. "[Beyond performance: the role of task demand, effort, and individual differences in ab initio pilots](#)". In *Nature Sci Rep* 13, 14035 (2023).
56. Irene Tenison, Sai Aravind Sreeramadas, Vaikkunth Mugunthan, Edouard Oyallon, Irina Rish, Eugene Belilovsky. "[Gradient Masked Averaging for Federated Learning](#)". In *Transactions on Machine Learning Research* (2023).
57. Gagnon-Audet, Jean-Christophe, Kartik Ahuja, Mohammad-Javad Darvishi-Bayazi, Guillaume Dumas, and Irina Rish. "[WOODS: Benchmarks for Out-of-Distribution Generalization in Time Series Tasks](#)" In *Transactions on Machine Learning Research* (2023). arXiv preprint arXiv:2203.09978
58. Lam, Guillaume, Irina Rish, and Philippe C. Dixon. "[Estimating individual minimum calibration for deep-learning with predictive performance recovery: An example case of gait surface classification from wearable sensor gait data.](#)" *Journal of Biomechanics* (2023): 111606.
59. Ardavan S Nobandegani, Irina Rish, Thomas Shultz, "[Decision-Making Paradoxes in Humans vs Machines: The case of the Allais and Ellsberg Paradoxes](#)", The 45th Annual Conference of the Cognitive Science Society (**CogSci**), 2023.
60. Ardavan S. Nobandegani, Thomas R. Shultz. "[Neural Network Modeling of Pure Reasoning in Preverbal Infants](#)". The 45th Annual Conference of the Cognitive Science Society (**CogSci**), 2023.
61. Thomas R. Shultz, Ardavan S. Nobandegani, Zilong Wang. "[A Neural Model of Number Comparison with Robust Generalization](#)". The 45th Annual Conference of the Cognitive Science Society (**CogSci**), 2023.
62. Ardavan S. Nobandegani, Irina Rish, Thomas R. Shultz. "[AI Agents Learn to Trust](#)". [Abstract]. The 45th Annual Conference of the Cognitive Science Society (**CogSci**), 2023. (poster presentation) *arXiv preprint arXiv:2312.12868*.
63. Arnav Kumar Jain, Lucas Lehnert, Irina Rish and Glen Berseth, "[Maximum State Entropy Exploration using Predecessor and Successor Representations](#)", in Proc of NeurIPS 2023.
64. Bouneffouf, Djallel, Mayank Agarwal, and Irina Rish. "[Dialogue System with Missing Observation.](#)" *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023.
65. Timothée Lesort, Oleksiy Ostapenko, Diganta Misra, Md Rifat Arefin, Pau Rodríguez, Laurent Charlin, Irina Rish. "[Challenging Common Assumptions about Catastrophic Forgetting.](#)" Second Conference on Lifelong Learning Agents (CoLLAs 2023)
66. Caballero, Ethan; Gupta, Kshitij; Rish, Irina; Krueger, David (2022). "[Broken Neural Scaling Laws](#)". International Conference on Learning Representations (ICLR 2023).
67. Alexis Roger, Esma Aimeur, Irina Rish, "[Towards Ethical Multimodal Systems](#)", in the First Workshop on AI meets Moral Philosophy and Moral Psychology (MP2) at NeurIPS 2023.
68. Kshitij Gupta, Benjamin Therien, Adam Ibrahim, Mats Leon Richter, Quentin Anthony, Eugene Belilovsky, Irina Rish, Timothee Lesort, "[Continual Pre-Training of Large Language Models: How to \(re\)warm your model?](#)" (new version), in Efficient Natural Language and Speech Processing (ENLSP-III) workshop at NeurIPS 2023.

69. Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Arian Khorasani, George Adamopoulos, Rishika Bhagwatkar, Marin Biloš, Hena Ghonia, Nadhir Hassen, Anderson Schneider, Sahil Garg, Alexandre Drouin, Nicolas Chapados, Yuriy Nevmyvaka, Irina Rish, "[Lag-Llama: Towards Foundation Models for Time Series Forecasting](#)", in R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models workshop at NeurIPS 2023.
70. Nizar Islah, Diganta Misra, Timothy Nest, Matthew D Riemer, Irina Rish, Eilif Benjamin Muller, "[Mitigating Mode Collapse in Sparse Mixture of Experts](#)", in New in Machine Learning Workshop at NeurIPS 2023.
71. Germán Abrevaya, Mahta Ramezani-Panahi, Jean-Christophe Gagnon-Audet, Pablo Polosecki, Irina Rish, Silvina Ponce Dawson, Guillermo Cecchi, Guillaume Dumas, "[Effective Latent Differential Equation Models via Attention and Multiple Shooting](#)". The Symbiosis of Deep Learning and Differential Equations III workshop at NeurIPS 2023.
72. Kshitij Gupta, Benjamin Thérien, Adam Ibrahim, Mats Leon Richter, Quentin Anthony, Eugene Belilovsky, Irina Rish, Timothee Lesort. "[Continual Pre-Training of Large Language Models: How to re-warm your model?](#)". In ES-FoMo: Efficient Systems for Foundation Models workshop at ICML-2023 (2023).
73. Nobandegani, Ardavan S., Thomas Shultz, and Irina Rish. "[Cognitive Models as Simulators: Using Cognitive Models to Tap into Implicit Human Feedback](#)" In Interactive Learning with Implicit Human Feedback (ILHF) workshop at ICML-2023 (2023).
74. Jain, Arnav Kumar, Lucas Lehnert, Irina Rish, and Glen Berseth. "[Maximum State Entropy Exploration using Predecessor and Successor Representations](#)" In New Frontiers in Learning, Control, and Dynamical Systems workshop at the International Conference on Machine Learning (ICML-2023). *arXiv preprint arXiv:2306.14808* (2023).
75. Khetarpal, Khimya, Matthew Riemer, Irina Rish, and Doina Precup. "[Towards Continual Reinforcement Learning: A Review and Perspectives](#)" *Journal of Artificial Intelligence Research* 75 (2022): 1401-1476.
76. Ramezani-Panahi, Mahta, Germán Abrevaya, Jean-Christophe Gagnon-Audet, Vikram Voleti, Irina Rish, and Guillaume Dumas. "[Generative Models of Brain Dynamics](#)." *Frontiers in Artificial Intelligence* (2022): 147. June 2022.
77. JC Layoun, A Roger, I Rish. "[Aligning MAGMA by Few-Shot Learning and Finetuning](#)". Presented at Montreal AI Symposium 2022 conference. *arXiv:2210.14161*, 2022.
78. Ardavan S. Nobandegani, Thomas Shultz, and Irina Rish. "[Cognitive Models as Simulators: The Case of Moral Decision-Making](#)". In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*. <https://arxiv.org/abs/2210.04121> (2022).
79. Riemer, Matthew, Sharath Chandra Raparthy, Ignacio Cases, Gopeshh Subbaraj, Maximilian Puelma Touzel, and Irina Rish. "[Continual learning in environments with polynomial mixing times](#)" *Advances in Neural Information Processing Systems* 35 (2022): 21961-21973.
80. Ernout, Maxence M., Fabrice Normandin, Abhinav Moudgil, Sean Spinney, Eugene Belilovsky, Irina Rish, Blake Richards, and Yoshua Bengio. "[Towards Scaling Difference Target Propagation by Learning Backprop Targets](#)". In International Conference on Machine Learning (ICML), pp. 5968-5987. July 2022.
81. Ostapenko, Oleksiy, Timothee Lesort, Pau Rodríguez, Md Rifat Arefin, Arthur Douillard, Irina Rish, and Laurent Charlin. "[Foundational Models for Continual Learning: An Empirical Study of Latent Replay](#)" *arXiv preprint arXiv:2205.00329* (2022). 1st CoLLAs conference, August 2022.
82. Gauthier, Shanel, Benjamin Thérien, Laurent Alsène-Racicot, Muawiz Chaudhary, Irina Rish, Eugene Belilovsky, Michael Eickenberg, and Guy Wolf. "[Parametric Scattering Networks](#)" In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5749-5758. June 2022.
83. Gagnon-Audet, Jean-Christophe, Soroosh Shahtalebi, Frank Rudzicz, and Irina Rish. "[A Remedy For Distributional Shifts Through Expected Domain Translation](#)". In ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 4523-4527. IEEE, May 2022.
84. Mittal, Sarthak, Sharath Chandra Raparthy, Irina Rish, Yoshua Bengio, and Guillaume Lajoie. "[Compositional Attention: Disentangling Search and Retrieval](#)" *arXiv preprint arXiv:2110.09419*, ICLR, 2022..

85. Ibrahim, Adam; Guille-Escuret, Charles; Mitliagkas, Ioannis; Rish, Irina; Krueger, David; Bashivan, Pouya. "[Towards Out-of-Distribution Adversarial Robustness](#)", In New Frontiers in Adversarial Machine Learning workshop at ICML-2022, July 2022.
86. Lesort, Timothée, Oleksiy Ostapenko, Diganta Misra, Md Rifat Arefin, Pau Rodríguez, Laurent Charlin, and Irina Rish. "[Scaling the Number of Tasks in Continual Learning](#)" arXiv preprint arXiv:2207.04543, CoLLAs workshop (2022).
87. Diganta Misra, Bharat Runwal, Tianlong Chen, Zhangyang Wang, and Irina Rish. "[APP: Anytime Progressive Pruning](#)" arXiv preprint arXiv:2204.01640 (2022). In 1st Workshop on Dynamic Neural Networks, ICML 2022 (DyNN) and 2nd Sparsity in Neural Networks workshop 2022 (SNN), July 2022.
88. Memarian, Amin, Maximilian Puelma Touzel, Matthew D. Riemer, Rupali Bhati, and Irina Rish. "[Summarizing Societies: Agent Abstraction in Multi-Agent Reinforcement Learning](#)" From Cells to Societies: Collective Learning across Scales. ICLR workshop, April 2022.
89. Nadhir Hassen, Irina Rish. "[Approximate Bayesian Optimisation for Neural Networks](#)", in AAAI-22 Workshop on Learning Network Architecture During Training, 2022.
90. Abrevaya, Germán, Guillaume Dumas, Aleksandr Y Aravkin, Peng Zheng, Jean-Christophe Gagnon-Audet, James Kozloski, Pablo Polosecki, Guillaume Lajoie, David Cox, Silvina Ponce Dawson, Guillermo Cecchi, Irina Rish. "[Learning Brain Dynamics With Coupled Low-Dimensional Nonlinear Oscillators and Deep Recurrent Networks](#)", Neural Computation 33, no. 8 (2021): 2087-2127.
91. Ahuja, Kartik, Ethan Caballero, Dinghuai Zhang, Jean-Christophe Gagnon-Audet, Yoshua Bengio, Ioannis Mitliagkas, and Irina Rish. "[Invariance Principle Meets Information Bottleneck for Out-of-Distribution Generalization](#)." Advances in Neural Information Processing Systems (NeurIPS) 34 (2021): 3438-3450.
92. Bashivan, Pouya, Reza Bayat, Adam Ibrahim, Kartik Ahuja, Mojtaba Faramarzi, Touraj Laleh, Blake Richards, and Irina Rish. "[Adversarial Feature Desensitization](#)" Advances in Neural Information Processing Systems (NeurIPS) 34 (2021): 10665-10677.
93. Bouneffouf, Djallel, Raphael Feraud, Sohini Upadhyay, Mayank Agarwal, Yasaman Khazaeni, and Irina Rish. "[Toward Skills Dialog Orchestration with Online Learning](#)" In ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 3600-3604. IEEE, 2021.
94. Bouneffouf, Djallel, Raphaël Féraud, Sohini Upadhyay, Yasaman Khazaeni, and Irina Rish. "[Double-Linear Thompson Sampling for Context-Attentive Bandits](#)." In ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 3450-3454. IEEE, 2021.
95. Bengio Y; Gupta P; Maharaj, T; Rahaman, N; Weiss, M; Deleu, T; Muller E.B; Qu, M; Schmidt, V; St-charles, P.L; Alsdurf, H; Bilaniuk, O; Buckeridge, D; Caron G; Carrier, P.L; Ghosn, J; Ortiz Gagne, S; Pal, C; Rish, I; Schölkopf, B; Sharma, A; Tang, J; Williams, A. [Predicting Infectiousness for Proactive Contact Tracing](#). 2021. International Conference on Learning Representations (ICLR).
96. Shanel Gauthier, Benjamin Thérien, Laurent Alsène-Racicot, Irina Rish, Eugene Belilovsky, Michael Eickenberg, Guy Wolf. "[Parametric Scattering Networks](#)", in Proc of 2021 Conference on the Mathematical Theory of Deep Learning (DeepMath 2021).
97. Dinghuai Zhang, Kartik Ahuja, Yilun Xu, Yisen Wang, Aaron Courville. "[Can Subnetwork Structure be the Key to Out-of-Distribution Generalization?](#)" In Proc of ICML 2021.
98. Kartik Ahuja, P. Sattigeri, K. Shanmugam, D. Wei, K.N. Ramamurthy, M. Kocaglu. "[Conditionally Independent Data Generation](#)." In Proc of UAI 2021.
99. Kartik Ahuja, Jun Wang, Amit Dhurandhar, Karthikeyan Shanmugam, Kush R. Varshney. "[Empirical or Invariant Risk Minimization? A Sample Complexity Perspective](#)." In Proc of ICLR 2021.
100. Kartik Ahuja, Karthikeyan Shanmugam, Amit Dhurandhar, "[Linear Regression Games: Convergence Guarantees to Approximate Out-of-Distribution Solutions](#)". In Proc of AISTATS, 2021.
101. Normandin, Fabrice, Florian Golemo, Oleksiy Ostapenko, Pau Rodriguez, Matthew D Riemer, Julio Hurtado, Khimya Khetarpal, Dominic Zhao, Ryan Lindeborg, Timothée Lesort, Laurent Charlin, Irina Rish, Massimo Caccia. "[Sequoia: A Software Framework to Unify Continual Learning Research](#)", arXiv preprint arXiv:2108.01005 (2021). Workshop on Theory and Foundation of Continual Learning (ICML Workshop), 2021.

102. Tenison, Irene, Sreya Francis, and Irina Rish. "[Gradient Masked Federated Optimization](#)" arXiv preprint arXiv:2104.10322 (2021). ICLR 2021 workshop on Distributed and Private Machine Learning (DPML) and NeurIPS 2021 Workshop on Federated Learning.
103. Francis, Sreya, Irene Tenison, and Irina Rish. "[Towards Causal Federated Learning For Enhanced Robustness and Privacy](#)" arXiv preprint arXiv:2104.06557 (2021). ICLR 2021 workshop on Distributed and Private Machine Learning (DPML), 2021
104. Lesort, Timothée, Thomas George, and Irina Rish. "[Continual Learning in Deep Networks: an Analysis of the Last Layer](#)" arXiv preprint arXiv:2106.01834 (2021).. International Conference of Machine Learning 2021 (ICML) Workshop on Theory and Foundation of Continual Learning.
105. Lesort, Timothée, Massimo Caccia, and Irina Rish. "[Understanding Continual Learning Settings with Data Distribution Drift Analysis](#)". arXiv preprint arXiv:2104.01678 (2021). International Conference of Machine Learning 2021 (ICML) Workshop on Theory and Foundation of Continual Learning.
106. Shahtalebi, Soroosh, Jean-Christophe Gagnon-Audet, Touraj Laleh, Mojtaba Faramarzi, Kartik Ahuja, and Irina Rish. "[SAND-mask: An Enhanced Gradient Masking Strategy for the Discovery of Invariances in Domain Generalization](#)" arXiv preprint arXiv:2106.02266 (2021). ICML 2021 workshop on Uncertainty and Robustness in ML.
107. Gabriele Prato, Simon Guiroy, Ethan Caballero, Irina Rish and Sarath Chandar. "[Scaling Laws for the Few-Shot Adaptation of Pre-trained Image Classifiers](#)", in Proc of ICML 2021 Workshop on Uncertainty & Robustness in Deep Learning.
108. Amin Mansouri, Sean Spinney, Amin Memarian, Patricia Conrod, and Irina Rish. "[Identifying Invariant and Sparse Predictors in High-dimensional Data](#)", in Proc of ICML 2021 Workshop on Uncertainty & Robustness in Deep Learning.
109. Lin, Baihan, Guillermo Cecchi, Djallel Bouneffouf, Jenna Reinen, and Irina Rish. "[Models of Human Behavioral Agents in Bandits, Contextual Bandits and RL](#)" In International Workshop on Human Brain and Artificial Intelligence, pp. 14-33. Springer, Singapore, 2021.
110. Polosecki, P; Castro, E; Rish, I; Pustina, D; Warner J.H; Wood, A; Sampaio, C; Cecchi, G. [Resting-state connectivity stratifies premanifest Huntington's disease by longitudinal cognitive decline rate](#). 2020. Nature Scientific Reports 10 (1), 1-15.
111. Garg, S; Rish, I; Cecchi G; Goyal, P; Ghazarian, S; Gao, S; Ver Steeg, G; Galstyan, A. "[Modeling Psychotherapy Dialogues with Kernelized Hashcode Representations: A Nonparametric Information-Theoretic Approach](#)", 2020. Association for the Advancement of Artificial Intelligence (AAAI).
112. Lin, B; Cecchi, G; Bouneffouf, D; Reinen, J; Rish, I. "[A Story of Two Streams: Reinforcement Learning Models from Human Behavior and Neuropsychiatry](#)". 2020. Autonomous Agents & Multiagent Systems (AAMAS).
113. Caccia, M; Rodriguez, P; Ostapenko, O; Normandin, F; Lin, M; Caccia, L; Laradji, I.H; Rish, I; Lacoste, A; Vazquez, D; Charlin L. "[Online Fast Adaptation and Knowledge Accumulation: A New Approach to Continual Learning](#)". 2020. NeurIPS.
114. Bouneffouf, D; Rish, I; Aggarwal, C. *Survey on Applications of Multi-Armed and Contextual Bandits*. 2020. IEEE Congress on Evolutionary Computation (CEC).
115. Touraj L; Raparthy C.S; Rish, I. *Chaotic Continual Learning*. 2020. International Conference on Machine Learning (ICML) Workshop on Lifelong Learning.
116. Lin, B; Cecchi, G; Bouneffouf, D; Reinen, J; Rish, I. "[Unified Models of Human Behavioral Agents in Bandits, Contextual Bandits and RL](#)", Knowledge Discovery and Data Mining (KDD) Workshop on *Designing AI in support of Good Mental Health*. 2020.
117. Anna Choromanska, Benjamin Cowen, Sadhana Kumaravel, Ronny Luss, Mattia Rigotti, Irina Rish, Brian Kingsbury, Paolo DiAchille, Viatcheslav Gurev, Ravi Tejawani, Djallel Bouneffouf. Beyond Backprop: Online Alternating Minimization with Auxiliary Variables. In Proceedings of Thirty-sixth International Conference on Machine Learning (ICML 2019), 2019.
118. Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to Learn without Forgetting by Maximizing Transfer and Minimizing Interference. In Proceedings of Seventh International Conference on Learning Representations (ICLR 2019), 2019.
119. Sahil Garg, Aram Galstyan, Irina Rish, Guillermo Cecchi, and Shuyang Gao. Efficient Representation for Natural Language Processing via Kernelized Hashcodes. In Proceedings of AAAI-2019, 2019.
120. J. McDonnell, W. Hord, J. Reinen, P. Polosecki, I. Rish and G. Cecchi. Predicting conversion to psychosis in clinical high risk patients using resting-state functional MRI features. In Medical Imaging

- 2019: Biomedical Applications in Molecular, Structural, and Functional Imaging (Vol. 10953, p. 109532A). International Society for Optics and Photonics, 2019.
121. Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesaro. Continual Learning by Maximizing Transfer and Minimizing Interference. In NIPS 2018 Workshop on Continual Learning, 2018.
122. Pouya Bashivan, Martin Schrimpf, Robert Ajemian, Irina Rish, Matthew Riemer and Yuhai Tu. Continual Learning with Self-Organizing Maps. In NIPS 2018 Workshop on Continual Learning, 2018.
123. Baihan Lin, Guillermo Cecchi, Djallel Bouneffouf and Irina Rish. Contextual Bandit with Adaptive Feature Extraction. Preprint arXiv:1802.00981 [cs.AI], in Proc. 6th ICDM Workshop on Data Science and Big Data Analytics (DSBDA-2018), 2018.
124. Sahil Garg, Guillermo Cecchi, Irina Rish, Shuyang Gao, Bhavana Bhaskar, Greg Ver Steeg, Palash Goyal, and Aram Galstyan. Dialogue Modeling Via Hash Functions: Applications to Psychotherapy. Preprint arXiv:1804.10188, in Proceedings of IJCAI-2019 workshop on Linguistic and Cognitive Approaches To Dialog Agents (LaCATODA), 2018.
125. E. Castro, P. Polosecki, I. Rish, D. Pustina, JH Warner, A. Wood, C. Sampaio, GA Cecchi. Baseline multimodal information predicts future motor impairment in premanifest Huntington's disease. *NeuroImage - Clinical*, 2018.
126. J. Dakka, P. Bashivan, M. Gheiratmand, I. Rish, S. Jha, R. Greiner. Learning Neural Markers of Schizophrenia Disorder Using Recurrent Neural Networks. In NIPS 2017 Workshop on Machine Learning for Health (ML4H), 2017.
127. Mina Gheiratmand, Irina Rish, Guillermo A Cecchi, Matthew RG Brown, Russell Greiner, Pablo Polosecki, Pouya Bashivan, Andrew J Greenshaw, Rajamannar Ramasubbu, Serdar M Dursun. Learning stable and predictive network-based patterns of schizophrenia and its clinical symptoms. *Nature pj Schizophrenia* 3(1), 22, Nature Publishing Group, 2017.
128. Djallel Bouneffouf, Irina Rish, Raphael Feraud and Guillermo Cecchi. Attentive Bandit: Contextual Bandit with Restricted Context. In Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-2017), 2017.
129. Djallel Bouneffouf, Irina Rish, Guillermo A. Cecchi, Bandit Models of Human Behavior: Reward Processing in Mental Disorders. In Proceedings of Artificial General Intelligence (AGI-2017), 2017.
130. Sahil Garg, Irina Rish, Guillermo A. Cecchi, Aurelie Lozano. Neurogenesis-Inspired Dictionary Learning: Online Model Adaption in a Changing World. In Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-2017). Extended version arXiv:1701.06106.
131. P. Polosecki, E. Castro, A. Wood, J. H. Warner, I. Rish and G. A. Cecchi. Computational psychiatry: Advancing predictive modeling of neurodegeneration with neuroimaging of Huntington's disease. *IBM Journal of Research and Development* 61(2/3), IEEE, 2017.
132. G. A. Cecchi, V. Gurev, S.J. Heisig, R. Norel, I. Rish, S. R. Schrecke. Computing the structure of language for neuropsychiatric evaluation. *IBM Journal of Research and Development* 61(2/3), IEEE, 2017.
133. I. Rish and G.A. Cecchi. Holographic brain: Distributed versus local activation patterns in fMRI. *IBM Journal of Research and Development* 61(2/3), IEEE, 2017.
134. Mina Gheiratmand, Irina Rish, Guillermo Cecchi, Matthew Brown, Russell Greiner, Pouya Bashivan, Pablo Polosecki, Serdar Dursun. Learning Discriminative Functional Network Features of Schizophrenia. In Proc. of SPIE Medical Imaging, 2017.
135. I. Rish, P. Bashivan, G.A. Cecchi, R.Z. Goldstein. Evaluating Effects of Methylphenidate on BrainActivity in Cocaine Addiction: A Machine-Learning Approach. Proc. of SPIE Medical Imaging, 2016.
136. P. Bashivan, I. Rish, M. Yeasin, N. Codella. Learning Representations from EEG with Deep RecurrentConvolutional Neural Networks. In Proc. of International Conference on Learning Representations (ICLR), 2016.
137. D.C. Haws, I. Rish, S. Teyssedre, D. He, A.C. Lozano, P. Kambadur, Z. Karaman, L. Parida. VariableSelection Emerges on Top in Empirical Comparison of Whole-Genome Complex-Trait Prediction Methods. *PloS one* 10(10), e0138903, Public Library of Science, 2015.
138. P. Bashivan, I. Rish, S. Heisig. Mental State Recognition via Wearable EEG. In Proc. of NIPS Workshop on Machine Learning and Interpretation in Neuroimaging (MLINI), 2015.

- 139.D. He, I. Rish, D. Haws, L. Parida. MINT: Mutual Information based Transductive Feature Selection For Genetic Trait Prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* pp. 99, 2015.
- 140.G.A. Cecchi and I. Rish. Turing a la Freud: Test for an Automated Psychiatrist. In *Proc. of AAAI2015 Workshop on Beyond the Turing Test*, 2015.
- 141.D. He, I. Rish, L. Parida. Transductive HSIC Lasso. *Proc. of SDM*, 2014.
- 142.S. Heisig, G. Cecchi, R. Rao and I. Rish. Augmented Human: Human OS for Improved Mental Function. *AAAI 2014 Workshop on Cognitive Computing and Augmented Human Intelligence*, 2014.
- 143.I. Rish. Machine Learning and Sparse Modeling in Brain Imaging. Invited paper at European Conference in Machine Learning (ECML-2013), Prague, 2013.
- 144.I. Rish, G.A. Cecchi, B. Thyreau, B. Thirion, M. Plaze, M.L. Paillere-Martinot, C. Martelli, J.L.Martinot, J.B. Poline. Schizophrenia as a Network Disease: Disruption of Emergent Brain Function in Patients with Auditory Hallucinations. *PloS ONE* 8(1), e50625, Public Library of Science, 2013.
- 145.G.A. Cecchi, L. Huang, J.A. Hashmi, M. Baliki, M.V. Centeno, I. Rish, A.V. Apkarian. Predictive Dynamics of Human Pain Perception. *PLoS Computational Biology* 8(10), e1002719, Public Library of Science, 2012.
- 146.I. Rish, G.A. Cecchi, K. Heuton, M.N. Baliki, A.V. Apkarian. Sparse regression analysis of task-relevant information distribution in the brain. *Proc. of SPIE Medical Imaging*, 2012.
- 147.J. Honorio, D. Samaras, I. Rish, G.A. Cecchi. Variable Selection for Gaussian Graphical Models. *Proc. of AISTATS*, 2012.
- 148.I. Rish, G. Cecchi, M.N. Baliki and A.V. Apkarian. Sparse Regression Models of Pain Perception, in *Proc. of Brain Informatics (BI-2010)*, Toronto, Canada, August 2010.
- 149.K. Scheinberg and I. Rish. Learning Sparse Gaussian Markov Networks using a Greedy CoordinateAscent Approach, in *Proc. of ECML 2010*, Barcelona, Spain, September 2010.
- 150.K. Scheinberg, I. Rish, N. Bani Asadi, Sparse Markov Net Learning with Priors on RegularizationParameters, in *Proc. of the Intl. Symp. on AI and Math (ISAIM-2010)*, Fort Lauderdale, FL, Jan. 2010.
- 151.M.K. Carroll, G.A.Cecchi, I. Rish, R. Garg, A.R. Rao. Prediction and Interpretation of Distributed Neural Activity with Sparse Models, *Neuroimage* 44(1):112-22, 2009.
- 152.G. Cecchi, I. Rish, B. Thyreau, B. Thirion, M. Plaze, ML Paillere-Martinot, C. Martelli, JL Martinot,JB Poline. Discriminative Network Models of Schizophrenia, in *Proc. of NIPS 2009*, Vancouver, BC, Canada, December 2009.
- 153.I. Rish, G. Grabarnik. Sparse Signal Recovery with Exponential-Family Noise, in *Proc. of Allerton2009*, Monticello, IL, September 2009.
- 154.N. Bani Asadi, I. Rish, K. Scheinberg, D. Kanevsky, B. Ramabhadran. A MAP Approach to LearningSparse Gaussian Markov Networks, in *Proceedings of ICASSP 2009*, Taipei, Taiwan, April 2009.
- 155.I. Rish, G. Grabarnik, G. Cecchi, F. Pereira, G. Gordon. Closed-form Supervised Dimensionality Reduction with Generalized Linear Models, in *Proceedings of ICML 2008*, Helsinki, Finland, July 2008.
- 156.G. Chandalia and I. Rish. Blind Source Separation Approach to Performance Diagnosis and Dependency Discovery, in *Proceedings of IMC-2007*, San Diego, CA, October 2007.
- 157.I. Rish and G. Tesauro. Estimating End-to-End Performance by Collaborative Prediction with Active Sampling, in *Proceedings of IM-2007*, Munich, Germany, May 2007.
- 158.A. Beygelzimer, J. Kephart and I. Rish. Evaluation of Optimization Methods for Network Bottleneck Diagnosis, in *Proceedings of ICAC 2007*.
- 159.A. Jakulin and I. Rish. Bayesian Learning of Markov Network Structure, in *Proceedings of ECML-2006*, Berlin, Germany, September 2006.
- 160.A. Beygelzimer, M. Brodie, J. Lenchner, I. Rish, Automated Knowledge Elicitation and Flowchart Optimization for Problem Diagnosis, in *UAI-06 Workshop on Applications of Bayesian Networks*.
- 161.I. Rish, Active Sampling Approaches in Systems Management Applications, *SysML Workshop at SIGMETRICS-06*.
- 162.I. Rish, M. Brodie, S. Ma, N. Odintsova, A. Beygelzimer, G. Grabarnik, K. Hernandez. Adaptive Diagnosis in Distributed Systems, in *IEEE Transactions on Neural Networks (special issue on Adaptive Learning Systems in Communication Networks)*, vol. 16, no. 5, pp. 1088-1109, September 2005.

163. A. Beygelzimer, R. Linsker, G. Grinstein, I. Rish. Improving Network Robustness by Edge Modification, *Physica A*, Vol: 357, No: 3-4, pp. 593-612, November 2005.
164. I. Rish, Distributed Systems Diagnosis Using Belief Propagation, to appear in Proc. of the 43-d Annual Allerton Conference on Communication, Control and Computing, 2005.
165. A. Zheng, I. Rish, A. Beygelzimer. Efficient Test Selection in Active Diagnosis via Entropy Approximation, in Proceedings of UAI-2005.
166. N. Odintsova, I. Rish, S. Ma. Multifault Diagnosis in Dynamic Systems, in Proceedings of IM 2005.
167. A. Beygelzimer, M. Brodie, S. Ma, I. Rish. Test-based Diagnosis: Tree and Matrix Representations, in Proceedings of IM 2005.
168. E. Erdogan, S. Ma, A. Beygelzimer, I. Rish. Statistical Models for Unequally Spaced Time Series, in Proceedings of SIAM DM 2005.
169. I. Rish, M. Brodie, N. Odintsova, S. Ma, G. Grabarnik. Real-time Problem Determination in Distributed Systems using Active Probing, in Proceedings of NOMS-2004, Seoul, Korea, April 2004.
170. A. Beygelzimer, G. Grinstein, R. Linsker, I. Rish, Improving Network Robustness, in Proceedings of ICAC-04.
171. Dechter, R. and Rish, I. Mini-buckets: A general scheme for bounded inference, *JACM* 50(2): 107-153 (2003).
172. R. Vilalta and I. Rish (2003). A Decomposition of Classes Via Clustering to Explain and Improve Naive Bayes, in Proceedings of 14th European Conference on Machine Learning (ECML/PKDD 2003). **Best paper award.**
173. M. Brodie, I. Rish, S. Ma, N. Odintsova (2003). Active Probing Strategies for Problem Diagnosis in Distributed Systems, in Proceedings of IJCAI-03, Acapulco, Mexico.
174. R. K. Sahoo, A. J. Oliner, I. Rish, M. Gupta, J.E. Moreira, S. Ma, R. Vilalta, A. Sivasubramaniam(2003). Critical Event Prediction for Proactive Management in Large-scale Computer Clusters, in Proceedings of KDD-2003 (Industrial Track).
175. R.K. Sahoo, I. Rish, A.J. Oliner, M. Gupta, J.E. Moreira, S. Ma, R. Vilalta and A. Sivasubramaniam. Autonomic Computing Features for Large-scale Server Management and Control, in the IJCAI-03 workshop on AI and Autonomic Computing, August 2003, Acapulco, Mexico.
176. I. Rish, M. Brodie, N. Odintsova, S. Ma, G. Grabarnik. Problem Diagnosis in Distributed Systems using Active Probing, in the UAI-2003 workshop on Bayesian Modeling Applications, August 2003, Acapulco, Mexico.
177. I. Rish, M. Brodie and S. Ma. Accuracy vs. Efficiency Trade-offs in Probabilistic Diagnosis, in *Proceedings of AAAI-2002*, Edmonton, Alberta, Canada, 2002.
178. Rish, I., Brodie, M., and Ma, S. 2002. Efficient fault diagnosis using probing, in Proceedings of 2002 AAAI Spring Symposium on Information Refinement and Revision for Decision Making: Modeling for Diagnostics, Prognostics, and Prediction, Stanford, Palo Alto.
179. Wang, H., Rish, I., and Ma, S. 2002. Using Sensitivity Analysis for Selective Parameter Update in Bayesian Network Learning, in Proceedings of 2002 AAAI Spring Symposium on Information Refinement and Revision for Decision Making: Modeling for Diagnostics, Prognostics, and Prediction, Stanford, Palo Alto.
180. Beygelzimer, A. and Rish, I. 2002. Inference complexity as a model-selection criterion for learning Bayesian networks, in Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR2002), Toulouse, France.
181. Brodie, M., Rish, I. and Ma, S. 2001. Intelligent probing: a Cost-Efficient Approach to Fault Diagnosis in Computer Networks, in IBM Systems Journal special issue on AI 41(3), pp. 372-385.
182. Brodie, M., Rish, I. and Ma, S. 2001. Optimizing Probe Selection for Fault Localization, in Proceedings of DSOM-01.
183. Vilalta, R., Brodie, M., Oblinger, D., and Rish, I. 2001. A unified framework for evaluation metrics in classification using decision trees, in Proceedings of ECML-01, pp. 503-514.
184. I. Rish. An empirical study of the naive Bayes classifier. In *Proc. of IJCAI-01 workshop on Empirical Methods in AI*, vol. 3, no. 22, pp. 41-46, 2001.
185. Rish, I. and Dechter, R. 2000. Resolution versus search; Two strategies for SAT. *Journal of Automated Reasoning*, 24(1/2), pp. 225-275.
186. Hellerstein, J., Jayram, T.S., and Rish, I. 2000. Recognizing end-user transactions in performance management, in Proceedings of AAAI-2000, pp. 596-602. Austin, Texas.

187. Rish, I. 2000. Advances in Bayesian Learning. Proceedings of the 2000 International Conference on Artificial Intelligence (IC-AI'2000). Las Vegas.
188. Rish, I., Kask, K., and Dechter, R. 1998. Empirical Evaluation of Approximation Algorithms for Probabilistic Decoding. Proceedings of UAI98, Madison, Wisconsin.
189. Rish, I. and Dechter, R. 1998. On the impact of causal independence, 1998 Stanford Spring Symposium on Interactive and Mixed-Initiative Decision-Theoretic Systems.
190. Rish, I. and Frost, D. 1997. Statistical Analysis of Backtracking on Inconsistent CSPs, in Proceedings of Principles and Practice of Constraint Programming (CP97), Linz, Austria.
191. Dechter, R. and Rish, I. 1997. A scheme for approximating probabilistic inference, in Proceedings of UAI-97, Providence, Rhode Island.
192. Frost, D., Rish, I. and Vila, L. 1997. Summarizing CSP hardness with continuous probability distributions, in Proceedings of AAAI-97, pp. 327-333.
193. Rish, I. and Dechter, R. 1996. To Guess or to Think? Hybrid Algorithms for SAT. In Proceedings of PPCP-96. Cambridge, Massachusetts.
194. Rish, I. and Dechter, R. 1996. Variable Independence in Markov Decision Problems. In Proceedings of AAAI-96 Workshop on Structural Issues in Planning and Temporal Reasoning. Portland, Oregon.
195. Dechter, R. and Rish, I. 1994. Directional Resolution: The Davis-Putnam Procedure, Revisited. In Proceedings of KR-94, Bonn, Germany.

SELECTED PATENTS (OUT OF 63)

1. Predicting states of subjects. US Patent 8346688, granted 2013-01-01.
2. Recovering the structure of sparse markov networks from high-dimensional data. US Patent 8326787, granted 2012-12-04.
3. Method and apparatus for authoring and optimizing flowcharts. US Patent 8341610, granted 2012-12-25.
4. Presentation of websites to a computer user. US Patent 8271517, granted 2012-09-18.
5. Method for performance bottleneck diagnosis and dependency discovery in distributed systems and computer networks. US Patent 7890449, granted 2011-02-15.
6. Systems, methods and computer program products for supervised dimensionality reduction with mixed type features and labels. US Patent 7996342, granted 2011-08-09.
7. Automated energy transfer calculation and compensation. US Patent 7831525, granted 2010-11-09.
8. Apparatus and methods for determining availability and performance of entities providing services in a distributed system using filtered service consumer feedback. US Patent 7782792, granted 2010-08-24.
9. Active sampling collaborative prediction method for end-to-end performance prediction. US Patent 7640224, granted 2009-12-29.
10. Active probing for real-time diagnosis. US Patent 7389347, granted 2008-06-17.

REFERENCES

Dr. Tom Mitchell

E. Fredkin University Professor
Machine Learning Department
School of Computer Science
Carnegie Mellon University
email: Tom.Mitchell@cmu.edu
phone: (412) 268-2611

Dr. Russell Greiner

Professor
PI, Alberta Machine Intelligence Institute
Department of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2E8
email: rgreiner@ualberta.ca
phone: (780) 492-5461

Dr. Terrence Sejnowski

Professor and Laboratory Head
Computational Neurobiology Laboratory
Francis Crick Chair
Salk Institute for Biological Studies
10010 N Torrey Pines Rd
La Jolla, CA 92037
email: sejnowski@salk.edu

Dr. Guillermo Cecchi

Principal Research Staff Member
Manager, Computational Psychiatry and Neuroimaging
Computational Biology Center
IBM T.J. Watson Research Center
Yorktown Heights, NY
email: gcecchi@us.ibm.com
phone: (914) 945-1815

Dr. Katya Scheinberg

Harvey E. Wagner Endowed Chair Professor
Department of Industrial and Systems Engineering
Co-director of Institute for Data, Intelligent Systems and Computation
Lehigh University
Harold S. Mohler Laboratory
Email: kattyas@lehigh.edu
200 West Packer Avenue

IRINA RISH

Bethlehem, PA 18015-1582