

# Advances in Financial AI: Opportunities, Innovations, and Responsible AI

**Workshop Summary.** The financial industry is experiencing a paradigm shift propelled by rapid advancements in artificial intelligence. From algorithmic trading and fraud detection to personalized banking and investment strategies, AI technologies are redefining financial services. Our workshop aims to convene researchers, industry professionals, and policymakers to explore the latest developments, discuss challenges, and chart a course for responsible AI integration in finance.

Topics of interest include, but not limit to Generative AI with applications to finance, time-series modeling, financial datasets, multi-agent systems, and practical financial applications such as forecasting, fraud detection, risk management, and quantitative finance, etc. By bringing together diverse perspectives from academia and industry, we seek to foster collaboration and drive forward advancements in the responsible use of AI in finance. Specifically, this workshop is focused on three primary areas:

- **Exploring Challenges and Opportunities.** Delving deeply into the multifaceted challenges and abundant opportunities of applying machine learning in financial services. Due to the scale and variety of financial applications, this include time-series forecasting in quantitative trading; Generative models for stress testing and synthetic financial data generation; anomaly detection for fraud detection; reinforcement learning for portfolio management and option pricing; domain adaptation for non-stationarity in capital market; explainable AI for credit decisions; federated learning and privacy-preserving ML for collaborative models across banks without sharing sensitive data, to name a few.
- **Showcasing Innovative Research.** Highlighting cutting-edge studies and breakthrough developments at the intersection of machine learning and finance. By showcasing pioneering research, we aim to inspire new ideas, foster collaborations, and drive innovations that can lead to more efficient, secure, and customer-centric financial services.
- **Promoting Responsible AI.** Discussing the significant social impact and critical importance of implementing responsible AI and trustworthy machine

learning practices in finance. Emphasizing ethical decision-making, transparency, fairness, and accountability, we will address topics such as bias mitigation, ethical data use, and the development of AI systems that are both effective and aligned with societal values.

Through knowledge exchange, we aim to tackle key challenges and promote advancements in the industry. The workshop will host a series of sessions intended to stimulate discussion on the mentioned topics. These will include invited talks, short paper presentations, panel discussion and poster sessions, all structured to cultivate dynamic conversations and facilitate the exchange of ideas among participants. In pursuit of this goal, we have invited a diverse array of speakers from academia and industry to share their perspectives on the subject matter and lead discussions on ongoing work and future directions.

## Invited Speakers and Panelists

- **[Confirmed] Greg Mori** is Vice President, RBC Fellow at Royal Bank of Canada's Borealis AI and an Adjunct Professor in the School of Computing Science at Simon Fraser University. He received a Ph.D. in Computer Science from UC Berkeley in 2004 and an Hon. B.Sc. in Computer Science and Mathematics from the University of Toronto in 1999. He was a Visiting Scientist at Google in 2014-2015. He served as Director of the School of Computing Science at Simon Fraser University from 2015-2018. Dr. Mori conducts research in computer vision and machine learning. He received the ICCV Helmholtz Prize in 2017. He was a Program Chair for CVPR 2020 and a General Chair for CVPR 2023. At Borealis AI his team builds AI-based products for financial services. These include the award-winning NOMI Forecast and numerous other industry-leading machine learning solutions.
- **[Confirmed] Ariel Neufeld** is a tenured associate professor in mathematics at the Nanyang Technological University in Singapore. He received his PhD in mathematics in May 2015 at ETH Zurich, where he spent half of his PhD at the Columbia University in the City of New York. Prior to joining NTU he was

a postdoctoral researcher at ETH Zurich. In 2021, he received the SIAM Activity Group on Financial Mathematics and Engineering Early Career Prize. His research focuses on machine learning algorithms and their applications in finance and insurance, model uncertainty in financial markets and distributionally robust optimization, as well as stochastic analysis and stochastic optimal control.

- **[Confirmed] Patrick Pun** is currently a tenured Associate Professor, Assistant Chair (MSc Programmes), and the Programme Director of Master of Science in Financial Technology at School of Physical and Mathematics Sciences, Nanyang Technological University, Singapore. Prior to NTU, Patrick obtained his Ph.D. in Statistics at the Chinese University of Hong Kong in 2016. His Ph.D. thesis won numerous awards, including Nicola Bruti Liberati Prize 2016 and the Young Scholars Thesis Award 2016. His research paper on high-dimensional portfolio selection won Best Student Research Paper (First Place) in INFORMS Financial Section in 2015. Patrick has strong research interests in Financial / Actuarial Mathematics, Big Data Analytics, and AI applications in Finance, as evidenced by his numerous top-tier publications in these fields.
- **[Tentative] Robert Wardrop** is a Professor in Management Practice at the University of Cambridge Judge Business School and the director of the Cambridge Centre of Alternative Finance. He is an economic sociologist interested in the development of channels and instruments of finance emerging outside of the incumbent financial system and the impact of these developments on market practices and regulatory change. He is the Academic Programme Director of the CCAF's Cambridge Fintech and Regulatory Innovation online programme which has enrolled financial services regulators based in more than 100 countries.
- **[Tentative] Kamalika Chaudhuri** is a Professor in Computer Science and Engineering department at University of California San Diego, and a Research Scientist in FAIR team at Meta AI. She received her Bachelor's degree in Computer Science and Engineering in 2002 from Indian Institute of Technology, Kanpur, and a PhD in Computer Science from University of California at Berkeley in 2007. After a postdoctoral at UCSD, she joined the CSE department at UC San Diego as an assistant professor in 2010. She received an NSF CAREER Award in 2013 and a Hellman Faculty Fellowship in 2012. She served as the program co-chair for AISTATS 2019 and ICML 2019. Her interests are in the foundations of trustworthy machine learning, which includes problems such as learning from sensitive data while preserving privacy, learning under sampling bias, and in the presence of an adversary. She is particularly interested in privacy-preserving machine learning, which addresses how to learn good models and predictors

from sensitive data, while preserving privacy.

- **[Tentative] James Zou** is an Associate Professor of Biomedical Data Science and, by courtesy, of Computer Science and Electrical Engineering at Stanford University. Prof. Zou works on making machine learning more reliable, human-compatible and statistically rigorous. Several of his algorithms are widely used in tech and biotech industries. Zou received a Ph.D from Harvard in 2014, and was a member of Microsoft Research, a Gates Scholar at Cambridge and a Simons fellow at UC Berkeley. He joined Stanford in 2016 and has been a two-time Chan-Zuckerberg Investigator and the faculty director of the university-wide Stanford Data4Health hub. Zou is also a member of the Stanford AI Lab.
- **[Tentative] John Ho** is the Global Head of Legal, Financial Markets for Standard Chartered Bank ("SCB"), overseeing and providing legal advisory, transactional and documentation support for Financial Markets (FM) business globally for Standard Chartered Bank (SCB), its branches and affiliates. Mr Ho is a member of Regulatory Working Group of the ACI FMA. Mr. Ho is also the co-chair of (1) the ISDA South East Asia Legal and Regulatory Committee and (2) the ISDA Asia Pacific ESG Working group. He is an active participant in the industry's Fintech events and is a public speaker on topics such as Central Bank Digital Currencies, Blockchain, Digital Assets, Financial Markets, Sustainability and Regulatory Reforms. He serves as a mentor to startups and co-authored a publication from the World Economic Forum called The Digital Currency Governance Consortium Whitepaper, published in November 2021. He is also a mentor for the UK FCA Sustainability TechSprint, FCA Digital Sandbox, R3 and LongHash.

**Workshop Format and Tentative Schedule.** The workshop announcement and call for papers will be made available on December 10, 2024, AoE. The paper submission deadline is January 31st, 2025, AoE, and the accept/reject notification will be on March 1st, 2025, AoE.

This one-day workshop will involve 4 invited talks, 3 oral presentations selected from the peer-reviewed submissions, one panel discussion, and two poster sessions for all accepted peer-reviewed submissions (all talks are in a large-attendance talk format).

- 08:45-09:00 Opening Remark
- 09:00-09:45 Invited Talk #1
- 09:45-10:30 Invited Talk #2
- 10:30-10:45 Coffee Break
- 10:45-11:00 Contributed Talk #1
- 11:00-11:15 Contributed Talk #2
- 11:15-11:30 Contributed Talk #3
- 11:30-12:30 Poster Session #1
- 12:30-13:30 Lunch Break

- 13:30-14:15 Invited Talk #3
- 14:15-15:00 Invited Talk #4
- 15:00-16:00 Poster Session #2
- 16:00-16:50 Panel Discussion
- 16:50-17:00 Closing Remark

Invited talks will be 30 minutes, followed by 15 minutes dedicated to Q&A and discussions. Each contributed talk will be 10 minutes, followed by 5 minutes dedicated to Q&A. Panel discussion will be 50 minutes, with the last 15 minutes dedicated to interactive discussion with participating audience. We further encourage attendees to connect with the authors of contributed papers during the poster sessions.

**Format and Logistics.** Given the rapid growth and increasing importance of the field, we anticipate approximately 60-70 number of attendees. We expect approximately 60 paper submissions, and we aim an acceptance rate of 50%. The workshop will strictly follow the ICLR guidelines, including the double-blind review protocol. Accepted papers and talk abstracts will be made available on the workshop website before the event.

In addition to the workshop paper track, a short paper track will be implemented to attract under-represented, under-resourced, and budding researchers who may not (yet) have the resources to submit full papers. We encourage submission of late-breaking developments that would most benefit from feedback at ICLR, such as an implementation and evaluation of an unpublished but simple idea, a modest but self-contained theoretical result, a follow-up experiment or re-analysis of a previously published paper, or a fresh perspective on an existing publication, etc. Both tracks will use OpenReview to manage and review the contributed papers.

We expect this workshop to be a **hybrid** event, with a heavy emphasis on in-person communication. Invited talks, contributed talks and panel discussion will all feature a dedicated slot for QA and interactive discussions. To facilitate any participant who cannot show up in person, we will leverage zoom video system to broadcast each sessions. All talk recordings will be made available on our workshop website afterwards as well. Most of our organizers will be onsite to facilitate and coordinate the workshop. We will also make sure we have organizers assisting online sections, including technical availability, collecting questions and feedbacks, etc.

**Fostering Mentorship and Collaboration.** We are excited to offer mentee and mentor sessions to pair juniors with seniors in the field during the coffee breaks and lunch break. These sessions will provide an opportunity for emerging talents to learn from experienced professionals, gain valuable insights, and build meaningful connections, both from academia and industry. To facilitate this, we will have a call before the event to group participants who are interested and arrange the pairs accordingly.

**Marketing Plan.** We plan to execute our marketing strategy through two main avenues to maximize outreach and engagement: online marketing and targeted outreach. Online, we'll develop a user-friendly workshop website with detailed information like the agenda, speaker profiles, and registration details. We'll leverage social media platforms such as Twitter, LinkedIn, and Facebook by creating event pages, sharing engaging content related to AI and finance, posting regular updates, and interacting with potential attendees to build a community around the event. For targeted outreach, we'll engage directly with finance-related industry companies, professional groups, and associations through personalized invitations, partnerships, and featuring the workshop in industry newsletters and internal communications to attract professionals working in finance and AI. We'll connect with universities, colleges, and research labs specializing in finance, artificial intelligence, and machine learning by disseminating information through academic mailing lists, departmental announcements, collaborations with faculty, and student organizations to encourage participation from scholars and students. By combining a robust online presence with targeted outreach efforts, we aim to effectively promote the workshop to a broad audience and specific groups interested in the intersection of AI and finance, maximizing attendance, fostering diverse participation, and enhancing the overall impact of the event.

**Diversity and Inclusion Statement** We are committed to ensuring that this workshop is an equal opportunity event, dedicated to diversity and inclusion. The workshop is organized by a diverse group of organizers and speakers from various sectors, affiliations, nationalities, geographic areas, and scientific backgrounds, with a clear intention to ensure that the organizers and speakers meet the highest standards of representation. Our inclusion principle has helped us identify organizers, speakers, and reviewers who represent a full spectrum of academic seniority, from PhD candidates to assistant and full professors, with research backgrounds in statistics, machine learning, applied research, and finance. The invited talks will cover a wide range of topics, from an introduction to challenges and opportunities to emerging research and responsible AI.

**Travel Funding for Diverse Attendees.** Borealis AI and JP Morgan AI Research have generously agreed to provide funding to support attendees from diverse backgrounds. This initiative aims to promote inclusivity, ensuring that individuals from various geographical, cultural, and socioeconomic backgrounds can participate in the workshop. By reducing financial barriers, Borealis AI and JP Morgan AI Research help create a more diverse and enriching environment where a wider range of perspectives can contribute to the discussions.

**Previous related workshops** While there has been no exact version of this workshop in past ICLR workshops, many similar workshops have been held at AI and

ML conferences. For example, the "AAAI 2024 Bridge on AI and Finance", the "AAAI 2024 Workshop on AI in Finance for Social Impact", "NeurIPS 2020 Workshop on Fair AI in Finance", "ICML 2021 Workshop on Representation Learning for Finance", and the "NeurIPS 2018 Workshop on Challenges and Opportunities for AI in Finance" have addressed related themes. However, this marks the first time that two major leading banks in AI for finance—J.P. Morgan and the Royal Bank of Canada, both recognized as AI innovation leaders in the banking industry according to the Evident AI ranking 2024—along with leading researchers from academia, have collaborated to organize a workshop to address key topics in machine learning for finance, delve into important challenges and opportunities, highlight innovative research, and discuss the need for responsible AI in the field.

### Organizers and Program Committee

- **Yongjae Lee** is an Associate Professor in the Department of Industrial Engineering and Artificial Intelligence Graduate School at Ulsan National Institute of Science and Technology (UNIST). Dr. Lee utilizes quantitative techniques such as ML/AI and optimization to analyze financial data and derive optimal decisions. He is particularly interested in analyzing individual and household financial activity to draw useful insights and design personalized financial services. Dr. Lee has published more than 30 papers in international journals and conferences including Quantitative Finance, European Journal of Operational Research, Annals of Operations Research, Journal of Portfolio Management, AISTATS and ICAIF. He is an advisory editorial board member for the Journal of Financial Data Science and will serve as an organizing committee (workshop chair) of ICAIF'24. He received his B.S. degree in computer science and mathematical sciences and Ph.D. degree in industrial and systems engineering from KAIST. E-mail: yongjaelee@unist.ac.kr
- **Bo An** is a President's Chair Professor in Computer Science, Head of Division of Artificial Intelligence, and Co-Director of Artificial Intelligence Research Institute (AI.R) at Nanyang Technological University, Singapore. He received the Ph.D degree in Computer Science from the University of Massachusetts, Amherst. His current research interests include artificial intelligence, multiagent systems, computational game theory, reinforcement learning, and optimization. Dr. An was the recipient of the 2010 IFAAMAS Victor Lesser Distinguished Dissertation Award, an Operational Excellence Award from the Commander, First Coast Guard District of the United States, the 2012 INFORMS Daniel H. Wagner Prize for Excellence in Operations Research Practice, 2018 Nanyang Research Award (Young Investigator), and 2022 Nanyang Research Award. His publications won the Best Innovative Application Paper Award at AAMAS'12, the Innovative Application Award at IAAI'16, and the best paper award at DAI'20. He was invited to give Early Career Spotlight talk at IJCAI'17. He led the team HogRider which won the 2017 Microsoft Collaborative AI Challenge. He was named to IEEE Intelligent Systems' "AI's 10 to Watch" list for 2018. He was PC Co-Chair of AAMAS'20 and General Co-Chair of AAMAS'23. He will be PC Chair of IJCAI'27. He is a member of the editorial board of JAIR and is the Associate Editor of AIJ, JAAMAS, IEEE Intelligent Systems, ACM TAAS, and ACM TIST. He was elected to the board of directors of IFAAMAS, senior member of AAAI, and Distinguished member of ACM. E-mail: boan@ntu.edu.sg
- **Eric Jiawei He** is a Research Director at RBC Borealis, Royal Bank of Canada. He and his team focuses on leveraging AI for financial services. His research expertise and interest lies in time series analysis and modelling, and more recently generative models and their applications. Eric received his Ph.D. in computer science from Simon Fraser University, Canada in 2020; Master of Engineering from Nanyang Technological University, Singapore in 2015; and Bachelor of Engineering from Zhejiang University, China in 2013. Eric has been regularly involved as program committee members in AI conferences such as NeurIPS, ICLR, ICML, AAAI, CVPR, ICCV, ECCV, etc. He organized numerous workshops and symposiums including *2nd Symposium on Advances in Approximate Bayesian Inference*, *Diversity and Inclusion Social Event*, *CVPR '23*, *Bridge event on AI for Financial Services*, *AAAI 2024*, etc. Eric is also a proud member of LGBTQ+ community in tech. **Main point of contact.** E-mail: eric.j.he@borealisai.com
- **Sharon Li** Sharon Li is an Assistant Professor in the Department of Computer Sciences at the University of Wisconsin-Madison. She received a Ph.D. from Cornell University in 2017, advised by John E. Hopcroft. Subsequently, she was a postdoctoral scholar in the Computer Science department at Stanford University. Her research focuses on the algorithmic and theoretical foundations of reliable machine learning. She is the recipient of the AFOSR Young Investigator Program (YIP) award, the NSF CAREER award, MIT Technology Review TR-35 Award, Forbes30Under30 in Science, and multiple faculty research awards from Google, Meta, and Amazon. Her works received a NeurIPS Outstanding Paper Award, and an ICLR Outstanding Paper Award Honorable Mention in 2022. She has served as Area Chair for ICLR, NeurIPS, ICML, and Program Chair for Workshop on Uncertainty and Robustness in Deep Learning (2019, 2020, 2021). She also co-organized multiple workshops and tutorials including NeurIPS Workshop on Robustness in Sequence Modeling in 2022, ICML Workshop on Distribution-free Uncertainty Quantification (2021 and 2022), ICCV Tutorial on Reliability of Deep Learning for Real-World Deployment in 2023,

and Data-centric Artificial Intelligence Workshop at WWW 2024. E-mail: sharonli@cs.wisc.edu

- **Alberto Pozanco** is a research lead at JP Morgan AI Research, where he joined after receiving his PhD in Computer Science from Universidad Carlos III de Madrid. The main focus of his current work includes the use of automated planning and optimization techniques to solve real world problems at JP Morgan. His research interests extend to other related Artificial Intelligence areas such as reinforcement learning, heuristic search and knowledge representation. He is regularly involved as PC member in conferences such as AAAI, IJCAI or ICAPS, and has co-organized the last two editions of the Workshop on Planning and Scheduling for Financial Services (FinPlan) at ICAPS, as well as the bridge event on "AI for Financial Services" at AAAI'24. E-mail: alberto.pozanco@jpmorgan.com

**Program Committee.** We are expecting a volume of 60+ submissions, and thus has the plan to reach out to approximately 100 individuals to assist with the paper review process. With the target of each submission having at least 3 reviews, we will ensure at least 60 confirmed reviewers.

Here is the list of **confirmed** program committee members to date. Please note that we are actively reaching out to additional experts to further strengthen our committee and ensure comprehensive coverage of all relevant areas.

1. **Julius Berner**, PhD, California Institute of Technology
2. **Akbar Rafiey**, PhD, University of California San Diego
3. **Nasimeh Heydaribeni**, PhD, University of California San Diego
4. **Amauri H. Souza**, PhD, Aalto University
5. **Giuseppe Canonaco**, PhD, J.P. Morgan AI Research
6. **Tomas de la Rosa**, PhD, J.P. Morgan AI Research
7. **Daniel Borrajo**, PhD, J.P. Morgan AI Research
8. **Marianela Morales**, PhD, J.P. Morgan AI Research
9. **Edoardo Vittori**, PhD, Intesa Sanpaolo
10. **Wei Deng**, PhD, Morgan Stanley
11. **Mohammed Suhail**, PhD, Google
12. **Ying Chen**, PhD, National University of Singapore
13. **Shih-Yang Su**, PhD, University of British Columbia
14. **John R.J. Thompson**, PhD, University of British Columbia
15. **Juba Ziani**, PhD, Georgia Institute of Technology
16. **Jaewook Song**, PhD, Hanyang University
17. **Yong Zheng**, PhD, Illinois Institute of Technology
18. **Mahmoud Khademi**, PhD, Microsoft Research
19. **Amir Abdi**, PhD, Microsoft Research
20. **Yoontae Hwang**, PhD, Ulsan National Institute of Science and Technology
21. **Dong-Young Lim**, PhD, Ulsan National Institute of Science and Technology
22. **Hyoungwoo Kong**, PhD, Ulsan National Institute of Science and Technology
23. **Kiarash Zahirnia**, PhD, Simon Fraser University
24. **Nazanin Mehrasa**, PhD, Borealis AI, Royal Bank of Canada
25. **Tristan Sylvain**, PhD, Borealis AI, Royal Bank of Canada
26. **Hossein Hajimirsadeghi**, PhD, Borealis AI, Royal Bank of Canada
27. **Thibaut Durand**, PhD, Borealis AI, Royal Bank of Canada