# Deep Generative Model in Machine Learning: Theory, Principle and Efficacy

Website: https://delta-workshop.github.io/

## 1. Workshop Description

Deep Generative Models (DGMs) have played a transformative role in advancing artificial intelligence (AI) over the past decade. Pioneering approaches such as variational autoencoders (VAEs) [1], generative adversarial networks (GANs) [2], flow-based models [3], and the more recent diffusion models [4,5] have pushed the boundaries of what is achievable with AI. Despite these significant advancements, however, substantial challenges persist, both in terms of deepening the theoretical foundations of DGMs and addressing practical implementation hurdles.

The success of Deep Generative Models (DGMs) hinges on a deeper theoretical understanding of their underlying mechanisms. However, current theoretical frameworks fall short of providing sustainable and practical principles for DGMs due to several key factors:

1. There exists a considerable gap between theoretical assumptions and real-world practice, making it difficult to reliably describe and predict the behavior of DGMs in applied settings.
2. The rapid development of generative AI paradigms, such as flow matching and diffusion models, outpaces the progress in theoretical research. As a result, this leaves significant gaps in how we comprehend their behaviors, generalization properties, and limitations.
3. Many of the current frameworks fail to provide rigorous principles that can help understand the optimization and training dynamics in DGMs.

In addition to the theoretical challenges, Deep Generative Models (DGMs) also face several significant algorithmic hurdles, including training instability, vulnerability to adversarial attacks, issues related to scalability, adaptation to new domains, etc. These challenges highlight the need for more principled, theory-driven solutions that can enhance the stability, robustness, and efficiency of DGMs in real-world applications. In particular:

1. Training instability, convergence issues, and vulnerability in adversarial settings, can result in unreliable model outputs.
2. Scaling DGMs to handle high-dimensional spaces or multimodal datasets remains computationally demanding, often resulting in a trade-off between model accuracy and the required training time.
3. Adaption to structured domains, such as discrete spaces, manifolds, meshes or graphs requires the models to account for complex geometric and domain-specific constraints, which remains an ongoing challenge.

This workshop, "**D**eep G**e**nerative Model in Machine **L**earning: **T**heories, Principles and Effic**a**cy (**DeLTa**) will center around the aforementioned challenges. We aim to bring together

experts from theory and applications, and foster discussions and collaborations on two broad and transformative questions in DGM:

***Q1****: How to develop theoretical frameworks to understand and design advanced generative models?*
***Q2****: How to develop principled strategies to improve the practical efficiency, reliability and transferability of DGMs in real-world applications?*

Our organizing team is geographically diverse and brings together a wide range of expertise, including substantial experience in hosting workshops. We have secured commitments from five prominent invited speakers who have confirmed their attendance. Drawing on our team's strong dedication, extensive expertise, and proven organizational proficiency, we are confident that the DeLTa workshop will be a standout event and a significant enhancement to ICLR 2025.

**Uniqueness.** Our workshop on deep generative models is a unique endeavor that bridges the gap between theory and practice. While many workshops focus on structured domains, specific applications, or particular DGM paradigms, our event emphasizes theoretical developments aimed at understanding the fundamental mechanisms of DGMs. By concentrating on these foundational principles, we aspire to provide a principled framework that guides better algorithm design and enhances the applicability of generative models across diverse fields. We aim to bring together experts from both theoretical and applied backgrounds and to foster interdisciplinary collaboration that leads to robust, efficient, and transferable generative models. More discussions can be found in Section 9.

**Urgency.** In addition, our workshop addresses an urgent need in the field of deep generative models. Despite their remarkable success, DGMs currently lack a principled and general theoretical framework to underpin their development. Moreover, as DGMs continue to grow in both model size and dimensionality—handling increasingly complex data and architectures—the challenges associated with scalability, optimization, and generalization become more pronounced. The rapid expansion in model complexity without a solid theoretical foundation exacerbates issues like training instability and unpredictability in performance. Our workshop seeks to tackle these pressing challenges by fostering discussions and collaborations aimed at establishing a comprehensive theoretical framework for DGMs.

## 2. Topics

We broadly classified the relevant topics for the workshop into *theory* as well as *algorithms and applications*.

**Theory**:
(1) Expressivity of deep generative models: investigate the expressivity of deep generative models and their performance variations across different datasets
(2) Optimization and generalization of deep generative models: study optimization strategies and generalization properties of generative models, emphasizing the performance of optimization algorithms in complex model architectures

(3) Solving stochastic processes for deep generative models: Delve into solving stochastic differential equations (SDEs) or partial differential equations (PDEs) for stochastic processes in generative models (e.g., diffusion models)
(4) Model Stability and Convergence Analysis in DGMs: Investigate the stability and convergence properties of generative models (e.g., GANs, VAEs, Flow-based Models, and Diffusion Models) during training
(5) Implicit Bias and Regularization in Generative Models: Explore implicit biases present in generative models and their impact on generalization. Study the effectiveness of explicit and implicit regularization techniques
(6) Robustness and Generalization Boundaries of Generative Models: Analyze the robustness boundaries of generative models and their behavior under out-of-distribution scenarios
(7) Latent Space Geometry and Manifold Learning: Analyze the geometric structure of the latent space in generative models and its relationship with the input data distribution. Explore how to balance diversity and generation quality in the latent space, and investigate the effectiveness and limitations of different manifold learning techniques in complex data scenarios

**Algorithms and Applications**:
(1) Improved sampling schemes: Develop and analyze efficient sampling algorithms for DGMs, including methods for faster and more accurate sample generation.
(2) Adversarial Robustness and Defense Mechanisms: Develop algorithmic methods to make DGMs robust against adversarial attacks, including adversarial training, robust optimization, and defense strategies that improve model reliability in hostile environments.
(3) Scalability and Efficiency in High-Dimensional Generative Modeling: Create algorithms that enable DGMs to efficiently handle high-dimensional data and large-scale multimodal datasets, optimizing computational resources without compromising model performance.
(4) Multimodal Generative Modeling Algorithms: Develop algorithms that effectively handle and integrate multimodal data (e.g., text, images, audio) in generative modeling, enhancing cross-modal generation and representation learning.
(5) Structured Data Modeling: Develop algorithms for adapting DGMs to structured domains such as discrete spaces, manifolds, meshes, or graphs, accounting for complex geometric and domain-specific constraints.
(6) Generative models for scientific discovery (AI4Science): Develop DGMS for applications in biology, physics, chemistry, material science and environmental science.

## 3. Submission Instructions and Tentative Deadlines

We invite submissions in two formats: **short papers** (up to 4 pages, excluding references and appendices) and **long papers** (up to 9 pages, excluding references and appendices). The main manuscript and any appendices must be submitted as a single PDF file through the designated submission portal. Papers previously published at other conferences will not be considered for acceptance.

(All dates are in 11.59pm AOE)

- Paper submission deadline: 1 Feb 2025
- Reviewer bidding dates: 1 - 5 Feb 2025
- Review Deadline: 25 Feb 2025

- Paper notification date: 27 Feb 2025
- Camera ready deadline: 4 Mar 2025

## 4. Workshop outlines

| Vancouver Time (GMT-7) | Event | Vancouver Time (GMT-7) | Event |
|---|---|---|---|
| 9:00-9:10 | Opening Remarks | 14:00-14:40 | Invited Talk 4 |
| 9:10-9:50 | Invited Talk 1 | 14:40-15:20 | Invited Talk 5 |
| 9:50-10:30 | Invited Talk 2 | 15:20-16:10 | Poster/Break |
| 10:30-11:20 | Poster/Break | 16:10-16:50 | Invited Talk 6 |
| 11:20-12:00 | Invited Talk 3 | 16:50-17:20 | Oral Presentations |
| 12:00-12:30 | Oral Presentations | 17:20-17:30 | Awards |
| 12:30-14:00 | Lunch | 17:30-17:40 | Closing Remarks |

We anticipate receiving 150 long paper submissions plus 100 short paper submissions, with an acceptance rate of 30%~40%. Of the accepted papers, 3% will be invited for oral presentations, 1% will be selected for outstanding paper awards. We expect 200~300 attendees.

## 5. Invited Speakers (A-Z by Last Name)

We invited 6 speakers, each specializing in key areas of deep generative models (DGMs), and 5 have confirmed their attendance. The speakers cover a range of important aspects of DGMs, ensuring comprehensive representation of the field.

- Theoretical analysis of generative models: Sinho Chewi, Valentin De Bortoli
- Algorithms for generative models: Yang Song, Valentin De Bortoli
- Multi-modal generative models: Nancy F. Chen
- Generative models for science: Michael Levin
- Safety in generative AI: Dawn Song

Our invited speakers represent a broad spectrum of constituencies within the research community, encompassing individuals from industry and academia, as well as both junior and senior researchers. We are committed to promoting the inclusion of early-career researchers and those from underrepresented groups. The workshop aims to foster an interdisciplinary environment that actively encourages and integrates diverse perspectives, backgrounds, and areas of expertise, both theoretical and applied. Five invited speakers have confirmed their participation, which reflect a rich diversity in terms of gender, race, and professional experience.

- [Confirmed] **Nancy F. Chen** (Group Leander & Principal Investigator, A*STAR)

Dr. Nancy F. Chen is a female group leader and principal investigator, leading the Artificial Intelligence for Education (AI4EDU) programme at I2R (Institute for Infocomm Research) and is a principal investigator at CFAR (Centre for Frontier AI Research), A*STAR in Singapore. Her research spans widely in natural language processing, including multilingual analysis, and multimodal generative models. She serves as Program Chair of NeurIPS 2025, Program Chair of ICLR 2023 and was awarded IEEE SPS Distinguished Lecturer 2023-2024.

- [Confirmed] **Sinho Chewi** (Assistant Professor, Yale University)
  Dr. Sinho Chewi is an assistant professor in the Department of Statistics and Data Science at Yale University. His research focuses on the mathematics of machine learning and statistics, particularly in log-concave sampling, and variational inference. He has made significant contributions to the theoretical understanding of sampling algorithms—foundational tools in the development of deep generative models. Currently, he is writing a book on the complexity of log-concave sampling. Dr. Chewi received his Ph.D. in Mathematics and Statistics from the MIT under Prof. Philippe Rigollet. He has held visiting positions at prestigious institutions such as the Simons Institute for the Theory of Computing, Microsoft Research, and the Institute for Advanced Study, highlighting his prominence in the theoretical community and his alignment with the goals of our workshop.

- [Confirmed] **Valentin De Bortoli** (Research Scientist, Google DeepMind)
  Dr. Valentin De Bortoli is a research scientist at Google DeepMind in London, the United Kingdom. His research interests are machine learning, information geometry, markov chain theory and optimal transport. His significant research work on diffusion model theories and applications is highly consistent with the topics which we aim to discuss in this workshop, especially his work on score matching, Schrodinger bridge, and dynamic regimes on diffusion models. He also contributed to the generative AI community with a website listing papers on diffusion models from which researchers have been benefited for their literature review. Dr. De Bortoli is also a CNRS researcher and was a postdoctoral researcher in Oxford University. He received his Ph.D. in ENS Paris Saclay.

- [Confirmed] **Michael Levin** (Distinguished Professor, Tufts University)
  Dr. Michael Levin is a distinguished professor in the Department of Biology at Tufts University, holding the Vannevar Bush endowed Chair and directing both the Allen Discovery Center and the Tufts Center for Regenerative and Developmental Biology. His research focuses on understanding the biophysical mechanisms that underpin decision-making during complex pattern regulation. A key aspect of his work involves creating next-generation AI tools to help scientists understand top-down control of pattern regulation—a focus that aligns closely with our workshop's emphasis on theoretical developments in deep generative models.

- [Confirmed] **Dawn Song** (Professor, UC Berkeley)
  Dr. Dawn Song is a professor in the Computer Science Division at the University of California, Berkeley. Her research encompasses deep learning, security, and blockchain, with a focus on secure deep learning and artificial intelligence. As the faculty co-director of the UC Berkeley Center on Responsible Decentralized

Intelligence (RDI) and a member of the Berkeley Artificial Intelligence Research (BAIR) Lab, she works on developing AI tools that are robust and secure. Dr. Song's numerous accolades include being named an ACM Fellow, IEEE Fellow, MacArthur Fellow, and Guggenheim Fellow. Her work directly addresses the challenges our workshop aims to tackle, such as enhancing the reliability and security of deep generative models, making her insights highly valuable to our discussions on improving the practical efficacy of DGMs in real-world applications.

- [Invited] **Yang Song** (Research Scientist, OpenAI)
  Dr. Yang Song leads the Strategic Explorations team at OpenAI and is an incoming assistant professor in Electrical Engineering and Computing + Mathematical Sciences at Caltech. His research focuses on building powerful AI models capable of understanding, generating, and reasoning with high-dimensional data across diverse modalities. Dr. Song has invented foundational concepts and techniques in score-based diffusion models, which are pivotal to the advancement of deep generative models—a central theme of our workshop. Dr. Song received his Ph.D. in Computer Science from Stanford University under Prof. Stefano Ermon and has been recognized with honors such as the ICLR Outstanding Paper Award and the Apple PhD Fellowship in AI/ML.

## 6. Plan to get an audience for a workshop

1). **Target Audience** We identify the specific groups that would benefit the most from our workshop: (1) Academia: Professors, researchers, and students in machine learning, generative models, and learning theory. (2) Industry Professionals: Practitioners working on AI development, generative models, or related applications. (3) Entrepreneurs: Those looking to understand the potential of generative models for creating innovative applications. (4) Interdisciplinary Researchers: Experts in fields like healthcare, biology, finance, and physics, who are interested in AI applications in their domains.

2). **Develop an Engaging Website** We have designed a professional and user-friendly website: https://delta-workshop.github.io/ .

3). **Leverage Academic and Professional Networks** We promote the workshop through: (1) Academic Departments: Reach out to relevant university departments in USA, Japan, Australia, Europe and China based on our organizers and collaborators. (2) Industry Partners: Collaborate with companies that have a vested interest in generative models or deep learning including Google, ByteDance, Microsoft and OpenAI.

4). **Use Social Media Campaigns** We target relevant communities on social media platforms: LinkedIn, Twitter/X, Reddit, and Google Groups (such as ML News).

## 7. Diversity commitment

In organizing the workshop, we recognize the importance of fostering a diverse and inclusive environment that reflects the global community of researchers, practitioners, and learners in

AI and machine learning. We believe that diversity enriches discussions, promotes the exchange of diverse perspectives, and leads to more innovative and impactful research. To ensure a welcoming and balanced environment for all participants, our diversity commitment will focus on the following key principles and actions:

- **Gender diversity**: We promote gender diversity in both selected speakers as well as the organizing committee. In particular, the gender ratio is 4:2 for speakers and 7:3 for organizers.
- **Geographic and institutional diversity**: The invited speakers and organizers represent a diverse range of geographic locations. This includes 1 speaker and 6 organizers from the Asia-Pacific region, 1 speaker and 1 organizer from Europe, and 4 speakers and 3 organizers from North America. Additionally, both academia and industry are well-represented, with a speaker ratio of 4:2 and an organizer ratio of 8:2.
- **Career stage diversity**: The speakers include both assistant professors and prominent figures in the field of research. Our organizing committee is similarly diverse, consisting of 4 postdoctoral researchers, 2 research scientists, 2 assistant professors and 2 professors.
- **Encouragement of Early-Career and Underrepresented Researchers**: We actively seeked contributions from early-career researchers, as well as individuals from underrepresented groups, ensuring that their voices and perspectives are included in the discussions.
- **Inclusive Call for Papers and Participation**: Our Call for Papers and workshop announcements will explicitly welcome submissions and attendance from all individuals, irrespective of gender, race, ethnicity, nationality, disability, or other factors. Furthermore, we encourage the short paper submissions in order to embrace the late-breaking discoveries and developments that would benefit from the exposures and feedbacks from our workshop. Notably, these short papers are welcomed to present unpublished, simple but novel ideas, a modest but self-contained theoretical result, a follow-up experiment or re-analysis of a previously published paper, a fresh perspective on an existing publication, or new methods with preliminary experiments, due to the limited computational resource or the time limit.
- **Code of Conduct**: The workshop will adhere to a clear code of conduct that promotes a safe, respectful, and welcoming environment for all attendees. We will outline behavioral expectations and provide mechanisms for reporting any violations.

## 8. Virtual access to workshop materials and outcome

In keeping with our commitment to broad participation and inclusivity, we will ensure that workshop materials and outcomes are accessible to a global audience. This approach allows individuals who may not be able to attend in person to still engage with the workshop content and contribute to the ongoing discussions. To achieve this, we will implement the following strategies:

1). **Live Streaming and Virtual Participation Options**: We will notify all registered participants about how to join our events either personally or virtually. In addition, we will prepare video and screen sharing, video recording, and pay attention to the interactions, to

ensure the presentation quality to both in-person and online audiences. In case of any last-minute mode transitions, we will notify all audiences multiple times, and ensure that they can receive the latest updates about the schedules and joining methods.

2). **Workshop materials**: We will make all workshop materials (speaker slides, workshop papers, and all posters) available on the workshop website. Furthermore, we will share video recordings of talks where possible. We will also design all workshop materials with accessibility and usability in mind.

## 9. Relevant Workshops

Our workshop, **"Deep Generative Models in Machine Learning: Theories, Principles, and Efficacy (DeLTa),"** distinguishes itself from other workshops by uniquely focusing on bridging the gap between theory and practical deployment of deep generative models (DGMs) across diverse domains. While existing workshops have explored specific aspects of generative modeling, DeLTa emphasizes theoretical advancements, practical principles, and real-world efficacy, aiming to establish a comprehensive framework unifying diverse generative paradigms. Below is a comparison with several relevant workshops:

1. Workshop on Diffusion Models @ NeurIPS2023 [https://diffusionworkshop.github.io/]. This workshop concentrated on the rapidly evolving field of diffusion models, covering their theoretical foundations, methodologies, and applications. While diffusion models are a key area of interest for *DeLTa*, our workshop does not restrict itself to this paradigm. Instead, we encompass a broader spectrum of generative models, including GANs, VAEs, flow-based models, and emerging architectures, aiming to provide a comparative analysis across these frameworks. *DeLTa* will explore connections between different generative models, linking theoretical principles with practical applications.

2. Safe Generative AI Workshop @NeurIPS2024 [https://safegenaiworkshop.github.io/]. This workshop focuses on the safety, security, and ethical implications of deploying generative models, covering topics such as bias, robustness, and adversarial vulnerabilities. While *DeLTa* also addresses robustness and reliability, our primary emphasis is on the theoretical foundations and practical principles that enhance the efficiency, performance, and scalability of generative models in real-world applications. Unlike the Safe Generative AI Workshop, which prioritizes ethical concerns, *DeLTa* focuses on balancing these with practical deployment strategies.

3. Structured Probabilistic Inference & Generative Modeling (SPIGM) @ ICML2024, [https://spigmworkshop2024.github.io/]. SPIGM centers around probabilistic inference and generative modeling, particularly addressing issues like amortization, sampling, and integration in graphical models. In contrast, *DeLTa* takes a more holistic view, incorporating both probabilistic and non-probabilistic approaches such as adversarial training and variational techniques. Our workshop will explore synergies across these paradigms, providing a unified framework that bridges diverse approaches to

generative modeling while also focusing on theoretical properties, which are not the main focus of SPIGM.

4. How Far Are We from AGI? @ ICLR'24, [https://agiworkshop.github.io/2024/] This workshop explores broader topics related to Artificial General Intelligence (AGI), including model design, applications, and theoretical analysis across various AI fields, such as multi-agent systems, expert systems, and symbolic AI. While AGI is a future-oriented goal, *DeLTa* is more focused on the current state of DGMs, offering in-depth discussions on advancing generative models through both theoretical and practical lenses. Our workshop zeroes in on deep generative models, providing a more focused platform for this domain.

5. Numerous other workshops have touched on generative models in specific contexts, such as Red Teaming GenAI: What Can We Learn from Adversaries? (NeurIPS 2024 https://redteaming-gen-ai.github.io/), which addresses security and risks, or EvalEval (NeurIPS 2024 https://evaleval.github.io/call-for-papers.html), which discusses model evaluations. Additionally, workshops such as CVG (ICML 2024 https://sites.google.com/view/cvgicml2024/home) and Generative Models for Decision Making (ICLR 2024 https://sites.google.com/view/genai4dm-iclr2024) focus on specific generative applications.

However, these workshops often lack discussions on the generative models in terms of theory, principle, and efficacy, which will be a central theme of *DeLTa*.

## 10. Organizers

Taiji Suzuki (Professor, University of Tokyo), Qibin Zhao (Team Leader, RIKEN AIP), Mingyuan Bai (Postdoctoral Researcher, RIKEN AIP), Maud Lemercier (Postdoctoral Research Associate, University of Oxford), Andi Han (Postdoctoral Researcher, RIKEN AIP), Ye Yuan (Research Scientist, ByteDance), Wei Huang (Research Scientist, RIKEN AIP), Denny Wu (Faculty Fellow, New York University & Flatiron Institute), Ernest K. Ryu (Assistant Professor, UCLA), Bamdev Mishra (Principal Applied Scientist, Microsoft India).

**Diversity, background and experience**

Our organizing committee brings together a diverse group of individuals with expertise from both academia and industry, spanning a wide range of fields. Our areas of specialization include both theory-oriented topics, such as optimization and deep learning theory, and application-focused areas, such as diffusion models and sampling. The team also represents various career stages, from postdoctoral researchers and research scientists to faculty members.

In addition to our team's diversity in background, we have extensive experience in organizing high-impact workshops. For example, Qibin Zhao has organized multiple workshops, including the 'Quantum Tensor Networks in Machine Learning' workshops at NeurIPS 2020 and 2021, the 'International Workshop on Tensor Network Representations in Machine Learning' at IJCAI 2020, the 'Workshop on Tensor Models for Machine Learning' at IEEE CAI 2024, and the '2nd RIKEN AIP – SJTU CS Joint Workshop on Machine Learning and Brain-like Intelligence.' Professor Taiji Suzuki has organized the 'Information-Based Induction

Sciences Workshop' (IBIS 2024), which is Japan's largest machine learning workshop. He also served as Program Co-Chair for ACML 2019 and as Program Chair for IBIS 2014.

This combination of diverse expertise and experience positions our committee to successfully deliver a high-quality workshop.

**Taiji Suzuki** (Email: taiji@mist.i.u-toyo.ac.jp )
Website: https://ibis.t.u-tokyo.ac.jp/suzuki/
Google Scholar: https://scholar.google.com/citations?hl=en&user=x8osrBsAAAAJ
Taiji Suzuki is currently a Professor in the Department of Mathematical Informatics at the University of Tokyo. He also serves as the team leader of “Deep learning theory” team in AIP-RIKEN. He received his Ph.D. degree in information science and technology from the University of Tokyo in 2009. He worked as an assistant professor in the department of mathematical informatics, the University of Tokyo between 2009 and 2013, and then he was an associate professor in the department of mathematical and computing science, Tokyo Institute of Technology between 2013 and 2017. He served as area chairs of premier conferences such as NeurIPS, ICML, ICLR and AISTATS, a program chair of ACML2019, and an action editor of the Annals of Statistics. He received the Outstanding Paper Award at ICLR in 2021, the MEXT Young Scientists’ Prize, and Outstanding Achievement Award in 2017 from the Japan Statistical Society. He is interested in deep learning theory, nonparametric statistics, high dimensional statistics, and stochastic optimization. In particular, he is mainly working on deep learning theory from several aspects such as representation ability, generalization ability and optimization ability. He also has devoted stochastic optimization to accelerate large scale machine learning problems including variance reduction methods, Nesterov’s acceleration, federated learning and non-convex noisy optimization.

**Qibin Zhao** (Email: qibin.zhao@riken.jp)
Website: https://qibinzhao.github.io/
Google scholar: https://scholar.google.com/citations?user=cSQGe3YAAAAJ&hl=en
Qibin Zhao received the Ph.D. degree in computer science from Shanghai Jiao Tong University, Shanghai, China, in 2009. He was a Research Scientist with RIKEN Brain Science Institute, Tokyo, Japan, from 2009 to 2017. He joined the RIKEN Center for Advanced Intelligence Project (AIP), Tokyo, as the Unit Leader, from 2017 to 2019. He is currently the Team Leader of the Tensor Learning Team at RIKEN AIP. He is also a Visiting Professor with the Tokyo University of Agriculture and Technology, Fuchu, Japan. He has published more than 150 scientific papers, and coauthored two monographs on tensor networks. His research interests include machine learning, tensor factorization and tensor networks, and brain signal processing. Dr. Zhao serves as the Area Chair for top-tier ML conferences of NeurIPS, ICML, and ACML, an Action Editor for Neural Networks and Transactions on Machine Learning Research.

**Mingyuan Bai** (E-mail: mingyuan.bai@riken.jp)
Google Scholar: https://scholar.google.com/citations?user=lo0_2rMAAAAJ&hl=en
Dr. Mingyuan Bai is a Postdoctoral Researcher in the Tensor Learning Team at RIKEN AIP, Tokyo, working with Professor Qibin Zhao. She completed her PhD in the University of Sydney, advised by Professor Junbin Gao. Her current research interests are generative models for adversarial machine learning, tensor learning and graph neural networks. She

published her research work in top machine learning conferences such as ICML and NeurIPS, and journals such as TNNLS and TMLR.

**Maud Lemercier** (Email: maud.lemercier@maths.ox.ac.uk)
Website: https://www.maths.ox.ac.uk/people/maud.lemercier
Google scholar: https://scholar.google.com/citations?user=zKSZJbsAAAAJ&hl=en
Maud Lemercier is a Postdoctoral Researcher at the Mathematical Institute of the University of Oxford. Previously, she was a PhD student at the University of Warwick and a visiting researcher at the Alan Turing Institute, working under the supervision of Prof. Theo Damoulas. She is part of the Oxford-Warwick CDT programme (OxWaSP) and is also a member of the Warwick Machine Learning Group and the DataSig team. Her research focuses on developing methodologies leveraging tools from rough path theory to perform inference on large and complex datasets of multivariate sequential data. Prior to joining OxWaSP, she completed an MSc in Machine Learning at Imperial College London and obtained a Degree of Engineering at IMT Atlantique in France. She co-organizes the Warwick Machine Learning Reading Group.

**Andi Han** (E-mail: andi.han@riken.jp)
Website: https://andihan3.github.io/
Google Scholar: https://scholar.google.com/citations?user=AKHQHs0AAAAJ&hl=en
Dr. Andi Han is currently a Postdoctoral Researcher in the Continuous Optimization Team at RIKEN AIP, Tokyo, working with Prof. Akiko Takeda. He completed his PhD in University of Sydney, advised by Prof. Junbin Gao. His research interest spans across optimization (on manifolds), deep learning theory, large foundation models, efficiency in machine learning and graph neural networks. He has publications in top machine learning conferences such as NeurIPS, ICML, AISTATS, IJCAI and journals, such as TPAMI, SIOPT, ML, TMLR.

**Ye Yuan** (Email: ye.yuan@tamu.edu)
Website: https://yyvettey.github.io/resume/
Google scholar: https://scholar.google.com/citations?user=L0W1LjMAAAAJ&hl=en
Ye Yuan is a research scientist with the Intelligent Creation team at ByteDance. She received her Ph.D. in Computer Science from Texas A&M University in 2020 and her bachelor's degree from the University of Science and Technology of China in 2015. Prior to joining ByteDance, Dr. Yuan completed internships at Walmart, ByteDance, and Adobe. Her research primarily focuses on video and image understanding, as well as AI-assisted editing. This includes developing cutting-edge algorithms for analyzing and interpreting visual content, advancing AI-driven creative tools, and improving the quality of automated image and video generation and editing.

**Wei Huang** (Email: wei.huang.vr@riken.jp)
Website: weihuang05.github.io
Google Scholar: https://scholar.google.com/citations?user=RZfDh4MAAAAJ&hl=en
Dr. Wei Huang is a Research Scientist in the Deep Learning Theory Team at RIKEN AIP, Tokyo, working with Professor Taiji Suzuki. He obtained his Ph.D. in Computer Science from the University of Technology Sydney, guided by Professor Richard Xu, and holds a Master's degree in Statistical Physics from the University of Science and Technology of China. Dr. Huang focuses on exploring the theoretical foundations of interpretability and transparency in deep learning's expressivity, optimization and generalization, as well as on developing

new algorithms, models, and methodologies that enhance the interpretability and improve their performance in graph neural networks, self-supervised learning and Auto-ML. His contributions to the field are documented in publications such as NeurIPS, ICLR, and ICML et al. Dr. Huang also keeps a comprehensive and popular blog on the latest Deep Learning Theory works on social media, with more than 10k followers.

**Denny Wu** (Email: dennywu@nyu.edu)
Website: https://dennywu1.github.io/
Google scholar: https://scholar.google.com/citations?user=x7Q3zj4AAAAJ&hl=en
Denny Wu is a faculty fellow at NYU CDS and the Flatiron Institute. He obtained his PhD from the University of Toronto and the Vector Institute, under the supervision of Jimmy Ba and Murat A. Erdogdu. Before that he was an undergraduate student at Carnegie Mellon University advised by Ruslan Salakhutdinov. His research focuses on developing a theoretical understanding (e.g., optimization and generalization) of modern machine learning systems such as neural networks, using tools from high-dimensional statistics.

**Ernest Ryu** (Email: eryu@math.ucla.edu)
Website: https://ernestryu.com/
Google scholar: https://scholar.google.com/citations?user=CNOqUZoAAAAJ&hl=en
Ernest Ryu is an assistant professor in the Department of Mathematics at UCLA. His current research focus is on applied mathematics, deep learning, and optimization. Professor Ryu received a B.S. degree in Physics and Electrical engineering with honors at the California Institute of Technology in 2010 and an M.S. in Statistics and a Ph.D. in Computational and Mathematical Engineering with the Gene Golub Best Thesis Award at Stanford University in 2016. In 2016, he joined the Department of Mathematics at UCLA, as an Assistant Adjunct Professor. In 2020, he joined the Department of Mathematical Sciences at Seoul National University as a tenure-track faculty. In 2024, returned to UCLA as an assistant professor.

**Bamdev Mishra** (Email: bamdevm@gmail.com)
Website: https://bamdevmishra.in/
Google scholar: https://scholar.google.com/citations?user=25IuNrMAAAAJ&hl=en&oi=ao
Bamdev Mishra is a Principal Applied Scientist at Microsoft India in the Office India Intelligence team. He develops machine learning (ML) solutions for Kaizala, Office Lens, Operations, and Information Protection for Office365, to name a few. Before this, he worked on various ML problems in the retail domain including competitive price prediction, review abuse detection, and style recommendations. Bamdev looks into problem domains that allow to use and build ML solutions for industrial applications. On the research side, his primary research interests include nonlinear optimization, stochastic learning, and matrix and tensor learning methods. He has published many technical papers in ML, NLP, and numerical optimization. He regularly reviews papers in many technical avenues and serves as Area Chair in NeurIPS, ICML, and AAAI. He is currently an Action Editor for the Transactions on Machine Learning Research (TMLR) journal. He received his BTech and MTech degrees in Electrical Engineering from the Indian Institute of Technology Bombay (IIT Bombay), India, in 2010, and his Ph.D. degree from the University of Liège, Belgium, in 2014. He was a Postdoctoral Researcher at the University of Liège and a Visiting Research Associate at the University of Cambridge from 2014 to 2015.

## 11. References

[1] Diederik P. Kingma, Max Welling. "Auto-Encoding Variational Bayes, arXiv preprint arXiv:1312.6114 (2013).

[2] Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial networks." Communications of the ACM 63, no. 11 (2020): 139-144.

[3] Rezende, Danilo, and Shakir Mohamed. "Variational inference with normalizing flows." In International conference on machine learning, pp. 1530-1538. PMLR, 2015.

[4] Song, Yang, and Stefano Ermon. "Generative modeling by estimating gradients of the data distribution." Advances in neural information processing systems 32 (2019).

[5] Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." Advances in neural information processing systems 33 (2020): 6840-6851.